Statistical integration of diverse omics data

Bouhaddani, S. el

4

Probabilistic Partial Least Squares: identifiability, estimation and application

4.1 Abstract

With a rapid increase in volume and complexity of data sets, there is a need for methods that can extract useful information, for example the relationship between two data sets measured for the same persons. The Partial Least Squares (PLS) method can be used for this dimension reduction task. Within life sciences, results across studies are compared and combined. Therefore, parameters need to be identifiable, which is not the case for PLS. In addition, PLS is an algorithm, while epidemiological study designs are often outcome-dependent and methods to analyze such data require a probabilistic formulation. Moreover, a probabilistic model provides a statistical framework for inference. To address these issues, we develop Probabilistic PLS (PPLS).

We derive maximum likelihood estimators that satisfy the identifiability conditions by using an EM algorithm with a constrained optimization in the M step. We show that the PPLS parameters are identifiable up to sign. A simulation study is conducted to study the performance of PPLS compared to existing methods. The PPLS estimates performed well in various scenarios, even in high dimensions. Most notably, the estimates seem to be robust against departures from normality. To illustrate our method, we applied it to IgG glycan data from two cohorts. Our PPLS model provided insight as well as interpretable results across the two cohorts.

4.2 Introduction

With the exponentially growing volume of data sets, multivariate methods for reducing dimensionality are an important research area in statistics. For combining two data sets, Partial Least Squares (PLS) regression [28] is a popular dimension reduction method [1]. PLS decomposes variation in each data set in a joint part and a residual part. The joint part is a linear projection of one data set on the other that best explains the covariance between the two data sets. These projections are obtained by iterative algorithms, such as NIPALS [28]. Partial Least Squares is popular in chemometrics [3]. In this field, the focus is on development of algorithms with good prediction performance, while the underlying model is less important. For applications in life sciences, interpretation of parameter estimates is necessary to gain understanding of the underlying molecular mechanisms.

For interpretation, a model needs to be identifiable. A model is said to be unidentifiable if the model corresponds to more than one set of parameter values. For PLS, rotation of the parameters does not change the model [26]. Hence, PLS does not provide an identifiable model. By constraining the parameter space, identifiability can be obtained. This involves solving a challenging optimization problem, since PLS requires estimating a structured covariance matrix [19].

For many problems in life sciences the study design needs to be accounted for, and algorithmic approaches such as PLS cannot be applied. Hence, a probabilistic formulation is necessary. Since likelihood method provides asymptotic standard errors of parameter estimates, computer-intensive resampling procedures can be avoided.

Also for other dimension reduction techniques, probabilistic methods have been developed. In 1999, Tipping and Bishop [23] developed the Probabilistic Principal Component Analysis (PPCA), in order to deal with missing data and dependent

samples. In 2005, Bach and Jordan [2] developed Probabilistic Canonical correlation analysis (PCCA). However, for both PPCA and PCCA the model parameters are not identifiable, since rotation of the parameters does not change the model [23, 2]. In addition, in 2015, simultaneous envelopes models have been developed [4] for 'low-dimensional' settings. Further, Probabilistic PLS Regression and Probabilistic PLS have been proposed [14, 30]. For all these approaches, the model parameters are not identifiable.

In this paper we propose the Probabilistic Partial Least Squares (PPLS) model and show that the model parameters are identifiable up to a sign. We propose to maximize the PPLS likelihood with an EM algorithm that decouples the likelihood into several factors involving distinct sets of parameters. In the M step, a constrained optimization problem is solved by using a matrix of Lagrange multipliers.

The rest of the paper is organized as follows: In section 4.3 we develop the PPLS model and establish identifiability of the model parameters. We develop an efficient algorithm for estimating the PPLS parameters. In section 4.4 we study the performance of the PPLS estimators via simulations. In section 4.5 we illustrate the PPLS model with two data matrices from two cohorts. We finish with a discussion.

4.3 Model and estimation

4.3.1 The PPLS model

Let x and y be two random row-vectors of dimension p and q, respectively. The Probabilistic Partial Least Squares (PPLS) model describes the two random vectors in terms of a joint part and a noise part. The joint part consists of correlated latent vectors, denoted by t and u, while the noise part consists of isotropic normal random vectors referred to as e, f and h. The dimension of t and u is denoted by r. The PPLS model describing the relationship between x, y and the joint and noise parts is

$$x = tW^{\mathrm{T}} + e, \quad y = uC^{\mathrm{T}} + f, \quad u = tB + h. \tag{4.1}$$

Specifically, $e = (e_1, \ldots, e_p)$, $f = (f_1, \ldots, f_q)$ and $h = (h_1, \ldots, h_r)$ are independent with zero mean and referred to as noise variables. The distributions of e, f and h are multivariate normal with positive definite covariance matrix proportional to the identity matrix,

$$e \sim \mathcal{N}\left(0, \sigma_e^2 I_p\right), \quad f \sim \mathcal{N}\left(0, \sigma_f^2 I_q\right), \quad h \sim \mathcal{N}\left(0, \sigma_h^2 I_r\right).$$

The latent vector $t = (t_1, \ldots, t_r)$ is an r-dimensional multivariate normal vector with with zero mean and diagonal positive definite covariance matrix $\Sigma_t = \mathrm{diag}(\sigma_{t_1}^2, \ldots, \sigma_{t_r}^2)$, so

$$t \sim \mathcal{N}\left(0, \Sigma_t\right). \tag{4.2}$$

The matrix $B = \mathrm{diag}\left(b_1, \ldots, b_r\right)$ is a diagonal matrix of size r, containing regression coefficients of u on t. Finally W $(p \times r)$ and C $(q \times r)$ are parameter matrices, referred to as loadings. The PPLS model for the random p-dimensional row-vector x and random q-dimensional row-vector y is given in Eq. (4.1). Let θ be the parameters of the PPLS model, i.e.,

$$\theta = (W, C, B, \Sigma_t, \sigma_e, \sigma_f, \sigma_h). \tag{4.3}$$

The PPLS model and its parameters are formulated conditional on the value of the dimension of the latent space r.

The PPLS model (4.1) assumes a multivariate normal distribution for the observable random vectors x and y. The covariance between x and y is modeled by the regression of the latent vector u on t. The distribution of (x, y) is $\mathcal{N}(0, \Sigma)$ with density given, for $x \in \mathbb{R}^p$ and $y \in \mathbb{R}^q$, by

$$f(x,y) = (2\pi)^{-(p+q)/2}|\Sigma|^{-\frac{1}{2}}\mathrm{e}^{(x,y)\Sigma^{-1}(x,y)^{\mathrm{T}}},$$

and covariance matrix

$$\Sigma = \begin{pmatrix} \Sigma_x & \Sigma_{x,y} \\ \Sigma_{y,x} & \Sigma_y \end{pmatrix} = \begin{pmatrix} W\Sigma_t W^{\mathrm{T}} + \sigma_e^2 I_p & W\Sigma_t BC^{\mathrm{T}} \\ CB\Sigma_t W^{\mathrm{T}} & C(B^2\Sigma_t + \sigma_h^2 I_r)C^{\mathrm{T}} + \sigma_f^2 I_q \end{pmatrix}. \tag{4.4}$$

This follows from the normality property and from computing the variances and covariances of the random vectors; see Appendix 4.7 for the details.

4.3.2 Identifiability of probabilistic PLS

To establish identifiability of the PPLS model, some assumptions about its parameters have to be made. First, we assume that $0 < r < \min(p, q)$. Second, we assume that the diagonal elements of B are positive, $b_k > 0$ for $k \in \{1, \ldots, r\}$. This will not restrict the model, since $t_k b_k$ is equal to $-t_k b_k$ in distribution. To identify the order of the loading vectors, the elements of $(\sigma_{t_k}^2 b_k)_{k=1}^r$ are assumed to be strictly decreasing with k. Finally, we assume that the loading matrices W and C are orthogonal, i.e., $W^{\mathrm{T}}W = C^{\mathrm{T}}C = I_r$. Together with the diagonality of Σ_t in (4.2), it implies identifiability of all parameters up to sign. This is shown in the following theorem.

Theorem 4.3.1. *Let r be fixed such that $0 < r < \min(p, q)$. Let $(x, y)_1$ and $(x, y)_2$ be generated by the PPLS model (4.1) having covariance matrix Σ_1 and Σ_2 with underlying parameters θ_1 and θ_2 (as defined in (4.3)), respectively. Then, we have that*

$$\Sigma_1 = \Sigma_2$$

implies that $W_1 = W_2\Delta$, $C_1 = C_2\Delta$ for some diagonal matrix Δ with on the diagonal elements $\delta_i \in \{-1, 1\}$, for $i \in \{1, \ldots, r\}$, and all other parameters in θ_1 and θ_2 are equal.

The formal proof is given in Appendix 4.7. Identifiability up to sign can be represented by a diagonal orthogonal matrix (a diagonal matrix with diagonal elements in $\{-1, 1\}$). For example, taking the model for x in (4.1), we may substitute W by WR_S and t by tR_S, where R_S is a diagonal orthogonal matrix, and get

$$x = tR_S R_S^{\mathrm{T}} W^{\mathrm{T}} + e = \sum_{j=1}^{r} t_j (R_S)_{jj}^2 w_j^{\mathrm{T}} + e.$$

Since $(R_S)_{jj}^2 = 1$ and the distribution of t_j and $-t_j$ is the same, the right-hand side reduces to the original model for x in (4.1). Note that the PPLS model is not invariant under general rotation matrices. Take a general rotation matrix R, then we still get

$$x = tW^{\mathrm{T}} + e = tRR^{\mathrm{T}}W^{\mathrm{T}} + e,$$

since $RR^{\mathrm{T}} = I_r$. Inspecting the covariance of TR we see that $\mathrm{Cov}(TR) = R^{\mathrm{T}}\Sigma_t R$, which is not diagonal if R is not diagonal, and violates the PPLS model assumption on Σ_t in Eq. (4.2).

4.3.3 Estimating the parameters

Unlike the iterative PLS methods, we propose a simultaneous approach for estimating the parameters, while taking the constraints in the PPLS model into account. Given the number of PPLS components, r, the log likelihood of an independent and identically distributed (iid) sample $(X, Y) = \{(X_1, Y_1)^{\mathrm{T}}, \ldots, (X_N, Y_N)^{\mathrm{T}}\}^{\mathrm{T}}$ of size N from (x, y) is

$$L(\theta) = -\frac{N(p+q)}{2} - \frac{N}{2}\ln|\Sigma| - \frac{N}{2}\mathrm{tr}\left(S\Sigma^{-1}\right) \tag{4.5}$$

with $S = N^{-1}\sum_{i=1}^{N}(X_i, Y_i)^{\mathrm{T}}(X_i, Y_i)$ and Σ as in Eq. (4.4). To ensure empirical identifiability, we assume that $r < N$. Note that the data dimensionality p and q may be larger than N. For estimation of θ, maximum likelihood is used.

The log likelihood (4.5) depends in a non-linear way on the theoretical covariance matrix Σ, which contains the loadings and variances. Optimizing this function directly is a non-trivial task, especially in high dimensions (i.e. when p and q are large). However, the PPLS model allows for a more simple (but iterative) optimization approach. Indeed, the maximum likelihood estimates for θ are a least squares type solution if the latent variables t and u are observed, as the model for x and y in (4.1) involves known t and u. In contrast, knowing θ allows for reconstruction of t and u by computing their conditional means given x and y. Alternating these two scenarios is actually an Expectation-Maximization (EM) [5] algorithm, with observed data (x, y) and missing data (t, u).

The EM algorithm The joint distribution of the complete data (x, y, t, u) can (with abuse of notation) be decomposed as

$$f(x, y, t, u) = f(x|t)f(y|u)f(u|t)f(t). \tag{4.6}$$

This follows from

$$f(x, y, t, u) = f(x, y|t, u)f(t, u) = f(x|t, u)f(y|t, u)f(t, u).$$

The second equation is implied by the fact that x and y are independent given t and u. The first two factors in the right-hand side can be rewritten as $f(x|t, u) = f(x|t)$ and $f(y|t, u) = f(y|u)$, since x and u are independent given t, and y and t are independent given u. The last factor can be rewritten as $f(u|t)f(t)$, yielding Eq. (4.6). The logarithm of the first three factors in the product in (4.6) can be written

as

$$\ln f(X|T) = -\frac{Np\ln 2\pi\sigma_e^2}{2} - \frac{1}{2\sigma_e^2}\sum_{i=1}^{N}||X_i - T_iW^{\mathrm{T}}||^2,$$

$$\ln f(Y|U) = -\frac{Nq\ln 2\pi\sigma_f^2}{2} - \frac{1}{2\sigma_f^2}\sum_{i=1}^{N}||Y_i - U_iC^{\mathrm{T}}||^2,$$

$$\ln f(U|T) = -\frac{Nr\ln 2\pi\sigma_h^2}{2} - \frac{1}{2\sigma_h^2}\sum_{i=1}^{N}||U_i - T_iB||^2.$$

Denote by $L_{\mathrm{Comp}} = \ln f(X, Y, T, U)$ the complete data log-likelihood, and define

$$Q(\theta) = \mathbb{E}\left\{L_{\mathrm{Comp}}(\theta)|X, Y, \theta'\right\},$$

where the expectation is taken conditional on the observed X and Y, and θ' is a fixed current estimate of the parameters. By optimizing Q over all allowed θ, we get a non-negative increase in the *observed* log-likelihood L. Moreover, by iterating this process of taking the expectation and maximizing over θ, the estimates in general converge to a stationary point or, in particular, a (possibly local) maximum of L [5, 29]. The expectation step calculates the conditional expectation of the missing data given the observed data given by $Q(\theta)$, which may in general involve intractable integration. However, for the exponential family, in particular the multivariate normal family, the complete likelihood depends on the complete data only via the sufficient statistics (called $t(\mathbf{x})$ in [5]), which are given in terms of the first and second moments of the complete data for the multivariate normal distribution. Computing $Q(\theta)$ implies computing the expected first and second moment of the latent variables: $\mathbb{E}\left(T|X, Y, \theta\right)$, $\mathbb{E}\left(T^{\mathrm{T}}T|X, Y, \theta\right)$, $\mathbb{E}\left(U|X, Y, \theta\right)$, $\mathbb{E}\left(U^{\mathrm{T}}U|X, X, \theta\right)$ and $\mathbb{E}\left(U^{\mathrm{T}}T|X, Y, \theta\right)$; see Appendix 4.7 for details. Moreover, the decomposition in (4.6) allows for optimization of $\mathbb{E}\{\ln f(X|T)\}$, $\mathbb{E}\{\ln f(Y|U)\}$ and $\mathbb{E}\{\ln f(U|T)\}$ separately, while only considering parameters involved in each factor. Maximizing Q over θ yields parameter estimates for the next iteration in the EM algorithm. This leads us to the following theorem.

Theorem 4.3.2. *Let X and Y be an observed data sample of size N, generated according to the PPLS model (4.1). Let r be fixed such that $0 < r < \min(N, p, q)$. The parameters in θ can be estimated with an EM algorithm, yielding the following*

iterative scheme in k with given starting values for $k = 0$.

$$
\begin{aligned}
W^{k+1} &= X^{\mathrm{T}}\,\mathbb{E}\left(T|X,Y,\theta^k\right)\left(L_W^{\mathrm{T}}\right)^{-1};\\
C^{k+1} &= Y^{\mathrm{T}}\,\mathbb{E}\left(U|X,Y,\theta^k\right)\left(L_C^{\mathrm{T}}\right)^{-1};\\
B^{k+1} &= \mathbb{E}\left(U^{\mathrm{T}}T|X,Y,\theta^k\right)\left\{\mathbb{E}\left(T^{\mathrm{T}}T|X,Y,\theta^k\right)\right\}^{-1}\circ I_r;\\
\Sigma_t^{k+1} &= \frac{1}{N}\mathbb{E}\left(T^{\mathrm{T}}T|X,Y,\theta^k\right)\circ I_r;\\
(\sigma_h^2)^{k+1} &= \frac{1}{Nr}\mathrm{tr}\mathbb{E}\left(H^{\mathrm{T}}H|X,Y,\theta^k\right);\\
(\sigma_e^2)^{k+1} &= \frac{1}{Np}\mathrm{tr}\mathbb{E}\left(E^{\mathrm{T}}E|X,Y,\theta^k\right);\\
(\sigma_f^2)^{k+1} &= \frac{1}{Nq}\mathrm{tr}\mathbb{E}\left(F^{\mathrm{T}}F|X,Y,\theta^k\right);
\end{aligned}
$$

where L_W and L_C are such that

$$
\begin{aligned}
L_W L_W^{\mathrm{T}} &= \mathbb{E}\left(T^{\mathrm{T}}|X,Y,\theta^k\right)\,X\,X^{\mathrm{T}}\,\mathbb{E}\left(T|X,Y,\theta^k\right),\\
L_C L_C^{\mathrm{T}} &= \mathbb{E}\left(U^{\mathrm{T}}|X,Y,\theta^k\right)\,Y\,Y^{\mathrm{T}}\,\mathbb{E}\left(U|X,Y,\theta^k\right).
\end{aligned}
$$

The proof for Theorem 4.3.2 and the expressions for the conditional expectations are given in Appendix 4.7. Note the dependency of W^{k+1} and C^{k+1} on the matrices L_W and L_C. These matrices ensure orthogonality of W^{k+1} and C^{k+1} in each iteration:

$$
\left(W^{k+1}\right)^{\mathrm{T}} W^{k+1} = L_W^{-1}\,\tilde{T}^{\mathrm{T}}\,X\,X^{\mathrm{T}}\,\tilde{T}\left(L_W^{\mathrm{T}}\right)^{-1} = L_W^{-1}\,L_W\,L_W^{\mathrm{T}}\,\left(L_W^{\mathrm{T}}\right)^{-1} = I_r,
$$

where $\tilde{T} = \mathbb{E}(T|X,Y,\theta^k)$. The exact forms of L_W and L_C are not unique. Two choices are the eigenvectors of $\mathbb{E}(T^{\mathrm{T}}|X,Y,\theta^k)\,X\,X^{\mathrm{T}}\,\mathbb{E}\left(T|X,Y,\theta^k\right)$ and the lower triangular matrix of $X^{\mathrm{T}}\,\mathbb{E}(T|X,Y,\theta^k)$ in the Cholesky decomposition. Note that these two orthogonalization matrices are straightforward to calculate with standard linear algebra tools. Since the PPLS model is identifiable, all choices for L_W and L_C will lead to the same optimum as the iteration number k tends to infinity.

Standard errors for PPLS Asymptotic standard errors for maximum likelihood estimators are found by inverting the observed Fisher information matrix. Following the reasoning of [16], the observed information may be given by

$$
\mathbb{E}\left\{B(\hat{\theta})|X,Y\right\} - \mathbb{E}\left\{S(\hat{\theta})S(\hat{\theta})^{\mathrm{T}}|X,Y\right\}.
$$

Here $S(\hat{\theta}) = \nabla\lambda(\hat{\theta})$ and $B(\theta) = -\nabla^2\lambda(\hat{\theta})$ are the gradient and negative of the second derivative of the log likelihood $\lambda(\theta)$ respectively evaluated in the MLE $\hat{\theta}$. The explicit form of the asymptotic covariance matrix of w_k is given in Appendix 4.7. The square root of the diagonal elements are the asymptotic standard errors for the corresponding loading estimates.

Finding the number of components r Available approaches to determine the number of PPLS components r are minimizing a cross-validated loss function [9], visually inspecting eigenvalues of a covariance matrix [17], and selecting the number of components needed to achieve a certain proportion of variance explained by the components. In this paper we apply the last approach.

The PLS and PPLS algorithms are available as R packages at `github.com/selbouhaddani` under repository OmicsPLS and PPLS, respectively.

4.4 Simulation study

To evaluate the performance of the PPLS estimates, a simulation study was conducted. The aim was (1) to investigate the performance of PPLS for various scenarios, (2) to evaluate robustness of the PPLS estimates against departures from the normality assumption, (3) to compare the performance of the loading estimates with other probabilistic approaches, and (4) to compare the asymptotic PPLS standard errors with the bootstrap standard errors.

The simulated data were generated according to the PPLS model (4.1). The number of components was chosen to be 3, both in the data generation and in the estimation. We considered combinations of small and large sample size ($N \in \{50, 500\}$), low and high dimensionality ($p \in \{20, 1000\}$), and small and large proportion of noise (denoted by $\alpha_n \in \{0.1, 0.5\}$). The robustness of PPLS was evaluated by considering four different continuous and discrete distributions for the latent variables t, u, e, f and h; we chose the normal distribution, the t distribution with two degrees of freedom, the Poisson distribution with rate 1, and the Binomial distribution with two trials and success probability 0.25. These distributions cover a wide range of characteristics typically observed in omics data: heavy tailed, skewed and/or discrete. The latent variables were scaled to have mean zero and variances as specified below. All scenarios are summarized in Table 4.1.

The true loading values per component were generated from the normal density function with parameters μ and σ, denoted by $N(x; \mu, \sigma^2)$, as follows

$$w_{j,k} = N\{j; (1/2 + 1/10j)p, 1/10p\}, \quad c_{j,k} = N\{j; (3/5 + 1/10j)q, 1/10q\}.$$

The second columns in W and C were orthonormalized with respect to the first columns, and the third columns were orthonormalized with respect to the first two columns; we used a Gram-Schmidt procedure for both operations. The elements of the diagonal matrix B were set to $b_k = \mathrm{e}^{\ln(1.5)-3(k-1)/10} = (1.5, 1.11, 0.82)$, for Σ_t we chose $\sigma_{t_k} = \mathrm{e}^{-(k-1)/10} = (1, 0.90, 0.82)$.

For comparing the parameter estimates with the true values θ, we computed the bias and the variance of the estimates. To deal with the identifiability up to sign, we multiplied each estimated loading vector by -1 if the inner product of the estimated loading vector and the true loading vector was negative. Moreover, we swapped columns in W and C to maintain the same ordering as the ordering in the true loadings. This was done to avoid inflation of the bias or variance due to a wrong sign or ordering of the individual components.

PPLS estimates were compared to PLS estimates (with orthogonal loadings, see [20] for an overview) for all scenarios above. For comparing PPLS with PPCA and

PCCA, we constructed a 'null model', i.e., $B = 0$, as well as $B \neq 0$. We used the same scenarios as above, but we only considered the normal distribution.

Table 4.1: **Overview of the simulation scenarios.** The noise level is defined as the proportion of variation in the noise matrices E, F and H relative to the total variation in X, Y and U respectively.

Sample size	$N = (50,\ 500)$
Dimensionality	$p = q = (20,\ 1000)$
Noise level	$\alpha_n = (0.1,\ 0.5)$
Distribution of t, u, e, f and h	$\{\mathcal{N}(0,1),\ t_2,\ \mathcal{P}(1),\ Binom(2, 0.3)\}$

Regarding standard errors for PPLS loadings, we compared asymptotic standard errors (as in Section 4.3) and bootstrap standard errors [27]. One set of two data matrices X and Y was simulated from a PPLS model with $p = q = 20$ normally distributed variables. Based on these data, asymptotic and bootstrap standard errors were calculated. The number of bootstrap replicates was 1000. Furthermore, simulation-based standard errors for the loadings (based on standard deviations over 1000 data matrices drawn from the PPLS model used to generate the original data) were included as reference. Low and high noise levels ($\alpha_n = 0.1$ resp. $\alpha_n = 0.5$), and small, large and 'extra large' sample sizes ($N = 50$, $N = 500$ and $N = 5000$ respectively) were considered. In the 'extra large' sample size scenario, no simulation-based reference was calculated. The PPLS estimation algorithm was considered to be converged when either the log-likelihood increment was below 10^{-6}, or 10^4 EM steps were made. For each scenario, 1000 replicates are used.

4.4.1 Results

Results for the loadings The biases and variances of the estimated first component W_1 for the low dimensional case for normally distributed latent variables are graphically depicted in Figure 4.1. A black dot represents the average estimated PPLS loading value across 1000 simulations, whereas the width of the black dashed vertical line equals two times the standard deviation across 1000 simulations. The red star and red dashed vertical line represent the average loading value and twice the standard deviation for the PLS estimates. The true loading values are represented by a step function with steps at each index $j \in \{1, \ldots, p\}$. Results for other components and scenarios are included in the Supplement.

Comparing the estimates for the *first* loading component W_1, a better performance of PPLS compared to PLS was observed in terms of bias. In all scenarios the bias of the PPLS estimators were about the same as or less than the bias of the PLS estimators. Both estimators showed larger bias towards zero for higher absolute loading values. The biases decreased with a larger sample size and lower noise level. The biases of both estimators were very similar across different distributions. In the scenario where there is 50% noise and few (50) samples the variance of the PPLS estimators tended to be slightly larger than the variances of the PLS estimators when the true loading values were larger. This was observed across all distributions. The variances of the PPLS estimates were about the same or lower than the PPLS estimates in all other scenarios, where either the noise level was less or more samples were available. For

both PPLS and PLS estimators the variances tended to increase with higher loading values. The variances decreased with larger sample size and lower noise level. The variances of bots estimators were very similar across different distributions. For the loading component C_1 and their PLS and PPLS estimators the same conclusions were obtained.

For the *second* loading component W_2 (shown in the Supplement), the biases of the PPLS loading estimates were as good as, and often better than the PLS loading estimates, especially at lower values. In the scenarios of 50% noise and a small sample size ($N = 50$) the bias was slightly larger for PPLS estimators compared to PLS estimates when the loading values were larger. Both estimators showed larger bias towards zero for higher loading values. The biases decreased with a larger sample size and lower noise level. For all distributions, the biases of both estimators were very similar. The variances of the PPLS estimators were as good as or lower than the PLS estimators, except in the scenario in which both the noise level was high (50%) and the sample size was small (50). In this scenario the variances of the PPLS estimators were still lower if the true loading values were close to zero, but higher for higher loading values. For both PPLS and PLS estimators the variances tended to increase with higher loading values. The variances decreased with larger sample size and lower noise level. The variances of both estimators were very similar across different distributions. For the loading component C_2 and their PLS and PPLS estimators the same conclusions were obtained.

For the *third* loading components W_3 and C_3 (shown in the Supplement), the same observations were made as for the first loading components W_1 and C_1, both for the biases as for the variances.

For the high and extra high-dimensional case, the same results were obtained for the loadings W and C. See the Supplement for more details.

With regard to the comparison of PPLS with PPCA and PCCA, PPLS performed better than PCCA and similar to PPCA in most scenarios. Details are given in the Supplement.

Results for the variance parameters The performance of the estimators of the variance parameters B, σ_t, σ_e, σ_f and σ_h were also evaluated, the results are shown in Figure 4.2. We did not compare with PLS as these model parameters are not present in the PLS framework. For sake of comparison, we calculated the relative biases and variances of the estimates with respect to the true corresponding parameter value. The biases of the PPLS estimators for all variance parameters were very small for large sample size ($N = 500$), regardless of the noise. For small sample size ($N = 50$), the first two diagonal elements of B and Σ_t were overestimated, while the last component was underestimated. The noise parameters σ_e and σ_f were underestimated in these scenarios, while the estimator for σ_h was unbiased, except in combination with a low noise level (10%). The relative biases decreased slightly with lower noise level, except for the earlier mentioned σ_h, and decreased more with larger sample size. The relative variances of the estimators of B, Σ_t and σ_h were larger than the variances of the estimators of σ_e and σ_f. For B, there was a slight increase in variance across the three components. The variances decreased slightly with lower noise and more with larger sample size. The variances slightly decreased in the scenario of high dimensionality and high noise level. The same observations were made across the

different distributions.

Ordering of the loadings We compared the ordering of the true loadings W and C with the ordering of the estimated loadings. This provides a proportion across the 1000 simulation replicates in which the ordering matched. In Table 4.2, the proportion of correct orderings of W for the scenario with normally distributed latent variables is shown for different scenarios. It can be seen that the proportion of correct orderings tends to be lower with smaller sample size and with higher noise level. Moreover, if the sample size is small, the proportion of correct orderings is much lower with higher noise. A higher dimensionality has a slightly negative impact on the correct ordering proportion when the sample size is larger, but a positive impact in the small sample size scenario. Especially, when also the noise level is high, this can be considerable. The same observations were made for the other distributions. Exactly the same proportions were observed for the loadings C.

Table 4.2: **Proportion of correct order of loadings W and C across 1000 simulation replicates.** These were obtained for different values of the dimensionality (high = 1000 variables, low = 20 variables), sample size (large = 500 subjects, small = 50 subjects) and noise level (high = 50% noise, low = 10% noise).

Dimensionality	Sample size	Noise level	Correct ordering proportion
low	large	low	1.000
		high	0.989
	small	low	0.932
		high	0.435
high	large	low	0.990
		high	0.985
	small	low	0.940
		high	0.665

Comparison of PPLS standard errors The results for low noise level are shown in Figure 4.3. In all scenarios, the asymptotic standard errors were smaller than the bootstrap standard errors for nearly all loading elements. In particular, for high loading values the difference between asymptotic and bootstrap standard errors tended to be large. This difference decreased with larger sample size: In the 'extra large' sample size, the bootstrap and asymptotic standard errors had similar magnitude. Similar observations were made for other distributions. For details, see the Supplement.

4.5 Data analysis

To illustrate the Probabilistic Partial Least Squares model, we apply it to IgG glycan datasets. Glycans, in particular IgG glycans, play an important role in the innate immune system, as well as in cell signaling. IgG2 glycans are less abundant than IgG1 glycans and more difficult to measure. Therefore, by using the relationships between

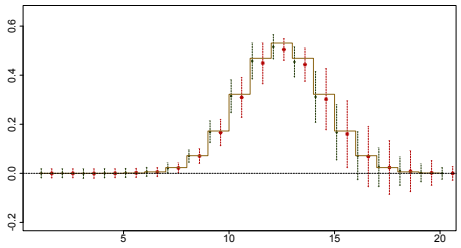

(a) Low noise (10%); small sample size ($N = 50$)

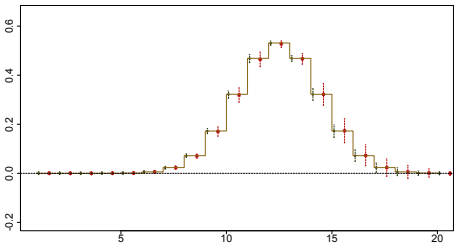

(b) Low noise (10%); large sample size ($N = 500$)

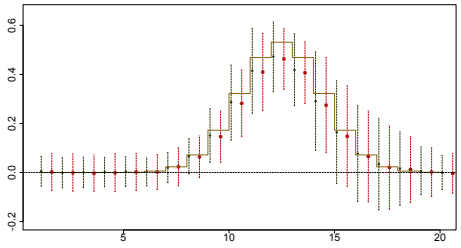

(c) High noise (50%); small sample size ($N = 50$)

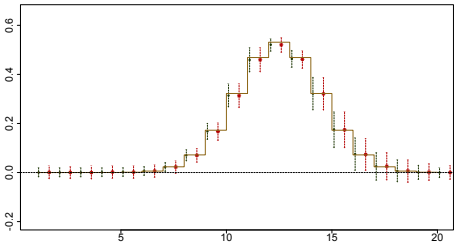

(d) High noise (50%); large sample size ($N = 500$)

Figure 4.1: **True and estimated loadings W_1 over 1000 simulation replications.** The black dots and dashed vertical lines (on the left of each pair) represent PPLS estimates, the red stars and dashed vertical lines (on the right of each pair) represent PLS estimates. The dots and stars are the average loading values across 1000 simulation replications; the width of the dashed lines are twice the standard deviations. The results are for normally distributed latent variables (t, e, f and h) and low dimensionality ($p = q = 20$ variables).

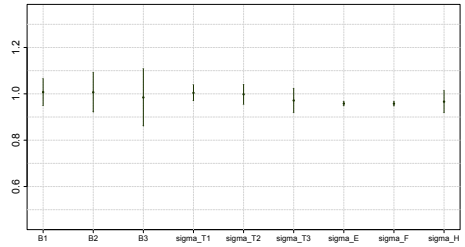

(a) Low noise (10%); small sample size ($N = 50$)

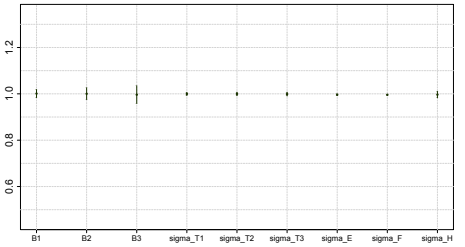

(b) Low noise (10%); large sample size ($N = 500$)

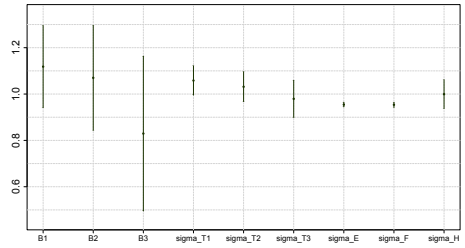

(c) High noise (50%); small sample size ($N = 50$)

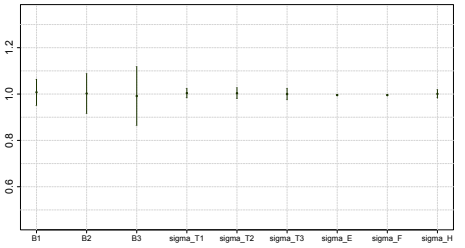

(d) High noise (50%); large sample size ($N = 500$)

Figure 4.2: **Relative estimated variance parameters B, Σ_t, σ_e, σ_f and σ_h (w.r.t. the truth) over 1000 simulation replications.** The dots are the average values across 1000 simulation replications; the width of the dashed lines are twice the standard deviations. The results are for normally distributed latent variables (t, e, f and h) and low dimensionality ($p = q = 20$ variables).

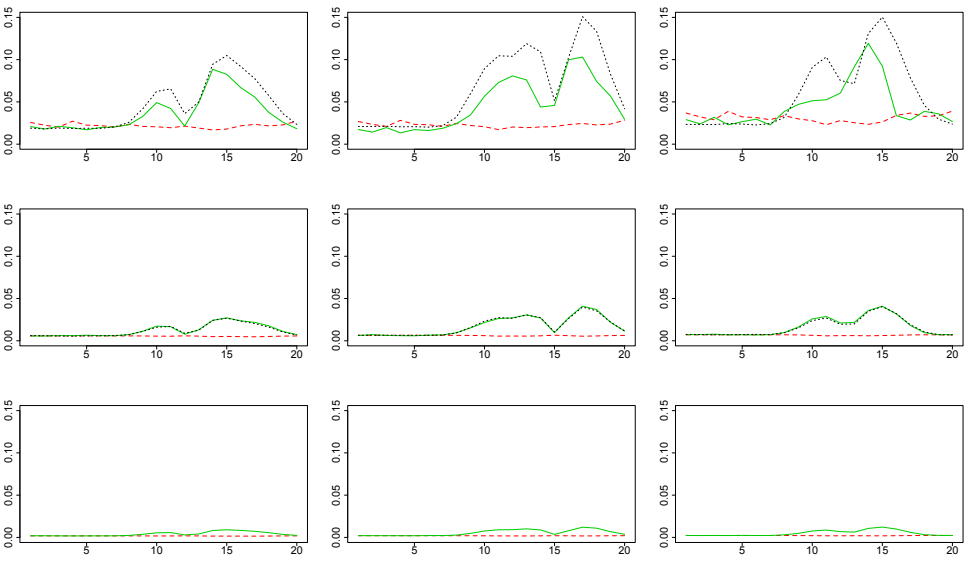

Figure 4.3: **Standard errors of the W loading elements per component.** Bootstrap standard errors (solid green line), asymptotic standard errors (dashed red line) and simulation-based standard errors (dotted black line) are plotted for the loading estimates in each component. Plots for the three sample sizes (small $N = 50$, large $N = 500$, 'extra large' $N = 5000$) are shown along the rows. The three loading components (W_1, W_3 and W_3) are plotted column wise. The last row does not include simulation-based standard errors. The data are generated from a normal distribution with $p = q = 20$ variables and low noise level ($\alpha_n = 0.1$).

IgG1 and IgG2 glycans, the characteristics of IgG2 can be better estimated. Hence, we will use IgG1 as X matrix, and IgG2 as Y matrix.

In total, 40 IgG glycans were measured, of which $p = 20$ are of subclass IgG1 and $q = 20$ are of subclass IgG2. These data were measured in two cohorts (CROATIA_Korcula with 951 participants and CROATIA_Vis with 796 participants) [12]. The data matrices containing IgG1 and IgG2 glycan variables are denoted by X_m and Y_m, with $m \in \{1, 2\}$, where $m = 1$ corresponds to CROATIA_Korcula and $m = 2$ corresponds to CROATIA_Vis. We apply PPLS to IgG1 and IgG2 glycans in each cohort *separately* and compare results.

In Figure 4.4, heatmaps of the correlations within and between the IgG1 and IgG2 glycans are shown, from which it is clear that there are many highly positive correlations between and within IgG1 and IgG2 in each data set. The RV coefficient [18], which generalizes the *squared* Pearson correlation coefficient to two matrices, was about 0.60 and 0.45 for CROATIA_Korcula and CROATIA_Vis cohorts respectively.

To determine the number of latent variables to use, we considered the total amount of variance explained by the latent space relative to the total amount of variation in the data: $||T_m||/||X_m||$ and $||U_m||/||Y_m||$ for $m \in \{1, 2\}$. By using four components, at least 89% of the total variation in each of the matrices X_1, X_2, Y_1 and Y_2 was accounted for.

For both cohorts, we fitted the PPLS models using $r = 4$ latent components. The amount of overlap in each cohort, estimated by $\mathrm{tr}\hat{\Sigma}_{x,y}/\mathrm{tr}\hat{\Sigma}_y$ given in (4.4), was 58% and 46% for CROATIA_Korcula and CROATIA_Vis cohorts, respectively. The PPLS loadings are inspected to identify which IgG glycans contribute most to this overlap. The estimated IgG1 loadings $w_{j,k}$, $j \in \{1, \ldots, p\}$ and $k \in \{1, \ldots, 4\}$, for both cohorts and both subclasses are depicted in Figure 4.5. The first joint component is proportional to the average glycan, as all glycans get about the same loading value. The second joint component involves especially G0 and G2 glycan subtypes, in which they are negatively correlated. Inspection of the loading values for the third component shows contibutions of fucosylated and non-fucosylated glycan subtypes. In the fourth component a pattern of positive and negative loading values is visible regarding the presence and absence of bisecting GlcNAc, respectively. The large loading value for G1NS is remarkable. The same conclusions hold for IgG2, as the estimated loading values were very similar. It is interesting to note that the observed patterns within components potentially reflect enzymatic synthesis where monosaccharides are added to glycan substrates [22]. Additionally, similar patterns are seen reflecting the inflammatory characteristics of glycans in aging and several different diseases [13]. Finally, the observed loading patterns were strikingly similar for both cohorts.

4.6 Discussion

We proposed PPLS to model the covariance matrix of two data sets. Maximum likelihood estimators for the model parameters were derived by solving a constrained optimization problem, and the parameter loadings were shown to be identifiable up to sign. This property ensures that PPLS estimates are comparable across several studies.

Our simulation study showed that the PPLS estimators had good performance

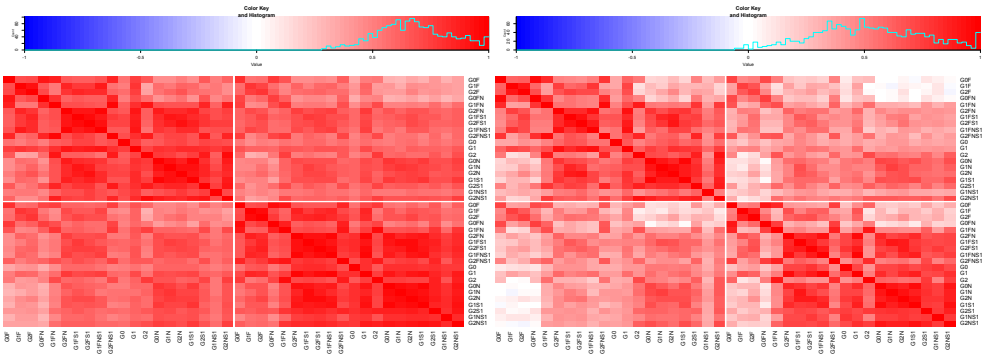

(a) The first (CROATIA_Korcula) cohort (b) The second (CROATIA_Vis) cohort

Figure 4.4: **Heatmaps of the correlations between IgG1 and IgG2 glycans.** In left the correlations of the CROATIA_Korcula cohort is shown. In right the CROATIA_Vis cohort is shown. The upper-left and lower-right block are the within subclass correlations, the off-diagonal block contains the correlations between IgG1 and IgG2 glycans. In both cohorts the glycans exhibit high positive correlations, especially between glycans within the IgG1 and IgG2 subclasses.

and lower bias compared to PLS. Most notably, the performance of PPLS was robust to misspecification of the distribution of the variables. A smaller sample size and high noise level had a negative effect on the accuracy of the estimates, but large loading values were still non-zero. Also, compared to Probabilistic CCA estimates, the PPLS estimates were less biased and more efficient. For high-dimensional data, PCCA estimates have larger bias and higher variance. This is likely to be caused by the unstable inverse sample covariance matrix calculated when using PCCA. Moreover, if the number of variables is larger than the sample size, PCCA estimates cannot be obtained. Therefore, especially in omics data analysis, PPLS provides more robust findings.

As an illustration of the PPLS model, we analyzed IgG glycomics data from two cohorts. The high correlations in the data (Figure 4.4) and the use of multiple cohorts illustrate the applicability of PPLS to facilitate combination of results derived from different experimental settings. We found that the estimated loading values were almost identical across the two cohorts (Figure 4.5).

When multiple cohorts are available, a meta-analysis on the parameter estimates can be performed. In ordinary regression models, this has been addressed for both low-dimensional [6] and high-dimensional [10] design matrices. When there is no access to all data, PPLS estimates can be combined by using standard meta-analysis approaches [6]. Such an approach requires that the PPLS parameter estimates are identifiable and asymptotically normally distributed. For the PLS framework, several approaches to combine estimates across cohorts were developed when there is access to all data. A group-PLS approach was considered [15] to incorporate several groups of subjects in the model. The authors showed that under certain assumptions this approach provided better predictions than a model without group effects. However their model is not identifiable and requires $N > p$. Another method is based on

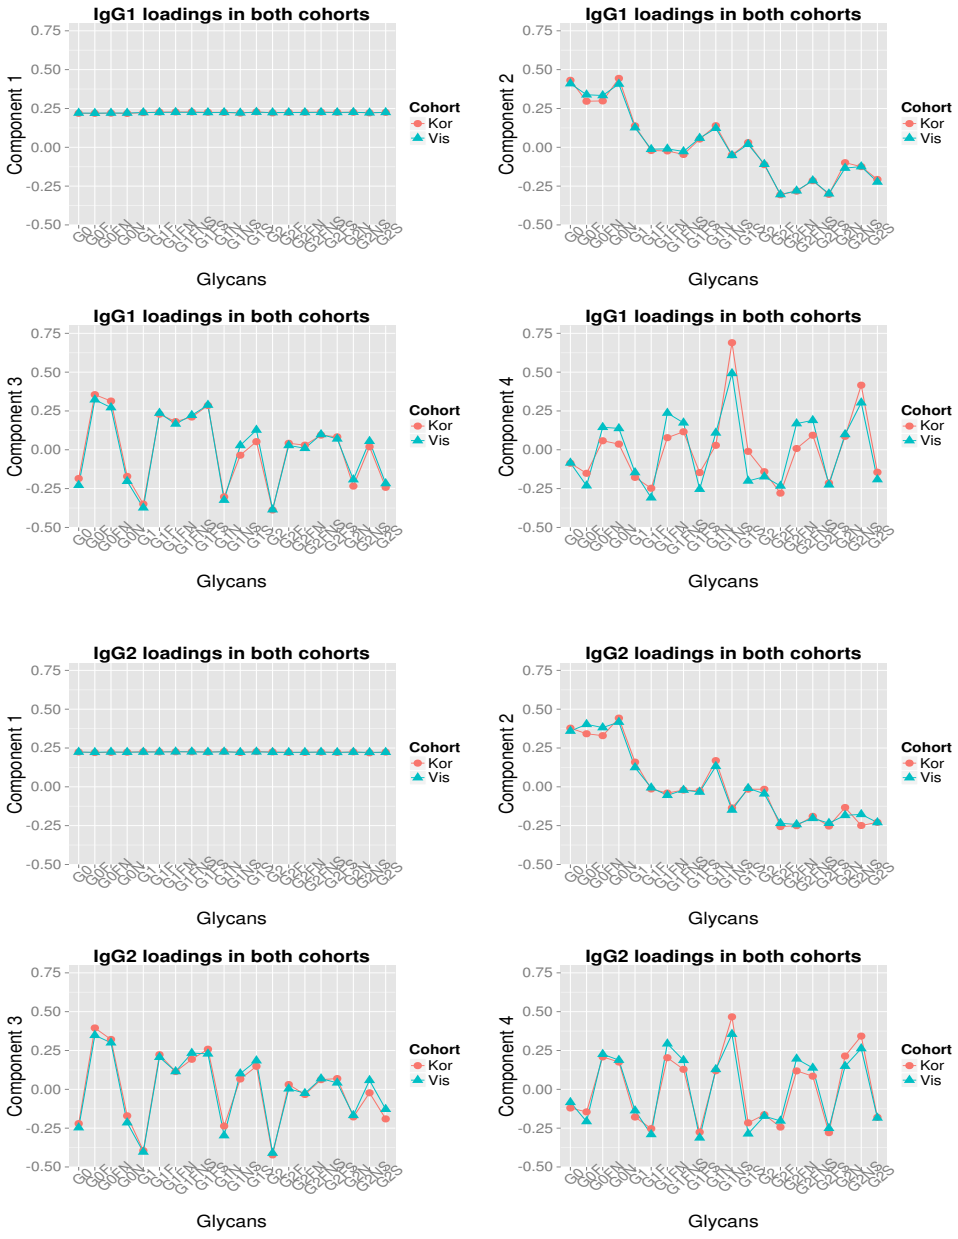

Figure 4.5: **Loadings per component for both cohorts.** In the top four plots loading values of IgG1 glycans (W) are plotted per glycan. The red dots connected by red lines are for the CROATIA_Korcula loadings. The four loading vectors are plotted left-to-right and top-to-bottom. The blue triangles and lines are for the CROATIA_Vis cohorts. In the bottom four plots the IgG2 glycan loading values are plotted in the same order and style.

weighted least squares to combine data from different studies with potentially different covariates [11]. An alternative method, when access to data is possible, is to estimate joint parts between the data sets and the studies simultaneously. This yields a joint space with variables that have high loading values in most studies. For example, in [25], a non-probabilistic approach is pursued in a least squares estimation method using PCA. Performing data integration across studies, while taking into account uncertainties within each study, is one of our topics for future research, and will lead to more powerful analysis of the IgG glycans across cohorts.

To assess the statistical significance of loadings, the probabilistic framework provides alternative approaches to jackknifing and bootstrapping [27]. The observed Fisher information matrix can be used to estimate standard errors for individual loading parameters. For small sample sizes, bootstrap approaches appears to better reflect the uncertainty of the parameters. For large enough sample sizes, the asymptotic standard errors are close to the simulation-based standard errors. Typically, in epidemiological studies, the sample size is large enough to use asymptotic standard errors.

In this paper we ignored the fact that different biological 'omics' measurements have different error structures. An extension of Partial Least Squares was proposed to correct for systematic variation (variation induced by latent variables uncorrelated to the other data set) in the data sets, named Two-Way Orthogonal PLS (O2PLS) [8, 24]. Such an extension can be pursued for PPLS by adding for both X and Y in (4.1) a set of independent latent variables multiplied by their loading parameters. We are currently working on exploring the possibilities of a Probabilistic O2PLS for heterogeneous data sets.

4.7 Appendices for Chapter 4

Appendix A. Variances and covariances

The covariance matrix of (x, y) is given in (4.4). First note that $\mathrm{Var}(u) = \mathrm{Var}(tB + h) = B^2\Sigma_t + \sigma_h^2 I_r$, then compute

$$\begin{aligned}
\mathrm{Var}(x) &= \mathrm{Var}\left(tW^{\mathrm{T}} + e\right) = W\mathrm{Var}(t)W^{\mathrm{T}} + \mathrm{Var}(e) = W\Sigma_t W^{\mathrm{T}} + \sigma_e^2 I_p,\\
\mathrm{Var}(y) &= \mathrm{Var}\left(UC^{\mathrm{T}} + f\right) = C\mathrm{Var}(u)C^{\mathrm{T}} + \mathrm{Var}(f) = C(B^2\Sigma_t + \sigma_h^2 I_r)C^{\mathrm{T}} + \sigma_f^2 I_q,\\
\mathrm{Cov}(x, y) &= \mathrm{Cov}\left(tW^{\mathrm{T}} + e, uC^{\mathrm{T}} + f\right) = W\mathrm{Cov}(t, tB)C^{\mathrm{T}} = WB\Sigma_t C^{\mathrm{T}}.
\end{aligned}$$

The covariances between the observed and latent variables are as follows

$$\begin{aligned}
\mathrm{Cov}(x, t) &= \mathrm{Cov}\left(tW^{\mathrm{T}} + e, t\right) = W\mathrm{Var}(t) = W\Sigma_t;\\
\mathrm{Cov}(x, u) &= \mathrm{Cov}\left(tW^{\mathrm{T}} + e, tB + h\right) = W\mathrm{Var}(t)B = W\Sigma_t B;\\
\mathrm{Cov}(y, t) &= \mathrm{Cov}\left(uC^{\mathrm{T}} + f, t\right) = C\mathrm{Cov}(tB + h, t) = C\Sigma_t B;\\
\mathrm{Cov}(y, u) &= \mathrm{Cov}\left(uC^{\mathrm{T}} + f, u\right) = C\mathrm{Cov}(tB + h, tB + h) = C\left(\Sigma_t B^2 + \sigma_h^2 I_r\right).
\end{aligned}$$

See, e.g., [21] for more details.

Appendix B. Identifiability of PPLS

For establishing identifiability of the PPLS model, we need to prove that if the distribution of (x, y) is given, there is only one corresponding set of parameters yielding this distribution. Since (x, y) follows a zero mean normal distribution, identifiability is equivalent to

$$\Sigma = \tilde{\Sigma} \quad \Leftrightarrow \quad \theta = \tilde{\theta},$$

where $\Sigma, \tilde{\Sigma}$ is defined, through $\theta, \tilde{\theta}$, in (4.4). The following lemma will be very useful in establishing identifiability.

Lemma 4.7.1. *Singular Value Decomposition. Let $W, \tilde{W}$ be $p \times r$ and $C, \tilde{C}$ be $q \times r$, all orthogonal matrices. Let $D, \tilde{D}$ be $r \times r$ diagonal with r distinct positive elements on the diagonal. Then $WDC^{\mathrm{T}} = \tilde{W}\tilde{D}\tilde{C}^{\mathrm{T}}$ (B.1) implies $W = \tilde{W}\Delta$, $C = \tilde{C}\Delta$ for some diagonal matrix Δ of size $r \times r$ with on the diagonal elements $\delta_i \in \{-1, 1\}$ and $D = \tilde{D}$.*

Proof. Let $A_1 = WDC^{\mathrm{T}}$ and $A_2 = \tilde{W}\tilde{D}\tilde{C}^{\mathrm{T}}$. Consider $A_i A_i^{\mathrm{T}}$ and $A_i^{\mathrm{T}} A_i$, $i \in \{1, 2\}$. The assertion (B.1) then implies the following.

$$A_1 A_1^{\mathrm{T}} = WD^2W^{\mathrm{T}} = \tilde{W}\tilde{D}^2\tilde{W}^{\mathrm{T}} = A_2 A_2^{\mathrm{T}}; \quad A_1^{\mathrm{T}} A_1 = CD^2C^{\mathrm{T}} = \tilde{C}\tilde{D}^2\tilde{C}^{\mathrm{T}} = A_2^{\mathrm{T}} A_2.$$

Note that both WD^2W^{T} and $\tilde{W}\tilde{D}^2\tilde{W}^{\mathrm{T}}$ are eigenvalue decompositions, as D^2 and $\tilde{D}^2$ are diagonal and $W, \tilde{W}$ and $C, \tilde{C}$ are orthogonal. The spectral theorem for matrices [7] then implies that whenever the elements in $D^2, \tilde{D}^2$ are distinct, the corresponding columns in $W, \tilde{W}$ and $C, \tilde{C}$ are equal up to multiplication with the same sign. We thus get $W = \tilde{W}\Delta$, $C = \tilde{C}\Delta$ and $D = \tilde{D}$. □

Using this Lemma, we show identifiability of the off-diagonal block part of the covariance matrix as given in (4.4).

Lemma 4.7.2. *If for matrices $W, \tilde{W}$, $C, \tilde{C}$ and diagonal $B, \tilde{B}$ and $\Sigma_t, \tilde{\Sigma}_t$, given as in the PPLS model, $W\Sigma_t BC^{\mathrm{T}} = \tilde{W}\tilde{\Sigma}_t\tilde{B}\tilde{C}^{\mathrm{T}}$, then $W = \tilde{W}\Delta$, $C = \tilde{C}\Delta$ and $\Sigma_t B = \tilde{\Sigma}_t\tilde{B}$.*

Proof. Applying Lemma 4.7.1 with $D = \Sigma_t B$ and $\tilde{D} = \tilde{\Sigma}_t\tilde{B}$ gives the desired result, since $\Sigma_t B$ and $\tilde{\Sigma}_t\tilde{B}$ are diagonal matrices with distinct ordered elements. □

Given $\Sigma_{x,y}$ we can identify W and C up to sign and the product $\Sigma_t B$. We now show that in particular also the individual parameters Σ_t and B are identified from the upper diagonal block matrix Σ_x.

Lemma 4.7.3. *If for matrix W, diagonal matrices Σ_t and $\tilde{\Sigma}_t$ and positive numbers $\sigma_e^2, \tilde{\sigma}_e^2$, given as in the PPLS model, $W\Sigma_t W^{\mathrm{T}} + \sigma_e^2 I_p = W\tilde{\Sigma}_t W^{\mathrm{T}} + \tilde{\sigma}_e^2 I_p$ (B.2), then $\sigma_e = \tilde{\sigma}_e$ and $\Sigma_t = \tilde{\Sigma}_t$.*

Proof. Suppose (B.2) holds. Since $r < p$ and $p > 1$, one can find a unit vector $w_\perp$ such that $W^{\mathrm{T}} w_\perp = 0$. Multiplying with such vector yields $\sigma_e^2 w_\perp = \tilde{\sigma}_e^2 w_\perp$. Multiplying again with $w_\perp^{\mathrm{T}}$ yields $\sigma_e^2 = \tilde{\sigma}_e^2$. It follows that we can identify σ_e^2. We can now reduce (B.2) to $W\Sigma_t W^{\mathrm{T}} = W\tilde{\Sigma}_t W^{\mathrm{T}}$. Pre-multiplying with W^{T} and post-multiplying with W on both sides yields $\Sigma_t = \tilde{\Sigma}_t$. □

We have seen in Theorem 4.7.2 that we can identify $\Sigma_t B$. Since we identified Σ_t we get identifiability of B. The remaining parameters σ_h^2 and σ_f^2 are now shown to be identified using the lower block diagonal Σ_y.

Lemma 4.7.4. *If for matrices C, B, Σ_t, $\sigma_f^2, \tilde{\sigma}_f^2$ and $\sigma_h^2, \tilde{\sigma}_h^2$, given as in the PPLS model, the assertion $\Sigma_y = \tilde{\Sigma}_y$ holds, i.e.,*

$$C(B^2\Sigma_t + \sigma_h^2 I_r)C^{\mathrm{T}} + \sigma_f^2 I_q = C(B^2\Sigma_t + \tilde{\sigma}_h^2 I_r)C^{\mathrm{T}} + \tilde{\sigma}_f^2 I_q,$$

then $\sigma_f^2 = \tilde{\sigma}_f^2$ and $\sigma_h^2 = \tilde{\sigma}_h^2$.

Proof. In Theorem 4.7.3 take W equal to C, σ_e^2 equal to σ_f^2, $\tilde{\sigma}_e^2$ equal to $\tilde{\sigma}_f^2$, and the diagonal covariance matrices Σ_t and $\tilde{\Sigma}_t$ equal to $\Sigma_t B^2 + \sigma_h^2 I_p$ and $\Sigma_t B^2 + \tilde{\sigma}_h^2 I_p$. We find that we can identify $\Sigma_t B^2 + \sigma_h^2$ and σ_f^2. Since we already identified Σ_t and B, we have also identifiability of σ_h^2. □

We conclude with the proof of Theorem 4.3.1.

Proof. Suppose $\Sigma = \tilde{\Sigma}$. This is true if and only if

$$\Sigma_{x,y} = \tilde{\Sigma}_{x,y}, \quad \Sigma_x = \tilde{\Sigma}_x, \quad \Sigma_y = \tilde{\Sigma}_y. \tag{4.7}$$

Applying Lemma 4.7.2 to the first equation, we identify W and C up to sign. Considering Lemma 4.7.3 together with Lemma 4.7.2, the second equation implies identifiability of Σ_t, B and σ_e. The three Lemmas 4.7.2, 4.7.3 and 4.7.4 together with the last equation imply identifiability of σ_h and σ_f.

□

Appendix C. An Expectation-Maximization algorithm for PPLS

To obtain parameter estimates in the PPLS model, maximum likelihood is used. The EM algorithm is an iterative procedure for maximizing the observed log-likelihood (4.5) and consists of an Expectation step and a Maximization step. The following Lemma is convenient to make the expectation step explicit.

Lemma 4.7.5. *Let the pair (z, x) be jointly multivariate normal row vectors with zero mean and covariance matrix*

$$\begin{pmatrix} \Sigma_z & \Sigma_{z,x} \\ \Sigma_{x,z} & \Sigma_x \end{pmatrix}.$$

Then $z|x$ is normally distributed with conditional mean $\mathbb{E}(z|x) = x\,\Sigma_x^{-1}\,\Sigma_{x,z}$, and conditional covariance matrix $\mathrm{Var}(z|x) = \Sigma_z - \Sigma_{z,x}\,\Sigma_x^{-1}\,\Sigma_{x,z}$. Secondly, if $z = (t, u)$, $\mathrm{Cov}(t, x) = \Sigma_{t,x}$ and $\mathrm{Cov}(x, u) = \Sigma_{x,u}$, then the conditional covariance between t and u is $\mathrm{Cov}(t, u|x) = \mathrm{Cov}(t, u) - \Sigma_{t,x}\,\Sigma_x^{-1}\,\Sigma_{x,u}$.

Proof. The proof for the first part of the Lemma is found in [21]. The second part follows from the off diagonal blocks of $\mathrm{Var}(z|x)$. □

Expectation The conditional first moments can be obtained by applying Lemma 4.7.5 while substituting t or u for z and (x, y) for x.

$$\mu_t = \mathbb{E}\,(t|x, y, \theta) = (x, y)\,\Sigma^{-1}\,\mathrm{Cov}\{(x, y), t\},$$
$$\mu_u = \mathbb{E}\,(u|x, y, \theta) = (x, y)\,\Sigma^{-1}\,\mathrm{Cov}\{(x, y), u\}.$$

The same substitution can be made for the conditional second moments. Using $\mathbb{E}(a^{\mathrm{T}}b|z) = \mathrm{Cov}(a, b|z) + \mathbb{E}(a|z)^{\mathrm{T}}\mathbb{E}(b|z)$, we get

$$\begin{aligned} C_{TT} = \mathbb{E}(t^{\mathrm{T}}t|x, y, \theta) &= I_r - \mathrm{Cov}\{t, (x, y)\}\,\Sigma^{-1}\,\mathrm{Cov}\{(x, y), t\} + \\ &\quad \mathrm{Cov}\{t, (x, y)\}\,\Sigma^{-1}S\Sigma^{-1}\,\mathrm{Cov}\{(x, y), t\}, \\ C_{UU} = \mathbb{E}(u^{\mathrm{T}}u|x, y, \theta) &= I_r - \mathrm{Cov}\{u, (x, y)\}\,\Sigma^{-1}\,\mathrm{Cov}\{(x, y), u\} + \\ &\quad \mathrm{Cov}\{u, (x, y)\}\,\Sigma^{-1}S\Sigma^{-1}\,\mathrm{Cov}\{(x, y), u\}, \end{aligned}$$

where S is the biased sample covariance matrix of (x, y). The conditional cross term equals

$$\begin{aligned} C_{UT} = \mathbb{E}(u^{\mathrm{T}}t|x, y, \theta) &= \Sigma_t B - \mathrm{Cov}\{u, (x, y)\}\,\Sigma^{-1}\,\mathrm{Cov}\{(x, y), t\} + \\ &\quad \mathrm{Cov}\{u, (x, y)\}\,\Sigma^{-1}S\Sigma^{-1}\,\mathrm{Cov}\{(x, y), t\} \end{aligned}$$

The covariances are given by

$$\mathrm{Cov}\{(x, y), t\} = \begin{pmatrix} W\Sigma_t \\ C\Sigma_t B \end{pmatrix}, \quad \mathrm{Cov}\{(x, y), u\} = \begin{pmatrix} W\Sigma_t B \\ C(\Sigma_t B + \sigma_h^2 I_r) \end{pmatrix}.$$

Although the the conditional expectations involve random variables and parameters, in the maximization step the calculated quantities are considered fixed and known.

Maximization The maximization step involves maximizing the complete-data likelihood (4.6), we have seen that it can be decomposed in distinct factors. This allows optimization of the expected complete data likelihood to be split into four submaximizations, given by the individual factors and their respective parameters in the following annotated product:

$$\underbrace{f(x|t)}_{W,\sigma_e}\,\underbrace{f(y|u)}_{C,\sigma_f}\,\underbrace{f(u|t)}_{B,\sigma_h}\,\underbrace{f(t)}_{\Sigma_t}$$

Moreover, it will become apparent that each parameter within each component can be decoupled, yielding a separate maximization per component per parameter. We focus on the part of $f(x|t)$ that depends on W, which is given by

$$\begin{aligned} \mathbb{E}\,\{\ln f(X|T)|X, Y\} &= -\mathbb{E}(||X - TW^{\mathrm{T}}||^2|X, Y) + \text{const.} \\ &= \mathrm{tr}(-X^{\mathrm{T}}X + 2X^{\mathrm{T}}\mu_t W^{\mathrm{T}} - WC_{TT}W^{\mathrm{T}}) + \text{const.} \end{aligned}$$

To take into account the constraints on W, namely $W^{\mathrm{T}}W = I_r$, we introduce a matrix of Lagrange multipliers Λ. We get as objective function

$$\mathrm{tr}(-X^{\mathrm{T}}X + 2X^{\mathrm{T}}\mu_t W^{\mathrm{T}} - WC_{TT}W^{\mathrm{T}}) - \mathrm{tr}\{(W^{\mathrm{T}}W - I_r)\Lambda\}.$$

Differentiating with respect to W yields $2X^{\mathrm{T}}\mu_t - 2WC_{TT} - 2W\Lambda = 2W\left(C_{TT}+\Lambda\right) - 2X^{\mathrm{T}}\mu_t$. One may choose Λ so that $C_{TT}+\Lambda$ is invertible. In a maximum W then satisfies $W = X^{\mathrm{T}}\mu_t\left(C_{TT}+\Lambda\right)^{-1}$. We want to find a Λ such that the constraint holds, i.e.,

$$I_r = W^{\mathrm{T}}W = \{(C_{TT}+\Lambda)^{-1}\}^{\mathrm{T}}\mu_t^{\mathrm{T}}XX^{\mathrm{T}}\mu_t\left(C_{TT}+\Lambda\right)^{-1},$$
$$\mu_t^{\mathrm{T}}XX^{\mathrm{T}}\mu_t = \left(C_{TT}+\Lambda\right)^{\mathrm{T}}\left(C_{TT}+\Lambda\right).$$

The last identity can be recognized as a Cholesky or Eigenvalue decomposition.

$$\mu_t^{\mathrm{T}}XX^{\mathrm{T}}\mu_t = \left(C_{TT}+\Lambda\right)^{\mathrm{T}}\left(C_{TT}+\Lambda\right) = L_tL_t^{\mathrm{T}}$$

with L_t the lower triangular matrix of a Cholesky decomposition of $\mu_t^{\mathrm{T}}XX^{\mathrm{T}}\mu_t$. Note that L_t exists, since $\mu_t^{\mathrm{T}}XX^{\mathrm{T}}\mu_t$ is always positive semi-definite. Choosing $\Lambda = L_t^{\mathrm{T}} - C_{TT}$, we get as update $W = X^{\mathrm{T}}\mu_t(L_t^{\mathrm{T}})^{-1}$. Following the same reasoning, we obtain for the $f(Y|U)$ part $C = Y^{\mathrm{T}}\mu_u(L_u^{\mathrm{T}})^{-1}$, where L_u is the lower triangular matrix from the Cholesky decomposition of $\mu_u^{\mathrm{T}}YY^{\mathrm{T}}\mu_u$.

The parameter B involves maximizing $\ln f(U|T)$, which is given by

$$-||U-TB||^2 = -\mathrm{tr}(U^{\mathrm{T}}U - 2U^{\mathrm{T}}TB + BT^{\mathrm{T}}TB) + \mathrm{const.}$$

Taking the conditional expectation with respect to (x,y) yields $-\mathrm{tr}\mathbb{E}(U^{\mathrm{T}}U - 2U^{\mathrm{T}}TB + BT^{\mathrm{T}}TB|X,Y)$. Differentiating with respect to B and equating to the zero matrix yields

$$B\mathbb{E}(T^{\mathrm{T}}T|X,Y) = \mathbb{E}(U^{\mathrm{T}}T|X,Y) \quad B = \mathbb{E}(U^{\mathrm{T}}T|X,Y)\{\mathbb{E}(T^{\mathrm{T}}T|X,Y)\}^{-1}$$

To incorporate the constraint that B should be diagonal, we set the diagonal elements to zero, yielding

$$B = \mathbb{E}(U^{\mathrm{T}}T|X,Y)\{\mathbb{E}(T^{\mathrm{T}}T|X,Y)\}^{-1} \circ I_r,$$

with $\circ$ the elementwise (Hadamart) product operator.

For the covariance matrix of Σ_t, we consider $\ln f(T)$ which is given by

$$2\ln f(T) = \mathrm{const.} - N\ln|\Sigma_T| - \mathrm{tr}(T^{\mathrm{T}}T\Sigma_T^{-1}) = \mathrm{const.} + N\ln|\Sigma_T^{-1}| - \mathrm{tr}(T^{\mathrm{T}}T\Sigma_T^{-1}).$$

After taking the conditional expectation of the last expression, it can be differentiated with respect to Σ_t^{-1}, which yields

$$2\frac{\partial}{\partial\,\Sigma_T^{-1}}\ln f(T) = N\Sigma_T - \mathbb{E}(T^{\mathrm{T}}T|x,y) = 0, \quad \Sigma_T = N^{-1}\mathbb{E}(T^{\mathrm{T}}T|x,y)\circ I_r$$

The last Hadamart product ensures Σ_t is diagonal.

To maximize over σ_e^2, we consider $\ln f(X|T)$ and note that $E = X - TW^{\mathrm{T}}$. Then $\ln f(X|T)$ is given by

$$2\ln f(X|T) = \mathrm{const.} - Np\ln|\sigma_e^2| - \sigma_e^{-2}\mathrm{tr}(E^{\mathrm{T}}E) = \mathrm{const.} + Np\ln\sigma_e^{-2} - \sigma_e^{-2}\mathrm{tr}(E^{\mathrm{T}}E)$$

After taking the conditional expectation of the last expression, we differentiate it with respect to σ_e^{-2}, yielding

$$2\frac{\partial}{\partial\,\sigma_e^{-1}}\ln f(X|T) = Np\sigma_e^2 - \mathbb{E}(E^{\mathrm{T}}E|X,Y) = 0, \quad \sigma_e^2 = (Np)^{-1}\mathbb{E}(E^{\mathrm{T}}E|X,Y)$$

The same derivation can be applied to $\ln f(y|u)$ and $\ln f(u|t)$ to find

$$\sigma_f^2 = (Nq)^{-1}\mathbb{E}(F^{\mathrm{T}}F|X,Y), \quad \sigma_h^2 = (Nr)^{-1}\mathbb{E}(H^{\mathrm{T}}H|X,Y)$$

Appendix D. Asymptotic standard errors for PPLS loadings

Using notation as in [16] we define

$$\lambda(W_k) = -\frac{1}{2\sigma_e^2}\mathrm{tr}(X^{\mathrm{T}}X - 2X^{\mathrm{T}}t_k w_k^{\mathrm{T}} + w_k t_k^{\mathrm{T}} t_k w_k^{\mathrm{T}})$$

to be the part of the log likelihood depending on w_k. We calculate the following first and second derivatives.

$$S(w_k) = \nabla\lambda = \sigma_e^{-2}\left(X^{\mathrm{T}}t_k - w_k t_k^{\mathrm{T}} t_k\right), \quad B(w_k) = -\nabla^2\lambda = \sigma_e^{-2}\left(t_k^{\mathrm{T}}t_k\right) I_p.$$

We obtain

$$\begin{aligned}
\sigma_e^4 S(w_k)S(w_k)^{\mathrm{T}} &= X^{\mathrm{T}}t_k t_k^{\mathrm{T}}X - 2X^{\mathrm{T}}t_k t_k^{\mathrm{T}}t_k w_k^{\mathrm{T}} + w_k t_k^{\mathrm{T}}t_k t_k^{\mathrm{T}}t_k w_k^{\mathrm{T}},\\
\sigma_e^4\mathbb{E}\left\{S(w_k)S(w_k)^{\mathrm{T}}|X,Y\right\} &= X^{\mathrm{T}}\mathbb{E}\left(t_k t_k^{\mathrm{T}}|X,Y\right)X - 2X^{\mathrm{T}}\mathbb{E}\left(t_k t_k^{\mathrm{T}}t_k|X,Y\right)w_k^{\mathrm{T}} +\\
&\quad w_k\mathbb{E}\left(t_k^{\mathrm{T}}t_k t_k^{\mathrm{T}}t_k|X,Y\right)w_k^{\mathrm{T}}\\
&= \sigma_k^2 X^{\mathrm{T}}X - 2X^{\mathrm{T}}\left(\mu_k||\mu_k||_2^2 + 3\mu_k\sigma_k^2\right)w_k^{\mathrm{T}} +\\
&\quad w_k\left(||\mu_k||_2^4 + 6||\mu_k||_2^2\sigma_k^2 + 3\sigma_k^4\right)w_k^{\mathrm{T}}.
\end{aligned}$$

Here $\mu_k = \mathbb{E}(t_k|X,Y)$ and $\sigma_k = \mathbb{E}(t_k^{\mathrm{T}}t_k|X,Y)$. For explicit expressions of these expectations, see Appendix 4.7. For the second derivative we get $\mathbb{E}\{B(w_k)|X,Y\} = \sigma_k^2 I_p/\sigma_e^2$. The observed Fisher information is now

$$I_{obs} = \mathbb{E}\{B(w_k)|X,Y\} - \mathbb{E}\{S(w_k)S(w_k)^{\mathrm{T}}|X,Y\},$$

and the asymptotic covariance matrix of w_k is $-I_{obs}^{-1}$. The square root of the diagonal elements are the standard errors of the corresponding loading elements.

4.8 Supplementary material for Chapter 4

Introduction

This supplement provides results, figures and tables referred to in the main article "*Probabilistic Partial Least Squares: identifiability, estimation and application*". In the main article some of the results are already discussed, in particular regarding the low-dimensional ($p = q = 20$) case with normally distributed variables. Here, figures and tables corresponding to all other scenarios are presented.

The rest of the supplement is organized as follows. In Section 4.8 figures for true and estimated parameter values are presented. The discussion of these figures can be found in the main article. In Section 4.8 a comparison is made between PPLS and PLS in the cases in which the data are very high-dimensional ($p = q = 10^4$ while $N = 50, 500$) and contain 50% noise. In Section 4.8 the performance of PPLS is compared to PLS, PCCA and PPCA. Here, the inner product between the true and estimated loading vectors is taken as performance indicator. An additional scenario was included in which the covariance between X and Y is zero. Only the normal distribution is considered in these simulations. In Section 4.8 the behavior of asymptotic standard errors of PPLS estimates is investigated and compared to a bootstrap approach. In

this simulation study an additional case is added in which the sample size was 'extra large' ($N = 5000$). We considered the scenarios in which the data were normally distributed and low-dimensional ($p = q = 20$).

Section 4.8 contains tables for Section 4.8. Section 4.8 contains figures for Section 4.8. Section 4.8 contains figures for Section 4.8. Section 4.8 contains figures for Section 4.8.

Extended simulation study

Simulation study as in main article

The results of this simulation study are presented in the main article. We refer to Figures 4.6-4.15 for the results.

Extra high dimensional scenario

A simulation study is conducted to evaluate the performance of the PPLS estimators. To compare PPLS to PLS, scenarios in which PLS is applied are chosen: both X and Y have $p = q = 10^4$ variables and the noise level is set to 50% of the total variation. All four distributions (normal, student-t Poisson and Binomial) and both large ($N = 500$) and small ($N = 50$) sample sizes are considered in the simulation study. A summary of these scenario's is presented in Table 4.3. All parameters are defined as in the simulation study in the main article. Results of all scenarios are

Sample size	$N = (50,\ 500)$
Dimensionality	$p = q = 10^4$
Noise level	$\alpha_n = 0.5$
Distribution	$\{\mathcal{N}(0,1),\ t_2,\ Pois(1),\ Bin(2, 0.3)\}$

Table 4.3: **Overview of the simulation scenarios.** The noise level is defined as the proportion of variation in the noise matrices E, F and H relative to the total variation in X, Y and U respectively. The last row contains the distributions from which the latent variables t, u, e, f and h were generated.

shown in Figures 4.18-4.23.

Comparing the estimates for the **first** loading component W_1, a better performance of PPLS is observed compared to PLS in terms of bias. The biases of the PPLS estimators are about the same or less than the biases of the PLS estimators, regardless of distribution or sample size. For higher absolute loading values, larger biases are observed. The variances of the PPLS estimators are similar to the variances of the PLS estimators. For the loading component C_1 and their PLS and PPLS estimators, the same observations are made.

For the **second** loading component W_2 similar conclusions can be drawn. The biases of the PPLS estimators are about the same or less than the biases of the PLS estimators regardless of distribution or sample size. For higher absolute loading values larger biases are observed. The variance of the PPLS estimates are less whenever the loadings are near zero. For the variables with high positive loading values in the third component, the variance of PPLS estimates for these loading values in the second

component is slightly higher than the variance of the PLS estimates. For the loading component C_2 and their PLS and PPLS estimators, the same observations are made.

In the **third** component, the same conclusions hold, except for the variables with high positive loading values: they v estimated with less bias and variance by PPLS. The difference between PLS and PPLS in bias and variance is even more prominent for the Y loadings C_3. These observations are similar across all four distributions.

The run-time of the PLS algorithm is less than the run-time of the PPLS algorithm. The median ratios of the PPLS run-time to the PLS run-time in each scenario is between 605 and 714 for small sample size and between 1137 and 1257 for large sample size.

Comparison to PCCA and PPCA

A simulation study is performed to compare PPLS, PLS, PPCA and PCCA. A 'null' scenario is included in which the data are not related, i.e. $B = 0$. However the latent structure TW^{T} and UC^{T} underlying each dataset is still present. Here PPCA is applied to each dataset separately to estimate W and C. In the scenarios for which $B \neq 0$ we did not consider PPCA, as it cannot estimate relationships between X and Y (which is fundamental to our research question). In the high dimensional case $p = q = 1000$, PCCA estimates cannot be obtained, as it requires $N > \max(p, q)$.

For each scenario, 1000 replicates are performed. For each method (PLS, PPLS, PPCA and PCCA) median inner products between the estimated and true loading vectors together with the Median Absolute Deviations (MAD) are calculated and shown in Table 4.4-4.6. For the null scenarios both PLS and PCCA fail to accurately recover the loading values (median inner products were between 0.449 and 0.772). Moreover, as PCCA breaks down if $\max(p, q) > N$, no comparison was possible in the high dimensional case. The PPLS and PPCA methods are both accurate in estimating the true loadings (median inner products were between 0.786 and 0.984). In addition, PPLS inner products are nearly as high as PPCA inner products. For the scenario in which T and U have non-zero correlation, both PLS and PPLS have good performance (median inner products were above 0.748). In this scenario, PPLS performs better. The PCCA method performs slightly worse than PLS (median inner product above 0.842), except for the high noise and small sample size scenario. Here the median inner product lies between 0.469 and 0.687 for the W loadings.

For the loading vectors in C similar observations are made. The median inner products are in general smaller than the median inner products with W, probably due to the extra noise component H in Y.

Asymptotic and bootstrap standard errors

To assess variability of PLS estimates, a bootstrap approach is often used. This approach is also available for PPLS. Additionally, asymptotic standard errors for the loading estimates can be calculated based on the observed Fisher information. To compare asymptotic and bootstrap standard errors, a dataset is simulated for several scenarios. In all cases the data are normally distributed and low dimensional ($p = q = 20$). We compare the standard errors of the elements of W and C under low and high noise levels ($\alpha_n = 0.1$ resp. $\alpha_n = 0.5$), and small, large and 'extra large' sample sizes ($N = 50$, $N = 500$ and $N = 5000$ respectively). Also simulation-based

standard errors (based on standard deviations of the 1000 replicates) are included as reference. The results are shown in Figures 4.24 and 4.25.

In all scenarios the asymptotic standard errors (SEs) are smaller than the bootstrap SEs for nearly all loading elements. In particular for high loading values the difference between asymptotic SEs and bootstrap SEs tends to be large. The bootstrap SEs are close to the simulation-based standard errors. All SEs tend to increase with increasing noise level, while a decrease is observed with increasing sample size. In the 'extra large' sample size the bootstrap SEs and asymptotic SEs have similar magnitude.

Tables

Table 4.4: **Median (MAD) inner product between true and estimated loading profiles.** Results for scenarios with zero correlation between T and U are shown, therefore PPCA was not included in the comparison. Data were simulated from a normal distribution and were low-dimensional ($p = q = 20$).

Method	Component W_1	Component W_2	Component W_3
PLS	0.771 (0.139)	0.755 (0.142)	0.713 (0.144)
PPLS	0.869 (0.087)	0.810 (0.135)	0.864 (0.093)
PPCA	0.875 (0.076)	0.842 (0.130)	0.872 (0.085)
PCCA	0.718 (0.168)	0.703 (0.161)	0.657 (0.197)

(a) Low noise level ($\alpha_n = 0.1$), small sample size ($N = 50$).

Method	Component W_1	Component W_2	Component W_3
PLS	0.772 (0.136)	0.753 (0.135)	0.726 (0.147)
PPLS	0.981 (0.019)	0.963 (0.035)	0.978 (0.023)
PPCA	0.982 (0.018)	0.968 (0.030)	0.981 (0.019)
PCCA	0.718 (0.158)	0.690 (0.177)	0.668 (0.171)

(b) Low noise level ($\alpha_n = 0.1$), large sample size ($N = 500$).

Method	Component W_1	Component W_2	Component W_3
PLS	0.680 (0.156)	0.641 (0.155)	0.585 (0.187)
PPLS	0.836 (0.111)	0.787 (0.110)	0.819 (0.107)
PPCA	0.840 (0.106)	0.786 (0.122)	0.820 (0.100)
PCCA	0.529 (0.194)	0.501 (0.201)	0.464 (0.213)

(c) High noise level ($\alpha_n = 0.5$), small sample size ($N = 50$).

Method	Component W_1	Component W_2	Component W_3
PLS	0.684 (0.151)	0.646 (0.160)	0.600 (0.171)
PPLS	0.970 (0.029)	0.947 (0.047)	0.966 (0.030)
PPCA	0.971 (0.027)	0.951 (0.044)	0.969 (0.026)
PCCA	0.535 (0.202)	0.494 (0.210)	0.449 (0.199)

(d) High noise level ($\alpha_n = 0.5$), large sample size ($N = 500$).

Table 4.5: **Median (MAD) inner product between true and estimated loading profiles.** Results for scenarios with zero correlation between T and U are shown. Data were simulated from a normal distribution and were high-dimensional ($p = q = 1000$). PCCA results are not available as it cannot cope with $\max(p, q) > N$ situations.

Method	Component W_1	Component W_2	Component W_3
PLS	0.762 (0.142)	0.726 (0.139)	0.699 (0.156)
PPLS	0.883 (0.078)	0.836 (0.126)	0.866 (0.076)
PPCA	0.881 (0.074)	0.842 (0.126)	0.870 (0.080)
PCCA	-	-	-

(a) Low noise level ($\alpha_n = 0.1$), small sample size ($N = 50$).

Method	Component W_1	Component W_2	Component W_3
PLS	0.760 (0.137)	0.721 (0.152)	0.706 (0.150)
PPLS	0.984 (0.017)	0.973 (0.027)	0.980 (0.020)
PPCA	0.984 (0.017)	0.972 (0.026)	0.982 (0.018)
PCCA	-	-	-

(b) Low noise level ($\alpha_n = 0.1$), large sample size ($N = 500$).

Method	Component W_1	Component W_2	Component W_3
PLS	0.679 (0.127)	0.617 (0.144)	0.548 (0.159)
PPLS	0.852 (0.085)	0.792 (0.125)	0.828 (0.087)
PPCA	0.863 (0.072)	0.822 (0.115)	0.836 (0.081)
PCCA	-	-	-

(c) High noise level ($\alpha_n = 0.5$), small sample size ($N = 50$).

Method	Component W_1	Component W_2	Component W_3
PLS	0.704 (0.127)	0.626 (0.140)	0.558 (0.171)
PPLS	0.981 (0.018)	0.968 (0.031)	0.976 (0.022)
PPCA	0.980 (0.018)	0.966 (0.029)	0.979 (0.017)
PCCA	-	-	-

(d) High noise level ($\alpha_n = 0.5$), large sample size ($N = 500$).

Table 4.6: **Median (MAD) inner product between true and estimated loading profiles.** Results for scenarios with non-zero correlation between T and U are shown. Data were simulated from a normal distribution and were low-dimensional ($p = q = 20$).

Method	Component W_1	Component W_2	Component W_3
PLS	0.964 (0.034)	0.940 (0.053)	0.955 (0.039)
PPLS	0.984 (0.014)	0.960 (0.041)	0.970 (0.029)
PPCA	-	-	-
PCCA	0.913 (0.086)	0.847 (0.121)	0.842 (0.118)

(a) Low noise level ($\alpha_n = 0.1$), small sample size ($N = 50$).

Method	Component W_1	Component W_2	Component W_3
PLS	0.996 (0.004)	0.993 (0.007)	0.995 (0.004)
PPLS	0.999 (0.001)	0.997 (0.003)	0.998 (0.002)
PPCA	-	-	-
PCCA	0.996 (0.004)	0.992 (0.008)	0.995 (0.005)

(b) Low noise level ($\alpha_n = 0.1$), large sample size ($N = 500$).

Method	Component W_1	Component W_2	Component W_3
PLS	0.878 (0.075)	0.784 (0.117)	0.748 (0.115)
PPLS	0.878 (0.089)	0.816 (0.116)	0.853 (0.103)
PPCA	-	-	-
PCCA	0.687 (0.155)	0.584 (0.183)	0.469 (0.218)

(c) High noise level ($\alpha_n = 0.5$), small sample size ($N = 50$).

Method	Component W_1	Component W_2	Component W_3
PLS	0.986 (0.009)	0.971 (0.018)	0.961 (0.018)
PPLS	0.989 (0.008)	0.977 (0.021)	0.983 (0.014)
PPCA	-	-	-
PCCA	0.984 (0.011)	0.961 (0.025)	0.938 (0.033)

(d) High noise level ($\alpha_n = 0.5$), large sample size ($N = 500$).

Table 4.7: **Median (MAD) inner product between true and estimated loading profiles.** Results for scenarios with zero correlation between T and U are shown, therefore PPCA was not included in the comparison. Data were simulated from a normal distribution and were low-dimensional ($p = q = 20$).

Method	Component C_1	Component C_2	Component C_3
PLS	0.770 (0.131)	0.770 (0.135)	0.773 (0.129)
PPLS	0.811 (0.125)	0.818 (0.127)	0.820 (0.125)
PPCA	0.743 (0.096)	0.741 (0.096)	0.737 (0.081)
PCCA	0.712 (0.169)	0.717 (0.168)	0.710 (0.173)

(a) Low noise level ($\alpha_n = 0.1$), small sample size ($N = 50$).

Method	Component C_1	Component C_2	Component C_3
PLS	0.770 (0.131)	0.770 (0.135)	0.773 (0.129)
PPLS	0.811 (0.125)	0.818 (0.127)	0.820 (0.125)
PPCA	0.743 (0.096)	0.741 (0.096)	0.737 (0.081)
PCCA	0.712 (0.169)	0.717 (0.168)	0.710 (0.173)

(b) Low noise level ($\alpha_n = 0.1$), large sample size ($N = 500$).

Method	Component C_1	Component C_2	Component C_3
PLS	0.666 (0.160)	0.670 (0.151)	0.674 (0.152)
PPLS	0.806 (0.116)	0.799 (0.122)	0.804 (0.119)
PPCA	0.750 (0.101)	0.757 (0.100)	0.748 (0.102)
PCCA	0.494 (0.212)	0.527 (0.209)	0.522 (0.213)

(c) High noise level ($\alpha_n = 0.5$), small sample size ($N = 50$).

Method	Component C_1	Component C_2	Component C_3
PLS	0.683 (0.145)	0.687 (0.128)	0.688 (0.144)
PPLS	0.822 (0.120)	0.819 (0.118)	0.828 (0.113)
PPCA	0.776 (0.099)	0.772 (0.097)	0.783 (0.109)
PCCA	0.509 (0.204)	0.520 (0.205)	0.511 (0.214)

(d) High noise level ($\alpha_n = 0.5$), large sample size ($N = 500$).

Table 4.8: **Median (MAD) inner product between true and estimated loading profiles.** Results for scenarios with zero correlation between T and U are shown. Data were simulated from a normal distribution and were high-dimensional ($p = q = 1000$). PCCA results are not available as it cannot cope with $\max(p, q) > N$ situations.

Method	Component C_1	Component C_2	Component C_3
PLS	0.767 (0.122)	0.766 (0.120)	0.767 (0.133)
PPLS	0.721 (0.094)	0.725 (0.099)	0.726 (0.100)
PPCA	0.714 (0.120)	0.725 (0.095)	0.722 (0.099)
PCCA	-	-	-

(a) Low noise level ($\alpha_n = 0.1$), small sample size ($N = 50$).

Method	Component C_1	Component C_2	Component C_3
PLS	0.767 (0.118)	0.767 (0.114)	0.781 (0.120)
PPLS	0.762 (0.076)	0.752 (0.117)	0.758 (0.109)
PPCA	0.721 (0.101)	0.723 (0.099)	0.719 (0.106)
PCCA	-	-	-

(b) Low noise level ($\alpha_n = 0.1$), large sample size ($N = 500$).

Method	Component C_1	Component C_2	Component C_3
PLS	0.704 (0.109)	0.696 (0.110)	0.699 (0.114)
PPLS	0.745 (0.107)	0.747 (0.109)	0.742 (0.112)
PPCA	0.703 (0.091)	0.709 (0.098)	0.706 (0.087)
PCCA	-	-	-

(c) High noise level ($\alpha_n = 0.5$), small sample size ($N = 50$).

Method	Component C_1	Component C_2	Component C_3
PLS	0.724 (0.113)	0.715 (0.109)	0.729 (0.113)
PPLS	0.767 (0.084)	0.746 (0.113)	0.761 (0.104)
PPCA	0.727 (0.106)	0.723 (0.081)	0.722 (0.097)
PCCA	-	-	-

(d) High noise level ($\alpha_n = 0.5$), large sample size ($N = 500$).

Table 4.9: **Median (MAD) inner product between true and estimated loading profiles.** Results for scenarios with non-zero correlation between T and U are shown. Data were simulated from a normal distribution and were low-dimensional ($p = q = 20$).

Method	Component C_1	Component C_2	Component C_3
PLS	0.964 (0.034)	0.940 (0.053)	0.955 (0.039)
PPLS	0.984 (0.014)	0.960 (0.041)	0.970 (0.029)
PPCA	-	-	-
PCCA	0.913 (0.086)	0.847 (0.121)	0.842 (0.118)

(a) Low noise level ($\alpha_n = 0.1$), small sample size ($N = 50$).

Method	Component C_1	Component C_2	Component C_3
PLS	0.996 (0.004)	0.993 (0.007)	0.995 (0.004)
PPLS	0.999 (0.001)	0.997 (0.003)	0.998 (0.002)
PPCA	-	-	-
PCCA	0.996 (0.004)	0.992 (0.008)	0.995 (0.005)

(b) Low noise level ($\alpha_n = 0.1$), large sample size ($N = 500$).

Method	Component C_1	Component C_2	Component C_3
PLS	0.878 (0.075)	0.784 (0.117)	0.748 (0.115)
PPLS	0.878 (0.089)	0.816 (0.116)	0.853 (0.103)
PPCA	-	-	-
PCCA	0.687 (0.155)	0.584 (0.183)	0.469 (0.218)

(c) High noise level ($\alpha_n = 0.5$), small sample size ($N = 50$).

Method	Component C_1	Component C_2	Component C_3
PLS	0.986 (0.009)	0.971 (0.018)	0.961 (0.018)
PPLS	0.989 (0.008)	0.977 (0.021)	0.983 (0.014)
PPCA	-	-	-
PCCA	0.984 (0.011)	0.961 (0.025)	0.938 (0.033)

(d) High noise level ($\alpha_n = 0.5$), large sample size ($N = 500$).

Figures of simulation study in main article

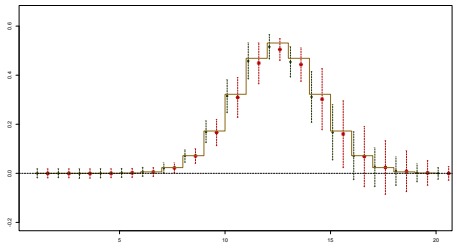

(a) Low noise (10%); small sample size ($N = 50$)

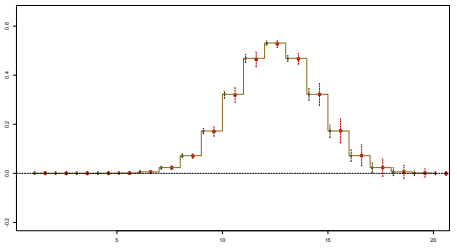

(b) Low noise (10%); large sample size ($N = 500$)

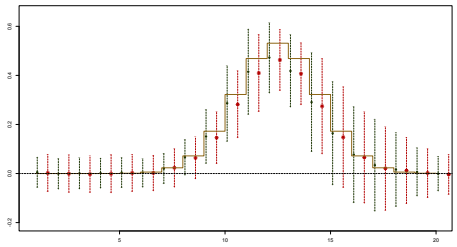

(c) High noise (50%); small sample size ($N = 50$)

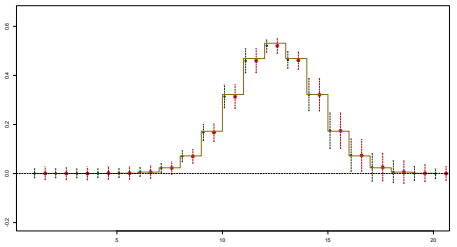

(d) High noise (50%); large sample size ($N = 500$)

Figure 4.6: **True and estimated loadings W_1 over 1000 simulation replications.** The black dots and dashed vertical lines (on the left of each pair) represent PPLS estimates, the red stars and dashed vertical lines (on the right of each pair) represent PLS estimates. The dots and stars are the average loading values across 1000 simulation replications; the width of the dashed lines are twice the standard deviations. The results are for normally distributed latent variables (t, e, f and h) and low dimensionality ($p = q = 20$ variables).

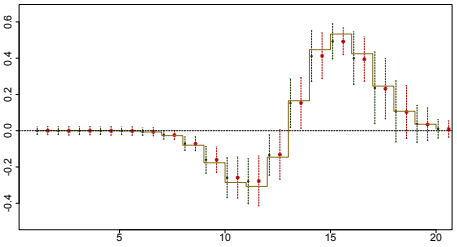

(a) Low noise (10%); small sample size ($N = 50$)

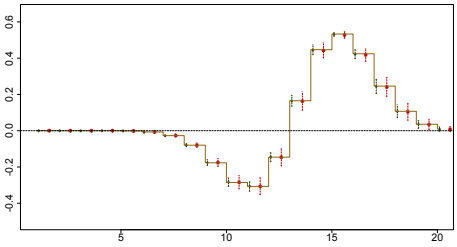

(b) Low noise (10%); large sample size ($N = 500$)

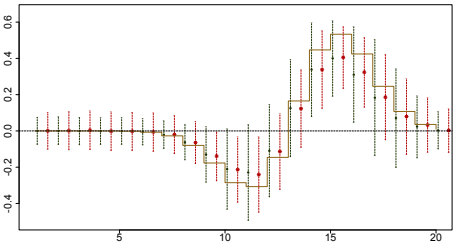

(c) High noise (50%); small sample size ($N = 50$)

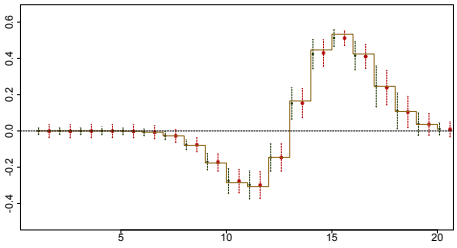

(d) High noise (50%); large sample size ($N = 500$)

Figure 4.7: **Bias and variance of the PPLS and PLS estimates of the loadings W_2 over 1000 simulation replications.** The black dots and dashed vertical lines (on the left of each pair) represent PPLS estimates, the red stars and dashed vertical lines (on the right of each pair) represent the PLS estimates. The dots and stars are the average loading values across 1000 simulation replications; the width of the dashed lines are twice the standard deviations. The results are for normally distributed latent variables (t, e, f and h) and low dimensionality ($p = q = 20$ variables).

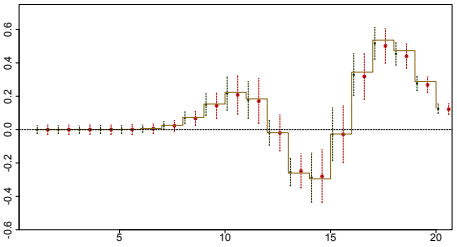

(a) Low noise (10%); small sample size ($N = 50$)

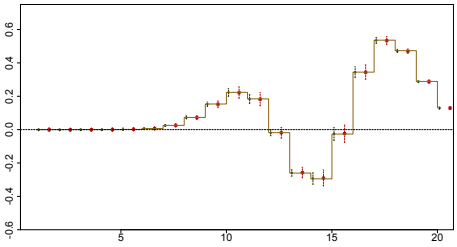

(b) Low noise (10%); large sample size ($N = 500$)

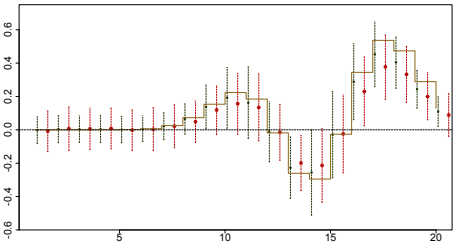

(c) High noise (50%); small sample size ($N = 50$)

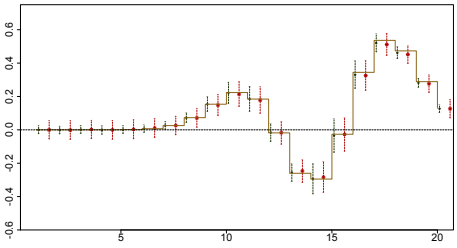

(d) High noise (50%); large sample size ($N = 500$)

Figure 4.8: **True and estimated loadings W_3 over 1000 simulation replications.** The black dots and dashed vertical lines (on the left of each pair) represent PPLS estimates, the red stars and dashed vertical lines (on the right of each pair) represent the PLS estimates. The dots and stars are the average loading values across 1000 simulation replications; the width of the dashed lines are twice the standard deviations. The results are for normally distributed latent variables (t, e, f and h) and low dimensionality ($p = q = 20$ variables).

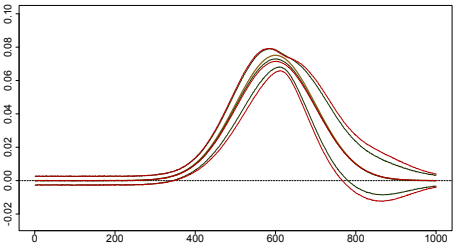

(a) Low noise (10%); small sample size ($N = 50$)

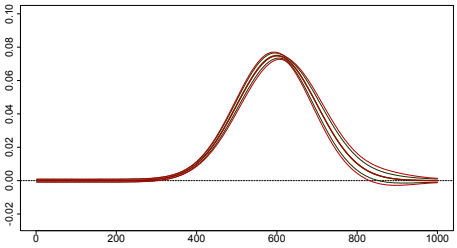

(b) Low noise (10%); large sample size ($N = 500$)

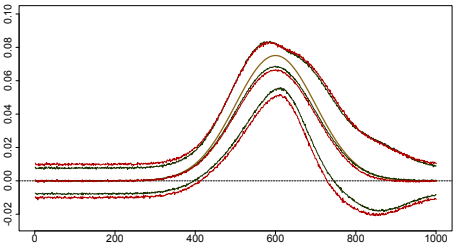

(c) High noise (50%); small sample size ($N = 50$)

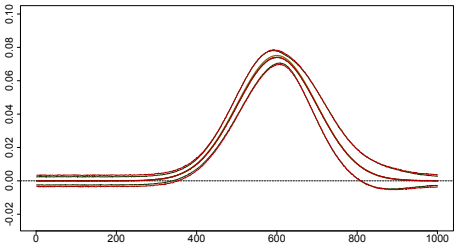

(d) High noise (50%); large sample size ($N = 500$)

Figure 4.9: **True and estimated loadings W_1 over 1000 simulation replications.** The black lines represent PPLS estimates, the red lines represent PLS estimates. The middle lines are the average loading values across 1000 simulation replications; the width of the two outer lines are twice the standard deviations. The results are for normally distributed latent variables (t, e, f and h) and low dimensionality ($p = q = 1000$ variables).

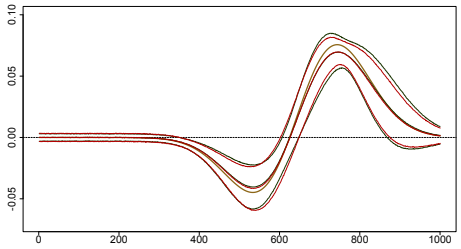

(a) Low noise (10%); small sample size ($N = 50$)

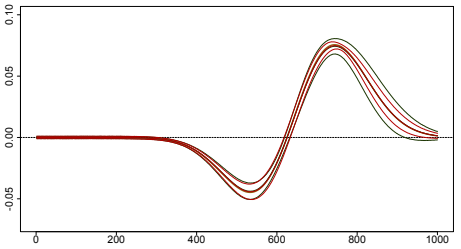

(b) Low noise (10%); large sample size ($N = 500$)

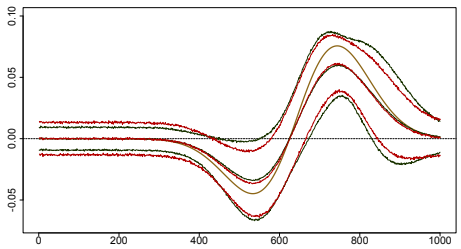

(c) High noise (50%); small sample size ($N = 50$)

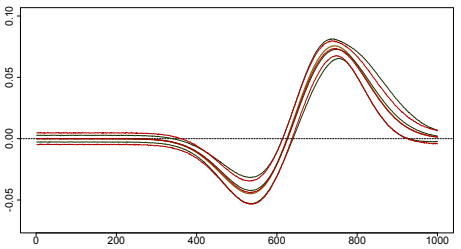

(d) High noise (50%); large sample size ($N = 500$)

Figure 4.10: **True and estimated loadings W_2 over 1000 simulation replications.** The black lines represent PPLS estimates, the red lines represent PLS estimates. The middle lines are the average loading values across 1000 simulation replications; the width of the two outer lines are twice the standard deviations. The results are for normally distributed latent variables (t, e, f and h) and low dimensionality ($p = q = 1000$ variables).

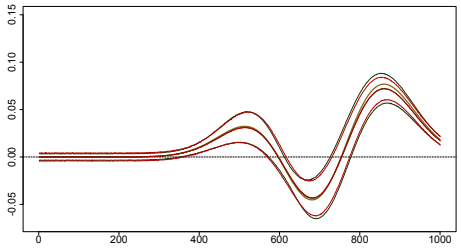

(a) Low noise (10%); small sample size ($N = 50$)

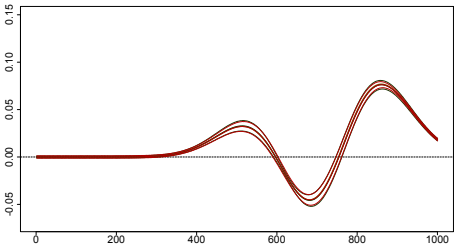

(b) Low noise (10%); large sample size ($N = 500$)

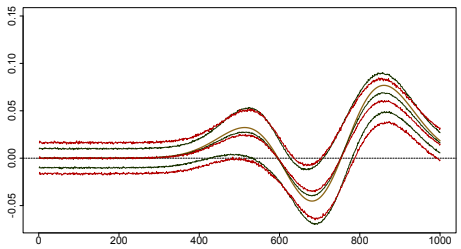

(c) High noise (50%); small sample size ($N = 50$)

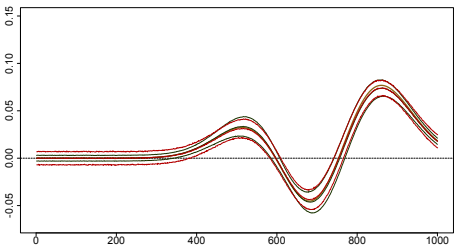

(d) High noise (50%); large sample size ($N = 500$)

Figure 4.11: **True and estimated loadings W_3 over 1000 simulation replications.** The black lines represent PPLS estimates, the red lines represent PLS estimates. The middle lines are the average loading values across 1000 simulation replications; the width of the two outer lines are twice the standard deviations. The results are for normally distributed latent variables (t, e, f and h) and low dimensionality ($p = q = 1000$ variables).

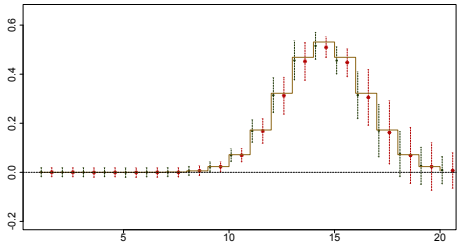

(a) Low noise (10%); small sample size ($N = 50$)

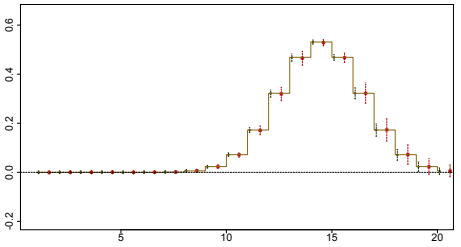

(b) Low noise (10%); large sample size ($N = 500$)

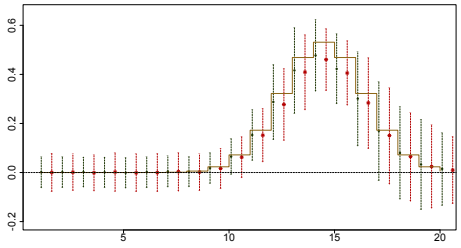

(c) High noise (50%); small sample size ($N = 50$)

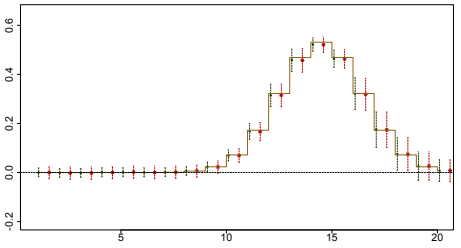

(d) High noise (50%); large sample size ($N = 500$)

Figure 4.12: **True and estimated loadings C_1 over 1000 simulation replications.** The black dots and dashed vertical lines (on the left of each pair) represent PPLS estimates, the red stars and dashed vertical lines (on the right of each pair) represent PLS estimates. The dots and stars are the average loading values across 1000 simulation replications; the width of the dashed lines are twice the standard deviations. The results are for normally distributed latent variables (t, e, f and h) and low dimensionality ($p = q = 20$ variables).

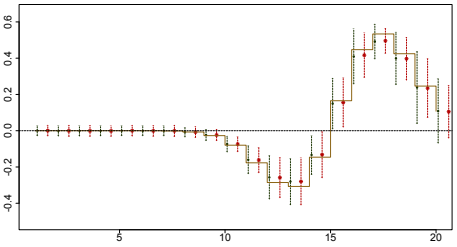

(a) Low noise (10%); small sample size ($N = 50$)

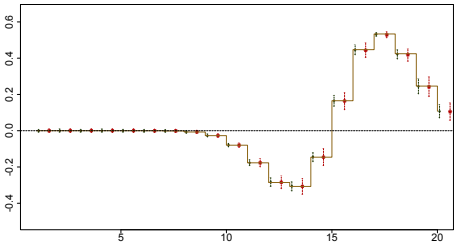

(b) Low noise (10%); large sample size ($N = 500$)

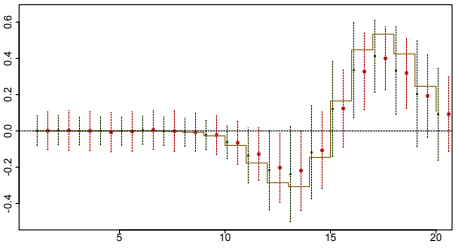

(c) High noise (50%); small sample size ($N = 50$)

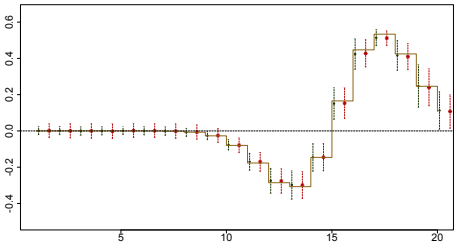

(d) High noise (50%); large sample size ($N = 500$)

Figure 4.13: **True and estimated loadings C_2 over 1000 simulation replications.** The black dots and dashed vertical lines (on the left of each pair) represent PPLS estimates, the red stars and dashed vertical lines (on the right of each pair) represent the PLS estimates. The dots and stars are the average loading values across 1000 simulation replications; the width of the dashed lines are twice the standard deviations. The results are for normally distributed latent variables (t, e, f and h) and low dimensionality ($p = q = 20$ variables).

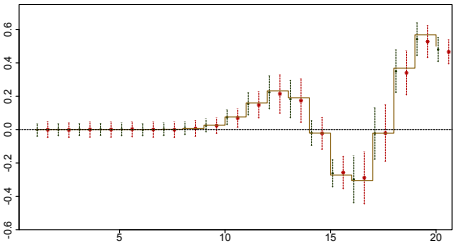

(a) Low noise (10%); small sample size ($N = 50$)

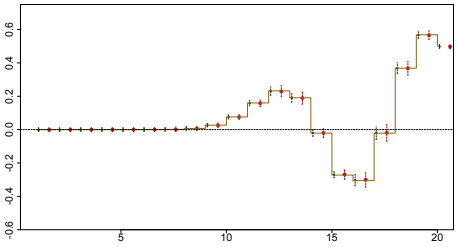

(b) Low noise (10%); large sample size ($N = 500$)

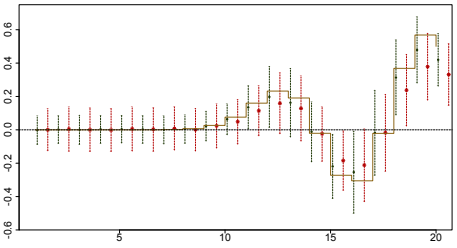

(c) High noise (50%); small sample size ($N = 50$)

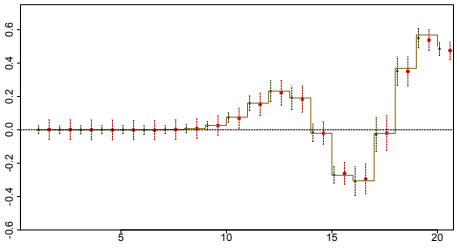

(d) High noise (50%); large sample size ($N = 500$)

Figure 4.14: **True and estimated loadings C_3 over 1000 simulation replications.** The black dots and dashed vertical lines (on the left of each pair) represent PPLS estimates, the red stars and dashed vertical lines (on the right of each pair) represent the PLS estimates. The dots and stars are the average loading values across 1000 simulation replications; the width of the dashed lines are twice the standard deviations. The results are for normally distributed latent variables (t, e, f and h) and low dimensionality ($p = q = 20$ variables).

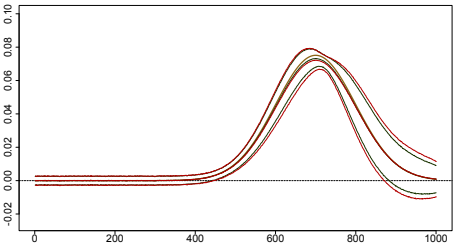

(a) Low noise (10%); small sample size ($N = 50$)

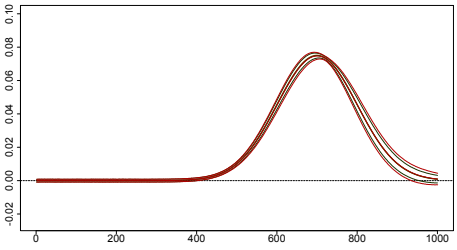

(b) Low noise (10%); large sample size ($N = 500$)

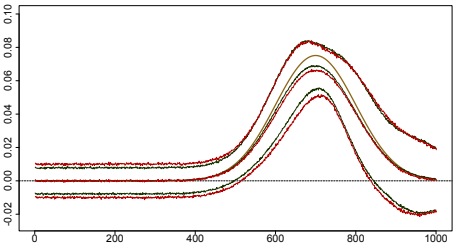

(c) High noise (50%); small sample size ($N = 50$)

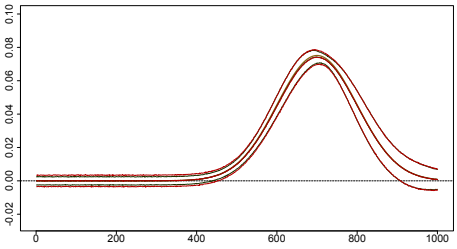

(d) High noise (50%); large sample size ($N = 500$)

Figure 4.15: **True and estimated loadings C_1 over 1000 simulation replications.** The black lines represent PPLS estimates, the red lines represent PLS estimates. The middle lines are the average loading values across 1000 simulation replications; the width of the two outer lines are twice the standard deviations. The results are for normally distributed latent variables (t, e, f and h) and low dimensionality ($p = q = 1000$ variables).

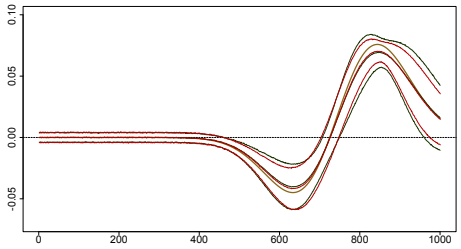

(a) Low noise (10%); small sample size ($N = 50$)

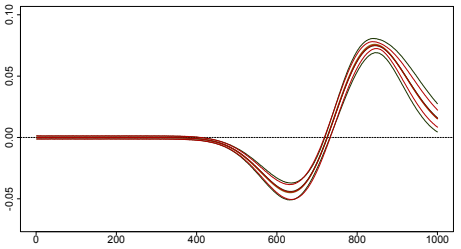

(b) Low noise (10%); large sample size ($N = 500$)

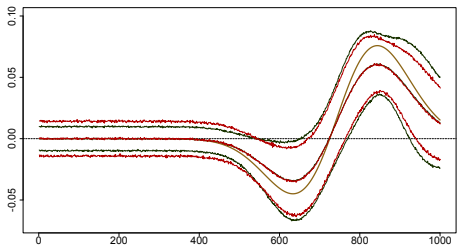

(c) High noise (50%); small sample size ($N = 50$)

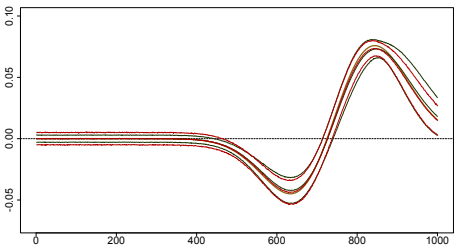

(d) High noise (50%); large sample size ($N = 500$)

Figure 4.16: **True and estimated loadings C_2 over 1000 simulation replications.** The black lines represent PPLS estimates, the red lines represent PLS estimates. The middle lines are the average loading values across 1000 simulation replications; the width of the two outer lines are twice the standard deviations. The results are for normally distributed latent variables (t, e, f and h) and low dimensionality ($p = q = 1000$ variables).

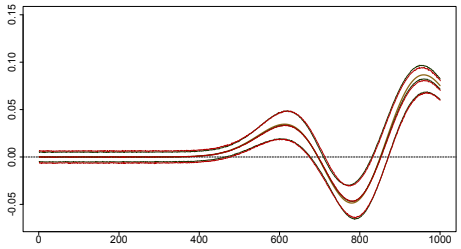

(a) Low noise (10%); small sample size ($N = 50$)

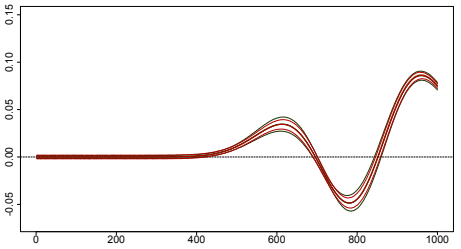

(b) Low noise (10%); large sample size ($N = 500$)

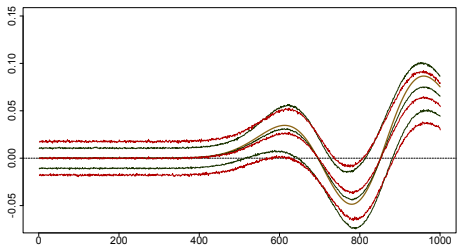

(c) High noise (50%); small sample size ($N = 50$)

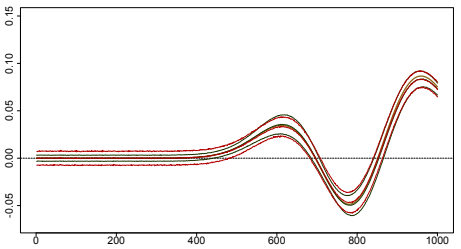

(d) High noise (50%); large sample size ($N = 500$)

Figure 4.17: **True and estimated loadings C_3 over 1000 simulation replications.** The black lines represent PPLS estimates, the red lines represent PLS estimates. The middle lines are the average loading values across 1000 simulation replications; the width of the two outer lines are twice the standard deviations. The results are for normally distributed latent variables (t, e, f and h) and low dimensionality ($p = q = 1000$ variables).

Figures for extra high dimensionality

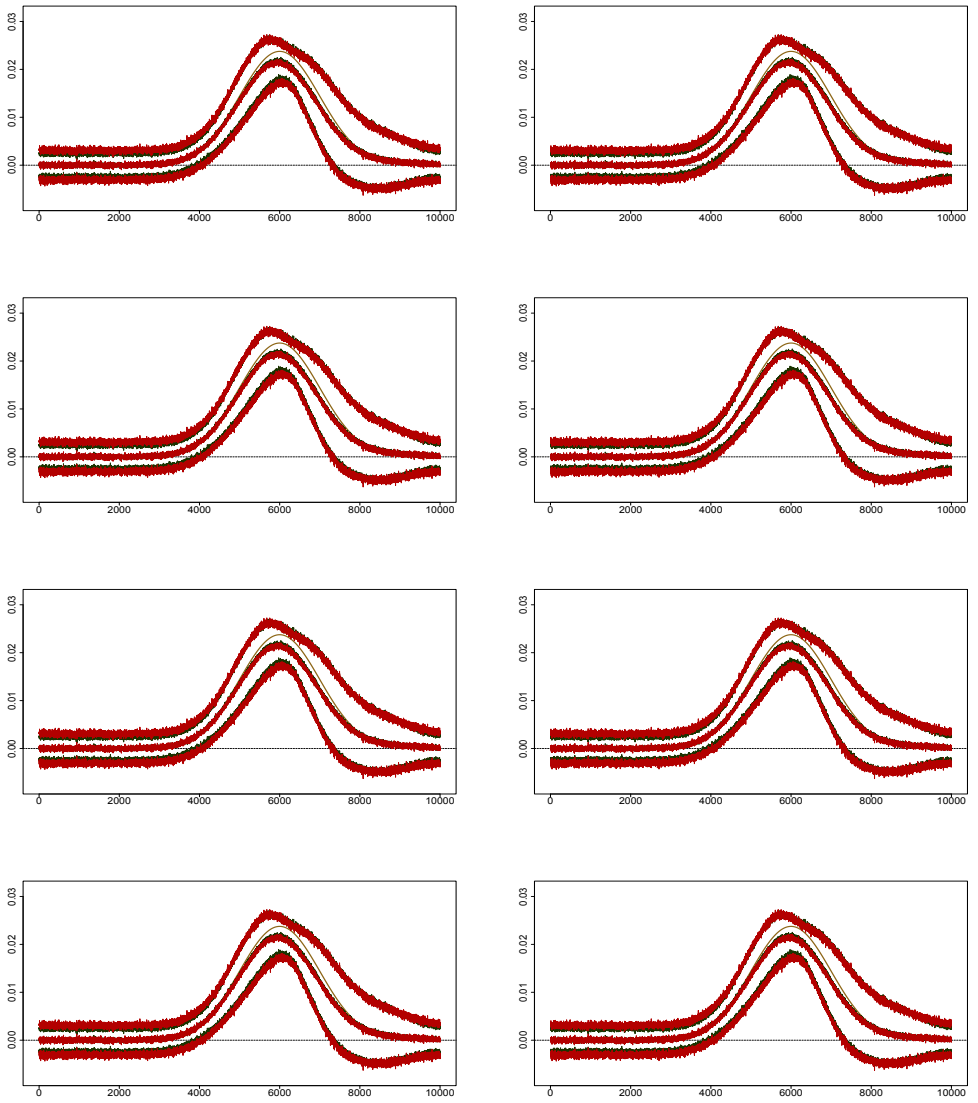

Figure 4.18: **True and estimated loadings W_1 over 1000 simulation replications.** The black and red lines represent PPLS and PLS estimates, respectively. The middle lines represent the average loading values. The distance between the upper and lower lines is twice the standard deviation for each loading value. The smooth golden line is the true loading profile. In the first column, the sample size is small ($N = 50$), the second column corresponds to large sample size ($N = 500$). The four rows correspond to a normal, student-t, Poisson and binomial distribution, respectively. Furthermore, the dimensionality is high ($p = q = 1000$) as well as the noise (equal to 50%).

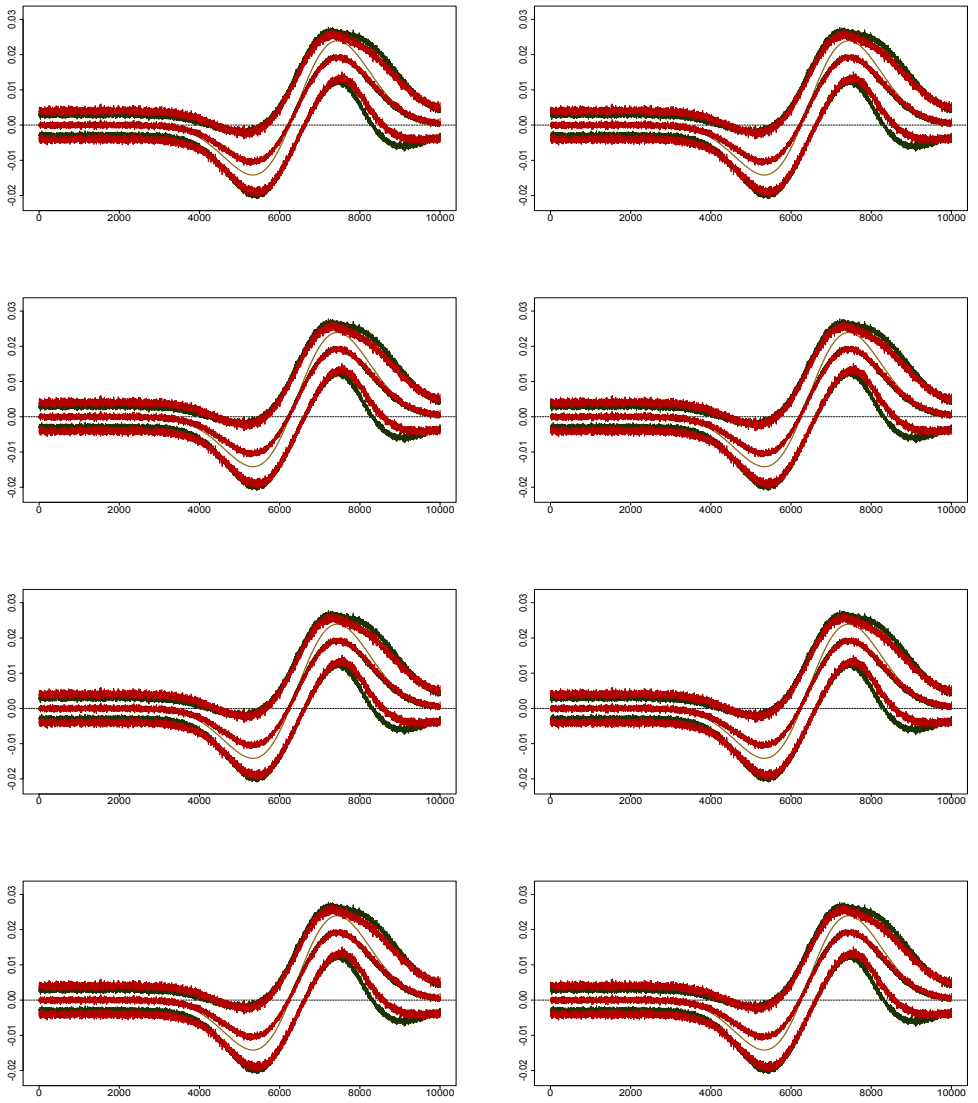

Figure 4.19: **True and estimated loadings W_2 over 1000 simulation replications.** The black and red lines represent PPLS and PLS estimates, respectively. The middle lines represent the average loading values. The distance between the upper and lower lines is twice the standard deviation for each loading value. The smooth golden line is the true loading profile. In the first column, the sample size is small ($N = 50$), the second column corresponds to large sample size ($N = 500$). The four rows correspond to a normal, student-t, Poisson and binomial distribution, respectively. Furthermore, the dimensionality is high ($p = q = 1000$) as well as the noise (equal to 50%).

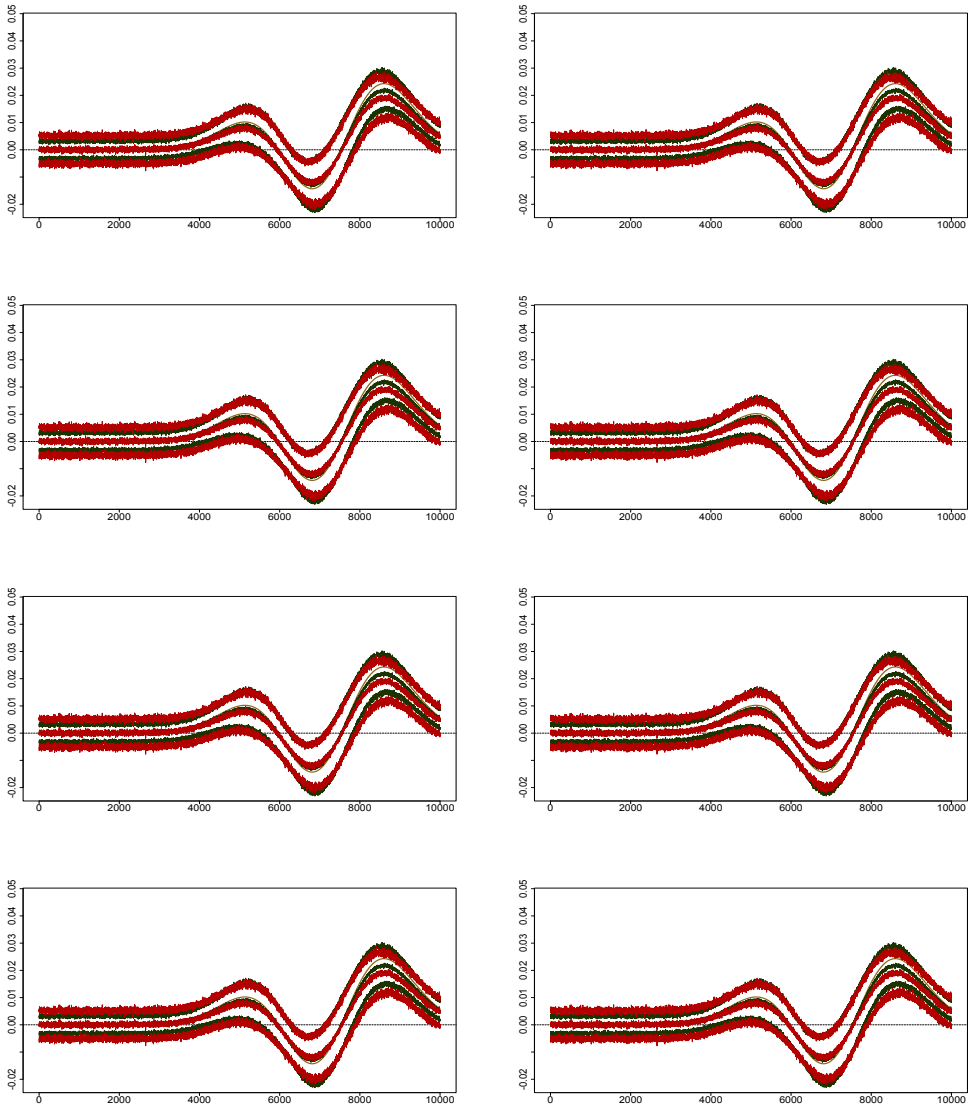

Figure 4.20: **True and estimated loadings W_3 over 1000 simulation replications.** The black and red lines represent PPLS and PLS estimates, respectively. The middle lines represent the average loading values. The distance between the upper and lower lines is twice the standard deviation for each loading value. The smooth golden line is the true loading profile. In the first column, the sample size is small ($N = 50$), the second column corresponds to large sample size ($N = 500$). The four rows correspond to a normal, student-t, Poisson and binomial distribution, respectively. Furthermore, the dimensionality is high ($p = q = 1000$) as well as the noise (equal to 50%).

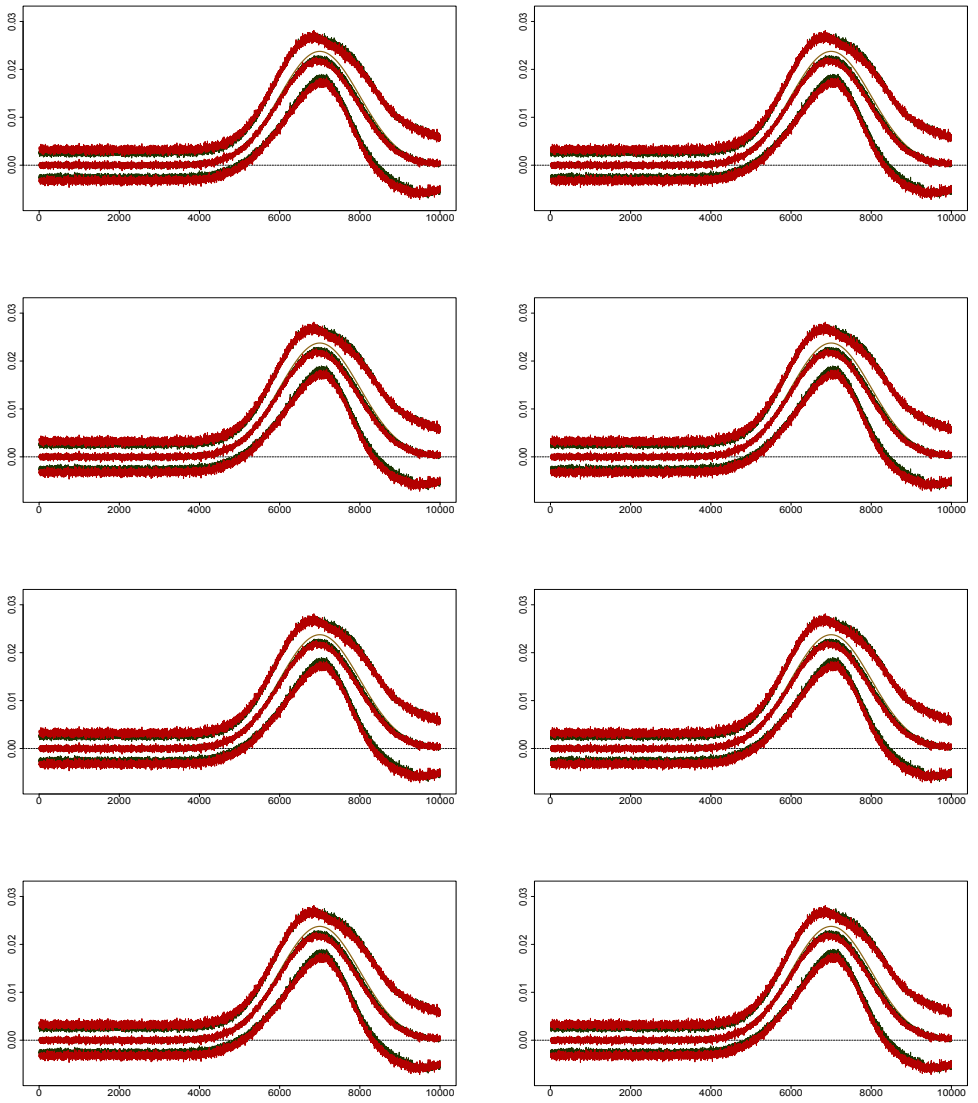

Figure 4.21: **True and estimated loadings C_1 over 1000 simulation replications.** The black and red lines represent PPLS and PLS estimates, respectively. The middle lines represent the average loading values. The distance between the upper and lower lines is twice the standard deviation for each loading value. The smooth golden line is the true loading profile. In the first column, the sample size is small ($N = 50$), the second column corresponds to large sample size ($N = 500$). The four rows correspond to a normal, student-t, Poisson and binomial distribution, respectively. Furthermore, the dimensionality is high ($p = q = 1000$) as well as the noise (equal to 50%).

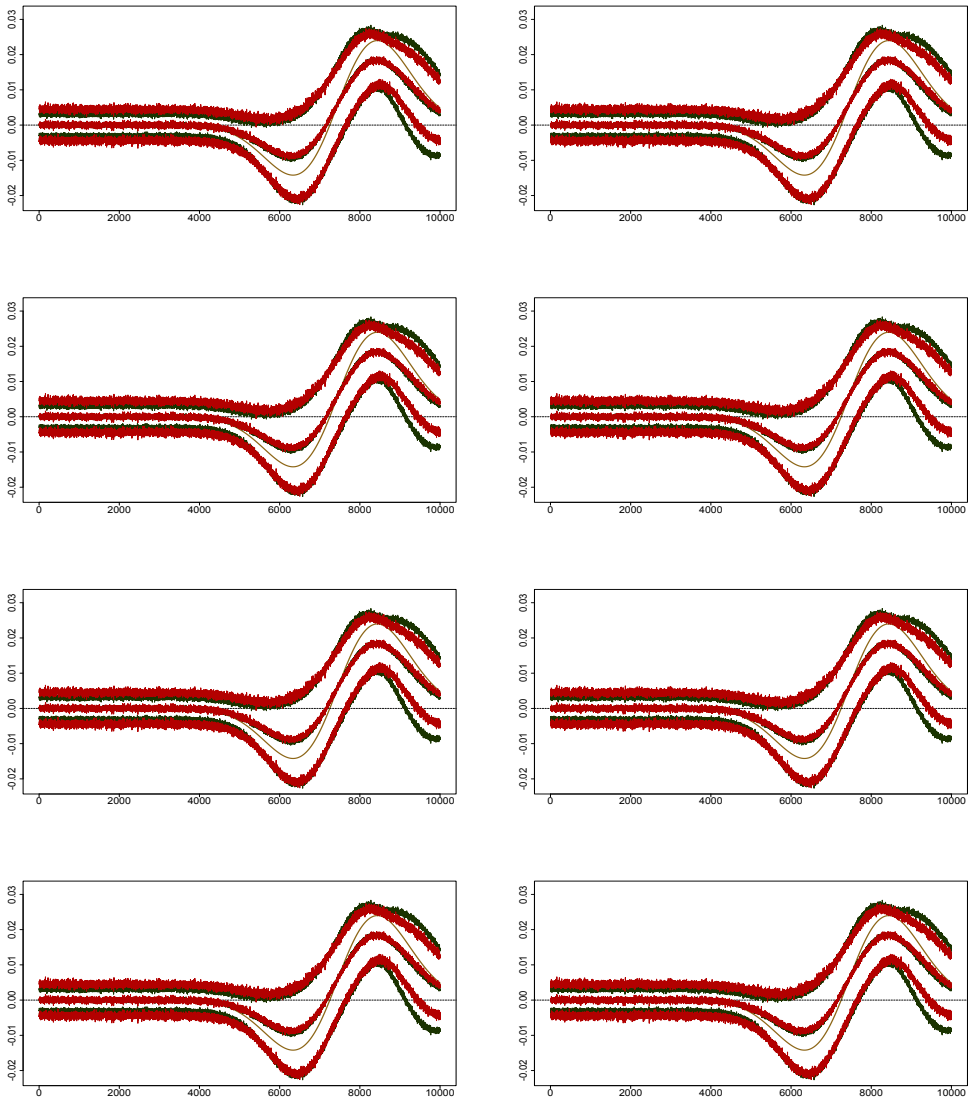

Figure 4.22: **True and estimated loadings C_2 over 1000 simulation replications.** The black and red lines represent PPLS and PLS estimates, respectively. The middle lines represent the average loading values. The distance between the upper and lower lines is twice the standard deviation for each loading value. The smooth golden line is the true loading profile. In the first column, the sample size is small ($N = 50$), the second column corresponds to large sample size ($N = 500$). The four rows correspond to a normal, student-t, Poisson and binomial distribution, respectively. Furthermore, the dimensionality is high ($p = q = 1000$) as well as the noise (equal to 50%).

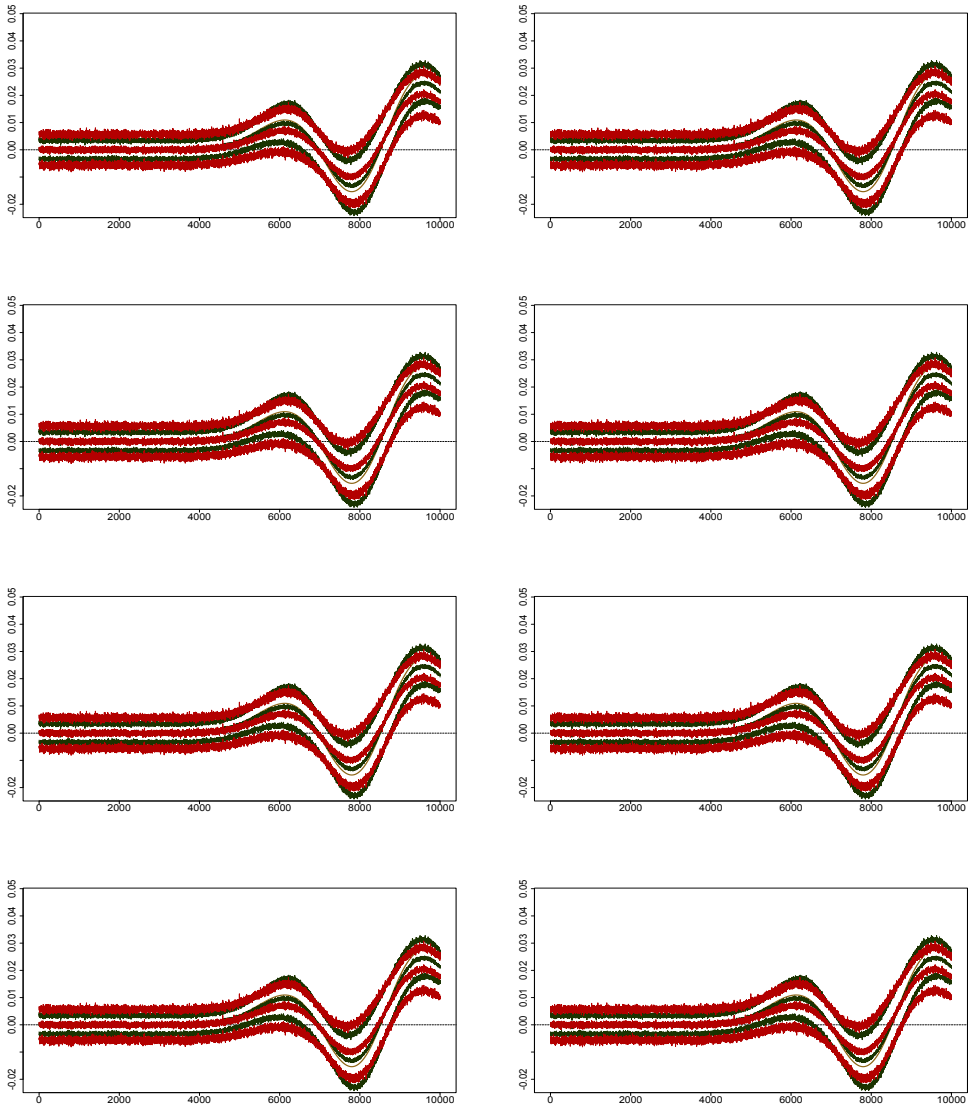

Figure 4.23: **True and estimated loadings C_3 over 1000 simulation replications.** The black and red lines represent PPLS and PLS estimates, respectively. The middle lines represent the average loading values. The distance between the upper and lower lines is twice the standard deviation for each loading value. The smooth golden line is the true loading profile. In the first column, the sample size is small ($N = 50$), the second column corresponds to large sample size ($N = 500$). The four rows correspond to a normal, student-t, Poisson and binomial distribution, respectively. Furthermore, the dimensionality is high ($p = q = 1000$) as well as the noise (equal to 50%).

Figures for standard errors

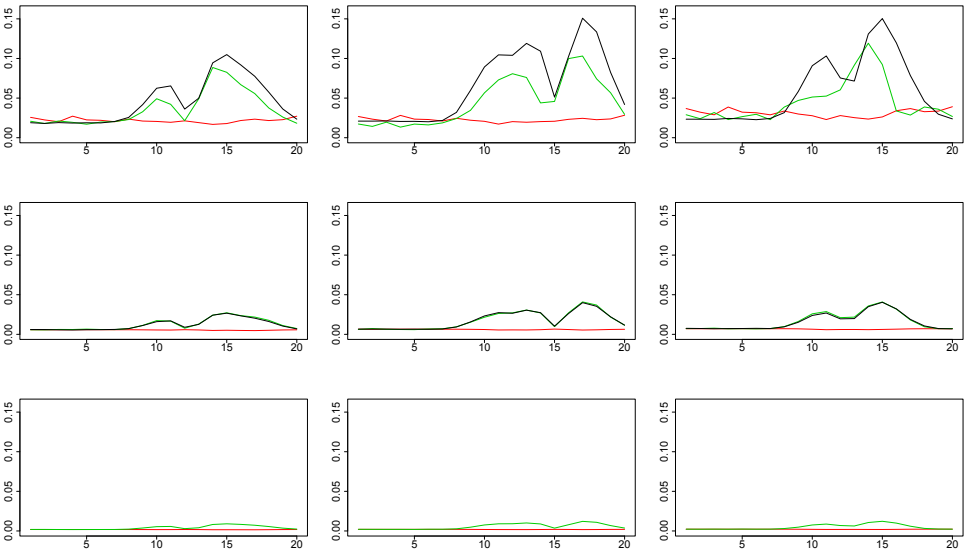

(a) Low noise level ($\alpha_n = 0.1$)

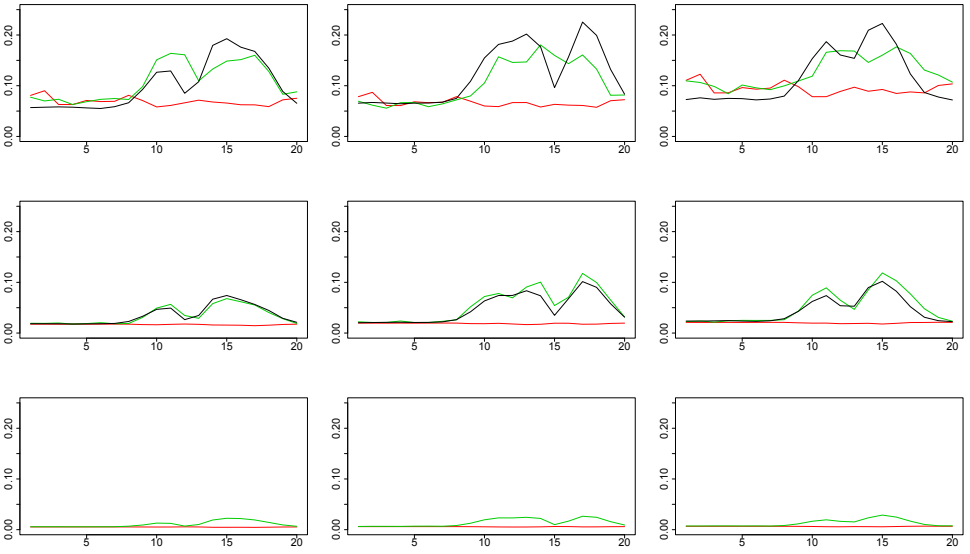

(b) High noise level ($\alpha_n = 0.5$)

Figure 4.24: **Standard errors of the W loading elements per component.** Bootstrap standard errors (shown in green), asymptotic standard errors (shown in red) and simulation-based standard errors (shown in black) are plotted for the loading estimates in each component. Plots for the three sample sizes (small $N = 50$, high $N = 500$, 'extra high' $N = 5000$) are shown along the rows. The three loading components (W_1, W_3 and W_3) are plotted column wise. The last row does not include simulation-based standard errors, as they are unavailable.

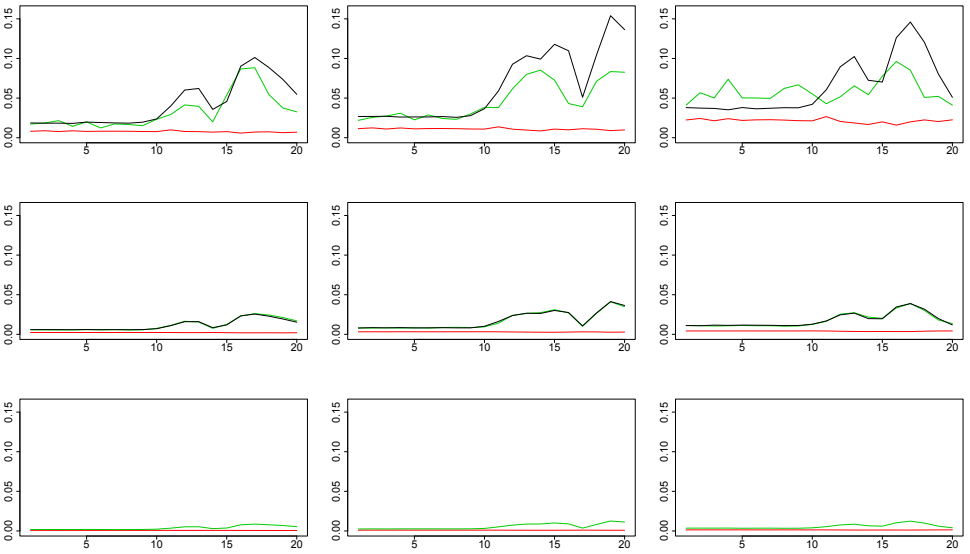

(a) Low noise level ($\alpha_n = 0.1$)

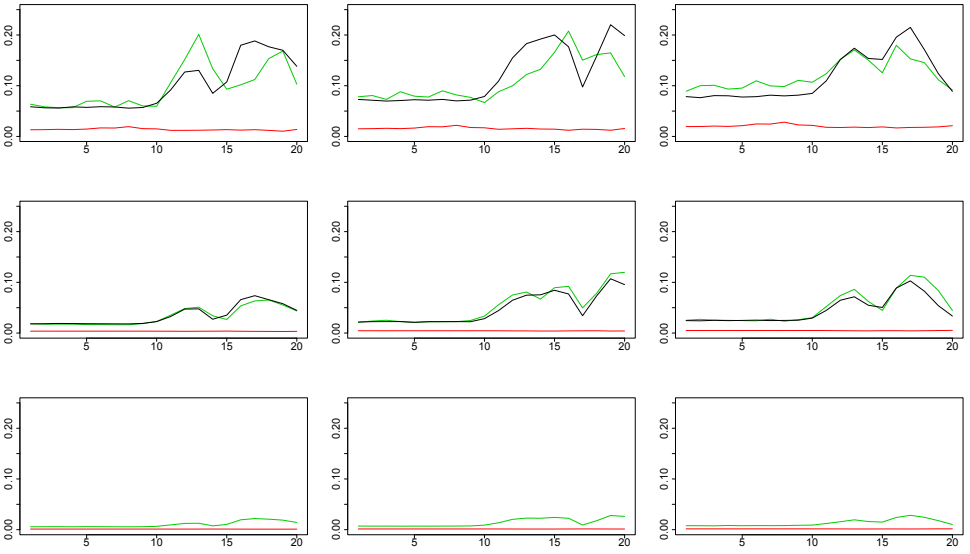

(b) High noise level ($\alpha_n = 0.5$)

Figure 4.25: **Standard errors of the C loading elements per component.** Bootstrap standard errors (shown in green), asymptotic standard errors (shown in red) and simulation-based standard errors (shown in black) are plotted for the loading estimates in each component. Plots for the three sample sizes (small $N = 50$, high $N = 500$, 'extra high' $N = 5000$) are shown along the rows. The three loading components (C_1, C_3 and C_3) are plotted column wise. The last row does not include simulation-based standard errors, as they are unavailable.

Bibliography

[1] H. Abdi. Partial least squares regression and projection on latent structure regression (PLS Regression). *Wiley Interdiscip. Rev. Comput. Stat.*, 2(1):97–106, jan 2010.

[2] F. R. Bach and M. I. Jordan. A Probabilistic Interpretation of Canonical Correlation Analysis. Technical Report 12, Department of Statistics, University of California, Berkeley, 2005.

[3] A. L. Boulesteix and K. Strimmer. Partial least squares: a versatile tool for the analysis of high-dimensional genomic data. *Br. Bioinform*, 8(1):32–44, 2007.

[4] R. D. Cook and X. Zhang. Simultaneous envelopes for multivariate linear regression. *Technometrics*, 57(1):11–25, 2015.

[5] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum Likelihood from Incomplete Data via the EM Algorithm. *J. R. Stat. Soc. Ser. B*, 39(1):1–38, 1977.

[6] R. DerSimonian and N. Laird. Meta-analysis in clinical trials. *Control. Clin. Trials*, 7(3):177–188, 1986.

[7] M. Eaton. *Multivariate statistics : a vector space approach*. Wiley, New York, 1983.

[8] S. el Bouhaddani, J. Houwing-Duistermaat, P. Salo, M. Perola, G. Jongbloed, and H.-W. Uh. Evaluation of O2PLS in Omics data integration. *BMC Bioinformatics*, 17(S2):S11, dec 2016.

[9] S. Geisser. Predictive Inference. *Philos. Sci.*, 24:180, 1993.

[10] Q. He, H. H. Zhang, C. L. Avery, and D. Y. Lin. Sparse meta-analysis with high-dimensional data. *Biostatistics*, 17(2):205–220, 2016.

[11] X. Huang, W. Pan, X. Han, Y. Chen, L. W. Miller, and J. Hall. Borrowing information from relevant microarray studies for sample classification using weighted partial least squares. *Comput. Biol. Chem.*, 29(3):204–211, 2005.

[12] G. Lauc, J. E. Huffman, M. Pučić, L. Zgaga, B. Adamczyk, A. Mužinić, M. Novokmet, O. Polašek, O. Gornik, J. Krištić, T. Keser, V. Vitart, B. Scheijen, H.-W. Uh, M. Molokhia, A. L. Patrick, P. McKeigue, I. Kolčić, I. K. Lukić, O. Swann, F. N. van Leeuwen, L. R. Ruhaak, J. J. Houwing-Duistermaat, P. E. Slagboom, M. Beekman, A. J. M. de Craen, A. M. Deelder, Q. Zeng, W. Wang, N. D. Hastie, U. Gyllensten, J. F. Wilson, M. Wuhrer, A. F. Wright, P. M. Rudd, C. Hayward, Y. Aulchenko, H. Campbell, and I. Rudan. Loci Associated with N-Glycosylation of Human Immunoglobulin G Show Pleiotropy with Autoimmune Diseases and Haematological Cancers. *PLoS Genet.*, 9(1):e1003225, jan 2013.

[13] G. Lauc, M. Pezer, I. Rudan, and H. Campbell. Mechanisms of disease: The human N-glycome. *Biochim. Biophys. Acta - Gen. Subj.*, 1860(8):1574–1582, aug 2016.

[14] S. Li, J. O. Nyagilo, D. P. Dave, W. Wang, B. Zhang, and J. Gao. Probabilistic partial least squares regression for quantitative analysis of Raman spectra. *Int. J. Data Min. Bioinform.*, 11(2):223–243, 2015.

[15] Y. Li, P. Uden, and D. V. Rosen. A two-step PLS inspired method for linear prediction with group effect. *Sankhya Indian J. Stat.*, 75:96–117, 2013.

[16] T. A. Louis. Finding the observed information matrix when using the EM algorithm. *J. Roy. Stat. Soc. Ser. B*, 44:226–233, 1982.

[17] K. V. Mardia, J. T. Kent, and J. M. Bibby. *Multivariate analysis*. Academic Press, 1979.

[18] P. Robert and Y. Escoufier. A Unifying Tool for Linear Multivariate Statistical Methods: The RV- Coefficient. *Appl. Stat.*, 25(3):257, 1976.

[19] B. Roś, F. Bijma, J. C. de Munck, and M. C. M. M. de Gunst. Existence and uniqueness of the maximum likelihood estimator for models with a Kronecker product covariance structure. *J. Multivar. Anal.*, 143:345–361, jan 2016.

[20] R. Rosipal and N. Krämer. Overview and Recent Advances in Partial Least Squares. In C. Saunders, M. Grobelnik, S. Gunn, and J. Shawe-Taylor, editors, *Subspace, Latent Struct. Featur. Sel.*, volume 3940, pages 34–51. Springer, Berlin, Heidelberg, 2006.

[21] G. A. F. Seber and A. J. Lee. *Linear regression analysis*. Wiley Series in Probability and Statistics. Wiley-Interscience [John Wiley & Sons], Hoboken, NJ, second edition, 2003.

[22] N. Taniguchi, K. Honke, M. Fukuda, H. Narimatsu, Y. Yamaguchi, and T. Angata, editors. *Handbook of Glycosyltransferases and Related Genes*. Springer Japan, Tokyo, 2014.

[23] M. E. Tipping and C. M. Bishop. Probabilistic principal component analysis. *J. R. Stat. Soc. Ser. B*, 61(3):611–622, 1999.

[24] J. Trygg and S. Wold. O2-PLS, a two-block (X–Y) latent variable regression (LVR) method with an integral OSC filter. *J. Chemom.*, 17(1):53–64, 2003.

[25] K. Van Deun, A. K. Smilde, M. J. van der Werf, H. a. L. Kiers, and I. Van Mechelen. A structured overview of simultaneous component based data integration. *BMC Bioinformatics*, 10:246, 2009.

[26] H. Wang, Q. Liu, and Y. Tu. Interpretation of partial least-squares regression models with VARIMAX rotation. *Comput. Stat. Data Anal.*, 48(1):207–219, jan 2005.

[27] R. Wehrens and W. E. van der Linden. Bootstrapping Principal Component Regression Models. *J. Chemom.*, 11(April 1996):157–171, 1997.

[28] H. Wold. Nonlinear iterative partial least squares (NIPALS) modelling: some current developments. In *Multivar. Anal. III (Proc. Third Internat. Symp. Wright State Univ., Dayton, Ohio, 1972)*, pages 383–407. Academic Press, New York, 1973.

[29] C. F. J. Wu. On the convergence properties of the EM algorithm. *Ann. Stat.*, 11(1):95–103, 1983.

[30] J. Zheng, Z. Song, and Z. Ge. Probabilistic learning of partial least squares regression model: Theory and industrial applications. *Chemom. Intell. Lab. Syst.*, 158:80–90, 2016.