



Structured Latent Factor Analysis for Large-scale Data: Identifiability, Estimability, and Their Implications: identifiability, estimability, and their implications

LSE Research Online URL for this paper: <http://eprints.lse.ac.uk/101122/>

Version: Accepted Version

Article:

Chen, Yunxiao, Li, Xiaou and Zhang, Siliang (2019) Structured Latent Factor Analysis for Large-scale Data: Identifiability, Estimability, and Their Implications: identifiability, estimability, and their implications. *Journal of the American Statistical Association*. ISSN 0162-1459

<https://doi.org/10.1080/01621459.2019.1635485>

Reuse

Items deposited in LSE Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the LSE Research Online record for the item.

Structured Latent Factor Analysis for Large-scale Data: Identifiability, Estimability, and Their Implications

Yunxiao Chen

Department of Statistics, London School of Economics and Political Science

Xiaou Li

School of Statistics, University of Minnesota

Siliang Zhang

Shanghai Center for Mathematical Sciences, Fudan University

Abstract

Latent factor models are widely used to measure unobserved latent traits in social and behavioral sciences, including psychology, education, and marketing. When used in a confirmatory manner, design information is incorporated as zero constraints on corresponding parameters, yielding structured (confirmatory) latent factor models. In this paper, we study how such design information affects the identifiability and the estimation of a structured latent factor model. Insights are gained through both asymptotic and non-asymptotic analyses. Our asymptotic results are established under a regime where both the number of manifest variables and the sample size diverge, motivated by applications to large-scale data. Under this regime, we define the structural identifiability of the latent factors and establish necessary and sufficient conditions that ensure structural identifiability. In addition, we propose an estimator which is shown to be consistent and rate optimal when structural identifiability holds. Finally, a non-asymptotic error bound is derived for this estimator, through which the effect of design information is further quantified. Our results shed lights on the design of

large-scale measurement in education and psychology and have important implications on measurement validity and reliability.

KEY WORDS: High-dimensional latent factor model, confirmatory factor analysis, identifiability of latent factors, structured low-rank matrix, large-scale psychological measurement

1 Introduction

Latent factor models are one of the major statistical tools for multivariate data analysis that have received many applications in social and behavioral sciences (Anderson, 2003; Bartholomew et al., 2011). Such models capture and interpret the common dependence among multiple manifest variables through the introduction of low-dimensional latent factors, where the latent factors are often given substantive interpretations. For example, in educational and psychological measurement, the manifest variables may be a test-taker's responses to test items, and the latent factors are often interpreted as his/her cognitive abilities and psychological traits, respectively. In marketing, the manifest variables may be an audience's ratings on movies, and the latent factors may be interpreted as his/her preferences on multiple characteristics of movies.

In many applications, latent factor models are used in a confirmatory manner, where design information on the relationship between manifest variables and latent factors is specified a priori. This is translated into a zero constraint on a loading parameter of the model. We call such a model a *structured latent factor model*. For example, consider a mathematics test consisting of J items that measures students' abilities on calculus and algebra. Then a structured latent factor model with two factors may be specified to model its item response data, where the two factors may be interpreted as the calculus and algebra factors, respectively. The relationship between the manifest variables and the two factors is coded by a $J \times 2$ binary matrix. The j th row being $(1, 0)$, $(0, 1)$, and $(1, 1)$ means that the j th items can be solved using only calculus skill, only algebra skill, and both, respectively.

Identifiability is an important property of a structured latent factor model for ensuring the substantive interpretation of latent factors. When the model is not identifiable, certain latent factors cannot be uniquely extracted from data and thus their substantive interpretations may not be valid. In particular, if no design information is incorporated, then it is well known that a latent factor model is not identifiable due to rotational indeterminacy, in which case one can simultaneously rotate the latent factors and the loading matrix of the model, without changing the distribution of data. It is thus important to study the relationship between design information of manifest variables and identifiability of latent factors. That is, when is a latent factor identifiable and when is not? A related problem is the estimation of a structured latent factor model. Under what notion of consistency, can a latent factors be consistently estimated? In that case, what is the convergence rate? These problems will be investigated in this paper. Similar identifiability problems have been considered in the context of linear factor models (Anderson and Rubin, 1956; Shapiro, 1985; Grayson and Marsh, 1994) and restricted latent class models (Xu and Zhang, 2016; Xu, 2017; Gu and Xu, 2018).

This paper considers the identifiability and estimability of structured latent factor models. We study this problem under a generalized latent factor modeling framework, which includes latent factor models for continuous, categorical, and count data. Unlike many existing works that take an empirical Bayes framework treating the latent factors as random variables, we treat them as unknown model parameters. This treatment has been taken in many works on latent factor models (Haberman, 1977; Holland, 1990; Bai and Li, 2012; Owen and Wang, 2016). Under this formulation, the identification of a latent factor becomes equivalent to the identification of the direction of a vector in \mathbb{R}^∞ , where each entry of the infinite-dimensional vector corresponds to a sample (e.g., a test-taker or a costumer) in a population. As will be shown in the sequel, to identify such a direction, we need an infinite number of manifest variables (e.g., measurements). Under such an asymptotic regime, we provide a necessary and sufficient condition on the design of manifest variables

for the identifiability of latent factors. This condition can be very easily interpreted from a repeated measurement perspective. In the rest of the paper, this notation of identifiability is called *structural identifiability*. Comparing with existing works (e.g., Anderson and Rubin, 1956; Shapiro, 1985; Grayson and Marsh, 1994), the current development applies to a more general family of latent factor models, and avoids distribution assumptions on the latent variables by adopting the double asymptotic regime that the sample size N and the number of manifest variables J simultaneously grow to infinity.

Under the double asymptotic regime and when structural identifiability holds, we propose an estimator and establish its consistency and convergence rate. Under suitable conditions, this convergence rate is shown to be optimal through an asymptotic lower bound for this estimation problem. Finally, a non-asymptotic error bound is derived for this estimator, which complements the asymptotic results by quantifying the estimation accuracy under a given design with finite N and J . To establish these results, we prove useful probabilistic error bounds and develop perturbation bounds on the intersection of linear subspaces which are of independent value for the theoretical analysis of low-rank matrix estimation.

The rest of the paper is organized as follows. In Section 2, we introduce a generalized latent factor modeling framework, within which our research questions are formulated. In Section 3, we discuss the structural identifiability for latent factors, establish the relationship between structural identifiability and estimability, and provide an estimator for which asymptotic results and a non-asymptotic error bound are established. Further implications of our theoretical results on large-scale measurement are provided in Section 4 and extensions of our results to more complex settings are discussed in Section 5. A new perturbation bound on linear subspaces is presented in Section 6 that is key to our main results. Numerical results are presented in Section 7 that contain two simulation studies and an application to a personality assessment dataset. Finally, concluding remarks are provided in Section 8. The proofs of all the technical results are provided as supplementary material.

2 Structured Latent Factor Analysis

2.1 Generalized Latent Factor Model

Suppose that there are N individuals (e.g., N test-takers) and J manifest variables (e.g. J test items). Let Y_{ij} be a random variable denoting the i th individual’s value on the j th manifest variable and let y_{ij} be its realization. For example, in educational tests, Y_{ij} s could be binary responses from the examinees, indicating whether the answers are correct or not. We further assume that each individual i is associated with a K -dimensional latent vector, denoted as $\boldsymbol{\theta}_i = (\theta_{i1}, \dots, \theta_{iK})^\top$ and each manifest variable j is associated with K parameters $\mathbf{a}_j = (a_{j1}, \dots, a_{jK})^\top$. We give two concrete contexts. Consider an educational test of mathematics, with $K = 3$ dimensions of “algebra”, “geometry”, and “calculus”. Then θ_{i1} , θ_{i2} , and θ_{i3} represent individual i ’s proficiency levels on algebra, geometry, and calculus, respectively. In the measurement of Big Five personality factors (Goldberg, 1993), $K = 5$ personality factors are considered, including “openness to experience”, “conscientiousness”, “extraversion”, “agreeableness”, and “neuroticism”. Then $\theta_{i1}, \dots, \theta_{i5}$ represent individual i ’s levels on the continuums of the five personality traits. The manifest parameter \mathbf{a}_j s can be understood as the regression coefficients when regressing Y_{ij} s on $\boldsymbol{\theta}_i$ s, $i = 1, \dots, N$. In many applications of latent factor models, especially in psychology and education, the estimations of $\boldsymbol{\theta}_i$ s and \mathbf{a}_j s are both of interest.

Our development is under a generalized latent factor model framework (Skrondal and Rabe-Hesketh, 2004), which extends the generalized linear model framework (McCullagh and Nelder, 1989) to latent factor analysis. Specifically, we assume that the distribution of Y_{ij} given $\boldsymbol{\theta}_i$ and \mathbf{a}_j is a member of the exponential family with natural parameter

$$m_{ij} = \mathbf{a}_j^\top \boldsymbol{\theta}_i = a_{j1}\theta_{i1} + \dots + a_{jK}\theta_{iK}, \quad (1)$$

and possibly a scale (i.e. dispersion) parameter ϕ . More precisely, the density/probability

mass function takes the form:

$$f(y|\mathbf{a}_j, \boldsymbol{\theta}_i, \phi) = \exp\left(\frac{ym_{ij} - b(m_{ij})}{\phi} + c(y, \phi)\right), \quad (2)$$

where $b(\cdot)$ and $c(\cdot)$ are pre-specified functions that depend on the member of the exponential family. Given $\boldsymbol{\theta}_i$ and \mathbf{a}_j , $i = 1, \dots, N$ and $j = 1, \dots, J$, we assume that all Y_{ij} s are independent. Consequently, the likelihood function, in which $\boldsymbol{\theta}_i$ s and \mathbf{a}_j s are treated as fixed effects, can be written as

$$L(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_N, \mathbf{a}_1, \dots, \mathbf{a}_J, \phi) = \prod_{i=1}^N \prod_{j=1}^J \exp\left(\frac{y_{ij}m_{ij} - b(m_{ij})}{\phi} + c(y_{ij}, \phi)\right). \quad (3)$$

This likelihood function is known as the *joint likelihood* function in the literature of latent variable models (Skrondal and Rabe-Hesketh, 2004). We remark that in the existing literature of latent factor models, there is often an intercept term indexed by j in the specification of (1), which can be easily realized under our formulation by constraining $\theta_{i1} = 1$, for all $i = 1, 2, \dots, N$. In that case, a_{j1} serves as the intercept term. This framework provides a large class of models, including linear factor models for continuous data, as well as logistic and Poisson models for multivariate binary and count data, as special cases. These special cases are listed below.

1. Linear factor model:

$$Y_{ij}|\boldsymbol{\theta}_i, \mathbf{a}_j \sim N(\mathbf{a}_j^\top \boldsymbol{\theta}_i, \sigma^2),$$

where the scale parameter $\phi = \sigma^2$.

2. Multidimensional Item Response Theory (MIRT) model:

$$Y_{ij}|\boldsymbol{\theta}_i, \mathbf{a}_j \sim \text{Bernoulli}\left(\frac{\exp(\mathbf{a}_j^\top \boldsymbol{\theta}_i)}{1 + \exp(\mathbf{a}_j^\top \boldsymbol{\theta}_i)}\right),$$

where the scale parameter $\phi = 1$.

3. Poisson factor model:

$$Y_{ij}|\boldsymbol{\theta}_i, \mathbf{a}_j \sim \text{Poisson}(\exp(\mathbf{a}_j^\top \boldsymbol{\theta}_i)),$$

where the scale parameter $\phi = 1$.

The analysis of this paper is based on the joint likelihood function (3) where both $\boldsymbol{\theta}_i$ s and \mathbf{a}_j s are treated as fixed effects, though in the literature the person parameters $\boldsymbol{\theta}_i$ are often treated as random effects and integrated out in the likelihood function (Holland, 1990). This fixed effect point of view allows us to straightforwardly treat the measurement problem as an estimation problem. For ease of exposition, we assume the scale parameter is known in the rest of the paper, while pointing out that it is straightforward to extend all the results to the case where it is unknown.

2.2 Confirmatory Structure

In this paper, we consider a confirmatory setting where the relationship between the manifest variables and the latent factors is known a priori. Suppose that there are J manifest variables and K latent factors. Then the confirmatory information is recorded by a $J \times K$ matrix, denoted by $Q = (q_{jk})_{J \times K}$, whose entries take value zero or one. In particular, $q_{jk} = 0$ means that manifest variable j is not directly associated with latent factor k . Such design information is often available in many applications of latent factor models, including in education, psychology, economics, and marketing (Thompson, 2004; Reckase, 2009; Gatignon, 2003).

The design matrix Q incorporates domain knowledge into the statistical analysis by imposing zero constraints on model parameters. For the generalized latent factor model, the loading parameter a_{jk} is constrained to zero if q_{jk} is zero. The constraints induced by the design matrix play an important role in the identifiability and the interpretation of the latent factors. Intuitively, suitable zero constraints on loading parameters will anchor the latent factors by preventing rotational indeterminacy, a major problem for the identifiability

of latent factor models. In the rest of this paper, we formalize this intuition by studying how the design matrix Q affects the identifiability and estimability of generalized latent factor models.

2.3 A Summary of Main Results

In this paper, we investigate how design information given by \mathbf{q}_j s affects the quality of measurement. This problem is tackled through both asymptotic analysis and a non-asymptotic error bound, under the generalized latent factor modeling framework.

Our asymptotic analysis focuses on the identifiability and estimability of the latent factors, under a setting where both N and J grow to infinity. To define identifiability, consider a population of people where $N = \infty$ and a universe of manifest variables where $J = \infty$. A latent factor k is a hypothetical construct, defined by the person population. More precisely, it is determined by the individual latent factor scores of the entire person population, denoted by $(\theta_{1k}^*, \theta_{2k}^* \dots) \in \mathbb{R}^{\mathbb{Z}^+}$, where θ_{ik}^* denotes the true latent factor score of person i on latent factor k and $\mathbb{R}^{\mathbb{Z}^+}$ denotes the set of vectors with countably infinite real number components. The identifiability of the k th latent factor then is equivalent to the identifiability of a vector in $\mathbb{R}^{\mathbb{Z}^+}$ under the distribution of an infinite dimensional random matrix, $\{Y_{ij} : i = 1, 2, \dots, j = 1, 2, \dots\}$. This setting is natural in the context of large-scale measurement where both N and J are large.

Under the above setting, this paper addresses three research questions. First, how should the identifiability of latent factors be suitably formalized? Second, under what design are the latent factors identifiable? Third, what is the relationship between the identifiability and estimability? In other words, whether and to what extent can we recover the scores of an identifiable latent factor from data?

We further provide a non-asymptotic error bound to complement the asymptotic results. For finite N and J , the effect of design information on the estimation of a latent factor is reflected by a multiplying coefficient in the bound.

2.4 Preliminaries

In this section, we fix some notations used throughout this paper.

Notations.

- a. \mathbb{Z}_+ : the set of all positive integers.
- b. $\mathbb{R}^{\mathbb{Z}_+}$: the set of vectors with countably infinite real number components.
- c. $\mathbb{R}^{\mathbb{Z}_+ \times \{1, \dots, K\}}$: the set of all the real matrices with countably infinite rows and K columns.
- d. $\{0, 1\}^{\mathbb{Z}_+ \times \{1, \dots, K\}}$: the set of all the binary matrices with countably infinite rows and K columns.
- e. Θ : the parameter matrix for the person population, $\Theta \in \mathbb{R}^{\mathbb{Z}_+ \times \{1, \dots, K\}}$.
- f. A : the parameter matrix for the manifest variable population, $A \in \mathbb{R}^{\mathbb{Z}_+ \times \{1, \dots, K\}}$.
- g. Q : the design matrix for the manifest variable population, $Q \in \{0, 1\}^{\mathbb{Z}_+ \times \{1, \dots, K\}}$.
- h. $\mathbf{0}$: the vector or matrix with all components being 0.
- i. $P_{\Theta, A}$: the probability distribution of $(Y_{ij}, i, j \in \mathbb{Z}_+)$, given person and item parameters Θ and A .
- j. $\mathbf{v}_{[1:m]}$: the first m components of a vector \mathbf{v} .
- k. $W_{[S_1, S_2]}$: the submatrix of a matrix W formed by rows S_1 and columns S_2 , where $S_1, S_2 \subset \mathbb{Z}_+$.
- l. $W_{[1:m, k]}$: the first m components of the k -th column of a matrix W .
- m. $W_{[k]}$: the k -th column of a matrix W .
- n. $\|\mathbf{v}\|$: the Euclidian norm of a vector \mathbf{v} .

o. $\sin \angle(\mathbf{u}, \mathbf{v})$: the sine of the angle between two vectors,

$$\sin \angle(\mathbf{u}, \mathbf{v}) = \sqrt{1 - \frac{(\mathbf{u}^\top \mathbf{v})^2}{\|\mathbf{u}\|^2 \|\mathbf{v}\|^2}},$$

where $\mathbf{u}, \mathbf{v} \in \mathbb{R}^m$, $\mathbf{u}, \mathbf{v} \neq \mathbf{0}$. We point out that the angles between vectors are assumed to belong to $[0, \pi]$ and $\sin \angle(\mathbf{u}, \mathbf{v}) \geq 0$ for all vectors \mathbf{u}, \mathbf{v} .

p. $\|W\|_F$: the Frobenius norm of a matrix $W = (w_{ij})_{m \times n}$, $\|W\|_F \triangleq \sqrt{\sum_{i=1}^m \sum_{j=1}^n w_{ij}^2}$.

q. $\|W\|_2$: the spectral norm of matrix W , i.e., the largest singular value of matrix.

r. $\sigma_1(W) \geq \sigma_2(W) \geq \dots \geq \sigma_n(W)$: the singular values of a matrix $W \in \mathbb{R}^{m \times n}$, in a descending order.

s. $\gamma(W)$: a function mapping from $\mathbb{R}^{\mathbb{Z}_+ \times n}$ to \mathbb{R} , defined as

$$\gamma(W) \triangleq \liminf_{m \rightarrow \infty} \frac{\sigma_n(W_{[1:m, 1:n]})}{\sqrt{m}}. \quad (4)$$

t. $|S|$: the cardinality of a set S .

u. $N \wedge J$: the minimum value between N and J .

3 Main Results

3.1 Structural Identifiability

We first formalize the definition of structural identifiability. For two vectors with countably infinite components $\mathbf{w} = (w_1, w_2, \dots)^\top$, $\mathbf{z} = (z_1, z_2, \dots)^\top \in \mathbb{R}^{\mathbb{Z}_+}$, we define

$$\sin_+ \angle(\mathbf{w}, \mathbf{z}) = \limsup_{n \rightarrow \infty} \sin \angle(\mathbf{w}_{[1:n]}, \mathbf{z}_{[1:n]}), \quad (5)$$

which quantifies the angle between two vectors \mathbf{w} and \mathbf{z} in $\mathbb{R}^{\mathbb{Z}_+}$. In particular, we say the angle between \mathbf{w} and \mathbf{z} is zero when $\sin_+ \angle(\mathbf{w}, \mathbf{z})$ is zero.

Definition 1 (Structural identifiability of a latent factor). *Consider the k th latent factor, where $k \in \{1, \dots, K\}$, and a nonempty parameter space $\mathcal{S} \subset \mathbb{R}^{\mathbb{Z}_+ \times \{1, \dots, K\}} \times \mathbb{R}^{\mathbb{Z}_+ \times \{1, \dots, K\}}$ for (Θ, A) . We say the k -th latent factor is structurally identifiable in the parameter space \mathcal{S} if for any $(\Theta, A), (\Theta', A') \in \mathcal{S}$, $P_{\Theta, A} = P_{\Theta', A'}$ implies $\sin_+ \angle(\Theta_{[k]}, \Theta'_{[k]}) = 0$.*

We point out that the parameter space \mathcal{S} , which will be specified later in this section, is essentially determined by the design information q_{jk} s. As will be shown shortly, a good design imposes suitable constraints on the parameter space, which further ensures the structure identifiability of the latent factors. This definition of identifiability avoids the consideration of the scale of the latent factor, which is not uniquely determined as the distribution of data only depends on $\{\boldsymbol{\theta}_i^\top \mathbf{a}_j : i, j \in \mathbb{Z}_+\}$. Moreover, the sine measure is a canonical way to quantify the distance between two linear spaces that has been used in, for example, the well-known sine theorems for matrix perturbation (Davis, 1963; Wedin, 1972). As will be shown in the sequel, this definition of structural identifiability naturally leads to a relationship between identifiability and estimability and has important implications on psychological measurement.

We now characterize the structural identifiability under suitable regularity conditions. We consider a design matrix Q for the manifest variable population, where $Q \in \{0, 1\}^{\mathbb{Z}_+ \times \{1, \dots, K\}}$. Our first regularity assumption is about the stability of the Q matrix.

A1 The limit

$$p_Q(S) = \lim_{J \rightarrow \infty} \frac{|\{j : q_{jk} = 1, \text{ if } k \in S \text{ and } q_{jk} = 0, \text{ if } k \notin S, 1 \leq j \leq J\}|}{J}$$

exists for any subset $S \subset \{1, \dots, K\}$. In addition, $p_Q(\emptyset) = 0$.

The above assumption requires that the frequency of manifest variables associated with and only with latent factors in S converges to a limit proportion $p_Q(S)$. In addition, $p_Q(\emptyset) =$

0 implies that there are few irrelevant manifest variables. We point out that this assumption is adopted mainly to simplify the statements in Theorem 1 and Theorem 2. As discussed later, this assumption is automatically satisfied if \mathbf{q}_j s are generated under a stochastic design. This assumption will be further relaxed in Section 3.4 where a non-asymptotic error bound is established. We also make the following assumption on the generalized latent factor model, which is satisfied under most of the widely used models, including the linear factor model, MIRT model, and the Poisson factor model listed above.

A2 The natural parameter space $\{\nu : |b(\nu)| < \infty\} = \mathbb{R}$.

Under the above assumptions, Theorem 1 provides a necessary and sufficient condition on the design matrix Q for the structural identifiability of the k th latent factor. This result is established within the parameter space $\mathcal{S}_Q \subset \mathbb{R}^{\mathbb{Z}_+ \times \{1, \dots, K\}} \times \mathbb{R}^{\mathbb{Z}_+ \times \{1, \dots, K\}}$,

$$\mathcal{S}_Q = \mathcal{S}_Q^{(1)} \times \mathcal{S}_Q^{(2)}. \quad (6)$$

Here, we define

$$\mathcal{S}_Q^{(1)} = \{\Theta \in \mathbb{R}^{\mathbb{Z}_+ \times \{1, \dots, K\}} : \|\boldsymbol{\theta}_i\| \leq C \text{ and } \gamma(\Theta) > 0\} \quad (7)$$

and

$$\begin{aligned} \mathcal{S}_Q^{(2)} = \left\{ A \in \mathbb{R}^{\mathbb{Z}_+ \times \{1, \dots, K\}} : \|\mathbf{a}_j\| \leq C, A_{[R_Q(S), S^c]} = \mathbf{0} \text{ for all } S \subset \{1, \dots, K\}, \right. \\ \left. \text{and } \gamma(A_{[R_Q(S), S]}) > 0 \text{ for all } S, \text{ s.t. } p_Q(S) > 0 \right\}, \end{aligned} \quad (8)$$

where C is a positive constant, the γ function is defined in (4), and

$$R_Q(S) = \{j : q_{jk'} = 1, \text{ for all } k' \in S \text{ and } q_{jk'} = 0, \text{ for all } k' \notin S\} \quad (9)$$

denotes the set of manifest variables that are associated with and only with latent factors in S . Discussions on the parameter space are provided after the statement of Theorem 1.

Theorem 1. *Under Assumptions A1 and A2, the k -th latent factor is structurally identifiable in \mathcal{S}_Q if and only if*

$$\{k\} = \bigcap_{k \in S, p_Q(S) > 0} S, \quad (10)$$

where we define $\bigcap_{k \in S, p_Q(S) > 0} S = \emptyset$ if $p_Q(S) = 0$ for all S that contains k .

The following proposition guarantees that the parameter space is nontrivial.

Proposition 1. *For any Q satisfying A1, $\mathcal{S}_Q \neq \emptyset$.*

We further remark on the parameter space \mathcal{S}_Q . First, \mathcal{S}_Q requires some regularities on each θ_i and \mathbf{a}_j (i.e., $\|\theta_i\| \leq C, \|\mathbf{a}_j\| \leq C$) and the A -matrix satisfying the constraints imposed by Q ($A_{[R_Q(S), S^c]} = \mathbf{0}$) for all S . It further requires that there is enough variation among people, quantified by $\gamma(\Theta) > 0$, where the γ function is defined in (4). Note that this requirement is mild, in the sense that if θ_i s are independent and identically distributed (i.i.d.) with a strictly positive definite covariance matrix, then $\gamma(\Theta) > 0$ a.s., according to the strong law of large numbers. Furthermore, $\gamma(A_{[R_Q(S), S]}) > 0$ for S satisfying $p_Q(S) > 0$ requires that each type of manifest variables (categorized by S) contains sufficient information if appearing frequently ($p_Q(S) > 0$). Similar to the justification for Θ , $\gamma(A_{[R_Q(S), S]}) > 0$ can also be justified by considering that \mathbf{a}_j s are i.i.d. following a certain distribution for $j \in R_Q(S)$.

We provide an example to facilitate the understanding of Theorem 1. Suppose that assumptions A1 and A2 hold. If $K = 2$ and $p_Q(\{1\}) = p_Q(\{1, 2\}) = 1/2$, then the second latent factor is not structurally identifiable, even if it is associated with infinitely many manifest variables. In addition, having many manifest variables with a simple structure ensures the structural identifiability of a latent factor. That is, if $p_Q(\{k\}) > 0$, then the k th factor is structurally identifiable.

Finally, we briefly discuss Assumption A1. The next proposition implies that if we have a stochastic design where Q has i.i.d. rows, then Assumption A1 is satisfied almost surely.

Proposition 2. *Assume that the design matrix Q has i.i.d. rows and $\mathbf{q}_i \neq (0, \dots, 0)$ for all $i = 1, 2, \dots$. Then, Assumption A1 is satisfied almost surely. Moreover, let*

$$w_S = P(q_{1k} = 1 \text{ for all } k \in S, \text{ and } q_{1k} = 0 \text{ for all } k \notin S) \quad (11)$$

for each $S \subset \{1, \dots, K\}$. Then, Q satisfies Assumption A1 with $p_S(Q) = w_S$ for all $S \subset \{1, \dots, K\}$ almost surely.

3.2 Identifiability and Estimability

It is well known that for a fixed dimensional parametric model with i.i.d. observations, the identifiability of a model parameter is necessary for the existence of a consistent estimator. We extend this result to the infinite-dimensional parameter space under the current setting. We start with a generalized definition for the consistency of estimating a latent factor. An estimator given N individuals and J manifest variables is denoted by $(\hat{\Theta}^{(N,J)}, \hat{A}^{(N,J)})$, which only depends on $Y_{[1:N,1:J]}$ for all $N, J \in \mathbb{Z}_+$.

Definition 2 (Consistency for estimating latent factor k). *The sequence of estimators $\{(\hat{\Theta}^{(N,J)}, \hat{A}^{(N,J)}), N, J \in \mathbb{Z}_+\}$ is said to consistently estimate the latent factor k if*

$$\sin \angle(\hat{\Theta}_{[k]}^{(N,J)}, \Theta_{[1:N,k]}) \xrightarrow{P_{\Theta, A}} 0, \text{ as } N, J \rightarrow \infty, \quad (12)$$

for all $(\Theta, A) \in \mathcal{S}_Q$.

The next proposition establishes the necessity of the structural identifiability of a latent factor on its estimability.

Proposition 3. *If latent factor k is not structurally identifiable in \mathcal{S}_Q , then there does not exist a consistent estimator for latent factor k .*

Remark 1. *The above results of identifiability and estimability are all established under a double asymptotic regime that both N and J grow to infinity, which is suitable for large-*

scale applications where both N and J are large. It differs from the classical asymptotic setting (e.g., Anderson and Rubin, 1956; Bing et al., 2017) where J is fixed and N grows to infinity. The classical setting treats the latent factors as random effects and focuses on the identifiability of the loading parameters. There are several reasons for adopting the double asymptotic regime. First, the double asymptotic regime allows us to directly focus on the identifiability and the estimation of factor scores θ_i , which is of more interest in many applications (e.g, psychological and educational measurement) than the loading parameters. To consistently estimate the factor scores, naturally, we need the number of measurements J to grow to infinity. Second, having a varying J allows us to take the number of manifest variables into consideration in the research design and data collection when measuring certain traits with substantive meaning through a latent factor model. Third, even if we only focus on the loading parameters, the classical asymptotic regime does not lead to consistency on the estimation of the loading parameters. For example, for the MIRT model for binary data, one cannot consistently estimate the loading parameters, when J is fixed and N grows to infinity, unless we assume θ_i s to be i.i.d. samples from a certain parametric distribution and use this parametric assumption in the estimation procedure. Finally, this double asymptotic regime is not entirely new. In fact, this regime has been adopted in, for example, Haberman (1977), Bai and Li (2012), and Owen and Wang (2016) among others, for the asymptotic analysis of linear and nonlinear factor models.

3.3 Estimation and Its Consistency

We further show that the structural identifiability and estimability are equivalent under our setting. For ease of exposition, let $Q \in \{0, 1\}^{\mathbb{Z}_+ \times \{1, \dots, K\}}$ be a design matrix satisfying Assumption A1. In addition, let $(\Theta^*, A^*) \in \mathcal{S}_Q$ be the true parameters for the person and the manifest variable populations. We provide an estimator $(\hat{\Theta}^{(N, J)}, \hat{A}^{(N, J)})$ such that

$$\sin \angle(\hat{\Theta}_{[k]}^{(N, J)}, \Theta_{[1:N, k]}^*) \xrightarrow{P_{\Theta^*, A^*}} 0, \quad N, J, \rightarrow \infty,$$

when Q satisfies (10) which leads to the structural identifiability of latent factor k according to Theorem 1. Specifically, we consider the following estimator

$$\begin{aligned} (\hat{\Theta}^{(N,J)}, \hat{A}^{(N,J)}) &\in \arg \min -l(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_N, \mathbf{a}_1, \dots, \mathbf{a}_J), \\ \text{s.t. } \|\boldsymbol{\theta}_i\| &\leq C', \|\mathbf{a}_j\| \leq C', \\ \mathbf{a}_j &\in \mathcal{D}_j, i = 1, \dots, N, j = 1, \dots, J, \end{aligned} \quad (13)$$

where $l(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_N, \mathbf{a}_1, \dots, \mathbf{a}_J) = \sum_{i=1}^N \sum_{j=1}^J y_{ij}(\boldsymbol{\theta}_i^\top \mathbf{a}_j) - b(\boldsymbol{\theta}_i^\top \mathbf{a}_j)$, C' is any constant greater than C in the definition of \mathcal{S}_Q , and $\mathcal{D}_j = \{\mathbf{a} \in \mathbb{R}^K : a_{jk} = 0 \text{ if } q_{jk} = 0\}$ imposes the constraint on \mathbf{a}_j . Note that maximizing $l(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_N, \mathbf{a}_1, \dots, \mathbf{a}_J)$ is equivalent to maximizing the joint likelihood (3), due to the natural exponential family form. The next theorem provides an error bound on $(\hat{\Theta}^{(N,J)}, \hat{A}^{(N,J)})$.

Theorem 2. *Under assumptions A1-A2 and $(\Theta^*, A^*) \in \mathcal{S}_Q$, there exists a constant κ_1 (independent of N and J , depending on the function b , the constant C' , and K) such that,*

$$\frac{1}{\sqrt{NJ}} E \left\| \hat{\Theta}^{(N,J)} (\hat{A}^{(N,J)})^\top - \Theta_{[1:N,1:K]}^* A_{[1:J,1:K]}^{*\top} \right\|_F \leq \frac{\kappa_1}{\sqrt{N \wedge J}}. \quad (14)$$

Moreover, if Q satisfies (10) and thus latent factor k is structurally identifiable, then there exists $\kappa_2 > 0$ such that

$$E \sin \angle(\Theta_{[1:N,k]}^*, \hat{\Theta}_{[k]}^{(N,J)}) \leq \frac{\kappa_2}{\sqrt{N \wedge J}}. \quad (15)$$

Proposition 3, and Theorems 1 and 2 together imply that the structural identifiability and estimability over \mathcal{S}_Q are equivalent, which is summarized in the following corollary.

Corollary 1. *Under Assumptions A1 and A2, there exists an estimator $(\hat{\Theta}^{(N,J)}, \hat{A}^{(N,J)})$ such that $\lim_{N,J \rightarrow \infty} \sin \angle(\hat{\Theta}_{[k]}^{(N,J)}, \Theta_{[1:N,k]}) = 0$ in $P_{\Theta,A}$ for all $(\Theta, A) \in \mathcal{S}_Q$ if and only if the design matrix Q satisfies (10).*

Remark 2. *The estimator (13) is closely related to the joint likelihood estimator in the literature of econometrics and psychometrics. When J is fixed and N grows to infinity,*

this estimator is shown to be inconsistent, due to the simultaneous growth of the sample size and the parameter space (Neyman and Scott, 1948; Andersen, 1970; Ghosh, 1995). However, if N and J simultaneously grow to infinity, Haberman (1977) shows that the joint maximum likelihood estimator is consistent under the Rasch model, a simple unidimensional item response theory model. Then under a general family of multidimensional item response theory models and under an exploratory factor analysis setting, an estimator similar to (13) is considered in Chen et al. (2019) but without the zero constraints given by the Q -matrix. As a result, the estimator obtained in Chen et al. (2019) is rotationally indeterminant.

Remark 3. The error bound (14) holds even when one or more latent factors are not structurally identifiable. In particular, (14) holds when removing the constraint $\mathbf{a}_j \in \mathcal{D}_j$ from (13), which corresponds to the exploratory factor analysis setting where no design matrix Q is pre-specified (or in other words, $q_{jk} = 1$ for all j and k ; see the setting of Chen et al., 2019). In that case, the best one can achieve is to recover the linear space spanned by the column vectors of Θ^* and similarly the linear space spanned by the column vectors of A^* . To make sense of such exploratory factor analysis results, one needs an additional rotation step to find an approximately sparse estimate of A^* (Chen et al., 2019).

Remark 4. The proposed estimator (13) and its error bound are related to exact low-rank matrix completion (Bhaskar and Javanmard, 2015) and approximate low-rank matrix completion (e.g. Candès and Plan, 2010; Davenport et al., 2014; Cai and Zhou, 2013), where a bound similar to (14) can typically be derived. The key differences are (a) the research on matrix completion is only interested in the estimation of $\Theta^* A^{*\top}$, while the current paper focuses on the estimation of Θ^* that is a fundamental problem of psychological measurement and (b) our results are derived under a more general family of models.

To further evaluate the efficiency of the proposed estimator, we provide the following lower bounds.

Theorem 3. Suppose that Assumptions A1 and A2 hold. For $N, J \in \mathbb{Z}_+$, let $\bar{M}^{(N,J)}$ be

an arbitrary estimator which maps data $Y_{[1:N,1:J]}$ to $\mathbb{R}^{N \times J}$. Then there exists $\varepsilon_1 > 0$ and $N_0, J_0 > 0$ such that for $N \geq N_0, J \geq J_0$, there exists $(\Theta^*, A^*) \in \mathcal{S}_Q$ such that

$$P_{\Theta^*, A^*} \left(\frac{1}{\sqrt{NJ}} \|\bar{M}^{(N,J)} - \Theta_{[1:N,1:K]}^* (A_{[1:J,1:K]}^*)^\top\|_F \geq \frac{\varepsilon_1}{\sqrt{N \wedge J}} \right) \geq \frac{1}{2}. \quad (16)$$

Moreover, let $\bar{\Theta}_{[k]}^{(N,J)}$ be an arbitrary estimator which maps data $Y_{[1:N,1:J]}$ to \mathbb{R}^N . Then for each $A^* \in \mathcal{S}_Q^{(2)}$ defined in (8), there exists $\varepsilon_2 > 0$ and $N_0, J_0 > 0$ such that for $N \geq N_0, J \geq J_0$, there exists $\Theta^* \in \mathcal{S}_Q^{(1)}$ such that

$$P_{\Theta^*, A^*} \left(\sin \angle(\bar{\Theta}_{[k]}^{(N,J)}, \Theta_{[1:N,k]}^*) \geq \frac{\varepsilon_2}{\sqrt{J}} \right) \geq \frac{1}{2}. \quad (17)$$

Based on the asymptotic upper bounds and lower bounds obtained in Theorems 2 and 3, we have the following findings. First, the proposed estimator $\hat{\Theta}^{(N,J)}(\hat{A}^{(N,J)})^\top$ is rate optimal for the estimation of $\Theta_{[1:N,1:K]}^* (A_{[1:J,1:K]}^*)^\top$, because the upper bound (14) matches the lower bound (16). Second, for estimating $\Theta_{[1:N,k]}^*$, the proposed estimator $\hat{\Theta}_{[k]}$ is rate optimal when $\limsup_{N,J \rightarrow \infty} J/N < \infty$. This assumption on the rate J/N seems reasonable in many applications of confirmatory generalized latent factor models, as N is typically larger than J . Finally, if $N = o(J)$, then the asymptotic upper bound (15) and lower bound (17) do not match, in which case the proposed estimator may not be rate optimal. In particular, when $N = o(J)$, it is known that the lower bound for the estimation of $\Theta_{[1:N,k]}^*$ can be achieved under a linear factor model (see Bai and Li, 2012). However, this lower bound may not be achievable under a generalized latent factor model for categorical data, such as the MIRT model for binary data. This problem is worth further investigation.

We end this section by providing an alternating minimization algorithm (Algorithm 1) for solving the optimization program (13), which is computationally efficient through our parallel computing implementation using Open Multi-Processing (OpenMP; Dagum and Menon, 1998). Specifically, to handle the constraints, we adopt a projected gradient descent algorithm (e.g. Parikh and Boyd, 2014) for solving (18) and (19) in each iteration, where the

projections have closed-form solutions. Similar algorithms have been considered in other works, such as Udell et al. (2016), Zhu et al. (2016) and Bi et al. (2017), for solving optimization problems with respect to low-rank matrices. Following Corollary 2 of Grippo and Sciandrone (2000), we obtain Lemma 1 below on the convergence property of the proposed algorithm.

Lemma 1. *Any limit point of the sequence $\{\boldsymbol{\theta}_i^{(l)}, \mathbf{a}_j^{(l)}, i = 1, \dots, N, j = 1, \dots, J, l = 1, 2, \dots\}$ obtained from Algorithm 1 is a critical point of the optimization program (13).*

Algorithm 1: Alternating minimization algorithm

1. **Input:** Data $(y_{ij} : 1 \leq i \leq N, 1 \leq j \leq J)$, dimension K , constraint parameter C' , initial iteration number $l = 1$, and initial value $\boldsymbol{\theta}_i^{(0)}$ and $\mathbf{a}_j^{(0)} \in \mathcal{D}_j, i = 1, \dots, N, j = 1, \dots, J$.

2. **Alternating minimization:**

for $l = 1, 2, \dots$ **do**

for $i = 1, 2, \dots, N$ **do**

$$\boldsymbol{\theta}_i^{(l)} \in \arg \min_{\|\boldsymbol{\theta}\| \leq C'} -l_i(\boldsymbol{\theta}), \quad (18)$$

 where $l_i(\boldsymbol{\theta}) = \sum_{j=1}^J y_{ij}(\boldsymbol{\theta}^\top \mathbf{a}_j^{(l-1)}) - b(\boldsymbol{\theta}^\top \mathbf{a}_j^{(l-1)})$.

end

for $j = 1, 2, \dots, J$ **do**

$$\mathbf{a}_j^{(l)} \in \arg \min_{\mathbf{a} \in \mathcal{D}_j, \|\mathbf{a}\| \leq C'} -\tilde{l}_j(\mathbf{a}), \quad (19)$$

 where $\tilde{l}_j(\mathbf{a}) = \sum_{i=1}^N y_{ij}((\boldsymbol{\theta}_i^{(l)})^\top \mathbf{a}) - b((\boldsymbol{\theta}_i^{(l)})^\top \mathbf{a})$.

end

end

3. **Output:** Iteratively perform Step 2 until convergence. Output $\hat{\boldsymbol{\theta}}_i = \boldsymbol{\theta}_i^{(L)}$, $\hat{\mathbf{a}}_j = \mathbf{a}_j^{(L)}, i = 1, \dots, N, j = 1, \dots, J$, where L is the last iteration number.

3.4 A Non-asymptotic Error Bound

We further provide a non-asymptotic error bound as a complement to the asymptotic results. Through this error bound, the effect of design information on the estimation of latent factors is quantified for finite N and J , without requiring N and J to grow to infinity. We introduce the following definition on collections of manifest variable types.

Definition 3 (Feasible collection of subsets). *For a given $k \subset \{1, \dots, K\}$, we say a collection of subsets of $\{1, \dots, K\}$, \mathcal{A} , is feasible for the k th dimension if*

$$\{k\} = \bigcap_{S \in \mathcal{A}} S. \quad (20)$$

The concept of feasible collection is closely related to the necessary and sufficient condition (10) for the structural identifiability of the k th factor. Given the concept of feasible collection of manifest variable types, we then define an index as follows,

$$\sigma_{N,J} = \min \left\{ \frac{\max_{\mathcal{A}: \mathcal{A} \text{ is feasible}} \min_{S \in \mathcal{A}} \sigma_{|S|}(A_{[R_Q(S) \cap \{1, \dots, J\}, S]}^*)}{\sqrt{J}}, \frac{\sigma_K(\Theta_{[1:N, 1:K]}^*)}{\sqrt{N}} \right\}, \quad (21)$$

which serves as a measure of the signal strength on dimension k . We elaborate on this index. First, for a given feasible collection \mathcal{A} , the first quantity in the large brackets of (21), $\min_{S \in \mathcal{A}} \sigma_{|S|}(A_{[R_Q(S) \cap \{1, \dots, J\}, S]}^*)/\sqrt{J}$, measures the amount of information contained in manifest variables associated with the set of factors $S \in \mathcal{A}$. In particular, if for any feasible collection \mathcal{A} , there exists an $S \in \mathcal{A}$ such that manifest variables of type S do not exist, then this quantity becomes zero and thus $\sigma_{N,J} = 0$ as the second term in the brackets of (21) is nonnegative. In that case, the manifest variables essentially contain no information about the k th dimension. Second, the second term in the brackets, $\sigma_K(\Theta_{[1:N, 1:K]}^*)/\sqrt{N}$, measures the co-linearity of different dimensions of $\Theta_{[1:N, 1:K]}^*$. In summary, $\sigma_{N,J}$ is a nonnegative index, with a larger value suggesting that $\Theta_{[1:N, k]}^*$ is easier to estimate. An error bound is then established in the next theorem.

Theorem 4. *Suppose that Assumption A2 holds and $\|\boldsymbol{\theta}_i^*\| \leq C$, $\|\mathbf{a}_j^*\| \leq C$ for all $1 \leq i \leq N$ and $1 \leq j \leq J$. Then, there is a constant κ_3 , independent of N and J , such that*

$$\frac{1}{\sqrt{NJ}} E \left\| \hat{\Theta}^{(N,J)} (\hat{A}^{(N,J)})^\top - \Theta_{[1:N, 1:K]}^* A_{[1:J, 1:K]}^{*\top} \right\|_F \leq \frac{\kappa_3}{\sqrt{N \wedge J}}. \quad (22)$$

Moreover, if $\sigma_{N,J} > 4\kappa_3^{1/2}(N \wedge J)^{-1/2}$, then there exists a monotone decreasing function

$\kappa_4(\cdot) : R_+ \rightarrow R_+$ (independent of N, J and may depend on C', K and function b) such that for $N \geq N_0$ and $J \geq J_0$,

$$E \sin \angle(\Theta_{[1:N,k]}^*, \hat{\Theta}_{[k]}^{(N,J)}) \leq \frac{\kappa_4(\sigma_{N,J})}{\sqrt{N \wedge J}}. \quad (23)$$

The precise forms of the constant κ_3 and the function $\kappa_4(\cdot)$ are given in the Appendix.

The above theorem does not require N and J to grow to infinity and thus is referred to as a non-asymptotic result. Comparing with Theorem 2, the assumptions of Theorem 4 are weaker. That is, Theorem 4 does not require any limit-type assumptions as required in Theorem 2, including Assumption A1, the requirement $(\Theta^*, A^*) \in \mathcal{S}_Q$, and condition (10). Instead, it quantifies the effect of design information on the estimation of latent factors under a non-asymptotic setting, only requiring the signal level $\sigma_{N,J}$ is higher than the noise level $4\kappa_3^{1/2}(N \wedge J)^{-1/2}$. According to the proposition below, this assumption can be implied by the limit-type assumptions of Theorem 2 and thus is weaker.

Proposition 4. *If Assumption A1 and condition (10) are satisfied and $(\Theta^*, A^*) \in \mathcal{S}_Q$, then*

$$\liminf_{N, J \rightarrow \infty} \sigma_{N,J} > 0.$$

4 Further Implications

In this section, we discuss the implications of the above results in the applications of generalized latent factor models to large-scale educational and psychological measurement. In such applications, each manifest variable typically corresponds to an item in a test, each individual is a test-taker, and both N and J are large. The latent factors are often interpreted as individuals' cognitive abilities, psychological traits, etc., depending on the context of the test. It is often of interest to estimate the latent factor scores θ_i .

	1	2	3	4	5	6	...
Q^\top	1	1	0	1	1	0	...
	1	0	1	1	0	1	...
	0	1	1	0	1	1	...

Table 1: An example of the design matrix Q , which has infinite rows and $K = 3$ columns. The rows of Q are given by repeating the first 3×3 submatrix infinite times.

4.1 On the Design of Tests.

According to Theorems 1 and 2, the key to the structural identifiability and consistent estimation of factor k is

$$\{k\} = \bigcap_{k \in S, p_Q(S) > 0} S, \quad (24)$$

which provides insights on the measurement design. First, it implies that the “simple structure” design advocated in psychological measurement is a safe design. Under the simple structure design, each manifest variable is associated with one and only one factor. If each latent factor k is associated with many manifest variables that only measure factor k , or more precisely $p_Q(\{k\}) > 0$, (24) is satisfied.

Second, our result implies that a simple structure is not necessary for a good measurement design. A latent factor can still be identified even when it is always measured together with some other factors. For example, consider the Q -matrix in Table 1. Under this design, all three factors satisfy (24) even when there is no item measuring a single latent factor.

Third, (24) is not satisfied when there exists a $k' \neq k$ and $k' \in \bigcap_{k \in S, p_Q(S) > 0} S$. That is, almost all manifest variables that are associated with factor k are also associated with factor k' , in the asymptotic sense. Consequently, one cannot distinguish factor k from factor k' , making factor k structurally unidentifiable. We point out that in this case, factor k' may still be structurally identifiable, as $p_Q(\{k'\}) > 0$ is still possible.

Finally, (24) is also not satisfied when $\bigcap_{k \in S, p_Q(S) > 0} S = \emptyset$. It implies that the factor k is not structurally identifiable when the factor is not measured by a sufficient number of manifest variables.

4.2 Properties of Estimated Factor Scores

A useful result. Let $(\Theta^*, A^*) \in \mathcal{S}_Q$ be the true parameters for the person and the manifest variable populations. We start with a lemma connecting sine angle consistency and L_2 consistency.

Lemma 2. Let $\mathbf{w}, \mathbf{w}' \in R^n$ with $\mathbf{w} \neq \mathbf{0}$ and $\mathbf{w}' \neq \mathbf{0}$ and $c = \text{sign}(\cos \angle(\mathbf{w}, \mathbf{w}'))$. Then,

$$\left\| \frac{\mathbf{w}}{\|\mathbf{w}\|} - c \frac{\mathbf{w}'}{\|\mathbf{w}'\|} \right\|^2 = 2 - 2\sqrt{1 - \sin^2 \angle(\mathbf{w}, \mathbf{w}')}.$$
 (25)

Combining Theorem 2 and Lemma 2, we have the following corollary which establishes a relationship between the true person parameters and their estimates. This result is the key to the rest of the results in this section.

Corollary 2. Under Assumption A1-A2 and (10) is satisfied for some k , then there exists a sequence of random variables $c_{N,J} \in \{-1, 1\}$, such that

$$\left\| \frac{\Theta_{[1:N,k]}^*}{\|\Theta_{[1:N,k]}^*\|} - c_{N,J} \frac{\hat{\Theta}_{[k]}^{(N,J)}}{\|\hat{\Theta}_{[k]}^{(N,J)}\|} \right\| \xrightarrow{P_{\Theta^*, A^*}} 0, \quad N, J, \rightarrow \infty.$$
 (26)

Remark 5. Corollary 2 follows directly from (15). It provides an alternative view on how $\hat{\Theta}_{[k]}^{(N,J)}$ approximates $\Theta_{[1:N,k]}^*$. Since the likelihood function depends on $\Theta_{[1:N,1:K]}$ and $A_{[1:J,1:K]}$ only through $\Theta_{[1:N,1:K]} A_{[1:J,1:K]}^\top$, the scale of $\Theta_{[1:N,k]}$ is not identifiable even when it is structurally identifiable. This phenomenon is intrinsic to latent variable models (e.g. Skrondal and Rabe-Hesketh, 2004). Corollary 2 states that $\Theta_{[1:N,k]}^*$ and $\hat{\Theta}_{[k]}^{(N,J)}$ are close in Euclidian distance after properly normalized. The normalized vectors $\Theta_{[1:N,k]}^* / \|\Theta_{[1:N,k]}^*\|$ and $c_{N,J} \hat{\Theta}_{[k]}^{(N,J)} / \|\hat{\Theta}_{[k]}^{(N,J)}\|$ are both of unit length. The value of $c_{N,J}$ depends on the angle between $\Theta_{[1:N,k]}^*$ and $\hat{\Theta}_{[k]}^{(N,J)}$. Specifically, $c_{N,J} = 1$ if $\cos \angle(\Theta_{[1:N,k]}^*, \hat{\Theta}_{[k]}^{(N,J)}) > 0$ and $c_{N,J} = -1$ otherwise. In practice, especially in psychological measurement, $c_{N,J}$ can typically be determined by additional domain knowledge.

On the distribution of person population. In psychological measurement, the distribution of true factor scores is typically of interest, which may provide an overview of the population on the constructs being measured. Corollary 2 implies the following proposition on the empirical distribution of the factor scores.

Proposition 5. *Suppose assumptions A1-A2 are satisfied and furthermore (10) is satisfied for factor k . We normalize θ_{ik}^* and $\hat{\theta}_{ik}^{(N,J)}$ by*

$$v_i = \frac{\sqrt{N}\theta_{ik}^*}{\|\Theta_{[1:N,k]}^*\|} \text{ and } \hat{v}_i = \frac{c_{N,J}\sqrt{N}\hat{\theta}_{ik}^{(N,J)}}{\|\hat{\Theta}_{[k]}^{(N,J)}\|}, \quad i = 1, \dots, N, \quad (27)$$

where $c_{N,J}$ is defined and discussed in Corollary 2. Let F_N and $\hat{F}_{N,J}$ be the empirical measures of v_1, \dots, v_N and $\hat{v}_1, \dots, \hat{v}_N$, respectively. Then,

$$\text{Wass}(F_N, \hat{F}_{N,J}) \xrightarrow{P_{\Theta^*, A^*}} 0, \quad N, J, \rightarrow \infty,$$

where $\text{Wass}(\cdot, \cdot)$ denotes the Wasserstein distance between two probability measures

$$\text{Wass}(\mu, \nu) = \sup_{h \text{ is 1-Lipschitz}} \left| \int h d\mu - \int h d\nu \right|.$$

We point out that the normalization in (27) is reasonable. Consider a random design setting where θ_{ik}^* s are i.i.d. samples from some distribution with a finite second moment. Then F_N converges weakly to the distribution of $\eta/\sqrt{E\eta^2}$, where η is a random variable following the same distribution. Proposition 5 then implies that when factor k is structurally identifiable and both N and J are large, the empirical distribution of $\hat{\theta}_{1k}^{(N,J)}, \hat{\theta}_{2k}^{(N,J)}, \dots, \hat{\theta}_{Nk}^{(N,J)}$ approximates the empirical distribution of $\theta_{1k}^*, \theta_{2k}^*, \dots, \theta_{Nk}^*$ accurately, up to a scaling. Specifically, for any 1-Lipschitz function h , $\int h(x)\hat{F}_{N,J}(dx)$ is a consistent estimator for $\int h(x)F_N(dx)$ according to the definition of Wasserstein distance. Furthermore, Corollary 2 states that under the regularity conditions, $\lim_{N \rightarrow \infty} \sum_{i=1}^N (v_i - \hat{v}_i)^2/N = 0$, implying that $\sum_{i=1}^N \mathbf{1}_{\{(v_i - \hat{v}_i)^2 \geq \epsilon\}}/N = 0$, for all $\epsilon > 0$. That is, most of the \hat{v}_i s will fall into a small neighborhood of the corresponding

v_i s.

On ranking consistency. The estimated factor scores may also be used to rank individuals along a certain trait. In particular, in educational testing, the ranking provides an ordering of the students' proficiency in a certain ability (e.g., calculus, algebra, etc.). Our results also imply the validity of the ranking along a latent factor when it is structurally identifiable and N and J are sufficiently large. More precisely, we have the following proposition.

Proposition 6. *Suppose assumptions A1-A2 are satisfied and furthermore (10) is satisfied for factor k . Consider v_i and \hat{v}_i , the normalized versions of θ_{ik}^* and $\hat{\theta}_{ik}^{(N,J)}$ as defined in (27). In addition, assume that there exists a constant κ_R such that for any sufficiently small $\epsilon > 0$ and sufficiently large N ,*

$$\frac{\sum_{i \neq i'} I\{|v_i - v_{i'}| \leq \epsilon\}}{N(N-1)/2} \leq \kappa_R \epsilon. \quad (28)$$

Then,

$$\frac{\tau(\mathbf{v}, \hat{\mathbf{v}})}{N(N-1)/2} \xrightarrow{P_{(\Theta^*, A^*)}} 0, \quad N, J, \rightarrow \infty, \quad (29)$$

where $\tau(\mathbf{v}, \hat{\mathbf{v}}) = \sum_{i \neq i'} I(v_i > v_{i'}, \hat{v}_i < \hat{v}_{i'}) + I(v_i < v_{i'}, \hat{v}_i > \hat{v}_{i'})$ is the number of inconsistent pairs according to the ranks of $\mathbf{v} = (v_1, \dots, v_N)$ and $\hat{\mathbf{v}} = (\hat{v}_1, \dots, \hat{v}_N)$.

We point out that (28) is a mild regularity condition on the empirical distribution F_N . It requires that the probability mass under F_N does not concentrate in any small ϵ -neighborhood, which further implies that the pairs of individuals who are difficult to distinguish along factor k , i.e., (i, i') s that v_i and $v_{i'}$ are close, take only a small proportion among all the $(N-1)N/2$ pairs. In fact, it can be shown that (28) is true with probability tending to 1 as N grows to infinity, when θ_{ik}^* s are i.i.d. samples from a distribution with a bounded density function. Proposition 6 then implies that if we rank the individuals using \hat{v}_i (assuming $c_{N,J}$ can be consistently estimated based on other information), the proportion of incorrectly ranked pairs converges to 0. Note that $\tau(\mathbf{v}, \hat{\mathbf{v}})$ is known as the Kendall's tau

distance (Kendall and Gibbons, 1990), a widely used measure for ranking consistency.

On classification consistency. Another common practice of utilizing estimated factor scores is to classify individuals into two or more groups along a certain construct. For example, in an educational mastery test, it is of interest to classify examinees into “mastery” and “nonmastery” groups according to their proficiency in a certain ability (Lord, 1980; Bartroff et al., 2008). In measuring psychopathology, it is common to classify respondents into “diseased” and “non-diseased” groups based on a mental health disorder. We justify the validity of making classification based on the estimated factor score.

Proposition 7. *Suppose assumptions A1-A2 are satisfied and furthermore (10) is satisfied for factor k . Consider v_i and \hat{v}_i , the normalized versions of θ_{ik}^* and $\hat{\theta}_{ik}^{(N,J)}$ as defined in (27). Let $\tau_- < \tau_+$ be the classification thresholds, then*

$$\frac{\sum_{i=1}^N I\{\hat{v}_i \geq \tau_+, v_i \leq \tau_-\} + I\{\hat{v}_i \leq \tau_-, v_i \geq \tau_+\}}{N} \xrightarrow{P_{\Theta^*, A^*}} 0, \quad N, J, \rightarrow \infty, . \quad (30)$$

Considering two pre-specified thresholds τ_- and τ_+ is the well-known indifference zone formulation of educational mastery test (e.g. Bartroff et al., 2008). In that context, examinees with $v_i \geq \tau_+$ are classified into the “mastery” group and those with $v_i \leq \tau_-$ are classified into the “nonmastery” group. The interval (τ_-, τ_+) is known as the indifference zone, within which no decision is made. Proposition 7 then implies that when factor k is structurally identifiable, the classification error tends to 0 as both N and J grow to infinity.

5 Extensions

5.1 Generalized Latent Factor Models with Intercepts

As mentioned in Section 2.1, intercepts can be easily incorporated into the generalized latent factor model by restricting $\theta_{i1} = 1$. Then, a_{j1} s are the intercept parameters and $q_{j1} = 1$ for

all j . Consequently, for any S satisfying $p_Q(S) > 0$, $1 \in S$ and thus the latent factors 2 - K are not structurally identifiable according to Theorem 1. Interestingly, these factors are still structurally identifiable if we restrict to the following parameter space

$$\mathcal{S}_{Q,-} = \left\{ (\Theta, A) \in \mathcal{S}_Q : \lim_{N \rightarrow \infty} \frac{1}{N} \mathbf{1}_N^\top \Theta_{[1:N,m]} = 0 \text{ for } m \geq 2, \text{ and } \theta_{i1} = 1 \text{ for } i \in \mathbb{Z}_+ \right\}, \quad (31)$$

which requires that $\Theta_{[k]}$ and $\Theta_{[1]}$ are asymptotically orthogonal, for all $k \geq 2$.

Proposition 8. *Under Assumptions A1-A2, and assuming that $q_{j1} = 1$ for all $j \in \mathbb{Z}_+$ and $K \geq 2$, then the k th latent factor is structurally identifiable in $\mathcal{S}_{Q,-}$ if and only if*

$$\{1, k\} = \bigcap_{k \in S, p_Q(S) > 0} S, \quad (32)$$

for $k \geq 2$.

The next proposition guarantees that $\mathcal{S}_{Q,-}$ is also non-empty.

Proposition 9. *For all Q satisfying A1 and $q_{j1} = 1$ for all $j \in \mathbb{Z}_+$, and in addition $C > 1$, then $\mathcal{S}_{Q,-} \neq \emptyset$.*

Remark 6. *When having intercepts in the model, similar results of consistency and a non-asymptotic error bound can be established for the estimator*

$$\begin{aligned} (\hat{\Theta}^{(N,J)}, \hat{A}^{(N,J)}) &\in \arg \min -l(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_N, \mathbf{a}_1, \dots, \mathbf{a}_J), \\ \text{s.t. } &\|\boldsymbol{\theta}_i\| \leq C', \|\mathbf{a}_j\| \leq C', \mathbf{a}_j \in \mathcal{D}_j, \\ &\theta_{i1} = 1, \sum_{i'=1}^N \theta_{i'k} = 0, \\ &i = 1, \dots, N, j = 1, \dots, J, k = 2, \dots, K. \end{aligned} \quad (33)$$

5.2 Extension to Missing Values

Our estimator can also handle missing data which are often encountered in practice. Let $\Omega = (\omega_{ij})_{N \times J}$ be the indicator matrix of nonmissing values, where $\omega_{ij} = 1$ if Y_{ij} is observed and $\omega_{ij} = 0$ if Y_{ij} is missing. When data are completely missing at random, the joint likelihood function becomes

$$L^\Omega(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_N, \mathbf{a}_1, \dots, \mathbf{a}_J, \phi) = \prod_{i,j:\omega_{ij}=1} \exp\left(\frac{y_{ij}m_{ij} - b(m_{ij})}{\phi} + c(y_{ij}, \phi)\right)$$

and our estimator becomes

$$\begin{aligned} (\hat{\Theta}^{(N,J)}, \hat{A}^{(N,J)}) &\in \arg \min -l^\Omega(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_N, \mathbf{a}_1, \dots, \mathbf{a}_J), \\ \text{s.t. } \|\boldsymbol{\theta}_i\| &\leq C', \|\mathbf{a}_j\| \leq C', \\ \mathbf{a}_j &\in \mathcal{D}_j, i = 1, \dots, N, j = 1, \dots, J, \end{aligned} \tag{34}$$

where $l^\Omega(\Theta A^\top) = \sum_{i,j:\omega_{ij}=1} y_{ij}m_{ij} - b(m_{ij})$. Moreover, results similar to Theorem 2 can be established even when having missing data. Specifically, we assume

A3 ω_{ij} s in Ω are independent and identically distributed Bernoulli random variables with

$$P(\omega_{ij} = 1) = \frac{n}{NJ}.$$

This assumption implies that data are completely missing at random and only about n entries of $(Y_{ij})_{N \times J}$ are observed. We have the following result.

Proposition 10. *Under assumptions A1-A3 and $(\Theta^*, A^*) \in \mathcal{S}_Q$, there exists $\kappa_5 > 0$ such that,*

$$\begin{aligned} &\frac{1}{\sqrt{NJ}} E \left(\|\hat{\Theta}^{(N,J)}(\hat{A}^{(N,J)})^\top - \Theta_{[1:N,1:K]}^*(A_{[1:N,1:K]}^*)^\top\|_F \right) \\ &\leq \kappa_5 \max \left\{ \sqrt{\frac{N \vee J}{n}}, \frac{(NJ)^{1/2}}{n^{3/4}} \right\}. \end{aligned} \tag{35}$$

Moreover, if Q satisfies (10) and thus latent factor k is structurally identifiable, then there exists a constant κ_6 , such that

$$E \sin \angle(\Theta_{[1:N,k]}^*, \hat{\Theta}_{[k]}^{(N,J)}) \leq \kappa_6 \max \left\{ \sqrt{\frac{N \vee J}{n}}, \frac{(NJ)^{1/2}}{n^{3/4}} \right\}. \quad (36)$$

Remark 7. Results similar to (35) have also been derived in the matrix completion literature (e.g. Candès and Plan, 2010; Davenport et al., 2014; Cai and Zhou, 2013; Bhaskar and Javanmard, 2015) under specific statistical models with an underlying low rank structure. Proposition 10 extends the existing results on matrix completion to a generalized latent factor model.

6 A Useful Perturbation Bound on Linear Subspace

The standard approach (see, e.g., Davenport et al. (2014)) for bounding the error of the maximum likelihood estimator is by making use of the strong/weak convexity of the log-likelihood function. However, in the generalized latent factor model, the log-likelihood function is not convex in $(\Theta_{[1:N,1:K]}, A_{[1:J,1:K]})$. Thus, the standard approach is not applicable for proving (15) in Theorem 2 and similar results on the estimation of $\Theta_{[1:N,k]}$ in Theorem 4 and Proposition 10.

New technical tools are developed to handle this problem. Specifically, there are two major steps for proving Theorem 2. In the first step, we establish (14), which bounds the error on the estimation of $\Theta_{[1:N,1:K]}^* A_{[1:J,1:K]}^{*\top}$ as a whole. This step extends the error bound for exact low-rank matrix completion (Bhaskar and Javanmard, 2015) under the generalized latent variable model setting. A small estimation error of $\Theta_{[1:N,1:K]}^* A_{[1:J,1:K]}^{*\top}$ implies that the linear space spanned by the column vectors of $\Theta_{[1:N,1:K]}^*$ can be recovered up to a small perturbation. In the second step, given the result from the first step and design information, we show that the direction of the vector $\Theta_{[1:N,k]}^*$ can be recovered up to a small perturbation. This step is technically challenging and is tackled by a new perturbation bound for the

intersection of linear spaces. This perturbation bound, as introduced below, may be of independent value for the theoretical analysis of low-rank matrix estimation.

Let $\mathcal{R}(W)$ denote the column space of a matrix W . Under the conditions of Theorem 2, the result of (14) combined with the Davis-Kahan-Wedin sine theorem (see e.g. Stewart and Sun, 1990) allows us to bound $\sin \angle(\mathcal{R}(\Theta_{[1:N,S]}^*), \mathcal{R}(\hat{\Theta}_{[S]}^{(N,J)}))$, for any S satisfying $p_Q(S) > 0$, where $\angle(L, M)$ denotes the largest principal angle between two linear spaces L and M , i.e., $\sin \angle(L, M) = \max_{\mathbf{u} \in M, \mathbf{u} \neq 0} \min_{\mathbf{v} \in L, \mathbf{v} \neq 0} \sin \angle(\mathbf{u}, \mathbf{v})$. Our strategy is to bound

$$\sin \angle(\Theta_{[1:N,k]}^*, \hat{\Theta}_{[k]}^{(N,J)}) = \sin \angle(\mathcal{R}(\Theta_{[1:N,k]}^*), \mathcal{R}(\hat{\Theta}_{[k]}^{(N,J)}))$$

by $\sin \angle(\mathcal{R}(\Theta_{[1:N,1:K]}^*), \mathcal{R}(\hat{\Theta}^{(N,J)}))$ under the assumptions of Theorem 2. Note that $\mathcal{R}(\Theta_{[1:N,k]}^*) = \bigcap_{k \in S, p_Q(S) > 0} \mathcal{R}(\Theta_{[1:N,S]}^*)$ and similarly $\mathcal{R}(\hat{\Theta}_{[k]}^{(N,J)}) = \bigcap_{k \in S, p_Q(S) > 0} \mathcal{R}(\hat{\Theta}_{[S]}^{(N,J)})$. Consequently, it remains to show that if the linear spaces are perturbed slightly, then their intersection does not change much. To this end, we establish a new perturbation bound on the intersection of general linear spaces in the next proposition.

Proposition 11 (Perturbation bound for intersection of linear spaces). *Let L, M, L', M' be linear subspaces of a finite dimensional vector space. Then,*

$$\|\mathbf{P}_{L' \cap M'} - \mathbf{P}_{L \cap M}\| \leq 8 \max\{\alpha(\theta_{\min,+}(L, M)), \alpha(\theta_{\min,+}(L', M'))\}(\|\mathbf{P}_L - \mathbf{P}_{L'}\| + \|\mathbf{P}_M - \mathbf{P}_{M'}\|), \quad (37)$$

where we define $\theta_{\min,+}(L, M)$ as the smallest positive principal angle between L and M (defined as 0 if all the principal angles are 0), \mathbf{P}_M denotes the orthogonal projection onto a linear space M , and $\alpha(\theta) = 2(1 + \cos \theta)/(1 - \cos \theta)^3$. Here, the norm $\|\cdot\|$ could be any unitary invariant, uniformly generated and normalized matrix norm. In particular, if we take $\|\cdot\|$ to be the spectral norm $\|\cdot\|_2$, then we have

$$\begin{aligned} & \sin \angle(L' \cap M', L \cap M) \\ & \leq 8 \max\{\alpha(\theta_{\min,+}(L, M)), \alpha(\theta_{\min,+}(L', M'))\}(\sin \angle(L, L') + \sin \angle(M, M')). \end{aligned} \quad (38)$$

We refer the readers to Stewart and Sun (1990) for more details on principal angles between two linear spaces and on matrix norms. In particular, the spectral norm is a unitary invariant, uniformly generated and normalized matrix norm. The result in Proposition 11 holds for all linear subspaces. The right-hand side in (37) is finite if and only if $\theta_{\max,+}(L, M) \neq 0$ and $\theta_{\max,+}(L', M') \neq 0$. In our problem, $L = \mathcal{R}(\Theta_{[1:N, S_1]})$ and $M = \mathcal{R}(\Theta_{[1:N, S_2]})$ for $S_1, S_2 \subset \{1, \dots, K\}$ and $S_1 \neq S_2$. The next lemma further bounds $\alpha(\theta_{\min,+}(L, M))$ when L and M are column spaces of a matrix, which is a key step in proving (15).

Lemma 3. *Let $W \in \mathbb{R}^{N \times K} \setminus \{\mathbf{0}\}$ for some positive integer N and K , and $S_1, S_2 \subset \{1, \dots, K\}$ be such that $S_1 \setminus S_2 \neq \emptyset$ and $S_2 \setminus S_1 \neq \emptyset$, then*

$$\cos(\theta_{\min,+}(\mathcal{R}(W_{[S_1]}), \mathcal{R}(W_{[S_2]}))) \leq 1 - \frac{\sigma_{|S_1 \cup S_2|}^2(W_{[S_1 \cup S_2]})}{\|W\|_2^2}. \quad (39)$$

7 Numerical Experiments

7.1 Simulation Study I

We first verify Theorem 2 and its implications when all latent factors are structurally identifiable. Specifically, we consider $K = 5$ under the three models discussed in Section 2.1, including the linear, the MIRT and the Poisson models. Two design structures are considered, including (1) a simple structure, where $p_Q(\{k\}) = 1/5$, $k = 1, \dots, 5$ and (2) a mixed structure, where $p_Q(S) = 1/5$, $S = \{1, 2, 3\}, \{2, 3, 4\}, \{3, 4, 5\}, \{4, 5, 1\}$, and $\{5, 1, 2\}$. The true person parameters θ_i^* s and the true manifest parameters \mathbf{a}_j^* s are generated i.i.d. from distributions over the ball $\{\mathbf{x} \in \mathbb{R}^K : \|\mathbf{x}\| \leq 2.5\}$, respectively, (i.e., $C = 2.5$ in S_Q). Under these settings, all the latent factors are structurally identifiable.

For each model and each design structure, a range of J values are considered and we let $N = 25J$. Specifically, we consider $J = 100, 200, \dots, 1000$ for the linear, the MIRT, and the Poisson models. For each combination of a model, a design structure, and a J value, 50 independent datasets are generated. For each dataset, we apply Algorithm 1 to solve (13),




Figure 1: The value of $\|\hat{\Theta}^{(N,J)}(\hat{A}^{(N,J)})^\top - \Theta_{[1:N,1:K]}^* A_{[1:J,1:K]}^{*\top}\|_F^2 / (NJ)$ versus the number of manifest variables J under different simulation settings (solid line: simple structure; dashed line: mixed structure). The median, 25% quantile and 75% quantile based on the 50 independent replications are shown by the dot, lower bar, and upper bar, respectively.

where $C' = 1.2C$.

Results are shown in Figures 1-5. Figure 1 shows the trend of $\frac{1}{NJ} \|\hat{\Theta}^{(N,J)}(\hat{A}^{(N,J)})^\top - \Theta_{[1:N,1:K]}^* A_{[1:J,1:K]}^{*\top}\|_F^2$ (y -axis) when J increases (x -axis), where each panel corresponds to a model. This result verifies (14) in Theorem 2. According to these plots, the normalized squared Frobenius norm between $\Theta_{[1:N,1:K]}^* A_{[1:J,1:K]}^{*\top}$ and its estimate $\hat{\Theta}^{(N,J)}(\hat{A}^{(N,J)})^\top$ decays towards zero, as J and N increase. Figures 2-5 present results based on the first latent factor and we point out that the results based on the other latent factors are almost the same. Figure 2 is used to verify (15) in Theorem 2, showing the pattern that $\sin \angle(\Theta_{[1:N,k]}^*, \hat{\Theta}_{[k]}^{(N,J)})$ decreases as J and N increase.

Moreover, Figure 3 provides evidence on the result of Proposition 5. Displayed in Figure 3 are the histograms of v_i s and \hat{v}_i s, respectively, based on a randomly selected dataset when $J = 1000$ under the Poisson model and the simple structure. According to this figure, little difference is observed between the empirical distribution of v_i s and that of \hat{v}_i s. Similar results are observed for other datasets under all these three models when J and N are large.

Finally, Figures 4 and 5 show results of ranking and classification based on the estimated




Figure 2: The value of $\sin \angle(\Theta_{[1:N,1]}^*, \hat{\Theta}_{[1]}^{(N,J)})$ under different simulation settings (solid line: simple structure; dashed line: mixed structure). The median, 25% quantile and 75% quantile based on the 50 independent replications are shown by the dot, lower bar, and upper bar, respectively.







Figure 3: Comparison between the histogram of v_i s and that of \hat{v}_i s for the first latent factor under the Poisson Model and the simple structure.



(a) Linear Factor Model




(b) MIRT Model




(c) Poisson Factor Model


Figure 4: The Kendall's tau ranking error calculated from $\Theta_{[1:N,1]}^*$ and $\hat{\Theta}_{[1]}^{(N,J)}$, under different simulation settings (solid line: simple structure; dashed line: mixed structure). The median, 25% quantile and 75% quantile based on the 50 independent replications are shown by the dot, lower bar, and upper bar, respectively.



(a) Linear Factor Model



(b) MIRT Model



(c) Poisson Factor Model

Figure 5: The classification error calculated from $\Theta_{[1:N,1]}^*$ and $\hat{\Theta}_{[1]}^{(N,J)}$ with indifference zone $(0.13, 0.43)$ under different simulation settings (solid line: simple structure; dashed line: mixed structure). The median, 25% quantile and 75% quantile based on the 50 independent replications are shown by the dot, lower bar, and upper bar, respectively.




Figure 6: The value of $\|\hat{\Theta}^{(N,J)}(\hat{A}^{(N,J)})^\top - \Theta_{[1:N,1:K]}^* A_{[1:J,1:K]}^{*\top}\|_F^2 / (NJ)$ versus the number of manifest variables J under different simulation settings (solid line: Linear Model; dashed line: MIRT Model; dotted line: Poisson Model). The median, 25% quantile and 75% quantile based on the 50 independent replications are shown by the dot, lower bar, and upper bar, respectively.

factor scores, whose theoretical results are given in Propositions 6 and 7. The y -axes of the two figures show the normalized Kendall’s tau distance in (29) and the classification error in (30), respectively. Specifically, τ_- and τ_+ are chosen as 0.14 and 0.43 which are the 55% and 65% quantiles of the v_i s. From these plots, both the ranking and the classification errors tend to zero, as J and N grow large.

7.2 Simulation Study II

We then provide an example, in which a latent factor is not identifiable. Specifically, we consider $K = 2$ and the same latent factor models as in Study I. The design structure is given by $p_Q(\{1\}) = 1/2$ and $p_Q(\{1, 2\}) = 1/2$. The true person parameters θ_i^* s and the true manifest parameters \mathbf{a}_j^* s are generated i.i.d. from distributions over the ball $\{\mathbf{x} \in \mathbb{R}^K : \|\mathbf{x}\| \leq 3\}$, respectively, (i.e., $C = 3$ in S_Q). Under these settings, the first latent factor is structurally identifiable and the second factor is not. For each model, we consider $J = 100, 200, \dots, 1000$. The rest of the simulation setting is the same as Study I. Results are shown in Figures 6 and 7. First, Figure 6 presents the pattern that $\frac{1}{NJ} \|\hat{\Theta}^{(N,J)}(\hat{A}^{(N,J)})^\top - \Theta_{[1:N,1:K]}^* A_{[1:J,1:K]}^{*\top}\|_F^2$




Figure 7: The value of $\sin \angle(\Theta_{[1:N,1]}^*, \hat{\Theta}_{[1]}^{(N,J)})$ under different simulation settings (solid line: the first latent trait; dashed line: the second latent trait). The median, 25% quantile and 75% quantile based on the 50 independent replications are shown by the dot, lower bar, and upper bar, respectively.

decays to 0 when J increases, even when a latent factor is not structurally identifiable. This is consistent with the first part of Theorem 2. Second, Figure 7 shows the trend of $\sin \angle(\Theta_{[1:N,k]}^*, \hat{\Theta}_{[k]}^{(N,J)})$ as J increases. In particular, the value of $\sin \angle(\Theta_{[1:N,k]}^*, \hat{\Theta}_{[k]}^{(N,J)})$ stays above 0.3 for most of the data sets for the factor which is structurally unidentifiable, while it still decays towards 0 for the identifiable one.

7.3 Real Data Example

As an illustration, we apply the proposed method to analyze a personality assessment dataset based on an International Personality Item Pool (IPIP) NEO personality inventory (Johnson, 2014). This inventory is a public-domain version of the widely used NEO personality inventory (Costa and McCrae, 1985), which is designed to measure the Big Five personality factors ($K = 5$), including Neuroticism (N), Agreeableness (A), Extraversion (E), Openness to experience (O), and Conscientiousness (C). The dataset contains 20,993 individuals and 300 items. We use a subset of the dataset which contains 7,325 individuals who have answered all 300 items. The measurement design matrix has a simple structure, which is a




Figure 8: The histogram of the estimated factor scores $\hat{\Theta}_{[k]}^{(N,J)}$, $k = 1, \dots, 5$ for the IPIP-NEO dataset.

	N	A	E	O	C
N	1.00	-0.23	-0.35	-0.04	-0.34
A	-0.23	1.00	0.22	0.19	0.31
E	-0.35	0.22	1.00	0.20	0.18
O	-0.04	0.19	0.20	1.00	-0.01
C	-0.34	0.31	0.18	-0.01	1.00

Table 2: The correlation matrix between the estimated factor scores for the IPIP-NEO dataset.

safe design according to our identifiability theory. Under this design, each item only measures one personality factor and each factor is measured by 60 items. All the items are on a five-category rating scale, where reverse-worded items were reversely recorded ($1 \rightarrow 5$, $2 \rightarrow 4$, $4 \rightarrow 2$, $5 \rightarrow 1$) at the time the respondent completed the inventory. In this analysis we dichotomize them by merging categories $\{1, 2, 3\}$ and $\{4, 5\}$, respectively, and then fit the MIRT model.

The results are shown in Figures 8 through 9 and Tables 2 and Table 3. In Figure 8, the histograms of $\hat{\Theta}_{[k]}^{(N,J)}$ are given, for $k = 1, \dots, 5$, which correspond to the N, A, E, O, and C factors, respectively. As we can see, the estimated factor scores are quite normally distributed, especially for the first three factors. For the O and C factors, the distributions of the estimated factor scores are slightly right skewed. In Table 2, the correlations between the factors are calculated using the estimated factor scores. The correlations between the factors are relatively small, which are largely consistent with the existing findings in the literature




Figure 9: The boxplot of the 60 estimated unconstrained loadings for each of the five factors for the IPIP-NEO dataset.

Factor	Items	Loading	Content
N	23(N+)	2.81	Am often down in the dumps.
	4(N+)	2.66	Get stressed out easily.
	58(N-)	2.51	Know how to cope.
	25(N+)	2.38	Have frequent mood swings.
	13(N+)	2.31	Get upset easily.
A	79(A-)	2.02	Take advantage of others.
	86(A-)	1.88	Look down on others.
	84(A+)	1.82	Am concerned about others.
	98(A-)	1.80	Insult people.
	67(A-)	1.80	Distrust people.
E	123(E+)	2.93	Feel comfortable around people.
	128(E-)	2.54	Avoid contacts with others.
	124(E+)	2.48	Act comfortably with others.
	139(E-)	2.25	Avoid crowds.
	132(E+)	2.19	Talk to a lot of different people at parties.
O	196(O-)	2.28	Do not like art.
	191(O+)	2.25	Believe in the importance of art.
	226(O-)	2.11	Am not interested in abstract ideas.
	225(O+)	2.07	Enjoy thinking about things.
	229(O-)	2.05	Am not interested in theoretical discussions.
C	286(C-)	2.23	Find it difficult to get down to work.
	272(C+)	2.21	Work hard.
	283(C+)	2.15	Start tasks right away.
	289(C-)	2.11	Have difficulty starting tasks.
	285(C+)	2.10	Carry out my plans.

Table 3: The content of the top five items with highest estimated loadings on each factor for the IPIP-NEO dataset.

of Big Five personality (Digman, 1997). Figure 9 shows the box plot of the 60 unconstrained loadings for each of the five factors, according to which all the estimated loadings take value between 0 and 3 and the majority of them take value between 0.5 and 2. Table 3 lists the content of the top five items with the highest estimated loadings for each factor. These items are indeed representative of the Big Five factors.

8 Concluding Remarks

In this paper, we study how design information affects the identifiability and estimability of a general family of structured latent factor models, through both asymptotic and non-asymptotic analyses. In particular, under a double asymptotic regime where both the numbers of individuals and manifest variables grow to infinity, we define the concept of structural identifiability for latent factors. Then necessary and sufficient conditions are established for the structural identifiability of a given latent factor. Moreover, an estimator is proposed that can consistently recover all the structurally identifiable latent factors. A non-asymptotic error bound is developed to characterize the effect of design information on the estimation of latent factors, which complements the asymptotic results. In establishing these results, new perturbation bounds on the intersection of linear subspaces, as well as some other technical tools, are developed, which may be of independent theoretical interest. As shown in Section 4, our results have significant implications on the use of generalized latent factor models in large-scale educational and psychological measurement.

There are many future directions along the current work. First, it is of interest to develop methods, such as information criteria, for model comparison based on the proposed estimator. These methods can be used to select a design matrix Q that best describes the data structure when there are multiple candidates, or to determine the underlying latent dimensions. Second, the current results may be further generalized by considering more general latent factor models beyond the exponential family. For example, we may establish similar

identifiability and estimability results when the distribution of Y_{ij} is a more complicated function or even an unknown function of θ_i and \mathbf{a}_j .

Acknowledgement

We would like to thank the editors and two referees for their helpful and constructive comments. We would also like to thank Prof. Zhiliang Ying for his valuable comments. Xiaou Li's research is partially supported by the NSF grant DMS-1712657. Yunxiao Chen's research is supported in part by NAEd/Spencer postdoctoral fellowship.

References

- Andersen, E. B. (1970). Asymptotic properties of conditional maximum-likelihood estimators, *Journal of the Royal Statistical Society. Series B (Methodological)* **32**: 283–301.
- Anderson, T. W. (2003). *An introduction to multivariate statistical analysis*, Wiley, Hoboken, NJ.
- Anderson, T. W. and Rubin, H. (1956). Statistical inference in factor analysis, in J. Neyman (ed.), *Proceedings of the third Berkeley symposium on mathematical statistics and probability*, Vol. 5, University of California Press, Berkeley, CA, pp. 111–150.
- Bai, J. and Li, K. (2012). Statistical analysis of factor models of high dimension, *The Annals of Statistics* **40**: 436–465.
- Bartholomew, D. J., Knott, M. and Moustaki, I. (2011). *Latent variable models and factor analysis: A unified approach*, Wiley, Hoboken, NJ.
- Bartroff, J., Finkelman, M. and Lai, T. L. (2008). Modern sequential analysis and its applications to computerized adaptive testing, *Psychometrika* **73**: 473–486.

- Bhaskar, S. A. and Javanmard, A. (2015). 1-bit matrix completion under exact low-rank constraint, *49th Annual Conference on Information Sciences and Systems (CISS)*, IEEE, pp. 1–6.
- Bi, X., Qu, A., Wang, J. and Shen, X. (2017). A group-specific recommender system, *Journal of the American Statistical Association* **112**: 1344–1353.
- Bing, X., Bunea, F., Ning, Y. and Wegkamp, M. (2017). Adaptive estimation in structured factor models with applications to overlapping clustering, *arXiv preprint:1704.06977* .
- Cai, T. and Zhou, W.-X. (2013). A max-norm constrained minimization approach to 1-bit matrix completion, *The Journal of Machine Learning Research* **14**: 3619–3647.
- Candès, E. J. and Plan, Y. (2010). Matrix completion with noise, *Proceedings of the IEEE* **98**: 925–936.
- Chen, Y., Li, X. and Zhang, S. (2019). Joint maximum likelihood estimation for high-dimensional exploratory item factor analysis, *Psychometrika* **84**: 124–146.
- Costa, P. T. and McCrae, R. R. (1985). *The NEO personality inventory*, Psychological Assessment Resources, Odessa, FL.
- Dagum, L. and Menon, R. (1998). OpenMP: An industry standard API for shared-memory programming, *Computational Science & Engineering, IEEE* **5**: 46–55.
- Davenport, M. A., Plan, Y., van den Berg, E. and Wootters, M. (2014). 1-bit matrix completion, *Information and Inference* **3**: 189–223.
- Davis, C. (1963). The rotation of eigenvectors by a perturbation, *Journal of Mathematical Analysis and Applications* **6**: 159–173.
- Digman, J. M. (1997). Higher-order factors of the Big Five, *Journal of Personality and Social Psychology* **73**: 1246– 1256.

- Gatignon, H. (2003). *Statistical analysis of management data*, Springer, New York, NY.
- Ghosh, M. (1995). Inconsistent maximum likelihood estimators for the Rasch model, *Statistics & Probability Letters* **23**: 165–170.
- Goldberg, L. R. (1993). The structure of phenotypic personality traits., *American Psychologist* **48**: 26–34.
- Grayson, D. and Marsh, H. W. (1994). Identification with deficient rank loading matrices in confirmatory factor analysis: Multitrait-multimethod models, *Psychometrika* **59**: 121–134.
- Grippo, L. and Sciandrone, M. (2000). On the convergence of the block nonlinear Gauss–Seidel method under convex constraints, *Operations Research Letters* **26**: 127–136.
- Gu, Y. and Xu, G. (2018). The sufficient and necessary condition for the identifiability and estimability of the DINA model, *Psychometrika* . To appear.
- Haberman, S. J. (1977). Maximum likelihood estimates in exponential response models, *The Annals of Statistics* **5**: 815–841.
- Holland, P. W. (1990). On the sampling theory foundations of item response theory models, *Psychometrika* **55**: 577–601.
- Johnson, J. A. (2014). Measuring thirty facets of the Five Factor Model with a 120-item public domain inventory: Development of the IPIP-NEO-120, *Journal of Research in Personality* **51**: 78–89.
- Kendall, M. G. and Gibbons, J. D. (1990). *Rank correlation methods*, Oxford University Press, New York, NY.
- Lord, F. M. (1980). *Applications of item response theory to practical testing problems*, Routledge, New York, NY.

- McCullagh, P. and Nelder, J. A. (1989). *Generalized Linear Models*, CRC Press, Boca Raton, FL.
- Neyman, J. and Scott, E. L. (1948). Consistent estimates based on partially consistent observations, *Econometrica* **16**: 1–32.
- Owen, A. B. and Wang, J. (2016). Bi-cross-validation for factor analysis, *Statistical Science* **31**: 119–139.
- Parikh, N. and Boyd, S. (2014). Proximal algorithms, *Foundations and Trends® in Optimization* **1**: 127–239.
- Reckase, M. (2009). *Multidimensional item response theory*, Springer, New York, NY.
- Shapiro, A. (1985). Identifiability of factor analysis: Some results and open problems, *Linear Algebra and its Applications* **70**: 1–7.
- Skrondal, A. and Rabe-Hesketh, S. (2004). *Generalized latent variable modeling: Multilevel, longitudinal, and structural equation models*, CRC Press, Boca Raton, FL.
- Stewart, G. and Sun, J. (1990). *Matrix perturbation theory*, Academic Press, Cambridge, MA.
- Thompson, B. (2004). *Exploratory and confirmatory factor analysis: Understanding concepts and applications*, American Psychological Association, Washington, DC.
- Udell, M., Horn, C., Zadeh, R. and Boyd, S. (2016). Generalized low rank models, *Foundations and Trends® in Machine Learning* **9**: 1–118.
- Wedin, P. (1972). Perturbation bounds in connection with singular value decomposition, *BIT Numerical Mathematics* **12**: 99–111.
- Xu, G. (2017). Identifiability of restricted latent class models with binary responses, *The Annals of Statistics* **45**: 675–707.

Xu, G. and Zhang, S. (2016). Identifiability of diagnostic classification models, *Psychometrika* **81**: 625–649.

Zhu, Y., Shen, X. and Ye, C. (2016). Personalized prediction and sparsity pursuit in latent factor models, *Journal of the American Statistical Association* **111**: 241–252.