

Bayesian Gaussian Copula Factor Models for Mixed Data

Jared S. Murray David B. Dunson Lawrence Carin Joseph E. Lucas

Abstract

Gaussian factor models have proven widely useful for parsimoniously characterizing dependence in multivariate data. There is a rich literature on their extension to mixed categorical and continuous variables, using latent Gaussian variables or through generalized latent trait models accommodating measurements in the exponential family. However, when generalizing to non-Gaussian measured variables the latent variables typically influence both the dependence structure and the form of the marginal distributions, complicating interpretation and introducing artifacts. To address this problem we propose a novel class of Bayesian Gaussian copula factor models which decouple the latent factors from the marginal distributions. A semiparametric specification for the marginals based on the extended rank likelihood yields straightforward implementation and substantial computational gains. We provide new theoretical and empirical justifications for using this likelihood in Bayesian inference. We propose new default priors for the factor loadings and develop efficient parameter-expanded Gibbs sampling for posterior computation. The methods are evaluated through simulations and applied to a dataset in political science. The models in this paper are implemented in the R package `bfa`.

Keywords: Factor analysis; Latent variables; Semiparametric; Extended rank likelihood; Parameter expansion; High dimensional

1 Introduction

Factor analysis and its generalizations are powerful tools for analyzing and exploring multivariate data, routinely used in applications as diverse as social science, genomics

Jared S. Murray is PhD Student, Dept. of Statistical Science, Duke University, Durham, NC 27708 (E-mail: jared.murray@stat.duke.edu). David B. Dunson is Professor, Dept. of Statistical Science, Duke University Durham, NC 27708 (E-mail: dunson@stat.duke.edu). Lawrence Carin is William H. Younger Professor, Dept. of Electrical & Computer Engineering, Pratt School of Engineering, Duke University, Durham, NC 27708 (E-mail: lcarin@ee.duke.edu) Joseph E. Lucas is Assistant Research Professor, Institute for Genome Sciences and Policy, Duke University, Durham, NC 27710 (E-mail: joe@stat.duke.edu) We would like to acknowledge the support of the Measurement to Understand Re-Classification of Disease of Cabarrus and Kannapolis (MURDOCK) Study and the NIH CTSA (Clinical and Translational Science Award) 1UL1RR024128-01, without whom this research would not have been possible. The first author would also like to thank Richard Hahn, Jerry Reiter and Scott Schmidler for helpful discussion.

and finance. The typical Gaussian factor model is given by

$$\mathbf{y}_i = \mathbf{\Lambda}\boldsymbol{\eta}_i + \boldsymbol{\epsilon}_i \tag{1.1}$$

where \mathbf{y}_i is a $p \times 1$ vector of observed variables, $\mathbf{\Lambda}$ is a $p \times k$ matrix of factor loadings ($k < p$), $\boldsymbol{\eta}_i \sim N(\mathbf{0}, \mathbf{I})$ is a $k \times 1$ vector of latent variables or factor scores, and $\boldsymbol{\epsilon}_i \sim N(\mathbf{0}, \boldsymbol{\Sigma})$ are idiosyncratic noise with $\boldsymbol{\Sigma} = \text{diag}(\sigma_1^2, \dots, \sigma_p^2)$. Marginalizing out the latent variables, $\mathbf{y}_i \sim N(\mathbf{0}, \mathbf{\Lambda}\mathbf{\Lambda}' + \boldsymbol{\Sigma})$, so that the covariance in \mathbf{y}_i is explained by the (lower-dimensional) latent factors. The model in (1.1) may be generalized to incorporate covariates at the level of the observed or latent variables, or to allow dependence between the latent factors. For exposition we focus on this simple case.

This model has been extended to data with mixed measurement scales, often by linking observed categorical variables to underlying Gaussian variables which follow a latent factor model (e.g. Muthén (1983)). An alternative is to include shared latent factors in separate generalized linear models for each observed variable (Dunson, 2000, 2003; Moustaki and Knott, 2000; Sammel et al., 1997). Unlike in the Gaussian factor model the latent variables typically impact both dependence and the form of the marginal distributions. For example, suppose $y_i = (y_{i1}, y_{i2})'$ are bivariate counts assigned Poisson log-linear models: $\log E(y_{ij} | \eta_i) = \mu_j + \lambda\eta_i$. Then λ governs both the dependence between y_{i1} , y_{i2} and the overdispersion in each marginal distribution. This confounding can lead to substantial artifacts and misleading inferences. Additionally, computation in such models is difficult and requiring marginal distributions in the exponential family can be restrictive.

There is a growing literature on semiparametric latent factor models to address the latter problem. A number of authors have proposed mixtures of factor models (Chen et al., 2010; Ghahramani and Beal, 2000). Song et al. (2010) instead allow flexible error distributions in Eq. (1.1). Yang and Dunson (2010) proposed a broad class of semiparametric structural equation models which allow an unknown distribution for $\boldsymbol{\eta}_i$. When building such flexible mixture models there is a sacrifice to be made in terms of interpretation, parsimony and computation, and subtle confounding effects remain. It would be appealing to retain the simplicity, interpretability and computational scalability of Gaussian factor models while allowing the marginal distributions to be unknown and free of the dependence structure.

To accomplish these ambitious goals we propose a semiparametric Bayesian Gaussian copula model utilizing the extended rank likelihood of Hoff (2007) for the marginal distributions. This approximation avoids a full model specification and is in some sense not fully Bayesian, but in practice we expect that this rank-based likelihood discards only a mild amount of information while providing robust inference. An additional contribution of this paper is to provide new theoretical and empirical justification for this approach.

The paper proceeds as follows: In Section 2, we propose the Gaussian copula factor model for mixed scale data and discuss its relationship to existing models. In Section 3 we develop a Bayesian approach to inference, specifying prior distributions and outlining a straightforward and efficient Gibbs sampler for posterior computation. Section

4 contains a simulation study, and Section 5 illustrates the utility of our method in a political science application. Section 6 concludes with a discussion.

2 The Gaussian copula factor model

A p -dimensional copula \mathbb{C} is a distribution function on $[0, 1]^p$ where each univariate marginal distribution is uniform on $[0, 1]$. Any joint distribution F can be completely specified by its marginal distributions and a copula; that is, there exists a copula \mathbb{C} such that

$$F(y_1, \dots, y_p) = \mathbb{C}(F_1(y_1), \dots, F_p(y_p)) \quad (2.1)$$

where F_j are the univariate marginal distributions of F (Sklar, 1959). If all F_j are continuous then \mathbb{C} is unique, otherwise it is uniquely determined on $\text{Ran}(F_1) \times \dots \times \text{Ran}(F_p)$ where $\text{Ran}(F_j)$ is the range of F_j . The copula of a multivariate distribution encodes its dependence structure, and is invariant to strictly increasing transformations of its univariate margins. Here we consider the Gaussian copula:

$$\mathbb{C}(u_1, \dots, u_p) = \Phi_p(\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_p) \mid \mathbf{C}), \quad (u_1, \dots, u_p) \in [0, 1]^p \quad (2.2)$$

where $\Phi_p(\cdot \mid \mathbf{C})$ is the p -dimensional Gaussian cdf with correlation matrix \mathbf{C} and Φ is the univariate standard normal cdf. Combining (2.1) and (2.2) we have

$$F(y_1, \dots, y_p) = \Phi_p(\Phi^{-1}(F_1(y_1)), \dots, \Phi^{-1}(F_p(y_p)) \mid \mathbf{C}) \quad (2.3)$$

From (2.3) a number of properties are clear: The joint marginal distribution of any subset of \mathbf{y} has a Gaussian copula with correlation matrix given by the appropriate submatrix of \mathbf{C} , and $y_j \perp\!\!\!\perp y_{j'}$ if and only if $c_{jj'} = 0$. When $F_j, F_{j'}$ are continuous, $c_{jj'} = \text{Corr}(\Phi^{-1}(F_j(y_j)), \Phi^{-1}(F_{j'}(y_{j'})))$ which is an upper bound on $\text{Corr}(y_j, y_{j'})$ (attained when the margins are Gaussian) (Klaassen and Wellner, 1997). The rank correlation coefficients Kendall's tau and Spearman's rho are monotonic functions of $c_{jj'}$ (Hult and Lindskog, 2002). For variables taking finitely many values $c_{jj'}$ gives the polychoric correlation coefficient (Muthén, 1983).

If the margins are all continuous then zeros in $\mathbf{R} = \mathbf{C}^{-1}$ imply conditional independence, as in the multivariate Gaussian distribution. However this is generally not the case when some variables are discrete. Even in the simple case where $p = 3$, Y_3 is discrete and $c_{13}c_{23} \neq 0$, if $r_{12} = 0$ then Y_1 and Y_2 are in fact *dependent* conditional on Y_3 (a similar result holds when conditioning on several continuous variables and a discrete variable as well - details available in supplementary materials). Results like these suggest that sparsity priors for \mathbf{R} in Gaussian copula models (e.g. Dobra and Lenkoski (2011); Pitt et al. (2006)) are perhaps not always well-motivated when discrete variables are present, and should be interpreted only with great care.

A Gaussian copula model can be expressed in terms of latent Gaussian variables \mathbf{z} : Let $F_j^{-1}(t) = \inf\{t : F_j(y) \geq t, y \in \mathbb{R}\}$ be the usual pseudo-inverse of F_j and suppose $\mathbf{\Omega}$ is a covariance matrix with \mathbf{C} as its correlation matrix. If $\mathbf{z} \sim N(0, \mathbf{\Omega})$ and

$y_j = F_j^{-1}(\Phi(z_j/\sqrt{\omega_{jj}}))$ for $1 \leq j \leq p$ then $F(\mathbf{y})$ has a Gaussian copula with correlation matrix \mathbf{C} and univariate marginals F_j . We utilize this representation to generalize the Gaussian factor model to Gaussian copula factor models by assigning \mathbf{z} a latent factor model:

$$\begin{aligned} \boldsymbol{\eta}_i &\sim N(0, \mathbf{I}), \quad \mathbf{z}_i | \boldsymbol{\eta}_i \sim N(\boldsymbol{\Lambda} \boldsymbol{\eta}_i, \mathbf{I}) \\ y_{ij} &= F_j^{-1} \left(\Phi \left(\frac{z_{ij}}{\sqrt{1 + \sum_{h=1}^k \lambda_{jh}^2}} \right) \right) \end{aligned} \quad (2.4)$$

Inference takes place on the scaled loadings

$$\tilde{\lambda}_{jh} = \frac{\lambda_{jh}}{\sqrt{1 + \sum_{h=1}^k \lambda_{jh}^2}} \quad (2.5)$$

so that $c_{jj'} = \sum_{h=1}^k \tilde{\lambda}_{jh} \tilde{\lambda}_{j'h}$. Rescaling is important as the factor loadings are not otherwise comparable across the different variables - even though $\boldsymbol{\Lambda}$ is technically identified it is not easily interpreted. We also consider the *uniqueness* of variable j , given by

$$u_j = 1 - \sum_{h=1}^k \tilde{\lambda}_{jh}^2 = \frac{1}{1 + \sum_{h=1}^k \lambda_{jh}^2} \quad (2.6)$$

In the Gaussian factor model u_j is $\sigma_j^2/(\sigma_j^2 + \sum_{h=1}^k \lambda_{jh}^2)$, the proportion of variance unexplained by the latent factors. In the Gaussian *copula* factor model this exact interpretation does not hold, but u_j still represents a measure of dependence on common factors.

2.1 Relationship to existing factor models

The Gaussian factor model and probit factor models are both special cases of the Gaussian copula factor model. Probit factor models for binary or ordered categorical data parameterize each margin by a collection of ‘‘cutpoints’’ $\gamma_{j0}, \dots, \gamma_{jc_j}$ (taking $\gamma_{j0} = -\infty$ and $\gamma_{jc_j} = \infty$ without loss of generality) so that $F_j(c) = \Phi \left(\gamma_{jc} (1 + \sum_{h=1}^k \lambda_{jh}^2)^{-1/2} \right)$. Then F_j has the pseudoinverse

$$F_j^{-1}(u_{ij}) = \sum_{c=1}^{c_j} c \mathbf{1} \left(\Phi \left(\frac{\gamma_{jc-1}}{\sqrt{1 + \sum_{h=1}^k \lambda_{jh}^2}} \right) < u_{ij} \leq \Phi \left(\frac{\gamma_{jc}}{\sqrt{1 + \sum_{h=1}^k \lambda_{jh}^2}} \right) \right)$$

Plugging this into (2.4) and simplifying gives $y_{ij} = \sum_{c=1}^{C_j} c \mathbf{1}(\gamma_{jc-1} < z_{ij} \leq \gamma_{jc})$ where $\mathbf{z}_i \sim N(\mathbf{0}, \boldsymbol{\Lambda} \boldsymbol{\Lambda}' + \mathbf{I})$, the data augmented representation of an ordinal probit factor model. Naturally the connection extends to mixed Gaussian/probit margins as well.

Other factor models which have Gaussian/probit models as special cases include semiparametric factor models, which assume non-Gaussian latent variables $\boldsymbol{\eta}_i$ or errors $\boldsymbol{\epsilon}_i$, retaining the linear model formulation (1.1) so that marginally $Cov(\mathbf{y}_i) = \boldsymbol{\Lambda}Cov(\boldsymbol{\eta}_i)\boldsymbol{\Lambda}' + \boldsymbol{\Sigma}$. But $F(\mathbf{y}_i)$ no longer has a Gaussian copula, and since the joint distribution is no longer elliptically symmetric the covariance matrix is unlikely to be an adequate measure of dependence. Further, the dependence and marginal distributions are confounded since the implied correlation matrix will depend on the parameters of the marginal distributions.

Our model overcomes these shortcomings. In the Gaussian copula factor model $\tilde{\boldsymbol{\Lambda}}$ governs the dependence separately from the marginal distributions, representing a factor-analytic decomposition for the scale-free copula parameter \mathbf{C} rather than $Cov(\mathbf{y}_i)$. The Gaussian copula model is also invariant to strictly monotone transformations of univariate margins. Therefore it is consistent with the common assumption that there exist monotonic functions h_1, \dots, h_p such that $(h_1(y_1), \dots, h_p(y_p))'$ follows a Gaussian factor model, while existing semiparametric approaches are not. Researchers using our method are not required to consider numerous univariate transformations to achieve “approximate normality”.

2.2 Marginal Distributions

One way to deal with marginal distributions in a copula model is to specify a parametric family for each margin and infer the parameters simultaneously with \mathbf{C} (see e.g. Pitt et al. (2006) for a Bayesian approach). This is computationally expensive for even moderate p , and there is often no obvious choice of parametric family for every margin. Since our goal is not to learn the whole joint distribution but rather to characterize its dependence structure we would prefer to treat the marginal distributions as nuisance parameters.

When the data are all continuous a popular semiparametric method is a two-stage approach in which an estimator \hat{F}_j is used to compute “pseudodata” $\hat{z}_{ij} = \Phi^{-1}(\hat{F}_j(y_{ij}))$, which are treated as fixed to infer the copula parameters. A natural candidate is $\hat{F}_j(t) = \frac{n}{n+1} \sum_{i=1}^n \frac{1}{n} \mathbf{1}(y_{ij} \leq t)$, the (scaled) empirical marginal cdf. Klaassen and Wellner (1997) considered such estimators in the Gaussian copula and Genest et al. (1995) developed them in the general case. However, this method cannot handle discrete margins. To accommodate mixed discrete and continuous data Hoff (2007) proposed an approximation to the full likelihood called the extended rank likelihood, derived as follows: Since the transformation $y_{ij} = F_j^{-1}(\Phi(z_{ij}))$ is nondecreasing, when we observe $\mathbf{y}_j = (y_{1j}, \dots, y_{nj})$ we also observe a partial ordering on $\mathbf{z}_j = (z_{1j}, \dots, z_{nj})$. To be precise we have that

$$\mathbf{z}_j \in D(\mathbf{y}_j) \equiv \{\mathbf{z}_j \in \mathbb{R}^n : y_{ij} < y_{i'j} \Rightarrow z_{ij} < z_{i'j}\} \quad (2.7)$$

The set $D(\mathbf{y}_j)$ is just the set of possible $\mathbf{z}_j = (z_{1j}, \dots, z_{nj})$ which are consistent with the ordering of the observed data on the j^{th} variable. Let $D(\mathbf{Y}) = \{\mathbf{Z} \in \mathbb{R}^{p \times n} : \mathbf{z}_j \in$

$D(\mathbf{y}_j) \quad \forall 1 \leq j \leq p\}$. Then we have

$$\begin{aligned} P(\mathbf{Y}|\mathbf{C}, F_1, \dots, F_p) &= P(\mathbf{Y}, \mathbf{Z} \in D(\mathbf{Y})|\mathbf{C}, F_1, \dots, F_p) \\ &= P(\mathbf{Z} \in D(\mathbf{Y})|\mathbf{C}) \times P(\mathbf{Y}|\mathbf{Z} \in D(\mathbf{Y}), \mathbf{C}, F_1, \dots, F_p) \end{aligned} \quad (2.8)$$

where (2.8) holds because given \mathbf{C} the event $\mathbf{Z} \in D(\mathbf{Y})$ does not depend on the marginal distributions. Hoff (2007) proposes dropping the second term in (2.8) and using $P(\mathbf{Z} \in D(\mathbf{Y})|\mathbf{C})$ as the likelihood. Intuitively we would expect the first term to include most of the information about \mathbf{C} . Simulations in Section 4 provide further evidence of this. Hoff (2007) shows that when the margins are all continuous the marginal ranks satisfy certain relaxed notions of sufficiency for \mathbf{C} , although these fail when some margins are discrete. Unfortunately theoretical results for applications involving mixed data have been lacking.

To address this we give a new proof of strong posterior consistency for \mathbf{C} under the extended rank likelihood with nearly any mixture of discrete and continuous margins (barring pathological cases which preclude identification of \mathbf{C}). Posterior consistency will generally fail under Gaussian/probit models when any margin is misspecified. A similar result for continuous data and a univariate rank likelihood was obtained by Gu and Ghosal (2009). We replace \mathbf{Y} with $\mathbf{Y}^{(n)}$ for notational clarity below.

Theorem 1. *Let Π be a prior distribution on the space of all positive semidefinite correlation matrices \mathcal{C} with corresponding density $\pi(\mathbf{C})$ with respect to a measure ν . Suppose $\pi(\mathbf{C}) > 0$ almost everywhere with respect to ν and that F_1, \dots, F_p , are the true marginal cdfs. Then for \mathbf{C}_0 a.e. $[\nu]$ and any neighborhood \mathcal{A} of \mathbf{C}_0 we have that*

$$\lim_{n \rightarrow \infty} \Pi(\mathbf{C} \in \mathcal{A} \mid \mathbf{Z}^{(n)} \in D(\mathbf{Y}^{(n)})) = 1 \text{ a.s. } [G_{\mathbf{C}_0, F_1, \dots, F_p}^\infty] \quad (2.9)$$

where $G_{\mathbf{C}_0, F_1, \dots, F_p}^\infty$ is the distribution of $\{\mathbf{y}_i\}_{i=1}^\infty$ under $\mathbf{C}_0, F_1, \dots, F_p$.

The proof is in Appendix A. We assumed a prior $\pi(\mathbf{C})$ having full support on \mathcal{C} . Under factor-analytic priors fixing $k < p$ restricts the support of π , and posterior consistency will only hold if \mathbf{C}_0 has a factor analytic decomposition in k or fewer factors. But by setting k large (or inferring it) it is straightforward to define factor-analytic priors which have full-support on \mathcal{C} (further discussion in Section 6). In practice, many correlation matrices which do not exactly have a k -factor decomposition are still well-approximated by a k -factor model. The result also applies to posterior consistency for $\hat{\Lambda}$ if k is chosen correctly, given compatible identifying restrictions. Theorem 1 also implies that for two divergent priors on Λ , the predictive distributions of \mathbf{z}_i and η_i will still converge, which is perhaps more directly relevant when e.g. estimating factor scores as in Section 5.

The efficiency of semiparametric estimators such as ours is also an important issue. Hoff et al. (2011) give some preliminary results which suggest that pseudo-MLE's based on the rank likelihood for continuous margins may be asymptotically relatively efficient. However, it is unclear whether even these results apply to the case of mixed

continuous and discrete margins, which is our primary focus. Simulations of the efficiency of posterior means under the extended rank likelihood versus correctly specified parametric models appear in Section 4.1. These results give an indication of the worst-case scenario in terms of efficiency lost in using the likelihood approximation, and are quite favorable in general.

3 Prior Specification and Posterior Inference

3.1 Prior Specification

Since the factor model is invariant under rotation or scaling of the loadings and scores we assume that sufficient identifying conditions are imposed (by introducing sign constraints and fixed zeros in $\mathbf{\Lambda}$), or that inference is on \mathbf{C} which does not suffer from this indeterminacy. For brevity we also assume here that k is known and fixed. Suggestions for incorporating uncertainty in k are in the Discussion.

A common prior for the unrestricted factor loadings in Gaussian, probit or mixed factor models is $\lambda_{jh} \sim N(0, 1/b)$. However, these priors have some troubling properties outside the Gaussian factor model: When $\sigma_j \equiv 1$ as in probit or mixed Gaussian/probit factor models – or in our copula model – the implied prior on u_j is

$$\pi(u_j) = \frac{(b/2)^{-k/2}}{\Gamma(k/2)} \left(\frac{1}{u_j^2}\right) \left(\frac{1-u_j}{u_j}\right)^{k/2-1} \exp\left[-\frac{b}{2} \left(\frac{1-u_j}{u_j}\right)\right] \quad (3.1)$$

Figure 1 shows that these priors are quite informative on the uniquenesses, especially as k increases. When k is small they are particularly informative on the scaled loadings, shrinking $\tilde{\lambda}_{jh}$ toward *large* values, rather than toward zero. This effect becomes worse as the prior variance increases. The problem is that the normal prior puts insufficient mass near zero. Coupled with the normalization this results in a “smearing” of mass across the columns of $\tilde{\mathbf{\Lambda}}$, deflating u_j , inducing spurious correlations, and giving inappropriately high probability to values of the scaled loadings near ± 1 . Therefore the normal prior is a very poor default choice in these models.

To address these shortcomings we consider shrinkage priors on λ_{jh} which place significant mass at or near zero. Such parsimony is also desirable for more interpretable results. Shrinkage priors have been thoroughly explored in the regression context (see e.g. Polson and Scott (2010) and references therein). In that context heavy-tailed distributions are desirable. While somewhat heavy tails are appealing here (so that $\pi(\tilde{\lambda}_{jh})$ decays slowly to zero as $|\tilde{\lambda}_{jh}| \rightarrow 1$), *extremely* heavy tails are inappropriate. Very heavy tails imply that with significant probability a single unscaled loading (say λ_{jm}) in a row j will be much larger than the others so that $\tilde{\lambda}_{jh} \approx \lambda_{jh}/|\lambda_{jm}|$ for $1 \leq h \leq k$. The resulting joint prior on $\tilde{\mathbf{\lambda}}_j = (\tilde{\lambda}_{j1}, \dots, \tilde{\lambda}_{jk})$ will assign undesirably high probability to vectors with one entry near ± 1 and the rest near 0, yielding correlations which are approximately 0 or ± 1 . Applying these priors in this new setting requires extra care.




Figure 1: Induced priors on the scaled factor loadings (top row) and uniquenesses (bottom row) implied by different priors as K varies

As a default choice we recommend the generalized double Pareto (GDP) prior of (Armagan and Dunson, 2011) which has the density

$$\pi(\lambda_{jh}) = \frac{\alpha}{2\beta} \left(1 + \frac{|\lambda_{jh}|}{\beta} \right)^{-(\alpha+1)} \quad (3.2)$$

which we will refer to as $\lambda_{jh} \sim GDP(\alpha, \beta)$. The GDP is a flexible generalization of the Laplace distribution with a sharper peak at zero and heavier tails. It has the following scale-mixture representation: $\lambda_{jh} | \psi_{jh} \sim N(0, \psi_{jh})$, $\psi_{jh} | \xi_{jh} \sim Exp(\xi_{jh}^2/2)$, and $\xi_{jh} \sim Ga(\alpha, \beta)$ which leads to conditional conjugacy and a simple Gibbs sampler. The GDP’s tail behavior is determined by α , and β is a scale parameter. Armagan and Dunson (2011) handle the hyperparameters by either fixing them both at 1 or assigning them a hyperprior. Here taking $\alpha = 3$, $\beta = 1$ is a good default choice: The $GDP(3, 1)$ distribution has mean 0 and variance 1, and $Pr(|\lambda_{jh}| < 2) \approx 0.96$. Critically, taking $\alpha = 3$ leads to tails of $\pi(\lambda_{jh})$ light enough to induce a sensible prior on $\tilde{\lambda}_{jk}$.

Figure 1 shows draws from the implied prior on u_j and $\tilde{\lambda}_{jh}$ under the $GDP(3, 1)$ prior, which are much more reasonable than the current default Normal priors. Note that as K increases, the prior puts increasing mass near zero without changing a great deal in the tails. This is reasonable since we expect each variable to load highly on only a few factors, and is difficult to mimic with the light-tailed normal priors. The prior on the uniquenesses remains relatively flat under the GDP prior, while the normal prior increasingly favors lower values and less parsimony.

3.2 Parameter-Expanded Gibbs Sampling

For efficient MCMC inference we introduce a parameter-expanded (PX) version of our original model. The PX approach (Liu and Wu, 1999; Meng and Van Dyk, 1999) adds redundant (non-identified) parameters to reduce serial dependence in MCMC and

improve convergence and mixing behavior. Naive Gibbs sampling in our model suffers from high autocorrelation due to strong dependence between \mathbf{Z} and $\mathbf{\Lambda}$. We modify (2.4) by adding scale parameters $\mathbf{V} = \text{diag}(v_1^2, \dots, v_p^2)$ to weaken this dependence:

$$\begin{aligned} \mathbf{w}_i &\sim N(\mathbf{V}^{1/2}\mathbf{\Lambda}\boldsymbol{\eta}_i, \mathbf{V}) \\ y_{ij} &= F_j^{-1}\left(\Phi\left(\frac{w_{ij}}{v_j\sqrt{1+\sum_{h=1}^K\lambda_{jh}^2}}\right)\right) \end{aligned} \quad (3.3)$$

Since w_{ij}/v_j and z_{ij} are equal in distribution (3.3) is observationally equivalent to the original model. We assume that \mathbf{V} is independent of the inferential parameters *a priori* so that $\pi(\mathbf{\Lambda}, \mathbf{H}, \mathbf{V}|\mathbf{Y}) = \pi(\mathbf{\Lambda}, \mathbf{H}|\mathbf{Y})\pi(\mathbf{V})$ (where \mathbf{H}' is the $n \times k$ matrix with entries η_{ik}) and the marginal posterior distribution of the inferential parameters is unchanged.

We choose the conjugate PX prior $1/v_j^2 \sim Ga(n_0/2, n_0/2)$ (independently). The greatest benefits from PX are realized when the PX prior is most diffuse, which would imply sending $n_0 \rightarrow 0$ and an improper PX prior. The resulting posterior for $(\mathbf{\Lambda}, \mathbf{H}, \mathbf{V})$ is also improper, but we can prove that the samples of $(\mathbf{\Lambda}, \mathbf{H})$ from the corresponding Gibbs sampler still have the desired stationary distribution $\pi(\mathbf{\Lambda}, \mathbf{H}|\mathbf{Y})$ (Appendix B). The PX Gibbs sampler is implemented as follows:

PX parameters: Draw $1/v_j^2 \sim Ga(n/2, s_j/2)$ where $\boldsymbol{\Psi}_j = \text{diag}(\psi_{j1}^2/2, \dots, \psi_{jh_j}^2/2)$ and $s_j = \mathbf{z}_j'(\mathbf{I} - \mathbf{H}_j'\boldsymbol{\Psi}_j^{-1} + \mathbf{H}_j\mathbf{H}_j')^{-1}\mathbf{H}_j\mathbf{z}_j'$

Factor Loadings: We assume a lower triangular loadings matrix with a positive diagonal; the extension to other constraints is straightforward. Let $k_j = \min(k, j)$ and \mathbf{H}_j' be the $n \times k_j$ matrix with entries η_{ik} for $1 \leq k \leq k_j$ and $1 \leq i \leq n$. Update nonzero elements in row j of $\mathbf{\Lambda}$ as $\boldsymbol{\lambda}'_j \sim N(\hat{\boldsymbol{\lambda}}'_j/v_j, (\boldsymbol{\Psi}_j^{-1} + \mathbf{H}_j\mathbf{H}_j')^{-1})$ where $\hat{\boldsymbol{\lambda}}'_j = (\boldsymbol{\Psi}_j^{-1} + \mathbf{H}_j\mathbf{H}_j')^{-1}\mathbf{H}_j\mathbf{z}_j'$ and λ_{jj} is restricted to be positive if $j \leq k$.

Hyperparameters: Update $(1/\psi_{jh}|-) \sim \text{InvGauss}(|\xi_{jh}/\lambda_{jh}|, \xi_{jh}^2)$ and $(\xi_{jh}|-) \sim Ga(\alpha + 1, \beta + |\lambda_{jh}|)$ where $\text{InvGauss}(a, b)$ is the inverse-Gaussian distribution with mean a and scale b .

Factor scores: Draw $\boldsymbol{\eta}_i$ from $(\boldsymbol{\eta}_i|-) \sim N([\boldsymbol{\Lambda}'\boldsymbol{\Lambda} + \mathbf{I}]^{-1}\boldsymbol{\Lambda}'z_i, [\boldsymbol{\Lambda}'\boldsymbol{\Lambda} + \mathbf{I}]^{-1})$

Augmented Data: Update \mathbf{Z} elementwise from

$$(z_{ij}|-) \sim TN\left(\sum_{h=1}^k \lambda_{jh}\eta_{ki}, 1, z_{ij}^l, z_{ij}^u\right) \quad (3.4)$$

where $TN(m, v, a, b)$ denotes the univariate normal distribution with mean m and variance v truncated to (a, b) , $z_{ij}^l = \max\{z_{i'j} : y_{i'j} < y_{ij}\}$ and $z_{ij}^u = \min\{z_{i'j} : y_{i'j} > y_{ij}\}$. If y_{ij} is missing then $(z_{ij}|-) \sim N(\sum_{h=1}^k \lambda_{jh}\eta_{ki}, 1)$. Note that (3.4) doesn't require a matrix inversion since $(z_{ij} \perp\!\!\!\perp z_{ij'} | \mathbf{\Lambda}, \boldsymbol{\eta}_i, \mathbf{Y})$ for $j \neq j'$, a unique property of our factor analytic representation and a significant computational benefit as p grows.

The PX-Gibbs sampler has mixing behavior at least as good as Gibbs sampling under the original model (which fixes $\mathbf{V} = \mathbf{I}$) (Liu and Wu, 1999; Meng and Van Dyk,

1999), and the additional computation is negligible. The PX-Gibbs sampler often increases the smallest effective sample size (associated with the largest loadings) by an order of magnitude or more in both real and synthetic data. The improved mixing is also vital for the multimodal posteriors sometimes induced by shrinkage priors. To our knowledge this is the first application of PX to factor analysis of mixed data, but PX has previously been applied to Gibbs sampling in Gaussian factor models by Ghosh and Dunson (2009) who introduce scale parameters for $\boldsymbol{\eta}_i$ to reduce dependence between \mathbf{H} and $\boldsymbol{\Lambda}$. Since MCMC in our model suffers primarily from dependence between \mathbf{Z} and $\boldsymbol{\Lambda}$ our approach is more appropriate. Hoff (2007) and Dobra and Lenkoski (2011) also use priors on unidentified covariance matrices to induce a prior on correlation matrices in Gaussian copula models. But the motivation there is simply to derive tractable MCMC updates and dependence between the priors on \mathbf{C} and \mathbf{V} precludes our strategy of choosing an optimal PX prior, limiting the benefits of PX.

3.3 Posterior Inference

Given MCMC samples we can address a number of inferential problems. The posterior distribution of the factor scores $\boldsymbol{\eta}_i$ provide a measure of the latent variables for each data point, describing a projection of the observed data into the latent factor space, and the factors themselves are characterized by the variables which load highly on them. Even if the factors are not directly interpretable this is a very useful exploratory technique for mixed data which is robust to outliers and handles missing data automatically (unlike common alternatives such as principal component analysis).

We can also do inference on conditional or marginal dependence relationships in \mathbf{y}_i . Here there is no need for identifying constraints in $\boldsymbol{\Lambda}$ which simplifies model specification. Tests of independence like $H_0 : c_{jj'} \leq \epsilon$ versus $H_1 : c_{jj'} > \epsilon$ are simple to construct from MCMC output. When the variables are continuous the conditional dependence relationships are encoded in $\mathbf{R} = \mathbf{C}^{-1}$ which we can compute as

$$\mathbf{R} = (\tilde{\boldsymbol{\Lambda}}\tilde{\boldsymbol{\Lambda}}' + \mathbf{U})^{-1} = \mathbf{U}^{-1} - \mathbf{U}^{-1}\tilde{\boldsymbol{\Lambda}}[\mathbf{I} + \tilde{\boldsymbol{\Lambda}}'\mathbf{U}^{-1}\tilde{\boldsymbol{\Lambda}}]^{-1}\tilde{\boldsymbol{\Lambda}}'\mathbf{U}^{-1} \quad (3.5)$$

Eq. (3.5) requires calculating only k -dimensional inverses, rather than p -dimensional inverses, a significant benefit of our factor-analytic representation.

As discussed in Section 2 the presence of discrete variables complicates inference on conditional dependence. Additionally, two discrete variables may be effectively marginally independent even if $|c_{jj'}| > 0$ simply by virtue of their levels of discretization. For these reasons, and for more readily interpretable results, it can be valuable to consider aspects of the posterior predictive distribution $\pi(\mathbf{y}^*|\mathbf{Y})$. Under our semiparametric model this distribution is somewhat ill-defined, but we can sample from an approximation to $\pi(\mathbf{y}^*|\mathbf{Y})$ by drawing $\tilde{\boldsymbol{\Lambda}}$ via the PX-Gibbs sampler, drawing $\mathbf{z}^* \sim N(0, \tilde{\boldsymbol{\Lambda}}\tilde{\boldsymbol{\Lambda}}' + \mathbf{U})$ and setting $y_j^* = \hat{F}_j^{-1}(\Phi(z_j^*))$ where \hat{F}_j are estimators of each marginal distribution. This disregards some uncertainty when making predictions; Hoff (2007) provides an alternative based on the values of z_{ij}^l, z_{ij}^u from (3.4) but (in keeping with his observations) we find both approaches to perform similarly.

Posterior predictive sampling of conditional distributions is detailed in Section 2 of the supplement. Importantly, the factor-analytic representation of \mathbf{C} allows us to directly sample any conditional distributions of interest (rather than using rejection sampling from the joint posterior predictive) by reducing the problem of sampling a truncated multivariate normal distribution to that of sampling independent truncated univariate normals.

4 Simulation Study

When fitting models in the following simulations we used the $GDP(3, 1)$ prior for λ_{jk} and take $1/\sigma^2 \sim \text{Gamma}(2, 2)$ for the Gaussian factor model and uniform priors on the cutpoints in the probit model. The cutpoints in the probit model were updated using independence Metropolis-Hastings steps with a proposal derived from the empirical cdfs. All models were fit using our R package.

4.1 Relative efficiency

First we examine finite-sample relative efficiency of the extended rank likelihood in the “worst-case scenario” (for our method). We compare the posterior mean correlation matrix under the Gaussian copula factor model with the extended rank likelihood to that under 1) a Gaussian factor model, when the factor model is true and 2) a probit factor model, when the probit model is true. Both are special cases of the Gaussian copula factor model so we can directly compare the parameters.

The true (unscaled) factor loadings were sampled iid $GDP(3, 1)$. For the probit case each margin had five levels with probabilities sampled $\text{Dirichlet}(1/2, \dots, 1/2)$. We fix k at the truth; additional simulations suggest that the relative performance is similar under misspecified k . We performed 100 replicates for various $p/k/n$ combinations in Fig. 2. Each model was fit using 100,000 MCMC iterations after a 10,000 burn-in, keeping every 20th sample. MCMC diagnostics for a sample of the fitted models indicated no convergence issues. We assess the performance of each method by computing a range of loss functions: Average and maximum absolute bias ($\frac{2}{p(p-1)} \sum_{i < j < p} |\hat{c}_{ij} - c_{ij}|$ and $\max_{i < j < p} |\hat{c}_{ij} - c_{ij}|$ respectively), root squared error: $\left[2 \sum_{i < j < p} (\hat{c}_{ij} - c_{ij})^2\right]^{1/2}$ and Stein’s loss: $\text{tr}(\hat{\mathbf{C}}\mathbf{C}^{-1}) + \log \det(\hat{\mathbf{C}}\mathbf{C}^{-1}) - p$. Stein’s loss is (up to a constant) the KL divergence from the Gaussian copula density with correlation matrix \mathbf{C} to the Gaussian copula density with correlation matrix $\hat{\mathbf{C}}$ and is therefore natural to consider here.

Figure 2 shows that the two methods are more or less indistinguishable in the probit case. Our method slightly outperforms the probit model in many cases because we do not have to specify a prior for the cutpoints. There are also some computational benefits here since in the copula model we avoid Metropolis-Hastings steps for the marginal distributions. In the continuous case our model also does well, although the Gaussian model is sometimes substantially more efficient under Stein’s loss. But as p




Figure 2: Efficiency (ratio of the loss under our model to that under the Gaussian/probit model) of the posterior mean under a range of loss functions

grows our model is increasingly competitive.

4.2 Misspecification Bias and Consistency

The previous simulations suggest that the loss in efficiency in worst-case scenarios is quite often minimal. To illustrate the practical benefit of our model (and the impact of our consistency result) in a realistic scenario we simulated data from a one-factor Gaussian copula factor model using the marginal distributions from Section 5 (Fig. 4). For simplicity we take $\tilde{\Lambda} = \tilde{\lambda}\mathbf{1}$ and consider $\tilde{\lambda} = 0.7$ and 0.8 (although we did not constrain the loadings to be equal when fitting models to the simulated data). Results in Section 5 suggest that these are plausible values.

Figure 3 shows that factor loadings for the two continuous variables (Black.Mkt.Premium and GDP.Per.Worker) are underestimated by the Gaussian/probit model. When all the variables are dependent there is a “ripple” effect so that even factor loadings for discrete variables are subject to some bias. We should expect this behavior in general – the copula correlations bound the observed Pearson correlations from above (in absolute value), with the bounds obtained only under Gaussian margins. The difference between the Pearson and copula correlation parameters, and hence the asymptotic bias, depends entirely on the form of the marginal distributions. This makes proper choices of transformations critical in the parametric model. Our model relieves this concern entirely. Although the magnitude of these effects is relatively mild here there is little reason to suspect this is true in general, especially in more complex models




Figure 3: Posterior mean factor loadings using 100 simulated datasets generated with the margins in Section 5 using our model (GCFM) versus a mixed Gaussian/probit model (G/P).

with multiple factors and a larger number of observed variables.

5 Application: Political-Economic Risk

Quinn (2004) considers measuring political-economic risk, a latent quantity, using five proxy variables and a Gaussian/probit factor model. The author defines political-economic risk as the risk of a state “manipulat[ing] economic rules to the advantage of itself and its constituents” following (North and Weingast, 1989, pp. 808). The dataset includes five indicators recorded for 62 countries: independent judiciary, black market premium, lack of appropriation risk, corruption, and gross domestic product per worker (GDPW) (Fig. 4). Additional background on political-economic risk and on the variables in this dataset is provided by Quinn (2004), and the data are available in the R package MCMCpack. Quinn (2004) transforms the positive continuous variables GDPW and black market premium by $\log(x)$ and $\log(x + 0.001)$ (resp.). The disproportionate number of zeros in black market premium (14/62 observations) leaves a large spike in the left tail and the normality assumption is obviously invalid. Since Quinn (2004) has already implicitly assumed a Gaussian copula, our model is a natural alternative to the misspecified Gaussian/probit model used there.

To explore sensitivity to prior distributions we fit the copula model under several priors: $GDP(3, 1)$, $N(0, 1)$ and the $N(0, 4)$ priors used by Quinn (2004). We use 100,000 MCMC iterations and save every 10^{th} sample after a burn-in of 10,000 itera-




Figure 4: Distributions of 4 variables from the political risk dataset in [Quinn \(2004\)](#). The fifth, Ind.Jud, is binary with 34/62 ones.

tions. Standard MCMC diagnostics gave no indication of lack of convergence. Figure 5 shows posterior means and credible intervals for the scaled loadings under each prior. Note that the $N(0, 4)$ prior, intended to be noninformative, is actually very informative on the scaled loadings (Fig 1). It pulls the scaled loadings toward ± 1 , with most pronounced influence on the binary variable Ind.Jud and the other categorical variables. The GDP prior instead shrinks toward zero as we would expect.

We also compare our model to the Gaussian/probit model in [Quinn \(2004\)](#), but using the $GDP(3, 1)$ prior in both cases. Figure 6 shows posterior predictive means and credible intervals for Kendall’s tau, as well as the observed values and bootstrapped confidence intervals. Our model fits well, considering the limited sample size, and fits almost uniformly better than the Gaussian-probit model. Other posterior predictive checks on rank correlation measures and in subsets of the data show similar results.

Incorrectly assuming a normal distribution for log black market premium is especially damaging. The copula correlation between GDPW and black market premium (on which the data are most informative) is underestimated in the Gaussian/probit model: mean -0.33 and 95% HPD interval $(-0.46, -0.22)$ as opposed to -0.56 and $(-0.73, -0.40)$ under our model. This is also evident in the posterior predictive samples of Kendall’s τ in Fig. 6. Figure 7 shows density estimates of draws from the bivariate posterior predictive of black market premium and GDPW. The Gaussian/probit model is clearly not a good fit, assigning very little mass to the bottom-right corner (which contains almost 25% of the data). The Gaussian copula factor model assigns appropriately high density to this region. Estimates of the latent variables are impacted as well: Figure 8 plots the mean factor scores from each model (after shifting and scaling to a common range) for low-risk countries. The seven countries with the lowest risk have identical covariate values except on GDPW. Our model infers mean scores that are sorted by GDPW (higher GDPW yielding a lower score). The Gaussian/probit model instead assigns these countries almost identical scores.




Figure 5: Posterior means/90% HPD intervals for scaled factor loadings under the different priors. Differences due to priors are larger for discrete variables, and largest for Ind.Jud (which is binary)

Posterior predictive distributions for Kendall's Tau




Figure 6: Posterior predictive mean and 95% HPD intervals of Kendall's τ under our model (Cop) and the Gaussian-probit model (Mix) as well as the observed values and bootstrapped 95% confidence intervals.

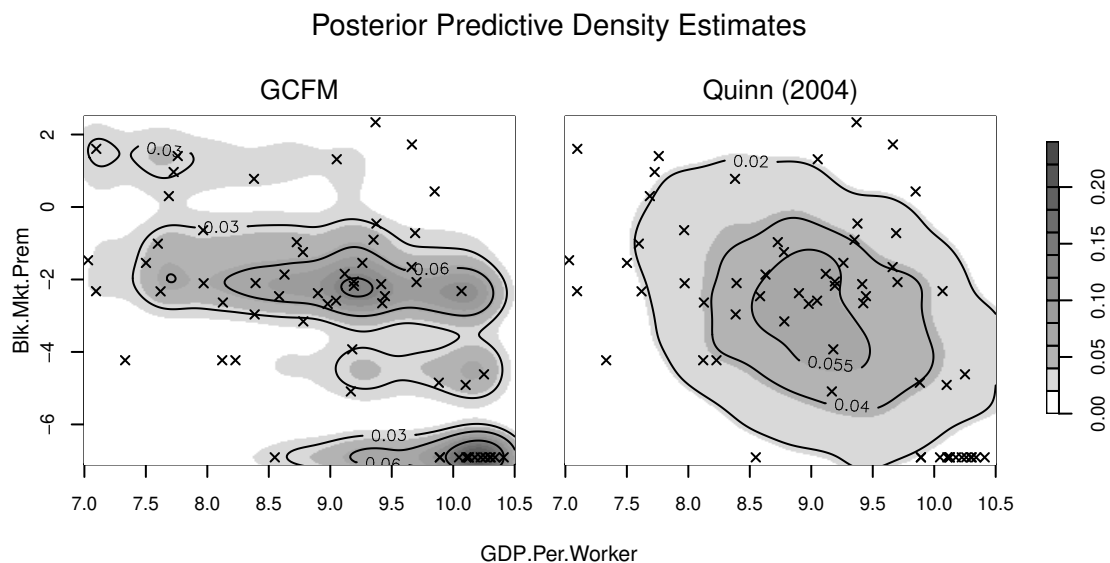


Figure 7: Posterior predictive distributions of log GDP and log black market premium, with observed data scatterplots. Note the cluster of points in the bottom-right corner; even though they represent over 20% of the sample the predictive density from the model in [Quinn \(2004\)](#) assigns very little mass to this area.




Figure 8: Comparison of the political-economic risk ranking obtained via our model and the mixed-data factor analysis of [Quinn \(2004\)](#). Points are posterior means and lines represent marginal 90% credible intervals.

6 Discussion

In this paper we have developed a new semiparametric approach to the factor analysis of mixed data which is both robust and efficient. We propose new default prior distributions for factor loadings which are more suited to routine use of this model (and similar models, such as probit factor models). As a byproduct we also induce attractive new priors on correlation matrices in Gaussian copula models; these are both more flexible and parsimonious than the inverse Wishart prior used by Hoff (2007), and much more efficient computationally than the graphical model priors of Dobra and Lenkoski (2011). They admit optimal parameter expansion schemes which are easy to implement, and are readily extended to informative specifications, to include covariates or to more complex latent variable models.

We have not considered the issue of uncertainty in the number of factors, but it is straightforward to do so by adapting existing methods for Gaussian factor models. In addition to posterior predictive checks, these include stochastic search (Carvalho et al., 2008), reversible jump MCMC (Lopes and West, 2004), Bayes factors (Ghosh and Dunson, 2009; Lopes and West, 2004) and nonparametric priors (Bhattacharya and Dunson, 2011; Paisley and Carin, 2009). The latter are especially promising when interest lies in \mathbf{C} since they preserve the computational advantages of factor-analytic priors while providing full support on correlation matrices (which fails for fixed $k < p$). Particularly when the plausible range of k is quite small, posterior predictive checking can be very effective.

A Proof of Theorem 1

Proof. We require a variant of Doob's theorem, presented in Gu and Ghosal (2009):

Doob's Theorem. *Let X_i be observations whose distributions depend on a parameter θ , both taking values in Polish spaces. Assume $\theta \sim \Pi$ and $X_i|\theta \sim P_\theta$. Let \mathcal{X}_N be the σ -field generated by X_1, \dots, X_N and $\mathcal{X}_\infty = \sigma(\bigcup_{i=1}^\infty \mathcal{X}_i)$. If there exists a \mathcal{X}_∞ measurable function f such that for $(\omega, \theta) \in \Omega^\infty \times \Theta$, $\theta = f(\omega)$ a.e. $[P_\theta^\infty \times \Pi]$ then the posterior is strongly consistent at θ for almost every θ $[\Pi]$.*

Therefore we must establish the existence of a consistent estimator of \mathbf{C} which is measurable with respect to the σ -field generated by the sequence $\{D(\mathbf{Y}^{(m)})\}_{m=1}^\infty$ (a coarsening of the σ -field generated by $\{Y^{(m)}\}_{m=1}^\infty$). Let $R_{nij} = \sum_{h=1}^n \mathbf{1}([y_{hj} \leq y_{ij}]) = n\hat{F}_j(y_{ij})$. Let $\mathbf{R}_{ni}(\mathbf{Y}^{(n)})$ be the p -vector with entry j given by R_{nij} and let $\mathbf{R}_n(\mathbf{Y}^{(n)}) = \{\mathbf{R}_{ni}\}_{i=1}^n$. Observe that the information contained in the extended rank likelihood (namely the boundary conditions in the definition of the set $D(\mathbf{Y}^{(n)})$) is equivalent to the information contained in $\mathbf{R}_n(\mathbf{Y}^{(n)})$. Hence a function that is measurable with respect to \mathcal{R}_n , the σ -field generated by $\{\mathbf{R}_m(\mathbf{Y}^{(m)})\}_{m=1}^n$, is also measurable with respect to the σ -field generated by $\{D(\mathbf{Y}^{(m)})\}_{m=1}^N$ and we may work exclusively with the former.

Let $\hat{U}_{nij} = \frac{R_{nij}}{n+1}$ and $\hat{\mathbf{U}}_{ni} = (\hat{U}_{ni1}, \dots, \hat{U}_{nip})'$. Then $\hat{U}_{nij} \xrightarrow{as} U_{ij}$ where $U_{ij} = F_j(y_{ij})$ by the SLLN, so $\hat{\mathbf{U}}_{ni} \xrightarrow{as} \mathbf{U}_i$ and therefore \mathbf{U}_i is $\mathcal{R}_\infty = \sigma(\bigcup_{i=1}^\infty \mathcal{R}_i)$ measurable. Note that if F_j is discrete U_{ij} is merely a relabeling of y_{ij} (each category/integer is “labeled” with its marginal cumulative probability). So \mathbf{U}_i is a sample from a Gaussian copula model with correlation matrix \mathbf{C}_0 where the continuous margins are all $U[0, 1]$ and the discrete marginal distributions are completely specified. The problem of estimating \mathbf{C} from \mathbf{U}_i reduces to estimating ordinary and polychoric/polyserial correlations with fixed marginals and it is straightforward to verify that the distribution of \mathbf{U}_i is a regular parametric family admitting a consistent estimator of \mathbf{C} , say $h_N(\mathbf{U}_1, \dots, \mathbf{U}_N)$. Therefore there exists a sequence of \mathcal{R}_∞ measurable functions $h_N(\mathbf{U}_1, \dots, \mathbf{U}_N) \rightarrow h(\mathbf{U}_1, \mathbf{U}_2, \dots) = \mathbf{C}_0$ almost surely and

$$\mathbf{C}_0 = h(\mathbf{U}_1, \mathbf{U}_2, \dots) = h^*(\{\mathbf{R}_{Ni} : N \geq 1, 1 \leq i \leq N\}) \text{ a.s. } [G_{\mathbf{C}_0, F_1, \dots, F_p}^\infty] \quad (\text{A.1})$$

where (A.1) holds because a null set under the measure induced by $\mathbf{R}_n(\mathbf{Y}^{(n)})$ is also null under $G_{\mathbf{C}_0, F_1, \dots, F_p}^\infty$. \square

B Validity of the PX Sampler

Let Θ be the inferential parameters and let $s_j = \mathbf{z}_j(\mathbf{I} - \mathbf{H}'_j(\Psi_j^{-1} + \mathbf{H}_j\mathbf{H}'_j)^{-1}\mathbf{H}_j)\mathbf{z}'_j$. Our working prior for (v_1, \dots, v_p) is $\prod_{j=1}^p IG(v_j^2; n_0/2, n_0/2)$. To verify that samples of Θ from the PX-Gibbs sampler have stationary distribution $\pi(\Theta|\mathbf{Y})$ we need to show that as $n_0 \rightarrow 0$ the transition kernels under the marginal sampling scheme (alternately drawing from $\pi(\mathbf{W}|\Theta, \mathbf{Y})$ and $\pi(\Theta|\mathbf{W})$) and the blocked sampling scheme (alternately drawing from $\pi(\mathbf{W}|\Theta, \mathbf{V}, \mathbf{Y})$ and $\pi(\Theta, \mathbf{V}|\mathbf{W})$) converge (Meng and Van Dyk, 1999). The t^{th} updates under the two schemes are as follows:

Scheme 1: Draw $1/v_{0j}^2 \sim Ga(n_0/2, n_0/2)$ and $1/v_{1j}^2 \sim Ga\left(\frac{n_0+n}{2}, \frac{n_0+v_{0j}^2 s_j}{2}\right)$. Set $r = v_{j0}/v_{j1}$ and draw $\lambda_j \sim N(r\hat{\lambda}_j', (\Psi_j^{-1} + \mathbf{H}\mathbf{H}')^{-1})$

Scheme 2: Draw $1/v_{tj}^2 \sim Ga\left(\frac{n_0+n}{2}, \frac{n_0+v_{(t-1)j} s_j}{2}\right)$. Set $r = v_{(t-1)j}/v_{tj}$ and draw $\lambda_j \sim N(r\hat{\lambda}_j', (\Psi_j^{-1} + \mathbf{H}\mathbf{H}')^{-1})$

Updates for the rest of Θ under both schemes are the same as in Section 3.2. As $n_0 \rightarrow 0$ under Scheme 1 the distribution of $1/v_{0j}^2$ approaches a point mass at 1 and Scheme 1 converges to Scheme 2 with $n_0 = 0$.

C Supplementary Material

C.1 Conditional independence

Assume $F(Y_1, Y_2, Y_3)$ has a Gaussian copula with correlation matrix \mathbf{C} , that Y_3 is discrete, and that $r_{12} = 0$. Let $(Z_1, Z_2, Z_3) \sim N(\mathbf{0}, \mathbf{C})$ and $\mathcal{B}_c = (F_3(c-1), F_3(c)]$

for c in the domain of Y_3 and define $g_j(z_3) = \Phi(F_j(y_j) - c_{j3}z_3)/(1 - c_{j3}^2)^{1/2}$. It is straightforward to show that

$$Pr(Y_1 \leq y_1 | Y_3 = c) Pr(Y_2 \leq y_2 | Y_3 = c) = E(g_1(z_3)) E(g_2(z_3)) \quad (\text{C.1})$$

$$Pr(Y_1 \leq y_1, Y_2 \leq y_2 | Y_3 = c) = E(g_1(z_3) g_2(z_3)) \quad (\text{C.2})$$

where the expectations are with respect to $\pi(z_3 | y_3 = c) = TN(0, 1, F_3(c - 1), F_3(c))$ and (C.2) holds because $\pi(z_1, z_2 | z_3) = \pi(z_1 | z_3) \pi(z_2 | z_3)$ when $r_{12} = 0$. Since g_1, g_2 are monotone it is well known that $E(g_1(z_3) g_2(z_3)) \neq E(g_1(z_3)) E(g_2(z_3))$ (and Y_1, Y_2 are dependent given Y_3) unless one or both functions are a.s. constant, which occurs only if one or both of Y_1, Y_2 are marginally independent of Y_3 ($c_{13}c_{23} = 0$). This result extends to conditioning on one discrete variable and any number of continuous variables since conditioning on a continuous variable $Y_4 = y_4$ implies that $Pr(z_4 = \Phi^{-1}(F(y_4))) = 1$, and $\pi(z_3 | y_3, z_4)$ is again univariate truncated normal (with a different mean and variance).

C.2 Posterior Predictive Conditional Distributions

To sample from conditional posterior predictive distributions such as $\pi(y_1^* | \mathbf{y}_{(-1)}^* = \mathbf{x}, \mathbf{Y})$ we could sample from $\pi(\mathbf{y}^* | \mathbf{Y})$ and discard draws where $y_j^* \neq x_j$ for any $2 \leq j \leq p$. This approach can be wasteful computationally since even in moderate dimensions most samples will be discarded. Instead we might prefer to estimate this distribution directly. We can write $Pr(y_1^* \leq y | \mathbf{y}_{(-1)}^* = \mathbf{x}, \mathbf{Y})$ as

$$\int_{\mathcal{C}} \int_{\mathbb{R}^{p-1}} \left(\int_{-\infty}^{\hat{F}_1(y)} \pi(z_1^* | \mathbf{z}_{(-1)}^*, \mathbf{C}) dz_1^* \right) \pi(\mathbf{z}_{(-1)}^* | \mathbf{y}_{(-1)}^* = \mathbf{x}, \mathbf{C}) \pi(\mathbf{C} | \mathbf{Y}) dz_{(-1)}^* d\mathbf{C} \quad (\text{C.3})$$

Assume that y_2, \dots, y_p are discrete, or that the empirical cdfs are used for \hat{F}_j (if y_j is continuous and \hat{F}_j is a smooth estimator then $z_j^* = \Phi^{-1}(\hat{F}_j(x_j))$ is fixed in the following). Then $\pi(\mathbf{z}_{(-1)}^* | \mathbf{y}_{(-1)}^* = \mathbf{x}, \mathbf{C})$ is the $(p - 1)$ -dimensional truncated normal distribution $N(\mathbf{0}, \mathbf{C}_{(-1)})$ where $\mathbf{C}_{(-1)}$ is obtained by dropping the first row and column of \mathbf{C} , restricted to the set $\mathcal{B}_x = \{\mathbf{z}_{(-1)}^*; \Phi^{-1}(\hat{F}_j(x_j^-)) < z_j^* \leq \Phi^{-1}(\hat{F}_j(x_j)) \forall 2 \leq j \leq p\}$ (where $F(x^-)$ is the lower limit of F at x). To estimate (C.3) from MCMC output we need to draw from this distribution (at least) once for every sample of \mathbf{C} . For a general \mathbf{C} this is prohibitive unless p is very small, but our factor-analytic representation allows us to efficiently draw from $\pi(\mathbf{z}_{(-1)}^* | \mathbf{y}_{(-1)}^* = \mathbf{x}, \mathbf{C})$ by sampling $(p - 1)$ univariate truncated normals: Let $\tilde{\Lambda}_{(-1)}$ be $\tilde{\Lambda}$ with the first row removed and $\mathbf{U}_{(-1)}$ be \mathbf{U} with the first row and column removed. Since $\mathbf{C}_{(-1)} = \tilde{\Lambda}_{(-1)} \tilde{\Lambda}'_{(-1)} + \mathbf{U}_{(-1)}$ we have

$$\begin{aligned} \pi(\mathbf{z}_{(-1)}^* | \mathbf{y}_{(-1)}^* = \mathbf{x}, \tilde{\Lambda}_{(-1)}) &\propto N(\mathbf{z}_{(-1)}^*; \mathbf{0}, \tilde{\Lambda}_{(-1)} \tilde{\Lambda}'_{(-1)} + \mathbf{U}_{(-1)}) \mathbf{1}((z_{(-1)}^* \in \mathcal{B}_x)) \\ &\propto \int_{\mathbb{R}^k} \prod_{j=2}^p \left(TN(\tilde{\lambda}_j \boldsymbol{\eta}, u_j, a_j, b_j) \right) N(\boldsymbol{\eta}; \mathbf{0}, \mathbf{I}) d\boldsymbol{\eta} \end{aligned}$$

where $a_j = \Phi^{-1}(\hat{F}_j(x_j^-))$, $b_j = \Phi^{-1}(\hat{F}_j(x_j))$ and $\boldsymbol{\eta}$ is an auxiliary variable. Therefore we can approximate (C.3) as follows:

1. Draw $\tilde{\boldsymbol{\Lambda}}$ via the PX-Gibbs sampler, and draw $\boldsymbol{\eta} \sim N(\mathbf{0}, \mathbf{I})$
2. Draw $z_j^* \sim TN(\tilde{\boldsymbol{\lambda}}_j \boldsymbol{\eta}, u_j, a_j, b_j)$ for $2 \leq j \leq p$
3. For each distinct value of \mathbf{y}_1 set $\tilde{F}^{(t)}(y_i) = \int_{-\infty}^{\hat{F}_1(y_i)} N(z_1^*; m, v) dz_1^*$ where

$$\begin{aligned} m &= \tilde{\boldsymbol{\lambda}}_1 \tilde{\boldsymbol{\Lambda}}'_{(-1)} [\tilde{\boldsymbol{\Lambda}}_{(-1)} \tilde{\boldsymbol{\Lambda}}'_{(-1)} + \mathbf{U}_{(-1)}]^{-1} \mathbf{z}_{(-1)}^* \\ v &= 1 - \tilde{\boldsymbol{\lambda}}_1 \tilde{\boldsymbol{\Lambda}}'_{(-1)} [\tilde{\boldsymbol{\Lambda}}_{(-1)} \tilde{\boldsymbol{\Lambda}}'_{(-1)} + \mathbf{U}_{(-1)}]^{-1} \tilde{\boldsymbol{\Lambda}}_{(-1)} \tilde{\boldsymbol{\lambda}}_1' \end{aligned} \quad (\text{C.4})$$

where again the matrix inverses in (C.4) can be computed efficiently as in (3.5). This procedure provides estimates of the conditional cdf at the observed data points. For a discrete response we can then directly compute conditional probabilities, odds ratios, and so on. When y_1 is continuous these can be interpolated to give a histogram estimate of $\pi(y_1 | \mathbf{y}_{(-1)} = \mathbf{x})$ with support on the range of the observed data. A number of modifications to this approach are possible; for example, to condition on a subset of $\mathbf{y}_{(-1)}$ we simply drop the irrelevant rows of $\boldsymbol{\Lambda}_{(-1)}$ and only perform step 3 for the j^{th} variable if we are conditioning on y_j .

This is a natural extension of factor regression models which posit a Gaussian factor model for $(y_i, \mathbf{x}'_i)'$, implying a linear regression model for $\pi(y_i | \mathbf{x}_i)$ (Carvalho et al., 2008; West, 2003). These are especially useful when $p > n$ as a model-based form of reduced rank regression (automatically selecting batches of correlated predictors by loading them highly on the same factor), or when there is missing data in \mathbf{X} . Here we have a flexible joint model which accommodates any ordered response or covariates while retaining the computational simplicity of factor regression models.

References

- Armagan, A. and Dunson, D. (2011). Generalized double Pareto shrinkage. *Arxiv preprint arXiv:1104.0861*.
- Bhattacharya, A. and Dunson, D. B. (2011). Sparse Bayesian infinite factor models. *Biometrika*, 98(2):291–306.
- Carvalho, C. M., Chang, J., Lucas, J. E., Nevins, J. R., Wang, Q., and West, M. (2008). High-Dimensional Sparse Factor Modeling: Applications in Gene Expression Genomics. *Journal of the American Statistical Association*, 103(484):1438–1456.
- Chen, M., Silva, J., Paisley, J., Wang, C., Dunson, D., and Carin, L. (2010). Compressive Sensing on Manifolds Using a Nonparametric Mixture of Factor Analyzers: Algorithm and Performance Bounds. *IEEE Transactions on Signal Processing*, 58(12):6140–6155.

- Dobra, A. and Lenkoski, A. (2011). Copula Gaussian graphical models and their application to modeling functional disability data. *The Annals of Applied Statistics*, 5(2A):969–993.
- Dunson, D. (2000). Bayesian latent variable models for clustered mixed outcomes. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, 62(2):355–366.
- Dunson, D. B. (2003). Dynamic Latent Trait Models for Multidimensional Longitudinal Data. *Journal of the American Statistical Association*, 98(463):555–563.
- Genest, C., Ghoudi, K., and Rivest, L.-P. (1995). A Semiparametric Estimation Procedure of Dependence Parameters in Multivariate Families of Distributions. *Biometrika*, 82(3):543.
- Ghahramani, Z. and Beal, M. (2000). Variational inference for Bayesian mixtures of factor analysers. *Advances in neural information processing systems*, 12:449–455.
- Ghosh, J. and Dunson, D. B. (2009). Default Prior Distributions and Efficient Posterior Computation in Bayesian Factor Analysis. *Journal of Computational and Graphical Statistics*, 18(2):306–320.
- Gu, J. and Ghosal, S. (2009). Bayesian ROC curve estimation under binormality using a rank likelihood. *Journal of Statistical Planning and Inference*, 139(6):2076–2083.
- Hoff, P. (2007). Extending the rank likelihood for semiparametric copula estimation. *Annals of Applied Statistics*, 1(1):265–283.
- Hoff, P., Niu, X., and Wellner, J. (2011). Information bounds for Gaussian copulas. *Arxiv preprint arXiv:1110.3572*, pages 1–24.
- Hult, H. and Lindskog, F. (2002). Multivariate extremes, aggregation and dependence in elliptical distributions. *Advances in Applied Probability*, 34(3):587–608.
- Klaassen, C. and Wellner, J. (1997). Efficient estimation in the bivariate normal copula model: normal margins are least favourable. *Bernoulli*, pages 55–77.
- Liu, J. S. and Wu, Y. N. (1999). Parameter Expansion for Data Augmentation. *Journal of the American Statistical Association*, 94(448):1264.
- Lopes, H. and West, M. (2004). Bayesian model assessment in factor analysis. *Statistica Sinica*, 14(1):41–68.
- Meng, X. and Van Dyk, D. (1999). Seeking Efficient Data Augmentation Schemes via Conditional and Marginal Augmentation. *Biometrika*, 86(2):301.
- Moustaki, I. and Knott, M. (2000). Generalized latent trait models. *Psychometrika*, 65(3):391–411.

- Muthén, B. (1983). Latent Variable Structural Equation Modeling with Categorical Data. *Journal of Econometrics*, 22(1-2):43–65.
- North, D. C. and Weingast, B. R. (1989). Constitutions and commitment: the evolution of institutions governing public choice in seventeenth-century England. *The journal of economic history*, 49(04):803–832.
- Paisley, J. and Carin, L. (2009). Nonparametric factor analysis with beta process priors. *Proceedings of the 26th Annual International Conference on Machine Learning - ICML '09*, pages 1–8.
- Pitt, M., Chan, D., and Kohn, R. (2006). Efficient Bayesian inference for Gaussian copula regression models. *Biometrika*, 93(3):537–554.
- Polson, N. and Scott, J. (2010). Shrink Globally, Act Locally: Sparse Bayesian Regularization and Prediction. *Bayesian Statistics*, 9.
- Quinn, K. M. (2004). Bayesian Factor Analysis for Mixed Ordinal and Continuous Responses. *Political Analysis*, 12(4):338–353.
- Sammel, M., Ryan, L., and Legler, J. (1997). Latent Variable Models for Mixed Discrete and Continuous Outcomes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 59(3):667–678.
- Sklar, A. (1959). Fonctions de répartition à n dimensions et leurs marges. *Publ. Inst. Statist. Univ. Paris*, 8:229–231.
- Song, X.-Y., Pan, J.-H., Kwok, T., Vandempuut, L., Ohlsson, C., and Leung, P.-C. (2010). A semiparametric Bayesian approach for structural equation models. *Biometrical journal. Biometrische Zeitschrift*, 52(3):314–32.
- West, M. (2003). Bayesian factor regression models in the large p, small n paradigm. *Bayesian statistics*, 7(2003):723–732.
- Yang, M. and Dunson, D. B. (2010). Bayesian Semiparametric Structural Equation Models with Latent Variables. *Psychometrika*, 75(4):675–693.