

# A Class of Conjugate Priors for Multinomial Probit Models which Includes the Multivariate Normal One

**Augusto Fasano**

AUGUSTO.FASANO@UNITO.IT

*Department of Economics and Statistics  
University of Torino and Collegio Carlo Alberto  
Corso Unione Sovietica 218/bis, 10134, Torino, Italy*

**Daniele Durante**

DANIELE.DURANTE@UNIBOCCONI.IT

*Department of Decision Sciences  
Bocconi Institute for Data Science and Analytics  
Bocconi University  
Via Röntgen 1, 20131, Milan, Italy*

**Editor:** David Blei

## Abstract

Multinomial probit models are routinely-implemented representations for learning how the class probabilities of categorical response data change with  $p$  observed predictors. Although several frequentist methods have been developed for estimation, inference and classification within such a class of models, Bayesian inference is still lagging behind. This is due to the apparent absence of a tractable class of conjugate priors, that may facilitate posterior inference on the multinomial probit coefficients. Such an issue has motivated increasing efforts toward the development of effective Markov chain Monte Carlo methods, but state-of-the-art solutions still face severe computational bottlenecks, especially in high dimensions. In this article, we show that the entire class of unified skew-normal (SUN) distributions is conjugate to several multinomial probit models. Leveraging this result and the SUN properties, we improve upon state-of-the-art solutions for posterior inference and classification both in terms of closed-form results for several functionals of interest, and also by developing novel computational methods relying either on independent and identically distributed samples from the exact posterior or on scalable and accurate variational approximations based on blocked partially-factorized representations. As illustrated in simulations and in a gastrointestinal lesions application, the magnitude of the improvements relative to current methods is particularly evident, in practice, when the focus is on high-dimensional studies.

**Keywords:** Bayesian inference, categorical data, classification, multinomial probit model, unified skew-normal distribution, variational Bayes

## 1. Introduction

Regression models for categorical data are ubiquitous in various fields of application and play a fundamental role in classification (e.g., Agresti, 2013). Within this framework, the overarching goal is to learn how a vector of  $L$  class probabilities  $\boldsymbol{\pi}(\mathbf{x}_i) = [\pi_1(\mathbf{x}_i), \dots, \pi_L(\mathbf{x}_i)]^\top = [\text{pr}(y_i = 1 \mid \boldsymbol{\beta}, \mathbf{x}_i), \dots, \text{pr}(y_i = L \mid \boldsymbol{\beta}, \mathbf{x}_i)]^\top$  changes with a set of  $p$  predictors  $\mathbf{x}_i$ , observed for every unit  $i = 1, \dots, n$ , where  $\boldsymbol{\beta}$  denotes a vector of coefficients controlling the predictors' effects. We refer to Maddala (1986); Greene (2003) and Agresti (2013) for a broad overview

of popular formulations to address such a goal, and focus in this article on the class of multinomial probit models. Indeed, such a broad set of formulations has gained vast popularity in social science, economics and machine learning applications, among others, due to their natural connection with Gaussian regression models, that act as latent predictor-dependent random utilities in a discrete choice setting and also ensure improved interpretability (Hausman and Wise, 1978; Daganzo, 1979). Moreover, expressing predictor-dependent class probabilities via correlated Gaussian latent utilities facilitates improved flexibility, thus avoiding restrictive assumptions, such as the *independence of irrelevant alternatives* (Hausman and Wise, 1978). These desirable properties have stimulated extensive implementations also in the machine learning context (e.g., Girolami and Rogers, 2006; Rogers and Girolami, 2007; Riihimäki et al., 2013; Johndrow et al., 2013; Agarwal et al., 2014; Kindo et al., 2016), while motivating several generalizations which extend the classical formulation in Hausman and Wise (1978) to incorporate class-specific predictor effects (Stern, 1992) and sequential discrete choices (Tutz, 1991).

The aforementioned benefits come, however, with computational difficulties in dealing with integrals of multivariate Gaussian densities (e.g., Genz, 1992; Horrace, 2005; Chopin, 2011; Botev, 2017; Genton et al., 2018; Cao et al., 2019, 2021). These challenges have stimulated an intensive research both in frequentist and in Bayesian settings. In this article, we aim to provide theoretical, methodological and computational advances for the second class of approaches to inference. Indeed, while the frequentist methods for estimation, inference and classification in multinomial probit models are relatively well-established (McFadden, 1989; Stern, 1992; Börsch-Supan and Hajivassiliou, 1993; Geweke et al., 1994; Natarajan et al., 2000), state-of-the-art Bayesian solutions rely either on Markov chain Monte Carlo (MCMC) methods (Albert and Chib, 1993; McCulloch and Rossi, 1994; Nobile, 1998; McCulloch et al., 2000; Albert and Chib, 2001; Chen and Kuo, 2002; Imai and Van Dyk, 2005; Zhang et al., 2006; Chan and Jeliaskov, 2009; Burgette and Nordheim, 2012; Johndrow et al., 2013) or on approximations of the posterior (Girolami and Rogers, 2006; Girolami and Zhong, 2007; Riihimäki et al., 2013; Knowles and Minka, 2011). Despite being widely implemented, both solutions still raise open questions in terms of accuracy, efficiency and computational tractability, especially in large  $p$  settings and in imbalanced situations where some classes are relatively less frequent than others. Recalling Chopin and Ridgway (2017); Johndrow et al. (2019); Durante (2019) and Fasano et al. (2022), these issues arise also in simple univariate probit models, and, as discussed in Section 2, are even more common in multinomial settings since the dimension of the parameters’ space often grows also with the number of classes  $L$ , due to the inclusion of class-specific effects (e.g., Stern, 1992; Tutz, 1991). In addition, MCMC and approximate methods are still sub-optimal relative to situations in which the posterior is analytically available from a tractable class of distributions.

In Sections 2 and 3, we generalize recent findings on univariate binary probit regression in Durante (2019) to prove that the entire class of unified skew-normal (SUN) distributions (Arellano-Valle and Azzalini, 2006)—which includes the classical Gaussian ones as a special case—is a conjugate prior for  $\beta$  in common multinomial probit models (Hausman and Wise, 1978; Stern, 1992; Tutz, 1991). Such a general class of distributions has been originally developed in seemingly unrelated contexts to introduce skewness in a multivariate Gaussian density through the cumulative distribution function of another Gaussian vector, thereby retaining several probabilistic properties of multivariate Gaussian variables (Arellano-Valle

and Azzalini, 2006; Azzalini and Capitanio, 2013). Leveraging such properties, we derive in Section 3 closed-form expressions for posterior predictive distributions and marginal likelihoods which facilitate classification, model selection and inference, also for the parameters regulating the dependence structure among the  $L$  alternatives. In fact, although the overarching focus of this article is to provide novel results that facilitate Bayesian inference for the  $\beta$  coefficients in multinomial probit models, the closed-form expression we derive for the marginal likelihood is also useful to develop improved methods for point estimation and full Bayesian inference (McCulloch and Rossi, 1994; Nobile, 1998; McCulloch et al., 2000; Imai and Van Dyk, 2005; Chan and Jeliazkov, 2009) also on the dependence structure between the different classes; see Section 3.1 for additional discussion and details.

The evaluation of more complex functionals of the posterior distribution for  $\beta$  proceeds instead via improved Monte Carlo methods which, unlike for state-of-the-art MCMC routines, rely on independent and identically distributed samples from the exact SUN posterior, thus avoid mixing issues and convergence diagnostics. As discussed in Section 3.2.1, such an improved strategy deals with multivariate truncated normals and cumulative distribution functions of multivariate Gaussians whose dimension grows with the sample size  $n$ . Hence, the proposed strategy is particularly useful, in practice, in small-to-moderate  $n$  settings, and massively improves state-of-the-art solutions in large  $p$  studies, a situation which occurs in various applications but is computationally impractical under the available implementations (Chopin and Ridgway, 2017). To address the scalability issues of the methods proposed in Section 3.2.1, we further improve and extend in Section 3.2.2 recent partially-factorized variational methods for univariate probit models (Fasano et al., 2022) to devise novel blocked partially-factorized approximations of the posterior distribution in multinomial probit regression which easily scale to large  $p$  and  $n$  datasets, and almost perfectly match the exact posterior, especially when  $p > n$ . These findings are further illustrated in a simulation study in Section 4, and in a gastrointestinal lesions application (Mesejo et al., 2016) in Section 5. Section 6 presents future directions of research which highlight how these novel results can motivate applied, methodological and computational advances in multinomial probit models. All proofs can be found in Appendix A, and extend conjugacy properties of Gaussian and SUN distributions in probit settings. Initial results on these properties are presented in Durante (2019), with a focus on Bayesian univariate binary probit regression. These results are a special case of our broader derivations which require novel extensions to incorporate classical multinomial probit models (Hausman and Wise, 1978), and related generalizations (Stern, 1992; Tutz, 1991). As clarified in Section 2, these formulations rely on more complex latent variable representations, typically based on the maximum of a multivariate vector of latent utilities that usually require a separate treatment relative to the univariate case.

## 2. Multinomial Probit Models

In this section we review three widely-implemented multinomial probit models that cover a large range of applications. These include the classical formulation presented in Hausman and Wise (1978), and two subsequent generalizations which account for class-specific predictor effects (Stern, 1992) and sequential discrete choices (Tutz, 1991). Despite providing different generative mechanisms for the class probability vector  $\boldsymbol{\pi}(\mathbf{x}_i) = [\pi_1(\mathbf{x}_i), \dots, \pi_L(\mathbf{x}_i)]^\top$ , all these representations rely on latent Gaussian random utilities and the associated likeli-

hood can be expressed via the cumulative distribution function of a multivariate Gaussian; see Sections 2.1–2.3. This facilitates the derivation of the conjugacy results for  $\boldsymbol{\beta}$  in Section 3. As mentioned in Section 1, estimation and inference for the parameters quantifying the dependence structure among the class-specific latent utilities is often of interest (e.g., McCulloch and Rossi, 1994; Nobile, 1998; McCulloch et al., 2000; Imai and Van Dyk, 2005; Chan and Jeliazkov, 2009). Although this goal goes beyond the scope of our contribution, in Section 3.1 we also discuss how the conjugacy results derived for  $\boldsymbol{\beta}$  can have direct consequences in improving estimation and inference on the dependence among the  $L$  alternatives.

## 2.1 Classical Discrete Choice Multinomial Probit Models

Let us first focus on the classical discrete choice model as originally formulated by Hausman and Wise (1978). Recalling Greene (2003), this representation expresses each class probability  $\pi_l(\mathbf{x}_i)$  via a random utility model in which every unit  $i$  chooses among  $L$  alternatives by maximizing a set of latent Gaussian utilities  $z_{i1}, \dots, z_{iL}$  that depend on  $p$ -dimensional vectors of class-specific attributes  $\mathbf{x}_{i1}, \dots, \mathbf{x}_{iL}$ —encoded in  $\mathbf{x}_i$ —as perceived by unit  $i$ . More specifically, each class probability  $\pi_l(\mathbf{x}_i)$  in  $\boldsymbol{\pi}(\mathbf{x}_i) = [\pi_1(\mathbf{x}_i), \dots, \pi_L(\mathbf{x}_i)]^\top$  can be written as

$$\text{pr}(y_i = l \mid \boldsymbol{\beta}, \mathbf{x}_i) = \text{pr}(z_{il} > z_{ik}, \forall k \neq l) = \text{pr}(\mathbf{x}_{il}^\top \boldsymbol{\beta} + \varepsilon_{il} > \mathbf{x}_{ik}^\top \boldsymbol{\beta} + \varepsilon_{ik}, \forall k \neq l), \quad (1)$$

for every  $l = 1, \dots, L$ , where  $\boldsymbol{\varepsilon}_i = (\varepsilon_{i1}, \dots, \varepsilon_{iL})^\top \sim \mathcal{N}_L(\mathbf{0}, \boldsymbol{\Sigma})$ , independently for  $i = 1, \dots, n$ ; see Greene (2003) for identifiability restrictions on the matrix  $\boldsymbol{\Sigma}$  regulating the dependence among the  $L$  alternatives.

In (1), the generic vector  $\mathbf{x}_{il} = (x_{il1}, \dots, x_{ilp})^\top$  of predictors has elements  $x_{ilj}$  measuring how the  $j$ th attribute of the  $l$ th alternative is perceived by unit  $i$ . For instance, in political studies (e.g., Dow and Endersby, 2004), each  $\mathbf{x}_{il}$  can include both information on voter  $i$  and attributes of candidate  $l$  as perceived by voter  $i$ . Hence, this specification assumes that to each individual  $i$  are associated  $L$  vectors of  $p$  observed predictors whose linear combinations  $\mathbf{x}_{i1}^\top \boldsymbol{\beta}, \dots, \mathbf{x}_{iL}^\top \boldsymbol{\beta}$  contribute to defining the  $L$  class-specific latent utilities  $z_{i1}, \dots, z_{iL}$ . Each individual  $i$  will then choose the alternative with the highest random utility  $z_{il} = \mathbf{x}_{il}^\top \boldsymbol{\beta} + \varepsilon_{il}$ , which is defined by a deterministic component  $\mathbf{x}_{il}^\top \boldsymbol{\beta}$  with  $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^\top$ , plus a Gaussian noise  $\varepsilon_{il}$ . This term accounts for deviations from the deterministic part due to potential unobserved attributes and, as stated in Proposition 1, it induces a joint likelihood for the observed response data  $\mathbf{y} = (y_1, \dots, y_n)^\top$  that coincides with the cumulative distribution function of an  $[n(L-1)]$ -variate Gaussian.

**Proposition 1** *Let  $\mathbf{v}_l$  denote an  $L \times 1$  vector having value 1 in position  $l$  and 0 elsewhere, for every  $l = 1, \dots, L$ . Moreover, for every  $l = 1, \dots, L$ , denote with  $\mathbf{V}_{[-l]}$  and  $\mathbf{X}_{i[-l]}$  the  $(L-1) \times L$  and  $(L-1) \times p$  matrices whose rows are obtained by stacking vectors  $(\mathbf{v}_k - \mathbf{v}_l)^\top$  and  $(\mathbf{x}_{il} - \mathbf{x}_{ik})^\top$ , respectively, for all  $k \neq l$ . Then, under the model in (1), with  $\boldsymbol{\varepsilon}_i \sim \mathcal{N}_L(\mathbf{0}, \boldsymbol{\Sigma})$  independently for every unit  $i = 1, \dots, n$ , we have*

$$p(\mathbf{y} \mid \boldsymbol{\beta}, \mathbf{X}) = \prod_{i=1}^n p(y_i \mid \boldsymbol{\beta}, \mathbf{x}_i) = \prod_{i=1}^n \Phi_{L-1}(\mathbf{X}_{i[-y_i]} \boldsymbol{\beta}; \mathbf{V}_{[-y_i]} \boldsymbol{\Sigma} \mathbf{V}_{[-y_i]}^\top) = \Phi_{n(L-1)}(\bar{\mathbf{X}} \boldsymbol{\beta}; \boldsymbol{\Lambda}), \quad (2)$$

where  $\bar{\mathbf{X}}$  is an  $[n(L-1)] \times p$  block matrix with  $(L-1) \times p$  row blocks  $\bar{\mathbf{X}}_{[i]} = \mathbf{X}_{i[-y_i]}$ , for each  $i = 1, \dots, n$ , whereas  $\boldsymbol{\Lambda}$  denotes an  $[n(L-1)] \times [n(L-1)]$  block diagonal covariance matrix

with  $(L - 1) \times (L - 1)$  diagonal blocks  $\mathbf{\Lambda}_{[ii]} = \mathbf{V}_{[-y_i]} \mathbf{\Sigma} \mathbf{V}_{[-y_i]}^\top$ , for every  $i = 1, \dots, n$ . In (2), the generic function  $\Phi_c(\mathbf{w}; \mathbf{S})$  denotes the cumulative distribution function, evaluated at  $\mathbf{w}$ , of a  $c$ -variate Gaussian with mean vector  $\mathbf{0}$  and covariance matrix  $\mathbf{S}$ .

The above results follow directly from (1) after noting that  $\text{pr}(y_i = l \mid \boldsymbol{\beta}, \mathbf{x}_i)$  can be written as  $\text{pr}(\varepsilon_{ik} - \varepsilon_{il} < (\mathbf{x}_{il} - \mathbf{x}_{ik})^\top \boldsymbol{\beta}, \forall k \neq l) = \text{pr}(\mathbf{V}_{[-l]} \boldsymbol{\varepsilon}_i < \mathbf{X}_{i[-l]} \boldsymbol{\beta}) = \Phi_{L-1}(\mathbf{X}_{i[-l]} \boldsymbol{\beta}; \mathbf{V}_{[-l]} \mathbf{\Sigma} \mathbf{V}_{[-l]}^\top)$ , where  $\boldsymbol{\varepsilon}_i \sim N_L(\mathbf{0}, \mathbf{\Sigma})$  and, hence,  $\mathbf{V}_{[-l]} \boldsymbol{\varepsilon}_i \sim N_{L-1}(\mathbf{0}, \mathbf{V}_{[-l]} \mathbf{\Sigma} \mathbf{V}_{[-l]}^\top)$ . The final equality in (2) is instead a direct consequence of the properties of multivariate Gaussian random variables. Indeed, since  $\mathbf{\Lambda}$  is a block diagonal covariance matrix and  $\bar{\mathbf{X}} \boldsymbol{\beta}$  is obtained by stacking sub-vectors  $\mathbf{X}_{i[-y_i]} \boldsymbol{\beta}$  for  $i = 1, \dots, n$ , it follows that  $\Phi_{n(L-1)}(\bar{\mathbf{X}} \boldsymbol{\beta}; \mathbf{\Lambda})$  factorizes as the product of  $n$  cumulative distribution functions of  $(L - 1)$ -variate Gaussians.

As mentioned previously, this formulation has been originally developed in social science and economic studies where there is a vector of predictors  $\mathbf{x}_{il}$  for each combination of unit  $i$  and class  $l$  (Hausman and Wise, 1978). This is, however, not always the case in general classification settings. Indeed, in these situations it is more common to observe only a single vector  $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})^\top$  of  $p$  predictors for each statistical unit  $i = 1, \dots, n$  and the focus is on modeling the vector  $\boldsymbol{\pi}(\mathbf{x}_i) = [\pi_1(\mathbf{x}_i), \dots, \pi_L(\mathbf{x}_i)]^\top$ , to ultimately predict the class  $y_i$  of unit  $i$ . In Sections 2.2 and 2.3 we focus on two widely-implemented representations (Stern, 1992; Tutz, 1991), which address this goal, while still relying on Gaussian latent utilities.

## 2.2 Discrete Choice Multinomial Probit Models with Class-Specific Effects

When a single vector  $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})^\top$  of  $p$  covariates is observed for each unit  $i = 1, \dots, n$ , an interpretable and common solution to model differences in the class probabilities within  $\boldsymbol{\pi}(\mathbf{x}_i) = [\pi_1(\mathbf{x}_i), \dots, \pi_L(\mathbf{x}_i)]^\top$  is to introduce class-specific predictors' effects  $\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_L$  as in Stern (1992), and define again  $\pi_l(\mathbf{x}_i)$  as a function of Gaussian utilities  $z_{i1}, \dots, z_{iL}$  via

$$\text{pr}(y_i = l \mid \boldsymbol{\beta}, \mathbf{x}_i) = \text{pr}(z_{il} > z_{ik}, \forall k \neq l) = \text{pr}(\mathbf{x}_i^\top \boldsymbol{\beta}_l + \varepsilon_{il} > \mathbf{x}_i^\top \boldsymbol{\beta}_k + \varepsilon_{ik}, \forall k \neq l), \quad (3)$$

for each  $l = 1, \dots, L$ , where  $\boldsymbol{\varepsilon}_i = (\varepsilon_{i1}, \dots, \varepsilon_{iL})^\top \sim N_L(\mathbf{0}, \mathbf{\Sigma})$ , independently for every unit  $i = 1, \dots, n$ , and  $\boldsymbol{\beta}_L = \mathbf{0}$  for identifiability purposes (Johndrow et al., 2013). Representation (3) and its interpretation are closely related to the classical discrete choice multinomial probit model in Section 2.1, with the only exception that the differences in the class-specific latent utilities  $z_{i1}, \dots, z_{iL}$ ,  $i = 1, \dots, n$ , are now driven by changes in the vectors of coefficients  $\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_L$ , rather than in the vectors of predictors as in model (1). For instance, recalling the political example discussed in Section 2.1, although the age is an attribute specific to voter  $i$ , it is reasonable to expect that such a covariate has a different effect in producing the utilities  $z_{i1} = \mathbf{x}_i^\top \boldsymbol{\beta}_1 + \varepsilon_{i1}, \dots, z_{iL} = \mathbf{x}_i^\top \boldsymbol{\beta}_L + \varepsilon_{iL}$  that voter  $i$  assigns to the different candidates  $l = 1, \dots, L$ . This property can be included by allowing the coefficient associated with the age attribute to change across classes, thus providing a formulation more similar to classical multinomial logit models (e.g., Greene, 2003), relative to (1). As stated in Proposition 2, also under this representation the likelihood for the observed response data  $\mathbf{y} = (y_1, \dots, y_n)^\top$  coincides with the cumulative distribution function of an  $[n(L - 1)]$ -variate Gaussian.

**Proposition 2** Denote with  $\mathbf{v}_l$  the  $L \times 1$  vector with value 1 in position  $l$  and 0 elsewhere, for each  $l = 1, \dots, L$ . Moreover, let  $\mathbf{x}_{il} = \bar{\mathbf{v}}_l \otimes \mathbf{x}_i$ , where  $\bar{\mathbf{v}}_l$  is the  $(L - 1) \times 1$  vector obtained

by the removing the  $L$ -th element from  $\mathbf{v}_l$ , whereas  $\otimes$  denotes the Kronecker product. Then, under model (3) with  $\varepsilon_i \sim \mathbf{N}_L(\mathbf{0}, \boldsymbol{\Sigma})$ , independently for each unit  $i = 1, \dots, n$ , we have

$$p(\mathbf{y} \mid \boldsymbol{\beta}, \mathbf{X}) = \prod_{i=1}^n p(y_i \mid \boldsymbol{\beta}, \mathbf{x}_i) = \prod_{i=1}^n \Phi_{L-1}(\mathbf{X}_{i[-y_i]}\boldsymbol{\beta}; \mathbf{V}_{[-y_i]}\boldsymbol{\Sigma}\mathbf{V}_{[-y_i]}^\top) = \Phi_{n(L-1)}(\bar{\mathbf{X}}\boldsymbol{\beta}; \boldsymbol{\Lambda}), \quad (4)$$

where  $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^\top, \dots, \boldsymbol{\beta}_{L-1}^\top)^\top$ , while  $\mathbf{X}_{i[-y_i]}$ ,  $\mathbf{V}_{[-y_i]}$ ,  $\bar{\mathbf{X}}$  and  $\boldsymbol{\Lambda}$  are defined as in Proposition 1, after setting  $\mathbf{x}_{il} = \bar{\mathbf{v}}_l \otimes \mathbf{x}_i$  for each  $i = 1, \dots, n$  and  $l = 1, \dots, L$ . Hence, in this case  $\mathbf{X}_{i[-y_i]}$  and  $\bar{\mathbf{X}}$  have dimension  $(L-1) \times [p(L-1)]$  and  $[n(L-1)] \times [p(L-1)]$ , respectively.

Proposition 2 follows as a directed consequence of Proposition 1, upon noticing that model (3) can be re-written as a particular case of model (1) with working covariates  $\mathbf{x}_{il}$  as defined in Proposition 2. Indeed, note that by setting  $\mathbf{x}_{il} = \bar{\mathbf{v}}_l \otimes \mathbf{x}_i$ ,  $i = 1, \dots, n$ ,  $l = 1, \dots, L$  and  $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^\top, \dots, \boldsymbol{\beta}_{L-1}^\top)^\top$ , the class probabilities in (3) can be re-expressed as  $\text{pr}(y_i = l \mid \boldsymbol{\beta}, \mathbf{x}_i) = \text{pr}(z_{il} > z_{ik}, \forall k \neq l) = \text{pr}(\mathbf{x}_i^\top \boldsymbol{\beta}_l + \varepsilon_{il} > \mathbf{x}_i^\top \boldsymbol{\beta}_k + \varepsilon_{ik}, \forall k \neq l) = \text{pr}(\mathbf{x}_{il}^\top \boldsymbol{\beta} + \varepsilon_{il} > \mathbf{x}_{ik}^\top \boldsymbol{\beta} + \varepsilon_{ik}, \forall k \neq l)$ , for  $l = 1, \dots, L$ , where the last quantity is the equation for the class probabilities in (1).

### 2.3 Sequential Discrete Choice Multinomial Probit Models

Before focusing on prior specification and posterior derivations, we consider also an extension of the sequential discrete choice multinomial probit model studied in Albert and Chib (2001) and originally proposed by Tutz (1991). Such a model still relies on a set of class-specific latent utilities but is conceptually different from those presented in Sections 2.1 and 2.2, since the choice among the  $L$  classes is modeled via a nested sequence of binary decisions where the generic step  $l$  of this sequential decision process is reached if individual  $i$  has not chosen classes  $1, \dots, l-1$ . At this step, the binary decision will be to either pick class  $l$  with probability  $\text{pr}(y_i = l \mid y_i > l-1, \boldsymbol{\beta}, \mathbf{x}_i) = \Phi(\mathbf{x}_i^\top \boldsymbol{\beta}_l)$  or to consider one of the subsequent alternatives  $l+1, \dots, L$  with complement probability  $\text{pr}(y_i > l \mid y_i > l-1, \boldsymbol{\beta}, \mathbf{x}_i) = 1 - \Phi(\mathbf{x}_i^\top \boldsymbol{\beta}_l)$ . Note that relative to the original formulations in Albert and Chib (2001) and Tutz (1991), here we consider a slightly different reparameterization and also allow the entire vector of coefficients, and not just the intercept, to change with the different labels, thus providing a more general representation. As discussed by Albert and Chib (2001) also this model has a latent utility representation which expresses each  $\pi_l(\mathbf{x}_i)$  in  $\boldsymbol{\pi}(\mathbf{x}_i) = [\pi_1(\mathbf{x}_i), \dots, \pi_L(\mathbf{x}_i)]^\top$  as

$$\text{pr}(y_i = l \mid \boldsymbol{\beta}, \mathbf{x}_i) = \text{pr}(z_{il} > 0) \prod_{k=1}^{l-1} \text{pr}(z_{ik} < 0) = \text{pr}(\mathbf{x}_i^\top \boldsymbol{\beta}_l + \varepsilon_{il} > 0) \prod_{k=1}^{l-1} \text{pr}(\mathbf{x}_i^\top \boldsymbol{\beta}_k + \varepsilon_{ik} < 0), \quad (5)$$

for  $l = 1, \dots, L-1$ , and  $\text{pr}(y_i = L \mid \boldsymbol{\beta}, \mathbf{x}_i) = \prod_{k=1}^{L-1} \text{pr}(\mathbf{x}_i^\top \boldsymbol{\beta}_k + \varepsilon_{ik} < 0)$ , where  $\varepsilon_{il} \sim \mathbf{N}(0, 1)$  independently for every unit  $i = 1, \dots, n$  and class  $l = 1, \dots, L-1$ . Model (5) provides a general representation in which each  $z_{il} = \mathbf{x}_i^\top \boldsymbol{\beta}_l + \varepsilon_{il}$  is the utility of choosing alternative  $l$  against the subsequent ones  $l+1, \dots, L$ , given that the classes  $1, \dots, l-1$  have not been selected in the previous steps of the sequential decision process. Proposition 3 shows that, although conceptually different from the models in Sections 2.1–2.2, also such a formulation admits a similar expression for the joint likelihood of the data  $\mathbf{y} = (y_1, \dots, y_n)^\top$ .

**Proposition 3** Define  $\bar{\mathbf{y}}_i = (\mathbf{0}_{y_i-1}^\top, 1)^\top$  if  $y_i \leq L-1$ , and  $\bar{\mathbf{y}}_i = \mathbf{0}_{L-1}$  if  $y_i = L$ , where the generic  $\mathbf{0}_c$  is a  $c \times 1$  vector of zeroes. Moreover, let  $\bar{n} = n_1 + \dots + n_n$  with  $n_i = \min(y_i, L-1)$ .

Then, under (5) with  $\varepsilon_{il} \sim \mathcal{N}(0, 1)$  independently for  $i = 1, \dots, n$ ,  $l = 1, \dots, L - 1$ , we have

$$p(\mathbf{y} \mid \boldsymbol{\beta}, \mathbf{X}) = \prod_{i=1}^n p(y_i \mid \boldsymbol{\beta}, \mathbf{x}_i) = \prod_{i=1}^n \Phi_{n_i}(\mathbf{X}_i \boldsymbol{\beta}; \mathbf{I}_{n_i}) = \Phi_{\bar{n}}(\bar{\mathbf{X}} \boldsymbol{\beta}; \boldsymbol{\Lambda}), \quad (6)$$

where  $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^\top, \dots, \boldsymbol{\beta}_{L-1}^\top)^\top$ ,  $\boldsymbol{\Lambda} = \mathbf{I}_{\bar{n}}$  and  $\bar{\mathbf{X}}$  denotes an  $\bar{n} \times [p(L-1)]$  matrix with  $n_i \times [p(L-1)]$  row blocks  $\bar{\mathbf{X}}_{[i]} = \mathbf{X}_i$  defined as  $\mathbf{X}_i = (\text{diag}(2\bar{\mathbf{y}}_i - \mathbf{1}) \otimes \mathbf{x}_i^\top, \mathbf{0}_{n_i \times [p(L-1)-n_i]})$ , for every statistical unit  $i = 1, \dots, n$ . In (6), the quantity  $\mathbf{I}_{n_i}$  refers to the  $n_i \times n_i$  identity matrix.

To clarify Proposition 3, it suffices to re-write  $\text{pr}(y_i = l \mid \boldsymbol{\beta}, \mathbf{x}_i)$ ,  $l = 1, \dots, L - 1$ , in (5), as  $\Phi(\mathbf{x}_i^\top \boldsymbol{\beta}_l) \prod_{k=1}^{l-1} [1 - \Phi(\mathbf{x}_i^\top \boldsymbol{\beta}_k)] = \prod_{k=1}^l \Phi[(2\bar{y}_{ik} - 1) \mathbf{x}_i^\top \boldsymbol{\beta}_k] = \Phi_l(\mathbf{X}_i \boldsymbol{\beta}; \mathbf{I}_l)$ , where  $\bar{\mathbf{y}}_i$  is defined as in Proposition 3. The above result leverages standard properties of multivariate Gaussians.

Combining Propositions 1–3 it is clear that, despite characterizing different utility-based decision mechanisms, models (1), (3) and (5) have a similar form for the joint likelihood. The only difference among such likelihoods is the dimension of the cumulative distribution functions and the definition of the known matrices  $\bar{\mathbf{X}}$  and  $\boldsymbol{\Lambda}$ , which change depending on the type of model. These results are fundamental for the novel conjugacy results in Section 3.

### 3. Conjugate Bayesian Inference for Multinomial Probit Models

Common Bayesian implementations of multinomial probit models consider a multivariate Gaussian prior  $\mathcal{N}_q(\boldsymbol{\xi}, \boldsymbol{\Omega})$  for the parameters in  $\boldsymbol{\beta}$ , where  $q$  is equal to  $p$  in model (1) and to  $p(L-1)$  in models (3) and (5), whereas  $\boldsymbol{\xi}$  and  $\boldsymbol{\Omega}$  denote the pre-specified prior mean vector and covariance matrix, respectively (Albert and Chib, 1993; McCulloch and Rossi, 1994; Nobile, 1998; McCulloch et al., 2000; Albert and Chib, 2001; Chen and Kuo, 2002; Imai and Van Dyk, 2005; Zhang et al., 2006; Burgette and Nordheim, 2012; Johndrow et al., 2013). Besides providing a default specification in various Bayesian regression models, this choice is also motivated by the Gaussian assumption for the latent utilities in (1), (3) and (5) which implies an augmented data representation facilitating the implementation of MCMC (e.g., Albert and Chib, 1993, 2001; Imai and Van Dyk, 2005; Holmes and Held, 2006; Chopin and Ridgway, 2017) and approximate methods (e.g., Girolami and Rogers, 2006; Girolami and Zhong, 2007; Riihimäki et al., 2013; Knowles and Minka, 2011) for inference and prediction.

As discussed in Section 1, the above strategies have computational drawbacks—especially in large  $p$  settings—and are motivated by the apparent absence of conjugacy between multinomial probit likelihoods and the Gaussian prior for  $\boldsymbol{\beta}$ . In Section 3.1, we show not only that the posterior in this setting is a SUN, but also that the whole SUN family is conjugate to multinomial probits, thereby obtaining closed-form posterior distributions under a broad variety of priors, which include also the default Gaussian one and, as a byproduct, Gaussian processes. Leveraging the novel results in Section 3.1, we develop in Section 3.2 improved Monte Carlo methods for full Bayesian inference and classification, along with scalable and accurate approximations of the SUN posterior in high-dimensional settings.

Before providing an overview of the SUN distribution (Arellano-Valle and Azzalini, 2006; Azzalini and Capitanio, 2013) and presenting our conjugacy results, we shall emphasize that some of the aforementioned contributions consider also priors for  $\boldsymbol{\Sigma}$  in models (1) and (3). Recalling Sections 1–2, our focus in this article is on the posterior for  $\boldsymbol{\beta}$  conditioned on  $\boldsymbol{\Sigma}$  and, therefore, we avoid additional identifiability and computational complications which

arise when including a prior also for  $\Sigma$ . Nonetheless, as discussed in Section 3.1, the closed-form expression for the marginal likelihood  $p(\mathbf{y} \mid \mathbf{X})$  presented in Corollary 1, and the i.i.d. sampler to generate values from the posterior  $p(\boldsymbol{\beta} \mid \mathbf{y}, \mathbf{X})$  outlined in Algorithm 1, can be useful to improve both point estimation and full Bayesian inference also on  $\Sigma$ .

### 3.1 Conjugacy via Unified Skew-Normal Priors

Consistent with Section 3, let us assume a  $\text{SUN}_{q,h}(\boldsymbol{\xi}, \boldsymbol{\Omega}, \boldsymbol{\Delta}, \boldsymbol{\gamma}, \boldsymbol{\Gamma})$  prior for  $\boldsymbol{\beta}$ , whose density

$$p(\boldsymbol{\beta}) = \phi_q(\boldsymbol{\beta} - \boldsymbol{\xi}; \boldsymbol{\Omega}) \frac{\Phi_h(\boldsymbol{\gamma} + \boldsymbol{\Delta}^\top \bar{\boldsymbol{\Omega}}^{-1} \boldsymbol{\omega}^{-1}(\boldsymbol{\beta} - \boldsymbol{\xi}); \boldsymbol{\Gamma} - \boldsymbol{\Delta}^\top \bar{\boldsymbol{\Omega}}^{-1} \boldsymbol{\Delta})}{\Phi_h(\boldsymbol{\gamma}; \boldsymbol{\Gamma})}, \quad (7)$$

is obtained by modifying the density function  $\phi_q(\boldsymbol{\beta} - \boldsymbol{\xi}; \boldsymbol{\Omega})$  of a  $q$ -variate Gaussian  $N_q(\boldsymbol{\xi}, \boldsymbol{\Omega})$ , via a skewness-inducing mechanism driven by the cumulative distribution function, computed at  $\boldsymbol{\gamma} + \boldsymbol{\Delta}^\top \bar{\boldsymbol{\Omega}}^{-1} \boldsymbol{\omega}^{-1}(\boldsymbol{\beta} - \boldsymbol{\xi}) \in \mathbb{R}^h$ , of an  $h$ -variate Gaussian with mean vector  $\mathbf{0}$  and  $h \times h$  covariance matrix  $\boldsymbol{\Gamma} - \boldsymbol{\Delta}^\top \bar{\boldsymbol{\Omega}}^{-1} \boldsymbol{\Delta}$ . The quantity  $\Phi_h(\boldsymbol{\gamma}; \boldsymbol{\Gamma})$  is instead the normalizing constant, which coincides with the cumulative distribution function, evaluated at  $\boldsymbol{\gamma} \in \mathbb{R}^h$ , of an  $h$ -variate Gaussian with mean vector  $\mathbf{0}$  and  $h \times h$  covariance matrix  $\boldsymbol{\Gamma}$ . As is clear from (7), the amount of skewness in the prior is mainly controlled by the  $q \times h$  matrix  $\boldsymbol{\Delta}$ . Indeed, when all the entries in  $\boldsymbol{\Delta}$  are 0, the prior  $p(\boldsymbol{\beta})$  in (7) coincides with the density of a  $q$ -variate Gaussian with mean vector  $\boldsymbol{\xi}$  and covariance matrix  $\boldsymbol{\Omega} = \boldsymbol{\omega} \bar{\boldsymbol{\Omega}} \boldsymbol{\omega}$  obtained via the quadratic combination among the correlation matrix  $\bar{\boldsymbol{\Omega}}$  and the diagonal scale matrix  $\boldsymbol{\omega} = (\boldsymbol{\Omega} \odot \mathbf{I}_q)^{1/2}$ , where  $\odot$  is the element-wise Hadamard product. Such a class of Gaussian priors can be also easily obtained by setting  $h = 0$  in (7). As discussed in Arellano-Valle and Azzalini (2006), the multivariate Gaussian case is just an example of a broad variety of distributions which can be obtained from prior (7) under suitable choices for its parameters. Additional priors of interest within this class are independent univariate skew-normals (Azzalini, 1985) for the coefficients in  $\boldsymbol{\beta}$  and classical multivariate skew-normals (Azzalini and Dalla Valle, 1996) for the entire vector  $\boldsymbol{\beta}$ . Therefore, our results allow tractable inference in Bayesian multinomial probit models under a broad class of priors that include Gaussian specifications along with asymmetric priors which may be useful in social science and econometric applications. Note that also non-linear effects modeled via Gaussian processes induce a multivariate Gaussian prior and, hence, our results can be directly applied to the flexible classification strategies discussed in Girolami and Rogers (2006) and Riihimäki et al. (2013), among others.

To further clarify the main roles of the parameters  $\boldsymbol{\xi}, \boldsymbol{\Omega}, \boldsymbol{\Delta}, \boldsymbol{\gamma}$  and  $\boldsymbol{\Gamma}$  note that, as shown in Arellano-Valle and Azzalini (2006), if  $\boldsymbol{\beta} \sim \text{SUN}_{q,h}(\boldsymbol{\xi}, \boldsymbol{\Omega}, \boldsymbol{\Delta}, \boldsymbol{\gamma}, \boldsymbol{\Gamma})$ , then

$$\boldsymbol{\beta} \stackrel{d}{=} \boldsymbol{\xi} + \boldsymbol{\omega}(\mathbf{V}_0 + \boldsymbol{\Delta} \boldsymbol{\Gamma}^{-1} \mathbf{V}_1), \quad \text{with } \mathbf{V}_0 \sim N_q(\mathbf{0}, \bar{\boldsymbol{\Omega}} - \boldsymbol{\Delta} \boldsymbol{\Gamma}^{-1} \boldsymbol{\Delta}^\top), \quad \mathbf{V}_1 \sim \text{TN}_h(-\boldsymbol{\gamma}; \mathbf{0}, \boldsymbol{\Gamma}), \quad (8)$$

where  $\text{TN}_h(-\boldsymbol{\gamma}; \mathbf{0}, \boldsymbol{\Gamma})$  denotes an  $h$ -variate Gaussian with zero mean, covariance matrix  $\boldsymbol{\Gamma}$  and truncation below  $-\boldsymbol{\gamma}$ . Recalling Arellano-Valle and Azzalini (2006), representation (8) also relates to a simple conditioning mechanism. In particular, if  $\mathbf{U}_1 \in \mathbb{R}^h$  and  $\mathbf{U}_0 \in \mathbb{R}^q$  are two random vectors jointly distributed as a  $N_{h+q}(\mathbf{0}, \boldsymbol{\Omega}^*)$ , where  $\boldsymbol{\Omega}^*$  denotes a correlation matrix with blocks  $\boldsymbol{\Omega}_{[11]}^* = \boldsymbol{\Gamma}$ ,  $\boldsymbol{\Omega}_{[22]}^* = \bar{\boldsymbol{\Omega}}$  and  $\boldsymbol{\Omega}_{[21]}^* = (\boldsymbol{\Omega}_{[12]}^*)^\top = \boldsymbol{\Delta}$ , then  $\boldsymbol{\beta} = \boldsymbol{\xi} + \boldsymbol{\omega} \bar{\boldsymbol{\beta}}$ , with  $\bar{\boldsymbol{\beta}} = (\mathbf{U}_0 \mid \mathbf{U}_1 + \boldsymbol{\gamma} > \mathbf{0})$ , has SUN distribution with density as in (7). Hence,  $\boldsymbol{\xi}$  and  $\boldsymbol{\omega}$  mostly regulate the location and the scale of the prior, while  $\boldsymbol{\Delta}, \boldsymbol{\Gamma}$ , and  $\bar{\boldsymbol{\Omega}}$  control dependence and skewness. Finally,  $\boldsymbol{\gamma}$  defines the truncation threshold in the conditioning part.



Model	Prior	SUN $_{q,h+m}(\boldsymbol{\xi}_{\text{pst}}, \boldsymbol{\Omega}_{\text{pst}}, \boldsymbol{\Delta}_{\text{pst}}, \boldsymbol{\gamma}_{\text{pst}}, \boldsymbol{\Gamma}_{\text{pst}})$		
		$q$	$h$	$m$
Model (1) (Proposition 1)		$p$	$h$	$n(L-1)$
Model (3) (Proposition 2)	SUN $_{q,h}(\boldsymbol{\xi}, \boldsymbol{\Omega}, \boldsymbol{\Delta}, \boldsymbol{\gamma}, \boldsymbol{\Gamma})$	$p(L-1)$	$h$	$n(L-1)$
Model (5) (Proposition 3)		$p(L-1)$	$h$	$\bar{n}$
Model (1) (Proposition 1)		$p$	0	$n(L-1)$
Model (3) (Proposition 2)	N $_q(\boldsymbol{\xi}, \boldsymbol{\Omega})$	$p(L-1)$	0	$n(L-1)$
Model (5) (Proposition 3)		$p(L-1)$	0	$\bar{n}$

Table 1: Dimension of the SUN $_{q,h+m}(\boldsymbol{\xi}_{\text{pst}}, \boldsymbol{\Omega}_{\text{pst}}, \boldsymbol{\Delta}_{\text{pst}}, \boldsymbol{\gamma}_{\text{pst}}, \boldsymbol{\Gamma}_{\text{pst}})$  posterior for  $\boldsymbol{\beta}$  in the multinomial probit models in Section 2, when considering either SUN or Gaussian priors.

Besides clarifying the role of the prior parameters, the additive representation (8) of the SUN random variable is useful also for posterior inference since, as we will discuss, it provides a direct strategy to sample i.i.d. values from the SUN distribution, thereby improving upon state-of-the-art MCMC methods for Bayesian multinomial probit models. Indeed, as shown in Theorem 1, the SUN prior in (7) is conjugate to the multinomial probit likelihoods reported in (2), (4) and (6), meaning that also the posterior  $(\boldsymbol{\beta} \mid \mathbf{y}, \mathbf{X})$  has a SUN distribution. In particular  $(\boldsymbol{\beta} \mid \mathbf{y}, \mathbf{X}) \sim \text{SUN}_{q,h+m}(\boldsymbol{\xi}_{\text{pst}}, \boldsymbol{\Omega}_{\text{pst}}, \boldsymbol{\Delta}_{\text{pst}}, \boldsymbol{\gamma}_{\text{pst}}, \boldsymbol{\Gamma}_{\text{pst}})$ ; see Table 1 for details on the specific dimensions of the SUN posterior under the three multinomial probit models discussed in Section 2, when considering either SUN $_{q,h}(\boldsymbol{\xi}, \boldsymbol{\Omega}, \boldsymbol{\Delta}, \boldsymbol{\gamma}, \boldsymbol{\Gamma})$  or N $_q(\boldsymbol{\xi}, \boldsymbol{\Omega})$  priors.

**Theorem 1** *Let  $p(\boldsymbol{\beta})$  denote the SUN prior density in (7), and define with  $\Phi_m(\bar{\mathbf{X}}\boldsymbol{\beta}; \boldsymbol{\Lambda})$  the generic multinomial probit likelihood reported in (2), (4) and (6), with  $m$ ,  $\bar{\mathbf{X}}$  and  $\boldsymbol{\Lambda}$  defined as in Propositions 1, 2 or 3 depending on whether the focus is on the model (1), (3) or (5), respectively; see also Table 1. Then, the posterior density  $p(\boldsymbol{\beta} \mid \mathbf{y}, \mathbf{X})$  of  $\boldsymbol{\beta}$  is*

$$p(\boldsymbol{\beta} \mid \mathbf{y}, \mathbf{X}) = \phi_q(\boldsymbol{\beta} - \boldsymbol{\xi}_{\text{pst}}; \boldsymbol{\Omega}_{\text{pst}}) \frac{\Phi_{h+m}(\boldsymbol{\gamma}_{\text{pst}} + \boldsymbol{\Delta}_{\text{pst}}^\top \bar{\boldsymbol{\Omega}}_{\text{pst}}^{-1} \boldsymbol{\omega}_{\text{pst}}^{-1} (\boldsymbol{\beta} - \boldsymbol{\xi}_{\text{pst}}); \boldsymbol{\Gamma}_{\text{pst}} - \boldsymbol{\Delta}_{\text{pst}}^\top \bar{\boldsymbol{\Omega}}_{\text{pst}}^{-1} \boldsymbol{\Delta}_{\text{pst}})}{\Phi_{h+m}(\boldsymbol{\gamma}_{\text{pst}}; \boldsymbol{\Gamma}_{\text{pst}})}, \quad (9)$$

with  $\boldsymbol{\xi}_{\text{pst}} = \boldsymbol{\xi}$ ,  $\boldsymbol{\Omega}_{\text{pst}} = \boldsymbol{\Omega}$ ,  $\boldsymbol{\Delta}_{\text{pst}} = (\boldsymbol{\Delta}, \bar{\boldsymbol{\Omega}}\boldsymbol{\omega}\bar{\mathbf{X}}^\top \mathbf{s}^{-1})$ ,  $\boldsymbol{\gamma}_{\text{pst}} = (\boldsymbol{\gamma}^\top, \boldsymbol{\xi}^\top \bar{\mathbf{X}}^\top \mathbf{s}^{-1})^\top$ , while  $\boldsymbol{\Gamma}_{\text{pst}}$  is an  $(h+m) \times (h+m)$  covariance matrix with blocks  $\boldsymbol{\Gamma}_{\text{pst}[11]} = \boldsymbol{\Gamma}$ ,  $\boldsymbol{\Gamma}_{\text{pst}[22]} = \mathbf{s}^{-1}(\bar{\mathbf{X}}\boldsymbol{\Omega}\bar{\mathbf{X}}^\top + \boldsymbol{\Lambda})\mathbf{s}^{-1}$  and  $\boldsymbol{\Gamma}_{\text{pst}[21]} = \boldsymbol{\Gamma}_{\text{pst}[12]}^\top = \mathbf{s}^{-1}\bar{\mathbf{X}}\boldsymbol{\omega}\boldsymbol{\Delta}$ , where  $\mathbf{s} = [(\bar{\mathbf{X}}\boldsymbol{\Omega}\bar{\mathbf{X}}^\top + \boldsymbol{\Lambda}) \odot \mathbf{I}_m]^{1/2}$ . Note that in (9), the dimension  $q$  is equal to  $p$  under model (1), and to  $p(L-1)$  under models (3) and (5).

**Remark 1** *As a consequence of Theorem 1, it follows that also the multivariate Gaussian prior—which provides a special case of unified skew-normal—yields to a SUN posterior when updated with the multinomial probit likelihoods in (2), (4) and (6). In particular, if  $p(\boldsymbol{\beta}) = \phi_q(\boldsymbol{\beta} - \boldsymbol{\xi}; \boldsymbol{\Omega})$  it immediately follows from Theorem 1 that the posterior distribution is a SUN having density as in (9), with  $h = 0$  and posterior parameters  $\boldsymbol{\xi}_{\text{pst}} = \boldsymbol{\xi}$ ,  $\boldsymbol{\Omega}_{\text{pst}} = \boldsymbol{\Omega}$ ,  $\boldsymbol{\Delta}_{\text{pst}} = \bar{\boldsymbol{\Omega}}\boldsymbol{\omega}\bar{\mathbf{X}}^\top \mathbf{s}^{-1}$ ,  $\boldsymbol{\gamma}_{\text{pst}} = \mathbf{s}^{-1}\bar{\mathbf{X}}\boldsymbol{\xi}$ ,  $\boldsymbol{\Gamma}_{\text{pst}} = \mathbf{s}^{-1}(\bar{\mathbf{X}}\boldsymbol{\Omega}\bar{\mathbf{X}}^\top + \boldsymbol{\Lambda})\mathbf{s}^{-1}$ , where  $\mathbf{s} = [(\bar{\mathbf{X}}\boldsymbol{\Omega}\bar{\mathbf{X}}^\top + \boldsymbol{\Lambda}) \odot \mathbf{I}_m]^{1/2}$ .*

Theorem 1 generalizes Corollary 4 in Durante (2019) to provide novel results with important implications in Bayesian inference for multinomial probit models. As discussed in Arellano-Valle and Azzalini (2006) SUN distributions share several common properties with

multivariate Gaussians. A relevant one is that this family is closed under marginalization, linear combinations and conditioning. Within our context, this means that the posterior for each single coefficient and linear combinations of interest—such as those defining the latent utilities—are still SUN and their parameters can be obtained via simple transformations of those in Theorem 1 (Arellano-Valle and Azzalini, 2006; Azzalini and Capitanio, 2013). According to (9), also the normalizing constant of the posterior is available in closed form and coincides with the cumulative distribution function  $\Phi_{h+m}(\boldsymbol{\gamma}_{\text{pst}}; \boldsymbol{\Gamma}_{\text{pst}})$ , evaluated at  $\boldsymbol{\gamma}_{\text{pst}}$ , of a multivariate Gaussian with  $\mathbf{0}$  mean and covariance matrix  $\boldsymbol{\Gamma}_{\text{pst}}$ . As outlined in Corollaries 1 and 2, this result is fundamental to obtain closed-form expressions for marginal likelihoods and predictive distributions, that are useful for model selection and classification.

**Corollary 1** *Under the settings in Theorem 1, the marginal likelihood can be expressed as*

$$p(\mathbf{y} \mid \mathbf{X}) = \frac{p(\mathbf{y}, \boldsymbol{\beta} \mid \mathbf{X})}{p(\boldsymbol{\beta} \mid \mathbf{y}, \mathbf{X})} = \frac{\Phi_{h+m}(\boldsymbol{\gamma}_{\text{pst}}; \boldsymbol{\Gamma}_{\text{pst}})}{\Phi_h(\boldsymbol{\gamma}; \boldsymbol{\Gamma})}, \quad (10)$$

with  $\boldsymbol{\gamma}_{\text{pst}}$  and  $\boldsymbol{\Gamma}_{\text{pst}}$  defined as in Theorem 1.

**Corollary 2** *Consider the expanded dataset in which, besides the original data  $\mathbf{y}$  and  $\mathbf{X}$ , we also have an additional unit with predictors  $\mathbf{x}_{\text{new}}$  and response  $y_{\text{new}} = l$ . Moreover, let  $m_l$ ,  $\bar{\mathbf{X}}_l$  and  $\boldsymbol{\Lambda}_l$  be defined analogously to  $m$ ,  $\bar{\mathbf{X}}$  and  $\boldsymbol{\Lambda}$  in Theorem 1, when the expanded dataset is considered. Then, under the settings of Theorem 1, we have that*

$$\text{pr}(y_{\text{new}} = l \mid \mathbf{y}, \mathbf{X}, \mathbf{x}_{\text{new}}) = \frac{\text{pr}(y_{\text{new}} = l, \mathbf{y} \mid \mathbf{X}, \mathbf{x}_{\text{new}})}{p(\mathbf{y} \mid \mathbf{X}, \mathbf{x}_{\text{new}})} = \frac{\Phi_{h+m_l}(\boldsymbol{\gamma}_{l\text{pst}}; \boldsymbol{\Gamma}_{l\text{pst}})}{\Phi_{h+m}(\boldsymbol{\gamma}_{\text{pst}}; \boldsymbol{\Gamma}_{\text{pst}})}, \quad (11)$$

for each  $l = 1, \dots, L$ , with  $\boldsymbol{\gamma}_{\text{pst}}$  and  $\boldsymbol{\Gamma}_{\text{pst}}$  as in Theorem 1, while  $\boldsymbol{\gamma}_{l\text{pst}}$  and  $\boldsymbol{\Gamma}_{l\text{pst}}$  coincide with  $\boldsymbol{\gamma}_{\text{pst}}$  and  $\boldsymbol{\Gamma}_{\text{pst}}$ , evaluated at  $\bar{\mathbf{X}}_l$  and  $\boldsymbol{\Lambda}_l$ , instead of  $\bar{\mathbf{X}}$  and  $\boldsymbol{\Lambda}$ .

Corollaries 1–2 facilitate closed-form Bayesian model selection and classification without the need to rely on MCMC. Moreover, although point estimation and full Bayesian inference on  $\boldsymbol{\Sigma}$  goes beyond the scope of the present contribution, as anticipated in Sections 1 and 2, Corollary 1 is practically relevant also for improving current solutions addressing this goal (e.g., McCulloch and Rossi, 1994; Nobile, 1998; McCulloch et al., 2000; Imai and Van Dyk, 2005; Chan and Jeliaskov, 2009). For example, estimation for the parameters in  $\boldsymbol{\Sigma}$  could proceed via direct maximization of the marginal likelihood in (10), after integrating out  $\boldsymbol{\beta}$  analytically. Also full Bayesian inference for  $\boldsymbol{\Sigma}$  can benefit from the closed-form marginal likelihood in Corollary 1, since it allows implementation of collapsed Metropolis–Hastings schemes that produce samples from the posterior of  $\boldsymbol{\Sigma}$  after integrating out  $\boldsymbol{\beta}$  analytically. This strategy is expected to yield gains in mixing relative to common MCMC methods that leverage full-conditional distributions depending both on  $\boldsymbol{\beta}$  and, potentially, on augmented data (Park and Van Dyk, 2009). Such advantages are practically relevant and notable in settings where  $\boldsymbol{\beta}$  is high-dimensional and the size of the distribution functions in Corollary 1 is small-to-moderate, thereby allowing accurate and rapid evaluation of (10) at different  $\boldsymbol{\Sigma}$  via recent strategies (Botev, 2017; Genton et al., 2018; Cao et al., 2019, 2021).

Exploiting the moment generating function of the SUN in Section 2.3 of Arellano-Valle and Azzalini (2006) and the additional derivations in Azzalini and Bacchieri (2010); Gupta

---

**Algorithm 1:** Strategy to sample i.i.d. from the SUN posterior in Theorem 1

---

**for**  $t = 1, \dots, T$  **do**

- [1] Sample  $\mathbf{V}_0^{(t)} \sim N_q(\mathbf{0}, \bar{\boldsymbol{\Omega}}_{\text{pst}} - \boldsymbol{\Delta}_{\text{pst}} \boldsymbol{\Gamma}_{\text{pst}}^{-1} \boldsymbol{\Delta}_{\text{pst}}^\top)$  [in R use `rmvnorm`]
- [2] Sample  $\mathbf{V}_1^{(t)} \sim \text{TN}_{h+m}(-\gamma_{\text{pst}}; \mathbf{0}, \boldsymbol{\Gamma}_{\text{pst}})$  [in R use `rtmvnorm` (Botev, 2017)]
- [3] Set  $\boldsymbol{\beta}^{(t)} = \boldsymbol{\xi}_{\text{pst}} + \boldsymbol{\omega}_{\text{pst}} (\mathbf{V}_0^{(t)} + \boldsymbol{\Delta}_{\text{pst}} \boldsymbol{\Gamma}_{\text{pst}}^{-1} \mathbf{V}_1^{(t)})$

**Output:** i.i.d. samples  $\boldsymbol{\beta}^{(1)}, \dots, \boldsymbol{\beta}^{(T)}$  from (9). Based on such samples, posterior functionals  $\mathbb{E}[g(\boldsymbol{\beta}) \mid \mathbf{y}, \mathbf{X}]$  can be computed, via Monte Carlo, as  $\sum_{t=1}^T g(\boldsymbol{\beta}^{(t)})/T$ .

---

et al. (2013) and Azzalini and Capitanio (2013) closed-form expressions can be derived also for the posterior mean, covariance matrix and cumulative distribution function of  $\boldsymbol{\beta}$ , thereby facilitating Bayesian point estimation, uncertainty quantification and classification. Such expressions require, however, the evaluation of multivariate Gaussian cumulative distribution functions and tedious derivations that do not facilitate calculation of more complex functionals, thus motivating the alternative computational methods presented in Section 3.2.

### 3.2 Computational Methods

This section provides new computational methods for Bayesian multinomial probit models that exploit results in Section 3.1 to improve upon state-of-the-art routines, especially in large  $q$  settings. In particular, in Section 3.2.1 we derive Monte Carlo methods that, unlike current MCMC solutions, rely on independent and identically distributed samples from the exact SUN posterior in (9). Such a strategy requires to sample from  $(h+m)$ -variate truncated normals with full covariance matrix and, hence, becomes impractical as  $h+m$  grows. To address this issue, we also propose in Section 3.2.2 a blocked partially-factorized variational Bayes that relaxes various factorization assumptions of classical mean-field families to obtain more accurate and computationally efficient approximations, that almost perfectly match the exact posterior in large  $q$  settings, especially when  $q > h+m$ ; see Table 1 for details on how  $q$ ,  $h$  and  $m$  relate to  $p$ ,  $n$  and  $L$  under the multinomial probit models in Sections 2.1–2.3.

#### 3.2.1 MONTE CARLO METHODS VIA INDEPENDENT SAMPLES FROM THE POSTERIOR

Complex functionals of the posterior can be effectively evaluated via Monte Carlo methods leveraging the additive representation of the SUN in (8). This allows to sample independent and identically distributed (i.i.d.) values from the posterior in Theorem 1, via linear combinations among samples from multivariate Gaussians and multivariate truncated normals. As outlined in Algorithm 1, this routine crucially avoids MCMC methods, thus circumventing convergence and mixing issues commonly seen in Bayesian multinomial probit (Johndrow et al., 2013), while allowing parallel implementations. A possible computational drawback in Algorithm 1 is sampling from  $\text{TN}_{h+m}(-\gamma_{\text{pst}}; \mathbf{0}, \boldsymbol{\Gamma}_{\text{pst}})$ . Recent advances based on minimax tilting methods (Botev, 2017) have made this task computationally feasible for multivariate truncated normals with a dimension of few hundreds, thereby making Algorithm 1 an efficient strategy in small-to-moderate  $h+m$  and large, potentially huge,  $q$  studies. Recalling Chopin and Ridgway (2017), these large  $q$  settings are actually those where state-of-the-art MCMC methods, including STAN implementations of Hamiltonian no-u-turn samplers (Hoff-

man and Gelman, 2014), are computationally unfeasible. The results in Botev (2017) are also useful to compute efficiently Gaussian cumulative distribution functions, and hence are practically relevant to evaluate (10) and (11) in small-to-moderate  $h + m$  settings.

### 3.2.2 BLOCKED PARTIALLY-FACTORIZED VARIATIONAL BAYES

As discussed in Section 3.2.1, when  $h + m$  is large, sampling from  $(h + m)$ -variate truncated normals with full covariance matrix becomes computationally unfeasible (Botev, 2017), thus making Algorithm 1 impractical in these settings. Typically,  $h$  is either 0—when Gaussian priors are considered—or is equal to a small value, whereas  $m$  depends on the sample size  $n$  and on the number of classes  $L$ ; see Table 1. Hence, it is necessary to devise more scalable methods, especially in common settings where  $n$  is larger than a few hundreds.

A possible solution to the above problem is to consider approximations of the posterior density, with variational Bayes providing a well-established procedure, especially in those models admitting simple augmented data representations (Blei et al., 2017). As clarified in Section 2, this is the case of multinomial probit models relying on Gaussian latent utilities. Such a property has motivated several variational strategies to approximate the joint posterior  $p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})$  of  $\boldsymbol{\beta}$  and the augmented data  $\bar{\mathbf{z}}$ , with a tractable density  $q^*(\boldsymbol{\beta}, \bar{\mathbf{z}})$ , which is the closest in Kullback–Leibler (KL) divergence (Kullback and Leibler, 1951) to  $p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})$ , among all the densities which belong to a pre-specified approximating family  $\mathcal{Q}$ . As for the development of simple Gibbs samplers relying on tractable full-conditionals (Albert and Chib, 1993), the inclusion of the augmented data facilitates the implementation of simple coordinate ascent variational inference (CAVI) routines (Bishop, 2006; Blei et al., 2017) to minimize, with respect to  $q(\boldsymbol{\beta}, \bar{\mathbf{z}})$ , the divergence  $\text{KL}[q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \parallel p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})]$ .

Clearly, the availability of simple optimization routines and strategies to derive the optimal marginal  $q^*(\boldsymbol{\beta})$  from  $q^*(\boldsymbol{\beta}, \bar{\mathbf{z}})$ , depend also on the choice of the family  $\mathcal{Q}$ . Common solutions in binary (e.g., Consonni and Marin, 2007) and multinomial (e.g., Girolami and Rogers, 2006) probit settings rely on mean-field families  $\mathcal{Q}_{\text{MF}} = \{q(\boldsymbol{\beta}, \bar{\mathbf{z}}) : q(\boldsymbol{\beta}, \bar{\mathbf{z}}) = q(\boldsymbol{\beta})q(\bar{\mathbf{z}})\}$  that assume independence between  $\boldsymbol{\beta}$  and  $\bar{\mathbf{z}}$ . These strategies come with simple CAVI algorithms which scale easily to high-dimensional settings and, due to the factorized form of  $q^*(\boldsymbol{\beta}, \bar{\mathbf{z}})$ , provide as a byproduct the approximating density  $q^*(\boldsymbol{\beta})$  of direct interest. However, recent theoretical and empirical studies on simple univariate probit models (Fasano et al., 2022), have shown that such a mean-field assumption often leads to a low-quality approximation in high-dimensional probit settings, which severely affects not only uncertainty quantification, but also estimation and classification. To address this issue in the context of basic univariate probit regression with Gaussian priors, Fasano et al. (2022) considered a partially-factorized mean-field approximating family  $\mathcal{Q}_{\text{PFM}} = \{q(\boldsymbol{\beta}, \bar{\mathbf{z}}) : q(\boldsymbol{\beta}, \bar{\mathbf{z}}) = q(\boldsymbol{\beta} \mid \bar{\mathbf{z}}) \prod_{r=1}^{h+m} q(\bar{z}_r)\}$  which avoids enforcing independence between  $\boldsymbol{\beta}$  and  $\bar{\mathbf{z}}$ , and only assumes that  $q(\bar{\mathbf{z}})$  factorizes as the product of its marginals. This novel class of approximating densities substantially improves the quality of the original mean-field approximation and almost perfectly matches the exact posterior in high-dimensional settings, especially when the number of predictors is higher than the sample size, without sacrificing computational tractability. Unfortunately, this strategy is only available for univariate binary probit models with Gaussian priors.

Motivated by the above discussion, we develop a new blocked partially-factorized mean-field approximation which extends the contribution of Fasano et al. (2022) in three main

important directions. In particular, we [i] allow the inclusion of SUN and not only Gaussian priors, [ii] generalize the methods to multinomial probit models, and [iii] further enlarge the class of approximating densities by replacing  $\prod_{r=1}^{h+m} q(\bar{z}_r)$  in  $\mathcal{Q}_{\text{PFM}}$  with  $\prod_{c=1}^C q(\bar{\mathbf{z}}_c)$ , where  $\bar{\mathbf{z}}_1, \dots, \bar{\mathbf{z}}_C$  are distinct sub-vectors of  $\bar{\mathbf{z}}$ , such that  $\bar{\mathbf{z}} = (\bar{\mathbf{z}}_1^\top, \dots, \bar{\mathbf{z}}_C^\top)^\top$ . Therefore, instead of enforcing independence among all the augmented data, we only make this assumption between pre-specified blocks. In fact, while in high-dimensional univariate binary settings the independence among all the augmented data does not seem to have a major impact on the quality of the approximation (Fasano et al., 2022), this may not be the case in multinomial probit models. For example, under the formulation presented in Section 2.2, every unit  $i$  enters the matrix  $\bar{\mathbf{X}}$  multiple times and, hence, it is reasonable to expect a relatively strong dependence among unit-specific augmented data, which cannot be accurately approximated by a fully factorized representation for  $q(\bar{\mathbf{z}})$ . Similar blocking ideas have been also considered by Chopin (2011); Genton et al. (2018) and Cao et al. (2019), to simulate from multivariate truncated normals and compute cumulative distribution functions of high-dimensional Gaussians. We adapt these ideas in the context of variational inference to obtain improved approximations of the posterior, without affecting computational performance.

To introduce the blocked partially-factorized mean-field approximation, first note that the kernel of the posterior density  $p(\boldsymbol{\beta} \mid \mathbf{y}, \mathbf{X})$  in (9) can be re-written as

$$p(\boldsymbol{\beta} \mid \mathbf{y}, \mathbf{X}) \propto \phi_q(\boldsymbol{\beta} - \boldsymbol{\xi}_{\text{pst}}; \boldsymbol{\Omega}_{\text{pst}}) \int \phi_{h+m}(\bar{\mathbf{z}} - (\boldsymbol{\eta}_{\text{pst}} + \mathbf{X}_{\text{pst}}\boldsymbol{\beta}); \boldsymbol{\Sigma}_{\text{pst}}) \mathbb{1}(\bar{\mathbf{z}} > \mathbf{0}) d\bar{\mathbf{z}}, \quad (12)$$

where  $\mathbf{X}_{\text{pst}} = \boldsymbol{\Delta}_{\text{pst}}^\top \bar{\boldsymbol{\Omega}}_{\text{pst}}^{-1} \boldsymbol{\omega}_{\text{pst}}^{-1}$ ,  $\boldsymbol{\eta}_{\text{pst}} = \boldsymbol{\gamma}_{\text{pst}} - \mathbf{X}_{\text{pst}} \boldsymbol{\xi}_{\text{pst}}$ , and  $\boldsymbol{\Sigma}_{\text{pst}} = \boldsymbol{\Gamma}_{\text{pst}} - \boldsymbol{\Delta}_{\text{pst}}^\top \bar{\boldsymbol{\Omega}}_{\text{pst}}^{-1} \boldsymbol{\Delta}_{\text{pst}}$ . To clarify the connection between (9) and (12) it suffices to note that the integral in (12) actually coincides with the multivariate Gaussian cumulative distribution function  $\Phi_{h+m}(\boldsymbol{\eta}_{\text{pst}} + \mathbf{X}_{\text{pst}}\boldsymbol{\beta}; \boldsymbol{\Sigma}_{\text{pst}})$  in the numerator of (9). Leveraging this alternative representation and Gaussian–Gaussian conjugacy, we can easily notice that

$$\begin{aligned} p(\boldsymbol{\beta} \mid \bar{\mathbf{z}}, \mathbf{y}, \mathbf{X}) &\propto \phi_q(\boldsymbol{\beta} - \boldsymbol{\xi}_{\text{pst}}; \boldsymbol{\Omega}_{\text{pst}}) \phi_{h+m}(\bar{\mathbf{z}} - (\boldsymbol{\eta}_{\text{pst}} + \mathbf{X}_{\text{pst}}\boldsymbol{\beta}); \boldsymbol{\Sigma}_{\text{pst}}), \\ &\propto \phi_q(\boldsymbol{\beta} - \mathbf{V}_{\text{pst}}[\mathbf{X}_{\text{pst}}^\top \boldsymbol{\Sigma}_{\text{pst}}^{-1}(\bar{\mathbf{z}} - \boldsymbol{\eta}_{\text{pst}}) + \boldsymbol{\Omega}_{\text{pst}}^{-1} \boldsymbol{\xi}_{\text{pst}}]; \mathbf{V}_{\text{pst}}), \end{aligned} \quad (13)$$

where  $\mathbf{V}_{\text{pst}} = (\mathbf{X}_{\text{pst}}^\top \boldsymbol{\Sigma}_{\text{pst}}^{-1} \mathbf{X}_{\text{pst}} + \boldsymbol{\Omega}_{\text{pst}}^{-1})^{-1}$ . Hence,  $(\boldsymbol{\beta} \mid \bar{\mathbf{z}}, \mathbf{y}, \mathbf{X}) \sim N_q(\mathbf{V}_{\text{pst}}[\mathbf{X}_{\text{pst}}^\top \boldsymbol{\Sigma}_{\text{pst}}^{-1}(\bar{\mathbf{z}} - \boldsymbol{\eta}_{\text{pst}}) + \boldsymbol{\Omega}_{\text{pst}}^{-1} \boldsymbol{\xi}_{\text{pst}}], \mathbf{V}_{\text{pst}})$ . On the other hand, according to (12), the conditional density  $p(\bar{\mathbf{z}} \mid \boldsymbol{\beta}, \mathbf{y}, \mathbf{X})$  of the augmented data  $\bar{\mathbf{z}}$  is a multivariate truncated normal with mean  $\boldsymbol{\eta}_{\text{pst}} + \mathbf{X}_{\text{pst}}\boldsymbol{\beta}$ , covariance matrix  $\boldsymbol{\Sigma}_{\text{pst}}$  and truncation below  $\mathbf{0}$ . Therefore, marginalizing out  $\boldsymbol{\beta}$  with density  $\phi_q(\boldsymbol{\beta} - \boldsymbol{\xi}_{\text{pst}}; \boldsymbol{\Omega}_{\text{pst}})$ , yields

$$\begin{aligned} p(\bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X}) &\propto \phi_{h+m}(\bar{\mathbf{z}} - (\boldsymbol{\eta}_{\text{pst}} + \mathbf{X}_{\text{pst}} \boldsymbol{\xi}_{\text{pst}}); \boldsymbol{\Sigma}_{\text{pst}} + \mathbf{X}_{\text{pst}} \boldsymbol{\Omega}_{\text{pst}} \mathbf{X}_{\text{pst}}^\top) \mathbb{1}(\bar{\mathbf{z}} > \mathbf{0}), \\ &\propto \phi_{h+m}(\bar{\mathbf{z}} - \boldsymbol{\gamma}_{\text{pst}}; \boldsymbol{\Gamma}_{\text{pst}}) \mathbb{1}(\bar{\mathbf{z}} > \mathbf{0}), \end{aligned} \quad (14)$$

since  $\boldsymbol{\eta}_{\text{pst}} = \boldsymbol{\gamma}_{\text{pst}} - \mathbf{X}_{\text{pst}} \boldsymbol{\xi}_{\text{pst}}$  and  $\boldsymbol{\Gamma}_{\text{pst}} = \boldsymbol{\Sigma}_{\text{pst}} + \mathbf{X}_{\text{pst}} \boldsymbol{\Omega}_{\text{pst}} \mathbf{X}_{\text{pst}}^\top$ . Combining (13)–(14) and recalling previous discussion, we aim to obtain an accurate approximation  $q^*(\boldsymbol{\beta}, \bar{\mathbf{z}})$  of the joint density

$$\begin{aligned} p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X}) &= p(\boldsymbol{\beta} \mid \bar{\mathbf{z}}, \mathbf{y}, \mathbf{X}) p(\bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X}), \\ &\propto \phi_q(\boldsymbol{\beta} - \mathbf{V}_{\text{pst}}[\mathbf{X}_{\text{pst}}^\top \boldsymbol{\Sigma}_{\text{pst}}^{-1}(\bar{\mathbf{z}} - \boldsymbol{\eta}_{\text{pst}}) + \boldsymbol{\Omega}_{\text{pst}}^{-1} \boldsymbol{\xi}_{\text{pst}}]; \mathbf{V}_{\text{pst}}) \phi_{h+m}(\bar{\mathbf{z}} - \boldsymbol{\gamma}_{\text{pst}}; \boldsymbol{\Gamma}_{\text{pst}}) \mathbb{1}(\bar{\mathbf{z}} > \mathbf{0}), \end{aligned} \quad (15)$$

such that  $q^*(\boldsymbol{\beta}, \bar{\mathbf{z}})$  minimizes the KL divergence  $\text{KL}[q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \parallel p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})]$  within the blocked partially-factorized mean-field family  $\mathcal{Q}_{\text{PFM-B}} = \{q(\boldsymbol{\beta}, \bar{\mathbf{z}}) : q(\boldsymbol{\beta}, \bar{\mathbf{z}}) = q(\boldsymbol{\beta} \mid \bar{\mathbf{z}}) \prod_{c=1}^C q(\bar{\mathbf{z}}_c)\}$ ,

where  $\bar{\mathbf{z}}_1, \dots, \bar{\mathbf{z}}_C$  are the pre-specified sub-vectors of  $\bar{\mathbf{z}}$ . Formulation (15) clarifies why  $\mathcal{Q}_{\text{PFM-B}}$  provides a particularly suitable family of approximating densities for  $p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})$ . In particular, since the exact conditional density  $p(\boldsymbol{\beta} \mid \bar{\mathbf{z}}, \mathbf{y}, \mathbf{X})$  has a tractable Gaussian form, assuming independence between  $\boldsymbol{\beta}$  and  $\bar{\mathbf{z}}$  as in classical mean-field variational Bayes seems an unnecessarily strong assumption. On the other hand, the main source of intractability in  $p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})$  arises from the high-dimensional truncated normal density  $p(\bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})$  with full covariance matrix  $\boldsymbol{\Gamma}_{\text{pst}}$ , thus motivating our attempt to approximate it via a set of  $C$  independent lower-dimensional truncated normal densities  $q^*(\bar{\mathbf{z}}_1) \cdots q^*(\bar{\mathbf{z}}_C)$ . Each of these blocks must be sufficiently small to allow tractable inference under the associated truncated normal approximation, and should be specified so as to group augmented data with strong correlations in  $\boldsymbol{\Gamma}_{\text{pst}}$ . Remark 2 discusses and motivates a possible default strategy to define the different blocks in multinomial probit models, when necessary.

**Remark 2** *In multinomial probit models, when necessary, it is typically sufficient to group augmented data associated with the same unit  $i$ , provided that there may be strong overlap in the rows of  $\bar{\mathbf{X}}$  referring to  $i$ , thereby leading to high correlation in  $\boldsymbol{\Gamma}_{\text{pst}}$ . This choice is further motivated by the fact that the optimal mean-field solution—which does not assume factorized forms for  $q(\bar{\mathbf{z}})$  in  $\mathcal{Q}_{\text{MF}}$ —is defined as  $q_{\text{MF}}^*(\boldsymbol{\beta}, \bar{\mathbf{z}}) = q_{\text{MF}}^*(\boldsymbol{\beta}) \prod_{i=1}^n q_{\text{MF}}^*(\bar{\mathbf{z}}_i)$  under Gaussian priors (Girolami and Rogers, 2006). Such a solution belongs also to  $\mathcal{Q}_{\text{PFM-B}}$  when blocking according to  $i$ . Therefore,  $\min_{q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \in \mathcal{Q}_{\text{PFM-B}}} \text{KL}[q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \mid p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})] \leq \min_{q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \in \mathcal{Q}_{\text{MF}}} \text{KL}[q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \mid p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})]$ . Moreover, since  $\mathcal{Q}_{\text{PFM}} \subset \mathcal{Q}_{\text{PFM-B}}$ , we also have that  $\min_{q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \in \mathcal{Q}_{\text{PFM-B}}} \text{KL}[q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \mid p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})] \leq \min_{q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \in \mathcal{Q}_{\text{PFM}}} \text{KL}[q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \mid p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})]$ . Hence, when blocking according to  $i$ , our solution is guaranteed to improve mean-field variational Bayes and recent partially-factorized extensions, under Gaussian priors. Similar arguments can be made under SUN priors.*

Besides providing a wider and more flexible class, the family  $\mathcal{Q}_{\text{PFM-B}}$  also allows straightforward optimization, as shown in Proposition 4.

**Proposition 4** *The KL divergence  $\text{KL}[q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \mid p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})]$  between  $q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \in \mathcal{Q}_{\text{PFM-B}}$  and  $p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})$  in (15), is minimized at  $q^*(\boldsymbol{\beta}, \bar{\mathbf{z}}) = q^*(\boldsymbol{\beta} \mid \bar{\mathbf{z}}) \prod_{c=1}^C q^*(\bar{\mathbf{z}}_c)$ , with*

$$q^*(\boldsymbol{\beta} \mid \bar{\mathbf{z}}) \propto \phi_q(\boldsymbol{\beta} - \mathbf{V}_{\text{pst}}[\mathbf{X}_{\text{pst}}^\top \boldsymbol{\Sigma}_{\text{pst}}^{-1}(\bar{\mathbf{z}} - \boldsymbol{\eta}_{\text{pst}}) + \boldsymbol{\Omega}_{\text{pst}}^{-1} \boldsymbol{\xi}_{\text{pst}}]; \mathbf{V}_{\text{pst}}), \quad (16)$$

$$q^*(\bar{\mathbf{z}}_c) \propto \phi_{n_c}(\bar{\mathbf{z}}_c - \boldsymbol{\gamma}_{\text{pst}[c]} - \mathbf{W}_{\text{pst}[c]}(\mathbb{E}_{q^*(\bar{\mathbf{z}}_{-c})}(\bar{\mathbf{z}}_{-c}) - \boldsymbol{\gamma}_{\text{pst}[-c]}); \boldsymbol{\Gamma}_{\text{pst}[c]}) \mathbb{1}(\bar{\mathbf{z}}_c > \mathbf{0}), \quad \forall c, \quad (17)$$

where  $\mathbf{W}_{\text{pst}[c]} = \boldsymbol{\Gamma}_{\text{pst}[c,-c]}(\boldsymbol{\Gamma}_{\text{pst}[-c,-c]})^{-1}$  and  $\boldsymbol{\Gamma}_{\text{pst}[c]} = \boldsymbol{\Gamma}_{\text{pst}[c,c]} - \boldsymbol{\Gamma}_{\text{pst}[c,-c]}(\boldsymbol{\Gamma}_{\text{pst}[-c,-c]})^{-1}\boldsymbol{\Gamma}_{\text{pst}[-c,c]}$ , with  $\boldsymbol{\Gamma}_{\text{pst}[c,c]}$ ,  $\boldsymbol{\Gamma}_{\text{pst}[-c,-c]}$ ,  $\boldsymbol{\Gamma}_{\text{pst}[-c,c]}$  and  $\boldsymbol{\Gamma}_{\text{pst}[c,-c]}$  denoting the four blocks of  $\boldsymbol{\Gamma}_{\text{pst}}$  when partitioned to highlight sub-vector  $\bar{\mathbf{z}}_c$  against all the others in  $\bar{\mathbf{z}}_{-c}$ . Similarly,  $\boldsymbol{\gamma}_{\text{pst}[c]}$  and  $\boldsymbol{\gamma}_{\text{pst}[-c]}$  are the sub-vectors of  $\boldsymbol{\gamma}_{\text{pst}}$  referring to block  $c$  and the remaining blocks, respectively. Finally,  $\mathbb{E}_{q^*(\bar{\mathbf{z}}_{-c})}(\bar{\mathbf{z}}_{-c}) = (\mathbb{E}_{q^*(\bar{\mathbf{z}}_1)}(\bar{\mathbf{z}}_1)^\top, \dots, \mathbb{E}_{q^*(\bar{\mathbf{z}}_{c-1})}(\bar{\mathbf{z}}_{c-1})^\top, \mathbb{E}_{q^*(\bar{\mathbf{z}}_{c+1})}(\bar{\mathbf{z}}_{c+1})^\top, \dots, \mathbb{E}_{q^*(\bar{\mathbf{z}}_C)}(\bar{\mathbf{z}}_C)^\top)^\top$ , where the expectations are taken with respect to the optimal truncated normal approximations.

The solution in (16) is a direct consequence of the chain rule for the KL divergence. In fact,  $\text{KL}[q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \mid p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})] = \text{KL}[q(\bar{\mathbf{z}}) \mid p(\bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})] + \mathbb{E}_{q(\bar{\mathbf{z}})}\{\text{KL}[q(\boldsymbol{\beta} \mid \bar{\mathbf{z}}) \mid p(\boldsymbol{\beta} \mid \bar{\mathbf{z}}, \mathbf{y}, \mathbf{X})]\}$ , and hence the non-negative second summand is exactly zero for every  $q(\bar{\mathbf{z}})$  only when  $q^*(\boldsymbol{\beta} \mid \bar{\mathbf{z}}) = p(\boldsymbol{\beta} \mid \bar{\mathbf{z}}, \mathbf{y}, \mathbf{X})$ . To clarify the result in (17), recall that the optimal solution for  $q(\bar{\mathbf{z}}_c)$  is proportional to  $\exp[\mathbb{E}_{q^*(\bar{\mathbf{z}}_{-c})}(\log[p(\bar{\mathbf{z}}_c \mid \bar{\mathbf{z}}_{-c}, \mathbf{y}, \mathbf{X})])]$  (Bishop, 2006; Blei et al., 2017). Hence,

---

**Algorithm 2:** CAVI for blocked partially-factorized approximation in Proposition 4
 

---

**for**  $t = 1$  *until convergence of the ELBO do*
**for**  $c = 1, \dots, C$  **do**

 Set  $\mathbb{E}_{q^{(t)}(\bar{\mathbf{z}}_c)}(\bar{\mathbf{z}}_c)$  equal to the expected value of an  $n_c$ -variate Gaussian with mean  $\boldsymbol{\gamma}_{\text{pst}[c]} + \mathbf{W}_{\text{pst}[c]}(\mathbb{E}_{q^{(t-1)}(\bar{\mathbf{z}}_{-c})}(\bar{\mathbf{z}}_{-c}) - \boldsymbol{\gamma}_{\text{pst}[-c]})$ , covariance matrix  $\boldsymbol{\Gamma}_{\text{pst}[c]}$  and truncation below  $\mathbf{0}$ , where  $\mathbb{E}_{q^{(t-1)}(\bar{\mathbf{z}}_{-c})}(\bar{\mathbf{z}}_{-c})$  is defined as

$$(\mathbb{E}_{q^{(t)}(\bar{\mathbf{z}}_1)}(\bar{\mathbf{z}}_1)^\top, \dots, \mathbb{E}_{q^{(t)}(\bar{\mathbf{z}}_{c-1})}(\bar{\mathbf{z}}_{c-1})^\top, \mathbb{E}_{q^{(t-1)}(\bar{\mathbf{z}}_{c+1})}(\bar{\mathbf{z}}_{c+1})^\top, \dots, \mathbb{E}_{q^{(t-1)}(\bar{\mathbf{z}}_C)}(\bar{\mathbf{z}}_C)^\top)^\top.$$

 [in R use function `MomTrunc` to compute the mean of truncated normals].

**Output:** Optimal truncated normal approximating densities  $q^*(\bar{\mathbf{z}}_1), \dots, q^*(\bar{\mathbf{z}}_C)$  from (17), which are combined with the closed-form solution for  $q^*(\boldsymbol{\beta} | \bar{\mathbf{z}})$  in (16), to provide the optimal joint approximate density  $q^*(\boldsymbol{\beta}, \bar{\mathbf{z}}) = q^*(\boldsymbol{\beta} | \bar{\mathbf{z}}) \prod_{c=1}^C q^*(\bar{\mathbf{z}}_c)$ .
 

---

since  $(\bar{\mathbf{z}} | \mathbf{y}, \mathbf{X})$  has a multivariate truncated Gaussian density in (14), it follows that also each  $p(\bar{\mathbf{z}}_c | \bar{\mathbf{z}}_{-c}, \mathbf{y}, \mathbf{X})$  is an  $n_c$ -variate truncated normal density, whose log-kernel is linear in  $\bar{\mathbf{z}}_{-c}$  and the remaining parameters are specified as in (17) (e.g., Horrace, 2005; Holmes and Held, 2006). According to Proposition 4, the only unknown parameters are  $\mathbb{E}_{q^*(\bar{\mathbf{z}}_c)}(\bar{\mathbf{z}}_c)$ ,  $c = 1, \dots, C$ , whose solution requires solving a non-linear system of equations. Algorithm 2 summarizes the steps of the CAVI to obtain such quantities via simple operations.

Once  $q^*(\boldsymbol{\beta} | \bar{\mathbf{z}})$  and  $q^*(\bar{\mathbf{z}}) = \prod_{c=1}^C q^*(\bar{\mathbf{z}}_c)$  are available, approximations of key functionals of  $\boldsymbol{\beta}$  can be easily derived leveraging the law of total expectation and results in Proposition 4. In particular, since  $\mathbb{E}_{q^*(\boldsymbol{\beta})}(\boldsymbol{\beta}) = \mathbb{E}_{q^*(\bar{\mathbf{z}})}[\mathbb{E}_{q^*(\boldsymbol{\beta}|\bar{\mathbf{z}})}(\boldsymbol{\beta})]$ , we have that

$$\mathbb{E}_{q^*(\boldsymbol{\beta})}(\boldsymbol{\beta}) = \mathbf{V}_{\text{pst}}[\mathbf{X}_{\text{pst}}^\top \boldsymbol{\Sigma}_{\text{pst}}^{-1}(\mathbb{E}_{q^*(\bar{\mathbf{z}})}(\bar{\mathbf{z}}) - \boldsymbol{\eta}_{\text{pst}}) + \boldsymbol{\Omega}_{\text{pst}}^{-1} \boldsymbol{\xi}_{\text{pst}}], \quad (18)$$

whereas, the equality  $\text{var}_{q^*(\boldsymbol{\beta})}(\boldsymbol{\beta}) = \mathbb{E}_{q^*(\bar{\mathbf{z}})}[\text{var}_{q^*(\boldsymbol{\beta}|\bar{\mathbf{z}})}(\boldsymbol{\beta})] + \text{var}_{q^*(\bar{\mathbf{z}})}[\mathbb{E}_{q^*(\boldsymbol{\beta}|\bar{\mathbf{z}})}(\boldsymbol{\beta})]$ , leads to

$$\text{var}_{q^*(\boldsymbol{\beta})}(\boldsymbol{\beta}) = \mathbf{V}_{\text{pst}} + \mathbf{V}_{\text{pst}} \mathbf{X}_{\text{pst}}^\top \boldsymbol{\Sigma}_{\text{pst}}^{-1} \text{var}_{q^*(\bar{\mathbf{z}})}(\bar{\mathbf{z}}) \boldsymbol{\Sigma}_{\text{pst}}^{-1} \mathbf{X}_{\text{pst}} \mathbf{V}_{\text{pst}}. \quad (19)$$

To evaluate (18) and (19), it is sufficient to compute  $\mathbb{E}_{q^*(\bar{\mathbf{z}}_c)}(\bar{\mathbf{z}}_c)$  and  $\text{var}_{q^*(\bar{\mathbf{z}}_c)}(\bar{\mathbf{z}}_c)$ , separately for each  $c = 1, \dots, C$ , since, due to the independence assumption among the  $C$  sub-vectors of  $\bar{\mathbf{z}}$ , the vector  $\mathbb{E}_{q^*(\bar{\mathbf{z}})}(\bar{\mathbf{z}})$  has blocks  $\mathbb{E}_{q^*(\bar{\mathbf{z}})}(\bar{\mathbf{z}})_{[c]} = \mathbb{E}_{q^*(\bar{\mathbf{z}}_c)}(\bar{\mathbf{z}}_c)$ , for each  $c = 1, \dots, C$ , whereas  $\text{var}_{q^*(\bar{\mathbf{z}})}(\bar{\mathbf{z}})$  is a block diagonal matrix with blocks  $\text{var}_{q^*(\bar{\mathbf{z}})}(\bar{\mathbf{z}})_{[cc]} = \text{var}_{q^*(\bar{\mathbf{z}}_c)}(\bar{\mathbf{z}}_c)$ . Recalling Remark 2, in multinomial probit models such blocks often refer to rows in the design matrix  $\bar{\mathbf{X}}$  corresponding to the same unit  $i$  and, hence, their dimensions are, by definition, equal or lower than the number of classes  $L$ , which is small in most applications. This allows rapid evaluation of  $\mathbb{E}_{q^*(\bar{\mathbf{z}}_c)}(\bar{\mathbf{z}}_c)$  and  $\text{var}_{q^*(\bar{\mathbf{z}}_c)}(\bar{\mathbf{z}}_c)$  via routine R functions such as `MomTrunc`.

Although (18) and (19) are typically the main quantities of interest, other generic functionals  $\mathbb{E}_{q^*(\boldsymbol{\beta})}[g(\boldsymbol{\beta})]$  can be easily derived via simple Monte Carlo methods based on samples from  $q^*(\boldsymbol{\beta})$ . Combining (16)–(17), such draws can be obtained by setting

$$\boldsymbol{\beta}^{(t)} = \mathbf{V}_{\text{pst}}[\mathbf{X}_{\text{pst}}^\top \boldsymbol{\Sigma}_{\text{pst}}^{-1}([\bar{\mathbf{z}}_1^{(t)\top}, \dots, \bar{\mathbf{z}}_C^{(t)\top}]^\top - \boldsymbol{\eta}_{\text{pst}}) + \boldsymbol{\Omega}_{\text{pst}}^{-1} \boldsymbol{\xi}_{\text{pst}}] + \boldsymbol{\varepsilon}^{(t)}, \quad t = 1, \dots, T, \quad (20)$$

where  $\boldsymbol{\varepsilon}^{(t)} \sim \mathcal{N}_q(\mathbf{0}, \mathbf{V}_{\text{pst}})$ , and  $\bar{\mathbf{z}}_c^{(t)} \sim \text{TN}_{n_c}(\mathbf{0}; \boldsymbol{\gamma}_{\text{pst}[c]} + \mathbf{W}_{\text{pst}[c]}(\mathbb{E}_{q^*(\bar{\mathbf{z}}_{-c})}(\bar{\mathbf{z}}_{-c}) - \boldsymbol{\gamma}_{\text{pst}[-c]}), \boldsymbol{\Gamma}_{\text{pst}[c]})$  for  $c = 1, \dots, C$ . Also in this case, since  $n_c$  is typically very small, samples from  $n_c$ -variate

truncated normals can be effectively obtained from common R functions, such as `rtmnorm`. This strategy is particularly useful to compute the predictive probabilities for a new unit with covariates  $\mathbf{x}_{\text{new}}$ . To accomplish this goal, it suffices to compute, for each sample  $\boldsymbol{\beta}^{(t)}$  of  $\boldsymbol{\beta}$ , the latent utilities  $z_{\text{new}l}^{(t)}$ ,  $l = 1, \dots, L$  defined either via (1), (3) or (5), depending on the multinomial probit model considered. Then, if the focus is on models (1) and (3), a Monte Carlo estimate for  $\text{pr}(y_{\text{new}} = l \mid \mathbf{y}, \mathbf{X}, \mathbf{x}_{\text{new}})$  can be obtained by computing the relative frequency of samples in which  $z_{\text{new}l}^{(t)} > z_{\text{new}k}^{(t)}$  for all  $k \neq l$ . If, instead, one considers the sequential representation in (5), the Monte Carlo estimate for  $\text{pr}(y_{\text{new}} = l \mid \mathbf{y}, \mathbf{X}, \mathbf{x}_{\text{new}})$  coincides with the relative frequency of samples in which  $z_{\text{new}l}^{(t)} > 0$  and  $z_{\text{new}k}^{(t)} < 0$ , for  $k < l$ .

#### 4. Simulation Study

To evaluate the performance of the computational methods presented in Section 3.2 relative to state-of-the-art competitors, we consider an extensive simulation study under different settings of  $p$  and  $L$ . The main objective is to assess empirical evidence of improved accuracy and computational scalability for Algorithms 1 and 2, over routinely-implemented mean-field approximations (e.g., Girolami and Rogers, 2006; Consonni and Marin, 2007), and Hamiltonian Monte Carlo sampling schemes (e.g., Hoffman and Gelman, 2014) under the `rstan` library. The latter MCMC strategy provides an accurate but expensive default solution in high dimensions, and, hence, is taken as a benchmark throughout the simulation study.

As discussed in Section 3.2, the gains provided by the proposed computational methods are valid for all the three multinomial probit models presented in Sections 2.1–2.3. Due to this and consistent with the application in Section 5, we consider the sequential multinomial probit model in Section 2.3 as a representative example to quantify empirically the magnitude of these gains at varying combinations of  $p \in \{75; 100; 200\}$  and  $L \in \{3; 5; 7; 9\}$ . The sample size is, instead, kept fixed and equal to  $n = 100$  consistent with the empirical results in Fasano et al. (2022) for basic univariate probit regression which show that the accuracy of variational strategies mainly depends on  $p$  and  $n$  through the ratio among such quantities. It is also worth noticing that the settings for  $n$ ,  $p$  and  $L$  considered in this simulation are much lower than those which could be handled under the blocked partially-factorized and mean-field approximations. Such moderate dimensions of  $n$ ,  $p$  and  $L$  are required to avoid possible computational bottlenecks in obtaining i.i.d. samples from Algorithm 1 and MCMC draws under `rstan`, respectively.

For each combination  $(p, L)$ , the predictors in  $\mathbf{X}$  are simulated from independent  $N(0, 1)$  variables and, as suggested in Gelman et al. (2008) and Chopin and Ridgway (2017), such predictors are subsequently standardized to have mean zero and standard deviation 0.5, for the training data. The coefficients  $\beta_{lj}$ ,  $l = 1, \dots, L - 1$ ,  $j = 1, \dots, p$  comprising the vector  $\boldsymbol{\beta}$  are, instead, generated independently from a uniform distribution in  $(-5, 5)$ . Leveraging the realizations of  $\mathbf{X}$  and  $\boldsymbol{\beta}$ , the categorical responses in  $\mathbf{y}$  are simulated from the sequential Bernoulli choice mechanism outlined in Section 2.3. To assess the quality in classification, we also simulate 20 test units following the same procedure presented for the  $n = 100$  training data. Consistent with common implementations of regression models for binary or categorical responses (e.g., Gelman et al., 2008; Chopin and Ridgway, 2017), Bayesian inference is performed under independent weakly informative Gaussian priors for the coefficients in  $\boldsymbol{\beta}$ , with zero mean and variance  $\omega^2 = 25$ .






Figure 1: For three functionals of main interest and different settings of  $L \in \{3; 5; 7; 9\}$ , trajectories for the median of the absolute differences, at varying  $p \in \{75; 100; 200\}$ , between an accurate but expensive Monte Carlo evaluation—via `rstan`—of such functionals and their estimates provided by the i.i.d. sampling scheme from the exact posterior presented in Section 3.2.1 (IID), the blocked partially-factorized variational Bayes proposed in Section 3.2.2 (PFM), and classical mean-field variational Bayes (MF), respectively.

To evaluate the performance of the proposed methods, we conduct posterior inference under the strategies developed in Sections 3.2.1–3.2.2 and compare the results against state-of-the-art alternatives comprising classical mean-field approximations (e.g., Consonni and Marin, 2007; Girolami and Rogers, 2006), and the `rstan` implementation of Hamiltonian Monte Carlo (e.g., Hoffman and Gelman, 2014). More specifically, we consider as benchmark posterior inference for selected functionals of interest computed from 5000 `rstan` samples, and compare such quantities with those resulting from 5000 i.i.d. samples from the exact SUN posterior under Algorithm 1, and the ones computed from the approximate densities provided by the blocked partially-factorized strategy in Algorithm 2, and the classical mean-field variational Bayes solution. Figure 1 summarizes the output of this comparison, with a focus on posterior means, standard deviations, and predictive probabilities for the 20 test

units. For these quantities, we display the median of the absolute differences between the `rstan` estimates and those arising from the other three strategies under analysis. In the first two panels, the medians are computed from the  $p(L - 1)$  absolute differences for the estimates of the posterior moments for every  $\beta_{lj}$ ,  $l = 1, \dots, L - 1$ ,  $j = 1, \dots, p$ , obtained under the different methods, while in the third panel such quantities are calculated for the  $20L$  predictive probabilities, estimated for each test unit and category.

As shown in Figure 1, the partially-factorized solution yields uniformly improved accuracy relative to the classical mean-field one, and the quality of the approximation increases with both  $p$  and  $L$ . Moreover, the error rapidly vanishes when the dimension of  $\beta$  exceeds  $n$ . These results are coherent with the empirical findings in Fasano et al. (2022) on classical univariate binary probit models. Such a superior performance comes at almost no expenses in computational budget, since the average runtime required to obtain the functionals of interest under blocked partially-factorized and mean-field approximations is, respectively, 2.9 and 2.5 seconds, with a maximum of 11.3 and 10.6 seconds, respectively, under the scenario  $(p, L) = (200, 9)$ . These runtimes are orders of magnitude faster than those of the `rstan` implementation of Hamiltonian Monte Carlo, which, on average, requires 489 seconds. Indeed, although `rstan` is efficient in low dimensions, such a method faces increasing computational difficulties as  $p$  and  $L$  grow, thereby yielding average runtimes of approximately 900 seconds in  $p > n$  scenarios. These are exactly the settings in which the i.i.d sampler described in Algorithm 1—that provides similarly accurate estimates relative to `rstan`—displays the highest computational advantages, with an average runtime of only 33 seconds. Therefore, Algorithm 1 addresses a gap in the literature regarding posterior inference in multinomial probit models with small-to-moderate sample size and high-dimensional coefficients vector, a setting where state-of-the-art MCMC are computationally inefficient. When  $n$  grows, the blocked partially-factorized approximation described in Section 3.2.2 and implemented in Algorithm 2 provides a practically feasible and effective solution which uniformly improves the accuracy of standard mean-field strategies, and yields almost the same estimates of state-of-the-art sampling methods when  $p(L - 1) > n$ , at massively lower runtimes. These results on accuracy remained consistent also when comparing other quantiles of the absolute differences.

## 5. Gastrointestinal Lesions Application

To confirm findings in Section 4 also on a real-world application, we consider a medical study by Mesejo et al. (2016) that focuses on 76 gastrointestinal lesions classified as **hyperplastic** ( $l = 1$ ), **serrated adenoma** ( $l = 2$ ) and **adenoma** ( $l = 3$ ) where the first is benign, whereas the others are malignant. For every individual lesion, a vector of 1396 features is available, and comprises 2D textural, 2D color, and 3D shape measurements, collected with white light and narrow band imaging. In our study we first remove the features that are always 0, and then standardize the remaining ones as suggested by Gelman et al. (2008) and Chopin and Ridgway (2017), thus obtaining  $p - 1 = 929$  predictors with mean 0 and standard deviation 0.5. To assess predictive performance, we also hold out 15 randomly chosen units from the calculation of the posterior, roughly corresponding to 20% of the observations.

As discussed in Section 1, Bayesian inference for such a high-dimensional study is computationally unfeasible under state-of-the-art MCMC methods (Chopin and Ridgway, 2017),




Figure 2: Comparison between the estimates of key functionals obtained under the methods discussed in Sections 3.2.1 and 3.2.2 ( $y$ -axis), against those provided by the `rstan` implementation of the Hamiltonian no-u-turn sampler ( $x$ -axis). Red squares refer to Monte Carlo estimates based on i.i.d. samples from the exact SUN posterior produced by Algorithm 1, whereas blue and green squares denote the estimates provided by classical mean-field variational Bayes and by the proposed blocked partially-factorized approximation, respectively.

and hence it provides a useful setting for quantifying to what extent our results in Section 3 can cover this gap. To address such a goal, we first focus on the sequential discrete choice multinomial probit model in Section 2.3 with Gaussian priors, and compare the computational performance of the methods developed in Section 3.2 with the `rstan` implementation of the Hamiltonian no-u-turn sampler in Hoffman and Gelman (2014). The choice of the sequential model is directly motivated by the type of response of interest in our study. Indeed, it is plausible to first model benign ( $l = 1$ ) against malignant ( $l > 1$ ) status, and then focus on comparing the two sub-categories  $l = 2$  and  $l = 3$  of malignant lesions. Under this model, the vector  $\beta$  has dimension 1860, corresponding to the two class-specific 929-dimensional parameter vectors plus a class-specific intercept term. Consistent with the simulation study in Section 4, we place a  $N_{1860}(\mathbf{0}, \omega^2 \mathbf{I}_{1860})$  prior on  $\beta$ , with  $\omega^2 = 25$  (Gelman et al., 2008).

Figure 2 compares the Monte Carlo estimates for selected functionals of interest based on 5000 MCMC samples from the Hamiltonian no-u-turn sampler (R package `rstan`), against

those provided by the Monte Carlo and approximate methods discussed in Sections 3.2.1–3.2.2. In particular, we compute such functionals using both 5000 i.i.d. samples from the exact SUN posterior provided by Algorithm 1, and also by leveraging the strategies associated with the blocked partially-factorized variational approximation in Algorithm 2. In computing such an approximation under the sequential discrete choice multinomial probit model, we follow the guidelines in Remark 2 and group those augmented data corresponding to the same unit  $i$ . We shall emphasize that when the coefficients are not shared across labels and have independent priors, the overlap among rows of  $\bar{\mathbf{X}}$  referring to the same unit  $i$  is absent in sequential discrete choice representations. Hence, in this very specific case, we have that  $\min_{q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \in \mathcal{Q}_{\text{PMF-B}}} \text{KL}[q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \| p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})] = \min_{q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \in \mathcal{Q}_{\text{PFM}}} \text{KL}[q(\boldsymbol{\beta}, \bar{\mathbf{z}}) \| p(\boldsymbol{\beta}, \bar{\mathbf{z}} \mid \mathbf{y}, \mathbf{X})]$ . As we will discuss in the following, this blocking approach is more crucial for the multinomial probit models in Sections 2.1–2.2. To highlight the benefits of the blocked partially-factorized approximation, we also compare results with classical mean-field variational Bayes enforcing independence between  $\boldsymbol{\beta}$  and  $\bar{\mathbf{z}}$  (Consonni and Marin, 2007; Girolami and Rogers, 2006).

As highlighted in Figure 2, the two sampling-based methods provide comparable results in terms of inference and prediction. However, Algorithm 1 produces almost 75 samples of  $\boldsymbol{\beta}$  per second, whereas the Hamiltonian no-u-turn sampler can only draw one sample every 3 seconds. This massive computational cost makes state-of-the-art MCMC methods rapidly unfeasible in large  $p$  settings. We shall highlight that by relying on i.i.d. samples, Algorithm 1 has also the advantage of avoiding the need of burn-in periods and convergence checks. However, as discussed in Sections 3.2 and 4, Algorithm 1 scales poorly with sample size and, hence, it becomes impractical in studies with  $n$  larger than a few hundreds. This motivates the blocked partially-factorized approximation in Section 3.2.2, that notably matches almost perfectly the Monte Carlo estimates in such a high-dimensional setting (see Figure 2), and requires only 0.25 seconds to converge and 16 seconds to compute the different functionals. Classical mean-field variational Bayes has comparable running times, but the independence assumption between  $\boldsymbol{\beta}$  and  $\bar{\mathbf{z}}$  induces notable overshrinkage of both the locations and scales, which massively affects the estimation of the predictive probabilities. These results confirm and further clarify the findings in Section 4.

Before concluding our analysis, we also implement the multinomial probit model with class-specific parameters presented in Section 2.2, assuming independent standard normal errors. Due to the form of the dataset, the classical discrete multinomial probit in Section 2.1 is not appropriate, since it would require a vector of covariates for each combination of unit  $i$  and lesion  $l$ , which is not the case for this study. Nonetheless, according to the results in Sections 2.1, 2.2, and 3, models (1) and (3) induce posteriors with comparable dimensions and, hence, the performance of the multinomial probit with class-specific coefficients is also indicative of the one associated with the classical specification outlined in Section 2.1. Here, we focus on comparing the computational and predictive performance between the already-implemented sequential formulation in (5) and the one having class-specific coefficients in (3), considering the Monte Carlo and variational estimates discussed in Section 3.2. Under model (3), blocking across units  $i$  was more crucial to obtain accurate variational inference. The Hamiltonian no-u-turn sampler faced, instead, severe mixing and convergence issues under model (3), further highlighting major issues of MCMC in such settings.

Figure 3 compares variational and Monte Carlo estimates of the predictive probabilities for all the units, under the two models. To estimate the predictive probabilities we split the




Figure 3: Comparison between the predictive probabilities provided by model (3) ( $y$ -axis) and (5) ( $x$ -axis). The dataset has been divided into 6 folds, and for each fold the predictive probabilities are computed using all the other available data as training set. Red squares refer to Monte Carlo estimates based on i.i.d. samples from the SUN posterior produced by Algorithm 1, whereas the green squares denote the estimates provided by the proposed blocked partially-factorized approximation.

dataset in six folds, four having 13 observations and two having 12 observations. Then, we compute the predictive probabilities for the observations in each fold, using the units in the remaining five folds to obtain the posterior distribution. As clarified in Figure 3, the two models provide similar, but not identical, predictive probabilities, whose values are almost the same when comparing the Monte Carlo and variational estimates. This result confirms the excellent performance of the proposed blocked partially-factorized approximation in high-dimensional settings, especially when the dimension of  $\beta$  is higher than the sample size. Indeed, by slightly increasing the dimension of the training set, the number of  $\beta$  samples per second produced by Algorithm 1 rapidly decreases from 75 to 50 in model (5), whereas the variational strategy still requires about 0.25 seconds to converge and 16 seconds to compute the functionals. The overall out-of-sample predictive accuracy under the two models is about 66.5%. Considering the simplicity of the multinomial probit models implemented, these values are quite satisfactory when compared with the 73.68% accuracy obtained under sophisticated black-box machine learning algorithms (Mesejo et al., 2016).

## 6. Discussion

This article provides novel conjugacy results and computational methods for a general class of multinomial probit models (Hausman and Wise, 1978; Stern, 1992; Tutz, 1991) with Gaussian priors, and extends such properties to the entire class of SUN (Arellano-Valle and Azzalini, 2006) priors. As discussed in Sections 3, 4 and 5, the availability of a SUN posterior allows major advances in terms of closed-form, Monte Carlo and approximate variational inference which cover a still unaddressed gap of MCMC methods in high-dimensional studies. These settings are common in a variety of fields, such as in medical applications collecting a huge number of predictors via state-of-the-art imaging technologies.

Our results open also several avenues for future research. For example, although Bayesian estimation and inference for the covariance matrix  $\Sigma$  goes beyond the scope of this article, as mentioned in Sections 1 and 3.1, the availability of a closed-form expression for the marginal likelihood in Corollary 1 motivates promising advances in point estimation and full Bayesian inference also on  $\Sigma$ , which deserve further exploration. The results in this article can be also included in more complex formulations. For instance, the sequential probit in (5) has been used within Bayesian nonparametric models for density regression based on probit stick-breaking process (Rodriguez and Dunson, 2011). Our findings could be useful in such settings to improve the computational performance and the theoretical treatment of predictor-dependent Bayesian nonparametric mixture models. Also extensions of our results to classification via Gaussian processes (Rasmussen and Williams, 2006; Girolami and Rogers, 2006; Cao et al., 2022) and state-space models (Fasano et al., 2021) are straightforward. Finally, it would be also interesting to exploit the strategies in Genton et al. (2018), Cao et al. (2019) and Cao et al. (2021) to identify suitable blocks of augmented data in a more data-driven way, which can be applied to perform accurate variational inference not only in multinomial but also in binary probit regression. Similarly, exploring other strategies for sampling from multivariate truncated normals, such as the sequential Monte Carlo method in Moffa and Kuipers (2014), could further increase the impact of our findings.

## Acknowledgments

We are grateful to the Editor and the two anonymous referees for their precious comments which helped us in improving the preliminary version of this paper. Daniele Durante acknowledges the support from MIUR-PRIN 2017 project 20177BRJXS.

## Appendix A. Proofs

**Proof of Theorem 1.** To prove Theorem 1, it suffices to first apply the Bayes rule and then recognize a SUN density in the kernel of  $p(\beta | \mathbf{y}, \mathbf{X})$ . In particular, note that  $p(\beta | \mathbf{y}, \mathbf{X}) \propto p(\beta)p(\mathbf{y} | \beta, \mathbf{X}) \propto \phi_q(\beta - \xi; \Omega)\Phi_h(\gamma + \Delta^\top \bar{\Omega}^{-1} \omega^{-1}(\beta - \xi); \Gamma - \Delta^\top \bar{\Omega}^{-1} \Delta)\Phi_m(\bar{\mathbf{X}}\beta; \Lambda)$ , and re-write  $\Phi_m(\bar{\mathbf{X}}\beta; \Lambda)$  as  $\Phi_m(\mathbf{s}^{-1}\bar{\mathbf{X}}\xi + (\bar{\Omega}\omega\bar{\mathbf{X}}^\top\mathbf{s}^{-1})^\top \bar{\Omega}^{-1}\omega^{-1}(\beta - \xi); \mathbf{s}^{-1}(\bar{\mathbf{X}}\Omega\bar{\mathbf{X}}^\top + \Lambda)\mathbf{s}^{-1} - (\bar{\Omega}\omega\bar{\mathbf{X}}^\top\mathbf{s}^{-1})^\top \bar{\Omega}^{-1}\bar{\Omega}\omega\bar{\mathbf{X}}^\top\mathbf{s}^{-1})$ . Replacing this quantity inside the kernel of the posterior and recalling proof of Corollary 4 in Durante (2019), it follows that

$$\begin{aligned} & \Phi_h(\gamma + \Delta^\top \bar{\Omega}^{-1} \omega^{-1}(\beta - \xi); \Gamma - \Delta^\top \bar{\Omega}^{-1} \Delta)\Phi_m(\bar{\mathbf{X}}\beta; \Lambda) \\ &= \Phi_{h+m}(\gamma_{\text{pst}} + \Delta_{\text{pst}}^\top \bar{\Omega}_{\text{pst}}^{-1} \omega_{\text{pst}}^{-1}(\beta - \xi_{\text{pst}}); \Gamma_{\text{pst}} - \Delta_{\text{pst}}^\top \bar{\Omega}_{\text{pst}}^{-1} \Delta_{\text{pst}}), \end{aligned}$$

with  $\xi_{\text{pst}}, \Omega_{\text{pst}}, \Delta_{\text{pst}}, \gamma_{\text{pst}}$  and  $\Gamma_{\text{pst}}$  as in Theorem 1. Leveraging this equality and recalling that  $\xi_{\text{pst}} = \xi, \Omega_{\text{pst}} = \Omega$ , it can be immediately noticed that  $p(\beta)p(\mathbf{y} | \beta, \mathbf{X})$  coincides with the kernel of the SUN in (9), thereby proving Theorem 1. To prove that  $\Omega_{\text{pst}}^*$  is a correlation matrix it suffices to replace  $\mathbf{I}_n$  with  $\Lambda$  in the proof of Corollary 4 in Durante (2019). ■

**Proof of Corollary 1.** To show that  $p(\mathbf{y} | \mathbf{X})$  can be written as in (10), note that from the proof of Theorem 1,  $p(\mathbf{y}, \beta | \mathbf{X}) = p(\beta)\Phi_m(\bar{\mathbf{X}}\beta; \Lambda) = p(\beta | \mathbf{y}, \mathbf{X})\Phi_{h+m}(\gamma_{\text{pst}}; \Gamma_{\text{pst}})/\Phi_h(\gamma; \Gamma)$ . Hence,  $p(\mathbf{y} | \mathbf{X}) = p(\mathbf{y}, \beta | \mathbf{X})/p(\beta | \mathbf{y}, \mathbf{X}) = [p(\beta | \mathbf{y}, \mathbf{X})\Phi_{h+m}(\gamma_{\text{pst}}; \Gamma_{\text{pst}})/\Phi_h(\gamma; \Gamma)]/p(\beta | \mathbf{y}, \mathbf{X})$ , which implies  $p(\mathbf{y} | \mathbf{X}) = \Phi_{h+m}(\gamma_{\text{pst}}; \Gamma_{\text{pst}})/\Phi_h(\gamma; \Gamma)$ . ■

**Proof of Corollary 2.** To prove Corollary 2 simply notice that (11) is the ratio between the marginal likelihoods of the new expanded dataset and the original one (i.e., the one without the additional unit with response  $y_{\text{new}} = l$  and covariates  $\mathbf{x}_{\text{new}}$ ). Hence, the expression for the predictive probabilities follows from Corollary 1 after noting that, due to the conditional independence assumption in (1), (3) or (5),  $p(\mathbf{y} \mid \mathbf{X}, \mathbf{x}_{\text{new}}) = p(\mathbf{y} \mid \mathbf{X})$ . ■

## References

- R. Agarwal, P. Ranjan, and H. Chipman. A new Bayesian ensemble of trees approach for land cover classification of satellite imagery. *Canadian Journal of Remote Sensing*, 39: 507–520, 2014.
- A. Agresti. *Categorical Data Analysis (Third Edition)*. Wiley, 2013.
- J.H. Albert and S. Chib. Bayesian analysis of binary and polychotomous response data. *Journal of the American Statistical Association*, 88:669–679, 1993.
- J.H. Albert and S. Chib. Sequential ordinal modeling with applications to survival data. *Biometrics*, 57:829–836, 2001.
- R.B. Arellano-Valle and A. Azzalini. On the unification of families of skew-normal distributions. *Scandinavian Journal of Statistics*, 33:561–574, 2006.
- A. Azzalini. A class of distributions which includes the normal ones. *Scandinavian Journal of Statistics*, 12:171–178, 1985.
- A. Azzalini and A. Bacchieri. A prospective combination of phase II and phase III in drug development. *Metron*, 68:347–369, 2010.
- A. Azzalini and A. Capitanio. *The Skew-Normal and Related Families*. Cambridge University Press, 2013.
- A. Azzalini and A. Dalla Valle. The multivariate skew-normal distribution. *Biometrika*, 83: 715–726, 1996.
- C.M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- D.M. Blei, A. Kucukelbir, and J.D. McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112:859–877, 2017.
- A. Börsch-Supan and V.A. Hajivassiliou. Smooth unbiased multivariate probability simulators for maximum likelihood estimation of limited dependent variable models. *Journal of Econometrics*, 58:347–368, 1993.
- Z.I. Botev. The normal law under linear restrictions: Simulation and estimation via minimax tilting. *Journal of the Royal Statistical Society: Series B*, 79:125–148, 2017.
- L.F. Burgette and E.V. Nordheim. The trace restriction: An alternative identification strategy for the Bayesian multinomial probit model. *Journal of Business & Economic Statistics*, 30:404–410, 2012.

- J. Cao, M.G. Genton, D.E. Keyes, and G.M. Turkiyyah. Hierarchical–block conditioning approximations for high-dimensional multivariate normal probabilities. *Statistics and Computing*, 29:585–598, 2019.
- J. Cao, M.G. Genton, D.E. Keyes, and G.M. Turkiyyah. Exploiting low-rank covariance structures for computing high-dimensional normal and student-t probabilities. *Statistics and Computing*, 31:2, 2021.
- J. Cao, D. Durante, and M.G. Genton. Scalable computation of predictive probabilities in probit models with Gaussian process priors. *Journal of Computational and Graphical Statistics*, (forthcoming), 2022.
- J. C.-C. Chan and I. Jeliaskov. MCMC estimation of restricted covariance matrices. *Journal of Computational and Graphical Statistics*, 18:457–480, 2009.
- Z. Chen and L. Kuo. Discrete choice models based on the scale mixture of multivariate normal distributions. *Sankhyā, Series B*, 64:192–213, 2002.
- N. Chopin. Fast simulation of truncated Gaussian distributions. *Statistics and Computing*, 21:275–288, 2011.
- N. Chopin and J. Ridgway. Leave Pima indians alone: Binary regression as a benchmark for Bayesian computation. *Statistical Science*, 32:64–87, 2017.
- G. Consonni and J.M. Marin. Mean-field variational approximate Bayesian inference for latent variable models. *Computational Statistics & Data Analysis*, 52:790–798, 2007.
- C. Daganzo. *Multinomial Probit*. Academic Press, 1979.
- J.K. Dow and J.W. Endersby. Multinomial probit and multinomial logit: a comparison of choice models for voting research. *Electoral Studies*, 23:107–122, 2004.
- D. Durante. Conjugate Bayes for probit regression via unified skew-normal distributions. *Biometrika*, 106:765–779, 2019.
- A. Fasano, G. Rebaudo, D. Durante, and S. Petrone. A closed-form filter for binary time series. *Statistics and Computing*, 31:47, 2021.
- A. Fasano, D. Durante, and G. Zanella. Scalable and accurate variational Bayes for high-dimensional binary regression models. *Biometrika*, (forthcoming), 2022.
- A. Gelman, A. Jakulin, M.G. Pittau, and Y.S. Su. A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics*, 2: 1360–1383, 2008.
- M.G. Genton, D.E. Keyes, and G.M. Turkiyyah. Hierarchical decompositions for the computation of high-dimensional multivariate normal probabilities. *Journal of Computational and Graphical Statistics*, 27:268–277, 2018.
- A. Genz. Numerical computation of multivariate normal probabilities. *Journal of Computational and Graphical Statistics*, 1:141–149, 1992.



- J. Geweke, M. Keane, and D. Runkle. Alternative computational approaches to inference in the multinomial probit model. *The Review of Economics and Statistics*, 76:609–632, 1994.
- M. Girolami and S. Rogers. Variational Bayesian multinomial probit regression with Gaussian process priors. *Neural Computation*, 18:1790–1817, 2006.
- M. Girolami and M. Zhong. Data integration for classification problems employing Gaussian process priors. In *Advances in Neural Information Processing Systems*, volume 20, pages 465–472, 2007.
- W.H. Greene. *Econometric Analysis*. Prentice Hall, 2003.
- A.K. Gupta, M.A. Aziz, and W. Ning. On some properties of the unified skew-normal distribution. *Journal of Statistical Theory and Practice*, 7:480–495, 2013.
- J.A. Hausman and D.A. Wise. A conditional probit model for qualitative choice: Discrete decisions recognizing interdependence and heterogeneous preferences. *Econometrica: Journal of the Econometric Society*, 46:403–426, 1978.
- M.D. Hoffman and A. Gelman. The No-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15:1593–1623, 2014.
- C.C. Holmes and L. Held. Bayesian auxiliary variable models for binary and multinomial regression. *Bayesian Analysis*, 1:145–168, 2006.
- W.C. Horrace. Some results on the multivariate truncated normal distribution. *Journal of Multivariate Analysis*, 94:209–221, 2005.
- K. Imai and D.A. Van Dyk. A Bayesian analysis of the multinomial probit model using marginal data augmentation. *Journal of Econometrics*, 124:311–334, 2005.
- J.E. Johndrow, D.B. Dunson, and K. Lum. Diagonal orthant multinomial probit models. In *Artificial Intelligence and Statistics*, volume 31, pages 29–38, 2013.
- J.E. Johndrow, A. Smith, N. Pillai, and D.B. Dunson. MCMC for imbalanced categorical data. *Journal of the American Statistical Association*, 114:1394–1403, 2019.
- B.P. Kindo, H. Wang, and E.A. Peña. Multinomial probit Bayesian additive regression trees. *Stat*, 5:119–131, 2016.
- D.A. Knowles and T. Minka. Non-conjugate variational message passing for multinomial and binary regression. In *Advances in Neural Information Processing Systems*, volume 24, pages 1701–1709, 2011.
- S. Kullback and R.A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22:79–86, 1951.
- G.S. Maddala. *Limited-Dependent and Qualitative Variables in Econometrics*. Cambridge University Press, 1986.

- R.E. McCulloch and P.E. Rossi. An exact likelihood analysis of the multinomial probit model. *Journal of Econometrics*, 64:207–240, 1994.
- R.E. McCulloch, N.G. Polson, and P.E. Rossi. A Bayesian analysis of the multinomial probit model with fully identified parameters. *Journal of Econometrics*, 99:173–193, 2000.
- D. McFadden. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica: Journal of the Econometric Society*, 57:995–1026, 1989.
- P. Mesejo, D. Pizarro, A. Abergel, O. Rouquette, S. Beorchia, L. Poincloux, and A. Bartoli. Computer-aided classification of gastrointestinal lesions in regular colonoscopy. *IEEE Transactions on Medical Imaging*, 35:2051–2063, 2016.
- G. Moffa and J. Kuipers. Sequential Monte Carlo EM for multivariate probit models. *Computational Statistics & Data Analysis*, 72:252–272, 2014.
- R. Natarajan, C.E. McCulloch, and N.M. Kiefer. A Monte Carlo EM method for estimating multinomial probit models. *Computational Statistics & Data Analysis*, 34:33–50, 2000.
- A. Nobile. A hybrid Markov chain for the Bayesian analysis of the multinomial probit model. *Statistics and Computing*, 8:229–242, 1998.
- T. Park and D.A. Van Dyk. Partially collapsed Gibbs samplers: Illustrations and applications. *Journal of Computational and Graphical Statistics*, 18:283–305, 2009.
- C.E. Rasmussen and C.K. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- J. Riihimäki, P. Jylänki, and A. Vehtari. Nested expectation propagation for Gaussian process classification with a multinomial probit likelihood. *Journal of Machine Learning Research*, 14:75–109, 2013.
- A. Rodriguez and D.B. Dunson. Nonparametric Bayesian models through probit stick-breaking processes. *Bayesian Analysis*, 6:145–178, 2011.
- S. Rogers and M. Girolami. Multi-class semi-supervised learning with the  $\epsilon$ -truncated multinomial probit Gaussian process. In *Journal of Machine Learning Research, Workshop & Proceedings*, volume 1, pages 17–32, 2007.
- S. Stern. A method for smoothing simulated moments of discrete probabilities in multinomial probit models. *Econometrica: Journal of the Econometric Society*, 60:943–952, 1992.
- G. Tutz. Sequential models in categorical regression. *Computational Statistics & Data Analysis*, 11:275–295, 1991.
- X. Zhang, W.J. Boscardin, and T.R. Belin. Sampling correlation matrices in Bayesian models with correlated latent variables. *Journal of Computational and Graphical Statistics*, 15:880–896, 2006.