

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Pairwise likelihood inference for multivariate ordinal responses with applications to customer satisfaction

This is a pre print version of the following article:

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1564327> since 2016-06-04T12:02:39Z

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Pairwise likelihood inference for multivariate ordinal responses with applications to customer satisfaction

Euloge Clovis Kenne Pagui¹ and Antonio Canale^{2 3}

¹ Department of Statistical Sciences, University of Padua, Padua, Italy

² Department of Economics and Statistics, University of Turin, Turin, Italy

³ Collegio Carlo Alberto, Moncalieri (TO), Italy

Address for correspondence: Antonio Canale
Dept. of Economics and Statistics
University of Turin
Corso Unione Sovietica 220 – 10100 Torino (TO), Italy

E-mail: antonio.canale@unito.it.

Phone: (+39) 011 670 5724.

Fax: (+39) 011 670 5783.

Abstract: A common practice in customer satisfaction analysis is to administer surveys where subjects are asked to express opinions on a number of statements, or satisfaction scales, by use of ordered categorical responses. Motivated by this application, we propose a pseudo-likelihood approach to estimate the dependence structure among multivariate categorical variables. As it is commonly done in this area, we assume that the responses are related to latent continuous variables which are truncated to induce categorical responses. A Gaussian likelihood is assumed for the latent variables leading to the so called ordered probit model. Since the calculation of the exact likelihood is computationally demanding, we adopt an approximate solution based on pairwise likelihood. To assess the performance of the approach, simulation studies are conducted comparing the proposed method with standard likelihood methods. A parametric bootstrap approach to evaluate the variance of the maximum pairwise likelihood estimator is proposed and discussed. An application to customer satisfaction survey is performed showing the effectiveness of the approach in the presence of covariates and under other generalisations of the model.

Key words: Composite likelihood; Latent continuous variables; Likert scales; Mutual information index; Polychoric correlation

1 Introduction

Surveys and questionnaires are a fundamental research tool in many application areas such as psychometry or marketing research among others. Responses in customer satisfaction

surveys, for example, are generally in the form of rating scale or Likert scale (Likert, 1932). When responding to a Likert questionnaire, customers specify their level of agreement or disagreement for each of the q questions or items. Usually each item consists of $K = 5$ ordered response levels, but sometimes also $K > 5$ can be adopted. Each item aims to measure one different aspect of the overall phenomenon under study (e.g. global customer satisfaction) and hence the responses are clearly correlated.

A single item, in Likert questionnaires, can be considered as an ordered-categorical outcome. An approach to model ordinal categorical responses which, albeit wrong, is commonly adopted, consists in treating them as continuous observations and to use standard methods such as linear regression to estimate covariate dependence, t -test or ANOVA models to test for differences among groups, and so forth. A more rigorous and elegant approach consists in considering that the observed categorical data are related to continuous latent variables. The relation between the latent continuous variables and the observed categorical variables is usually induced by use of thresholds partitioning the latent sample space into a series of regions corresponding to each ordinal category. If a Gaussian distribution is assumed for the latent variables, an ordered probit model is induced (Agresti, 2002). Ordered probit model has been extensively used under the Bayesian paradigm thanks to the celebrated strategy of Albert and Chib (1993) with applications ranging from political science and market research (see for exaple Rossi et al., 2012; Jackman, 2009).

Multivariate ordered probit models are the simple multivariate extension of ordered probit models where both the latent and the observed variables have dimension $q \geq 2$, as it is the case for customer satisfaction surveys. With regard to this class of models, we discuss the problem of the estimation of the dependence structure. The multivariate ordered probit model has clear computational problems related to the calculation of a q dimensional integral for each single likelihood contribution. Such a procedure is computationally demanding or even unfeasible also for moderate q . For that reason we adopt an approximate solution based on a simple pseudolikelihood belonging to the class of composite likelihoods (Lindsay, 1988): the pairwise likelihood (PL) (Cox and Reid, 2004). The use of the PL when the full likelihood is computationally unmanageable is becoming a commonly adopted procedure (Varin et al., 2011). Applications of PL to ordered probit models are not new, and significant contributions can be found in De Leon (2005), where the dependence among the ordered variables is evaluated in terms of polychoric correlation, i.e. the linear correlation of the latent variables, and in Varin and Vidoni (2006) and Varin and Czado (2010) in the context of longitudinal data and focusing on the mean evolution in time. Our contribution to the PL approach in the case of ordered probit model builds on De Leon (2005), but it differs for three main aspects. First, we estimate different measures of dependence and not only the polychoric correlation, i.e. the linear correlation coefficient and the L measure coefficient (Lu, 2011). The polychoric correlation is indeed based on the latent Gaussianity assumption which we consider more as a computational solution rather than an exact characteristic of the model. Second, the dimensions that we consider mimic real contexts and are greater than $q = 3$. This has clearly no theoretical difference, but it increases the computational burden. In addition, this causes the number of model parameters to expand rapidly, leading to the practical impossibility of evaluate the Fisher's and Godambe's information matrices

and forcing us to study other approaches to evaluate the variability of the estimators. Third, we apply the methods to a concrete customer satisfaction application, introducing subject-specific covariates and other complications. In the application, we discuss a bootstrap-based approach which allows us to perform classical hypothesis testing on each pairwise variable dependence and on the significance of the regression coefficients. This bootstrap-based approach is, to our knowledge, the first attempt to assess the variance of the maximum likelihood estimator for this class of problems.

The rest of the paper is organised as follows. In the next section we review the multivariate ordered probit model and introduce our solution based on pairwise likelihood. In Section 3 an extensive simulation experiment is conducted comparing the full and pairwise likelihoods both in terms of point estimates and computational time. In Section 4 the model is generalised including the effect of subject specific covariates, motivated by applications to customer satisfaction surveys. Data from a survey conducted by a company operating in the sector of information technology services and consulting are then analysed and discussed, introducing a bootstrap-based solution to assess the estimator standard errors. Section 5 summarises the main results.

2 Multivariate ordered probit model

Let $Y_i = (Y_{i1}, \dots, Y_{iq})^T$, $Y_{ij} \in \{1, 2, \dots, K\}$ for $j = 1, \dots, q$ be a q -dimensional ordinal categorical random vector with joint distribution $g(Y_i)$ depending on some unknown parameter θ , with $i = 1, \dots, n$ defining a collection of iid random vectors. The ordered probit model assumes that, for each $i = 1, \dots, n$, there exists a latent random vector $Z_i = (Z_{i1}, \dots, Z_{iq})^T$, with $Z_i \stackrel{iid}{\sim} N(0, \Sigma)$ where

$$\Sigma = \begin{pmatrix} 1 & \rho_{1,2} & \dots & \dots & \rho_{1,q} \\ & 1 & \dots & \dots & \rho_{2,q} \\ & & 1 & \rho_{rs} & \vdots \\ & & & 1 & \rho_{q-1,q} \\ & & & & 1 \end{pmatrix}$$

is an unknown q -dimensional positive definite correlation matrix. In observing y_{ij} , the following relation is assumed

$$y_{ij} = k \text{ if and only if } z_{ij} \in (a_{k-1}, a_k],$$

where $-\infty = a_0 < a_1 < \dots < a_K = \infty$ is a sequence of real numbers defining a disjoint partition of \mathbb{R} . The likelihood function for a single observation is then proportional to

$$\text{pr}(Y_{i1} = y_{i1}, \dots, Y_{iq} = y_{iq}) = \int_{a_{y_{i1}-1}}^{a_{y_{i1}}} \dots \int_{a_{y_{iq}-1}}^{a_{y_{iq}}} \phi_{\Sigma}(z_{i1}, \dots, z_{iq}) dz_{i1} \dots dz_{iq},$$

where $\phi_{\Sigma}(\cdot)$ denotes the multivariate Gaussian distribution with mean zero and variance Σ . The full log-likelihood for the n observations is therefore given by

$$\ell(\theta) = \sum_{i=1}^n \log \left\{ \int_{a_{y_{i1}-1}}^{a_{y_{i1}}} \dots \int_{a_{y_{iq}-1}}^{a_{y_{iq}}} \phi_{\Sigma}(z_{i1}, \dots, z_{iq}) dz_{i1} \dots dz_{iq} \right\}, \quad (2.1)$$

where the parameters of the model are the thresholds (a_1, \dots, a_{K-1}) and the latent correlations $\{\rho_{r,s}\}$ for $r, s = 1, \dots, q$, and $r < s$, and $\theta = (a_1, \dots, a_{K-1}, \rho_{1,2}, \dots, \rho_{q-1,q})$ is the joint vector of all parameters with $\hat{\theta}$ denoting its maximum likelihood estimator. Equation (2.1) involves a q -dimensional integral for each observation. Indeed, the likelihood is not easily tractable, since its direct evaluation requires the computation of n integrals of a q -dimensional Gaussian distribution.

2.1 Pairwise likelihood for multivariate probit model

The pairwise likelihood that we discuss in this paper, is a pseudolikelihood constructed from bivariate margins. Specifically, the pairwise log-likelihood for our ordered categorical data is

$$\begin{aligned} \ell^P(\theta) &= \sum_{i=1}^n \sum_{r=1}^{q-1} \sum_{s=r+1}^q \log \{ \text{pr}(Y_{ir} = y_{ir}, Y_{is} = y_{is}) \} \\ &= \sum_{i=1}^n \sum_{r=1}^{q-1} \sum_{s=r+1}^q \log \left\{ \int_{a_{y_{ir}-1}}^{a_{y_{ir}}} \int_{a_{y_{is}-1}}^{a_{y_{is}}} \phi_{\Sigma(\rho_{r,s})}(z_{ir}, z_{is}) dz_{ir} dz_{is} \right\}, \end{aligned}$$

where $\Sigma(\rho)$ denotes the 2×2 correlation matrix with off-diagonal entries equal to ρ . The PL approach is particularly attractive because it substitutes a log-likelihood involving high dimensional integrations with a sum of bivariate integrals, which can be easily evaluated with standard available software. Similar PL has been considered by [De Leon \(2005\)](#) for likelihood inference within group continuous models, that is models for multivariate ordinal data. While the latent model is the same, we are not estimating the polychoric correlations only (i.e. the linear correlation of the latent variables) but also directly a measure of relation among the ordered categorical responses. To this end, we use both the standard correlation coefficient and the L measure defined as

$$\text{Cor}(U, V) = \frac{\text{Cor}(U, V)}{\sqrt{\text{Var}(U)\text{Var}(V)}} \quad (2.2)$$

$$L(U, V) = \left[1 - \exp \left\{ \frac{-2I(U, V)}{1 - \frac{I(U, V)}{\min\{H(U)H(V)\}}} \right\} \right]^{1/2}, \quad (2.3)$$

where U and V are two discrete random variables, $I(U, V) = \sum_{u,v} \text{pr}(u, v) \log\{\text{pr}(u, v)/\text{pr}(u)\text{pr}(v)\}$ is the mutual information between U and V and $H(U) = -\sum_u \text{pr}(u) \log(\text{pr}(u))$ is the entropy of U . The L measure is a positive and symmetric index of dependence between any possible random variable. It equals 0 if and only if the random variables are independent

and 1 if there is a strict dependence between the random variables. The L measure is defined using the mutual information index $I(\cdot, \cdot)$, it is invariant under marginal one-to-one transformations of the random variables, and it equals the correlation coefficient if the two random variables are normally distributed. Note also that definition (2.3) exploits the fact that, in our case, each random variable is discrete. For further details see Lu (2011).

Generally, composite likelihood inferential procedures are similar to the ones based on the standard likelihood. For example, the pairwise score vector, $U^P(\theta) = \partial \ell^P(\theta) / \partial \theta$, is still unbiased being the sum of the score vector based on the likelihood contribution of each pair of observations. As for the standard likelihood, the maximum pairwise likelihood estimator $\hat{\theta}^P$ can be obtained either by maximizing $\ell^P(\theta)$ or by solving the pairwise score equation $U^P(\theta) = 0$. Under the usual regularity conditions (Molenberghs and Verbeke, 2005), the maximum pairwise likelihood estimator is consistent and asymptotically normal as $n \rightarrow \infty$. Specifically

$$\hat{\theta}^P \sim N(\theta, G(\theta)^{-1}),$$

where $G(\theta) = W(\theta)J(\theta)^{-1}W(\theta)$ is known as the Godambe information or sandwich information, $W(\theta) = E_{\theta}\{-\partial U^P(\theta) / \partial \theta\}$ is the sensitivity matrix and $J(\theta) = \text{Var}_{\theta}\{U^P(\theta)\}$ is the variability matrix.

The use of a pseudo-likelihood usually leads to some loss of efficiency (see, for example, Cox and Reid, 2004, Section 1; Zhao and Joe, 2005, Section 6 and Xu and Reid, 2011, Section 3 for some illustrations), but in our motivating examples this loss of efficiency is necessary if we want to be able to actually compute any likelihood-based point estimate. Moreover, for the underlying model, the simulations done in Section 3.2 and in the Supporting Information Online provide empirical evidence of a moderate loss of efficiency. Note that, however, there are few classes of models in which the maximum pairwise likelihood estimator is fully efficient (Mardia et al., 2009; Kenne Pagui et al., 2015).

With regard to the present model, both the standard and pairwise likelihood involve a numerical integration and thus the evaluation of the Fisher information, $W(\theta)$, and $J(\theta)$ may be a daunting task, particularly in the applications we have in mind which involve high dimensions q . To this end, in Section 4 we adopt a bootstrap-based approach.

3 Computation and simulation

To assess the performance of the procedure based on the PL we conduct a simulation study. Before describing and commenting the empirical results, we briefly discuss the computational issues related to the proposed model stressing when it is preferable to the full-likelihood approach.

3.1 Computational issues

Our code is implemented in the `PLordprod` package which can be downloaded from CRAN ([Kenne Pagui and Canale, 2014](#)). The core function is implemented in C. To evaluate the Gaussian integrals, the package uses the Fortran 77 subroutine `SADMVN` by [Genz \(1992\)](#). As discussed in [Varin and Vidoni \(2006\)](#), to assure the ordering of the thresholds $a_0, < \dots < a_K$, we reparametrise the model defining

$$\delta_k = \log(a_k - a_{k-1})$$

for $k = 2, \dots, K - 1$. Clearly, it is not possible to identify both the mean of the latent variables and the first threshold a_1 . To deal with this, we fix the latent mean to be zero. For optimisation, we consider as starting values for the thresholds $a_1 = 0$ and $\delta_k = 0$ for $k = 2, \dots, K - 1$. This is equivalent to assume that the thresholds are equally spaced with distance one. As for the covariance components, we consider as starting values the sample covariances of the observed categorical variables treated as continuous. Optimisation of the PL and of the full-likelihood functions is performed via quasi-Newton box-constrained optimisation algorithm, with a relative tolerance of 1×10^{-10} .

It is important to note that despite the admissible range for a correlation is $[-1, 1]$, we have here the additional constraint of the full covariance matrix Σ being a positive definite matrix. By breaking up the full likelihood into pairwise contributions, and keeping $[-1, 1]$ as the range for the quasi-Newton box-constrained optimisation, it may happen that a given configuration of correlations leads to a full correlation matrix Σ which is not positive definite. In order to avoid this, our numerical routines implement a check on the final eigenvalues of Σ and return a likelihood value equal to zero if one of these eigenvalue is not positive.

A comparison between the PL and the full-likelihood is reported in the next section. However such a comparison is feasible only for moderate q . Indeed our motivation to use the PL approach lies both in speeding up computations, and in avoiding the numerical instability intrinsically related to the calculation of q -dimensional integrals. [Figure 1](#) gives an idea of the former aspect. The left panel shows the executions time of a single evaluations of the likelihood functions for samples of size $n = 50$ and different dimensions $q = 2, 5, 10, 15, 20$. The computational burden of the full likelihood with respect to the PL is clearly comparable for $q = 2$ while it is double for $q = 5$ and even ten times bigger for $q > 5$. Note that in maximising the likelihood function numerically, several evaluation of the likelihood functions are needed leading to an increase of the differences between the two approaches. In particular, by our experience, we have noted that for $q > 5$ the maximisation of the full likelihood turns out to be extremely slow and presents enormous computational difficulties. An idea of the latter aspect is illustrated on the right panel of [Figure 1](#), where the average time of a complete maximisation is reported. Note that the differences are remarkable also for $q = 6$. The times are in seconds and calculated when running the algorithms in R version 3.1.1 on a 64bit Linux machine with a 2.70 GHz Intel quadri-core i7 processor with 8 Gb of RAM.




Figure 1: Execution time of a single evaluation (left panel) and of the whole maximisation procedure (right panel) of the full log-likelihood (continuous lines) and pairwise log-likelihood (dashed lines) for different dimensions of q . Sample size $n = 50$ and $K = 5$ for both situations.

3.2 Simulation experiment

To assess the performance of the proposed approach in terms of estimation we simulate data from different scenarios. Three scenarios consider $q = 5$ and $K = 5$ while a single scenario consider $q = 3$ and $K = 5$. As previously discussed the computations for the full likelihood when $q > 5$ are demanding and hence we decide to fix $q = 5$ as upper bound for sake of comparison. However, the PL proposal is feasible also for higher q as we show in the application of Section 4 and in the Supporting Information Online, where we perform a simulation with $q = 15$ for the PL only. This dimension is comparable to that of the application discussed in Section 4. In the first scenario with $q = 5$ we first simulate latent Z_i from a multivariate normal with zero mean vector and a full covariance matrix with different correlations and then round Z_i with thresholds $-\infty, -1, 0, 1, 2, \infty$. The only scenario with $q = 3$ also assume full covariance matrix. In the second scenario with $q = 5$ we generate latent Z_i from a constrained covariance matrix with all correlations equal to ρ while in the last one we use a sparse matrix with four zeros and five different non-zero entries. For each scenario we simulate 1,000 samples of sizes $n = 50, 100, 200$. From the estimated values of parameters, we derive the polychoric correlation coefficients and, with a plug-in approach, the pairwise correlation and L measure coefficients. The Monte Carlo mean bias and variance of the estimators for the polychoric correlations are reported in Table 1 and in the Supporting Information Online.

As expected, in all scenarios and considering all the three measures of dependence, both bias and variance decrease as the sample size increases, both in the full and in the pairwise

Table 1: Bias and variances of the polychoric correlation coefficients under the three scenarios. For each n , the lower and upper triangle refer to the bias and the variance, respectively.

Scenario	n	full likelihood					pairwise likelihood				
S1	50	—	0.0192	0.0191	0.0208	0.0166	—	0.0250	0.0243	0.0218	0.0183
		0.0177	—	0.0009	0.0221	0.0140	-0.0041	—	0.0008	0.0255	0.0152
		0.0070	0.0197	—	0.0207	0.0089	-0.0140	-0.0014	—	0.0243	0.0097
	100	-0.0017	0.0104	-0.0146	—	0.0100	-0.0013	-0.0025	-0.0050	—	0.0122
		0.0490	0.0422	-0.0026	-0.0078	—	0.0026	-0.0054	0.0012	-0.0071	—
		—	0.0086	0.0025	0.0107	0.0058	—	0.0121	0.0028	0.0116	0.0045
		0.0144	—	0.0046	0.0123	0.0091	0.0101	—	0.0058	0.0127	0.0130
		-0.0003	0.0228	—	0.0110	0.0026	0.0028	-0.0012	—	0.0110	0.0023
		0.0086	0.0271	0.0030	—	0.0108	0.0002	-0.0012	0.0047	—	0.0109
	200	0.0217	-0.0193	-0.0095	0.0263	—	0.0030	-0.0002	-0.0002	0.0041	—
		—	0.0027	0.0025	0.0053	0.0031	—	0.0060	0.0052	0.0060	0.0060
		0.0004	—	0.0001	0.0044	0.0046	0.0029	—	0.0008	0.0023	0.0059
		0.0018	0.0019	—	0.0049	0.0049	-0.0024	-0.0042	—	0.0059	0.0054
		-0.0005	0.0024	-0.0009	—	0.0015	-0.0007	-0.0046	0.0031	—	0.0063
		0.0079	0.0041	0.0002	-0.0025	—	-0.0048	-0.0069	-0.0044	-0.0009	—
S2	50	—	0.0052	0.0044	0.0047	0.0043	—	0.0046	0.0049	0.0047	0.0047
		-0.0088	—	0.0050	0.0050	0.0049	-0.0076	—	0.0044	0.0049	0.0048
		-0.0049	-0.0044	—	0.0044	0.0038	-0.0097	-0.0051	—	0.0050	0.0053
	100	-0.0060	-0.0051	-0.0051	—	0.0047	-0.0079	-0.0095	-0.0103	—	0.0050
		-0.0099	-0.0073	-0.0016	-0.0068	—	-0.0067	-0.0093	-0.0107	-0.0053	—
		—	0.0020	0.0020	0.0019	0.0020	—	0.0021	0.0022	0.0020	0.0021
		0.0014	—	0.0021	0.0022	0.0022	0.0005	—	0.0023	0.0023	0.0023
		0.0007	-0.0025	—	0.0022	0.0020	-0.0003	-0.0036	—	0.0024	0.0021
		-0.0020	-0.0015	-0.0010	—	0.0020	-0.0034	-0.0030	-0.0020	—	0.0021
	200	0.0014	-0.0009	-0.0045	0.0001	—	0.0008	-0.0017	-0.0054	-0.0006	—
		—	0.0010	0.0011	0.0011	0.0010	—	0.0010	0.0011	0.0011	0.0011
		-0.0005	—	0.0011	0.0011	0.0010	-0.0024	—	0.0011	0.0011	0.0010
		0.0002	-0.0004	—	0.0011	0.0010	-0.0010	-0.0022	—	0.0012	0.0010
		-0.0003	0.0001	-0.0017	—	0.0011	-0.0016	-0.0010	-0.0031	—	0.0012
		0.0005	-0.0023	-0.0010	-0.0015	—	-0.0009	-0.0036	-0.0018	-0.0029	—
S3	50	—	0.0255	0.0095	0.0199	0.0251	—	0.0247	0.0090	0.0211	0.0255
		-0.0042	—	0.0261	0.0232	0.0089	-0.0033	—	0.0254	0.0248	0.0079
		-0.0038	0.0133	—	0.0208	0.0270	-0.0018	0.0110	—	0.0222	0.0268
	100	-0.0064	-0.0068	0.0028	—	0.0171	-0.0080	-0.0057	-0.0031	—	0.0182
		0.0101	0.0101	0.0010	-0.0127	—	0.0095	0.0031	-0.0031	-0.0180	—
		—	0.0121	0.0084	0.0114	0.0106	—	0.0121	0.0083	0.0117	0.0108
		0.0124	—	0.0121	0.0111	0.0079	0.0119	—	0.0123	0.0126	0.0084
		0.0064	-0.0010	—	0.0117	0.0114	0.0037	-0.0020	—	0.0117	0.0124
		-0.0070	-0.0043	-0.0030	—	0.0022	-0.0079	-0.0058	-0.0024	—	0.0020
	200	0.0005	-0.0086	0.0053	0.0022	—	0.0012	-0.0096	0.0081	0.0022	—
		—	0.0061	0.0048	0.0062	0.0057	—	0.0061	0.0048	0.0062	0.0057
		0.0004	—	0.0062	0.0069	0.0058	0.0003	—	0.0062	0.0069	0.0057
		-0.0061	-0.0015	—	0.0049	0.0064	-0.0059	-0.0015	—	0.0049	0.0063
		-0.0056	0.0015	-0.0029	—	0.0060	-0.0058	0.0018	-0.0032	—	0.0060
		0.0015	-0.0094	0.0009	-0.0039	—	0.0022	-0.0095	0.0005	-0.0038	—

likelihood approach. In the second scenario, namely when we have all the polychoric correlations fixed to ρ the variances between the full likelihood and the PL approach are almost identical for $n = 200$ while it is evident that the PL has larger finite sample bias. In the last scenario the differences in terms of bias are smaller than in the first two.

Comparing the results of our approximate solution for the first scenario with $q = 3$ (reported in the Supporting Information Online) and $q = 5$ we note, as expected, a loss of efficiency in estimating the parameters as q increases.

4 Application to customer satisfaction analysis

Customer satisfaction surveys are useful tools to measure customers' opinions on products and services. For example, a typical surveys structure to measure the quality of services, follows the so called SERVQUAL structure (Parasuraman et al., 1988). The main idea behind this approach is that, in order to guarantee a good quality of service, it is necessary to go beyond a customer's expectations. Thus, it is important to measure the gap between expected services and experienced services. To this end, questionnaires are typically divided into two separate blocks of questions. In the first, customers are asked to say how important are some characteristics of a given service in general. In the second, customers are asked to report their actual satisfaction on the same characteristics, for the experienced service.

We focus on data from a questionnaires collected in 2009 by an Italian IT company producing and selling software and offering consulting service. The survey was carried out by an independent marketing research company specialized in such surveys. Similar data have been presented in Azzalini and Scarpa (2012). We focus on the responses of $n = 324$ customers to $q = 14$ questions, divided into two blocks of seven items each. In the first block, the question was "how important are the following aspects", while in the second block the question was "how satisfied are you regarding the following aspects". The seven aspects are: efficiency of the service, reliability of the service, flexibility of the service, velocity in giving solutions, ability to satisfy the customers' needs, ability to implement technological innovation, velocity in adapting to changing legislation. Each item is measured on a scale of $K = 10$ levels. In addition to the responses above, subject specific covariates are available, namely gender, age, and educational level (elementary school, high school, and university) of the customer. Let x_i the vector of the covariates for customer i .

From an application point of view, we are mainly interested in the mean responses, accounting for meaningful dependence structure between and within the two blocks. Specifically, we want to allow for different means for each different items in order to measure the discrepancy between experienced and expected services, while controlling for any effect due to the covariates. The company hopes to record similar means in the two block and possibly higher scores in the experienced quality block and positive correlation between an item of the first block and the relative item in the second block.

To this end we modify the model of Section 2, assuming

$$\mathbf{Z}_i = (Z_{i1}, \dots, Z_{iq})^T, \quad E[Z_{ij}] = \xi_j + x_i^T \beta, \quad \text{Var}[\mathbf{Z}_i] = \Sigma \quad (4.1)$$

for all $i = 1, \dots, n$ and $j = 1, \dots, q$, with Σ being a q -dimensional positive definite correlation matrix to be estimated. Let furthermore $\rho = (\rho_{1,2}, \dots, \rho_{q-1,q})^T$, the $q(q-1)/2$ vector containing all the upper diagonal elements of Σ and $\xi = (\xi_1, \dots, \xi_q)^T$, the vector of the item-specific means. We assume to observe $y_{ij} = k$ if $z_{ij} \in (a_{k-1}, a_k]$ and let $a = (a_2, \dots, a_{K-1})^T$, be the vector of $K-2$ thresholds. Note that it is not possible to identify both the mean of the latent variables and all the thresholds. Thus a_1 is fixed to zero. We fix the first thresholds rather than the means because the generalization (4.1) has a different mean for each item and subject. For the same reason, we cannot compute the actual correlations and L -measures, and hence we focus on the polychoric correlation estimation only.

The parameter vector of model (4.1) is thus made of $K-2$ thresholds, $q(q-1)/2$ polychoric correlations, q item-specific intercept levels, and p regression coefficients. The PL log-likelihood for the parameters is

$$\ell^P(\rho, \xi, \beta, a) = \sum_{i=1}^n \sum_{r=1}^{q-1} \sum_{s=r+1}^q \log \{ \text{pr}(Y_{ir} = y_{ir}, Y_{is} = y_{is}) \}$$

where

$$\text{pr}(Y_{ir} = y_{ir}, Y_{is} = y_{is}) = \int_{a_{y_{ir}-1}}^{a_{y_{ir}}} \int_{a_{y_{is}-1}}^{a_{y_{is}}} \phi_{\Sigma(\rho_{rs})}(z_{ir} - \mu_{ir}, z_{is} - \mu_{is}) dz_{ir} dz_{is},$$

and $\mu_{ik} = \xi_k + x_i^T \beta$. Clearly, the high number of parameters (117, for our dataset) makes the likelihood optimization a demanding task. By our experience, the quasi-Newton box-constrained optimization algorithm, greatly benefits from the following initialization. As before, we consider as starting values for the polychoric correlations, the sample covariances of the observed categorical variables. Then, for the β coefficients we first compute the average scores over the items for each subject, then we standardise them and after that, we perform a linear regression between the standardised-average-scores and the covariates. The least squares estimates for the regression coefficients act as starting values for the β parameters. For ξ , we simply consider the sample mean for each item over the subjects. For what concerns the initialization of the thresholds we first compute the sample mean (\bar{y}), standard deviation (s_y), and relative frequencies of the categories ($\hat{f}_1, \dots, \hat{f}_{10}$) of the whole sample. Then we calculate the quantiles of levels \hat{f}_j for $j = 2, \dots, 10$ of a Gaussian distribution with mean \bar{y} and standard deviation s_y and consider those values as starting values for the thresholds.

In this applied problem, we are facing a very complex model both from the analytical and computational viewpoints. Despite the reduction of the computational burden obtained using of the pairwise likelihood, it still remains to evaluate the unmanageable pairwise likelihood quantities $W(\theta)$ and $J(\theta)$. The nonparametric bootstrap methods reveal to be of great help to provide an alternative solution. The steps of the methodology are briefly summarised in what follows. Several independent resample from the original dataset are

carried out with replacement obtaining bootstrap parameter estimates, $\hat{\theta}_b^*$ for each $b = 1, \dots, B$, where B refers to the number of bootstrap replicates. The empirical distribution of $\hat{\theta}^*$ is used as a good estimate for the distribution of the estimator $\hat{\theta}$ from which we can derive for example bias, standard error, and so forth. Thanks to the independence among the bootstrap replicates, the procedure strongly benefits from a parallel implementation on multicore machines or clusters of CPUs. For this problem, we took $B = 1400$.

To calculate confidence intervals only for correlation parameters, we use the bias corrected and accelerated confidence intervals (BCa) which are recommended for complex models (Efron and Tibshirani, 1994; DiCiccio and Efron, 1996). The BCa method has the advantage to be transformation respecting and second order accurate. A $100(1 - \alpha)\%$ BCa percentile confidence interval is based upon the quantiles of the bootstrap distribution of the estimates and is given by $[\hat{\theta}^{*(\alpha_1)}, \hat{\theta}^{*(\alpha_2)}]$, where $\hat{\theta}^{*(\alpha)}$ is the empirical percentile of level α and

$$\begin{aligned}\alpha_1 &= \Phi \left(\hat{z}_0 + \frac{\hat{z}_0 + z^{(\alpha/2)}}{1 - \hat{a}(\hat{z}_0 + z^{(\alpha/2)})} \right) \\ \alpha_2 &= \Phi \left(\hat{z}_0 + \frac{\hat{z}_0 + z^{(1-\alpha/2)}}{1 - \hat{a}(\hat{z}_0 + z^{(1-\alpha/2)})} \right),\end{aligned}\tag{4.2}$$

where \hat{z} is a bias correction constant defined as

$$\hat{z}_0 = \Phi^{-1} \left(\frac{\#(\hat{\theta}_b^* < \hat{\theta})}{B} \right),$$

where $\#$ means “number of”. To calculate the accelerated constant \hat{a} we follow (Efron and Tibshirani, 1994, p. 186):

$$\hat{a} = \frac{\sum_{i=1}^n (\hat{\theta}_{(\cdot)} - \hat{\theta}_{(i)})^3}{6 \{ \sum_{i=1}^n (\hat{\theta}_{(\cdot)} - \hat{\theta}_{(i)})^2 \}^{3/2}},$$

where $\hat{\theta}_{(i)}$ is the estimate of θ based on the observed data with the i -th observation deleted and $\hat{\theta}_{(\cdot)} = \sum_{i=1}^n \hat{\theta}_{(i)} / n$.

To evaluate the performance of this bootstrap-based procedure, we conducted a small simulation study. We generated 1,000 samples under the setting of Scenario 1 described in Section 3 with both $q = 3, 5$. The results (reported in the Appendix) show that the actual coverage of 95% confidence intervals are very close to the nominal level for both cases. As expected, less precision is observed for the higher dimension.

Applying this procedure to our real data we obtain the estimates for the polychoric correlations and their 95% confidence intervals shown in Figure 2. The same results are reported in tabular form in the Supporting Information Online.

The results reported in Table 2 and 3 are obtained thanks to the normal approximation of estimators. From Table 2, we comment that age and sex are not linearly related to the global




Figure 2: Centipede plot for the (i, j) -th estimated polychoric correlations (index on the right) and 95 % accelerated bias corrected bootstrap confidence intervals for the quality of service dataset.




Figure 3: Maximum pairwise likelihood estimates for the polychoric correlations represented on a color scale ranging from dark (polychoric correlation equal to 0) to white (polychoric correlation equal to 1).

item scores. On the other side higher scores are more prominent among subjects with higher educational level (p -value lower than 5% level). To visualise the relations among the items, consider the plot in Figure 3 in which we represent the estimated polychoric correlations with a color ranging from white (correlation equal to one) to darker colours (correlation going to zero). Note that we do not have any negative estimate. Our estimates indicate that the correlations of the latent Gaussian variables are higher within the two blocks of questions and lower between the two blocks. This is a very important evidence, since it means that high rating in the first item, i.e. high expected importance for the first item, is strongly related to the importance of other items, but not so related to its experienced quality. In Table 3 the differences $\xi_{j+7} - \xi_j$, namely the gap between experienced and expected service for the j -th characteristic are reported. For all $j = 1, \dots, 7$, this gap is negative with high significance (p -value lower than 10^{-40}), meaning that the experienced quality is clearly below customers' expectations.

Table 2: Regression coefficients for the effect of the covariates

	Estimate	Std. Error	z value	$\Pr(> z)$
Male	0.0130	0.0086	1.5148	0.1298
Age	-0.2002	0.1645	-1.2170	0.2236
High school	1.2857	0.6051	2.1249	0.0336
University	1.4260	0.6222	2.2919	0.0219

Table 3: Differences between the item specific intercepts in the expected and experienced blocks

	Estimate	Std. Error	z value	$\Pr(< z)$
Gap 1	-2.0277	0.1587	-12.7737	<1e-40
Gap 2	-2.7384	0.1862	-14.7066	<1e-40
Gap 3	-2.7536	0.1736	-15.8601	<1e-40
Gap 4	-3.3125	0.1985	-16.6898	<1e-40
Gap 5	-2.9145	0.1907	-15.2857	<1e-40
Gap 6	-2.6237	0.1909	-13.7470	<1e-40
Gap 7	-2.9021	0.1892	-15.3385	<1e-40

Gap 1: efficiency of the service; Gap 2: reliability of the service; Gap 3: flexibility of the service; Gap 4: velocity in giving solutions; Gap 5: ability of satisfy the customer's needs; Gap 6: ability to implement technological innovation; Gap 7: velocity in adapting to changing legislation;

5 Discussion

In this paper we have studied the pairwise likelihood approach for multivariate ordinal probit models, showing that its computational gain is dramatic if compared to the standard likelihood approach. The method has been applied to customer satisfaction data, introducing several complications to the basic model, such as covariates dependence and different means parameter for each latent variable. Despite the considerable number of parameters, the likelihood optimisation is feasible in reasonable time. A bootstrap-based approach to evaluate the estimators' standard errors has been adopted. We hope that our contribution will encourage the use of pairwise likelihood for ordered probit models in several applied fields such as marketing, psychometry, and medical research.

Acknowledgement

The authors thank Bruno Scarpa for generously providing the data and Nicola Sartori and Francesca Bassi for fruitful comments. Suggestions made by two anonymous referees are gratefully acknowledged.

Supporting Information

Additional simulations results for this article is available online. These include additional tables for the Simulation of Section 3.2, additional simulation results for $q = 15$ and an additional table with the numerical values of Figures 2-3.

References

- Agresti, A. (2002). *Categorical data analysis*, volume 359. John Wiley & Sons.
- Albert, J. H. and Chib, S. (1993). Bayesian analysis of binary and polychotomous response data. *Journal of the American statistical Association*, **88**(422), 669–679.
- Azzalini, A. and Scarpa, B. (2012). *Data Analysis and Data Mining: An Introduction*. Oxford University Press, USA.
- Cox, D. R. and Reid, N. (2004). A note on pseudolikelihood constructed from marginal densities. *Biometrika*, **91**, 729–737.
- De Leon, A. (2005). Pairwise likelihood approach to grouped continuous model and its extension. *Statistics & probability letters*, **75**, 49–57.
- DiCiccio, T. J. and Efron, B. (1996). Bootstrap confidence intervals. *Statistical Science*, pages 189–212.
- Efron, B. and Tibshirani, R. J. (1994). *An introduction to the bootstrap*, volume 57. CRC press.
- Genz, A. (1992). Numerical computation of multivariate normal probabilities. *Journal of computational and graphical statistics*, **1**, 141–149.
- Jackman, S. (2009). *Bayesian analysis for the social sciences*, volume 846. John Wiley & Sons.
- Kenne Pagui, E. C. and Canale, A. (2014). *PLordprob: Multivariate Ordered Probit Model via Pairwise Likelihood*. R package version 1.0.
- Kenne Pagui, E. C., Salvan, A., and Sartori, N. (2015). On full efficiency of the maximum composite likelihood estimator. *Statistics & Probability Letters*, **97**, 120–124.
- Likert, R. (1932). A technique for the measurement of attitudes. *Archives of psychology*, **22**(140).
- Lindsay, B. (1988). Composite likelihood methods. *Contemporary Mathematics*, **80**, 221–240.
- Lu, S. (2011). Measuring dependence via mutual information. Master’s thesis, Queen’s University.
- Mardia, K. V., Kent, J. T., Hughes, G., and Taylor, C. C. (2009). Maximum likelihood estimation using composite likelihoods for closed exponential families. *Biometrika*, **96**, 975–982.
- Molenberghs, G. and Verbeke, G. (2005). *Models for discrete longitudinal data*. Springer Verlag.

- Parasuraman, A., Zeithaml, V., and L.L., B. (1988). SERVQUAL: a multiple-item scale for measuring consumer perceptions of service quality. *Journal of Retailing*, **64**, 12–37.
- Rossi, P. E., Allenby, G. M., and McCulloch, R. (2012). *Bayesian statistics and marketing*. John Wiley & Sons.
- Varin, C., Reid, N., and Firth, D. (2011). An overview of composite likelihood methods. *Statistica Sinica*, **21**, 5–42.
- Varin, C. and Czado, C. (2010). A mixed autoregressive probit model for ordinal longitudinal data. *Biostatistics*, **11**, 127–138.
- Varin, C. and Vidoni, P. (2006). Pairwise likelihood inference for ordinal categorical time series. *Computational Statistics & Data Analysis*, **51**, 2365–2373.
- Xu, X. and Reid, N. (2011). On the robustness of maximum composite likelihood estimate. *Journal of Statistical Planning and Inference*, **141**, 3047–3054.
- Zhao, Y. and Joe, H. (2005). Composite likelihood estimation in multivariate data analysis. *Canadian Journal of Statistics*, **33**, 335–356.

Appendix

The obtained coverage level obtained in the simulation study discussed in Section 4 are (ordered from the first row to the last row and considering only the upper triangular elements of the polychoric correlation matrix):

- for $q = 3$
 - 0.942, 0.946, and 0.940 for the polychoric correlation coefficients
 - 0.940, 0.944, and 0.942 for the Pearson’s correlation coefficients
 - 0.942, 0.942, and 0.940 for the L measure coefficients
- for $q = 5$
 - 0.936, 0.924, 0.922, 0.950, 0.909, 0.946, 0.946, 0.932, 0.942, and 0.934 for the polychoric correlation coefficients
 - 0.934, 0.922, 0.924, 0.948, 0.907, 0.944, 0.946, 0.932, 0.936, and 0.934 for the Pearson’s correlation coefficients
 - 0.934, 0.922, 0.926, 0.946, 0.907, 0.829, 0.946, 0.833, 0.887, and 0.932 for the L measure coefficients