Reflections on Analytical Choices in the Scaling Model for Test Scores in International
Large-Scale Assessment Studies

Alexander Robitzsch^{1,2} and Oliver Lüdtke^{1,2}

¹ IPN – Leibniz Institute for Science and Mathematics Education, Kiel, Germany

² Centre for International Student Assessment (ZIB), Germany

Author Note

Correspondence concerning this article should be sent to Alexander Robitzsch,
IPN – Leibniz Institute for Science and Mathematics Education, Olshausenstraße 62, 24118
Kiel, Germany. Email: robitzsch@leibniz-ipn.de

Abstract

International large-scale assessments (LSAs) such as the Programme for International Student Assessment (PISA) provide important information about the distribution of student proficiencies across a wide range of countries. The repeated assessments of these content domains offer policymakers important information for evaluating educational reforms and received considerable attention from the media. Furthermore, the analytical strategies employed in LSAs often define methodological standards for applied researchers in the field. Hence, it is vital to critically reflect the conceptual foundations of analytical choices in LSA studies. This article discusses methodological challenges in selecting and specifying the scaling model used to obtain proficiency estimates from the individual student responses in LSA studies. We distinguish design-based inference from model-based inference. It is argued that for the official reporting of LSA results, design-based inference should be preferred because it allows for a clear definition of the target of inference (e.g., country mean achievement) and is less sensitive to specific modeling assumptions. More specifically, we discuss five analytical choices in the specification of the scaling model: (1) Specification of the functional form of item response functions, (2) the treatment of local dependencies and multidimensionality, (3) the consideration of test-taking behavior for estimating student ability, and the role of country differential items functioning (DIF) for (4) cross-country comparisons, and (5) trend estimation. This article's primary goal is to stimulate discussion about recently implemented changes and suggested refinements of the scaling models in LSA studies.

Keywords: large-scale assessment, item response models, scaling, linking, differential item functioning, partial invariance, item response function, trend estimation, PISA, survey statistics, educational assessment

1. Introduction

In the last two decades, international large-scale assessments (LSAs) have provided important information about the distribution of student proficiencies across a wide range of countries and age groups. For example, every three years since 2000, the Programme for International Student Assessment (PISA) reported international comparisons of student performance in three content areas (reading, mathematics, and science; OECD, 2014). The repeated assessments of these content domains provide policymakers with important information for the evaluation of educational reforms and also received considerable attention from the media. Furthermore, LSAs provide unique research opportunities (Singer & Braun, 2018) that are increasingly used by researchers from different fields to investigate the relations between student proficiency and other cognitive and noncognitive variables such as students' gender, socioeconomic status, self-concept, or attitudes (e.g., Guo et al., 2021; Keller et al., 2020; Marsh, 2016). From the beginning, LSAs have been confronted with many methodological challenges (Rutkowski et al., 2013). In addition, it seems that the analytical strategies employed in LSAs often define methodological standards for applied researchers in the field. Hence, it is vital to critically reflect the conceptual foundations of analytical choices in LSA studies.

In the present article, we reflect on methodological challenges in the selection and specification of the scaling model used to obtain proficiency estimates from the individual student responses in LSA studies. In our discussion, we distinguish between design-based inference (based on sampling designs for specific populations of persons and test items) and model-based inference (based on specific assumptions of statistical models). It is argued that for the official reporting of LSA results, design-based inference should be preferred because it allows for a clear definition of the target of inference (e.g., country mean achievement) and is less sensitive to specific modeling assumptions. More specifically, we discuss five specific analytical choices for the scaling model that received considerable attention in the methodological literature, and that can affect the reporting of LSA results: (1) Specification of the functional form of item response functions, (2) the treatment of local dependencies and multidimensionality, (3) the consideration of test-taking behavior for estimating student ability, and the role of country differential items functioning (DIF) for (4) cross-country comparisons, and (5) trend estimation. The main goal of this article is to stimulate discussion about the role of recent changes that have been implemented in the scaling models of LSA studies (with a particular emphasis on PISA), or that were suggested by methodologists as further refinements of the currently used scaling models.

2. Model-Assisted Design-Based Inference

2.1. Model-Assisted Design-Based Inference for Persons

In the remainder of the article, we consider statistics (e.g., mean, standard deviation, quantiles) of the distribution of an ability variable (e.g., reading ability). Let θ_n denote a corresponding ability of person n . In the usual sampling design of LSA studies, not all students in a population (e.g., a country) are sampled. Frequently, stratified multi-stage sampling is employed in which schools are sampled in the first stage, and students within a school are sampled in the second stage (Meinck, 2020). Consequently, not all students within a country have the same probability of being sampled, and it is important to take into account the different selection probabilities when inferring from the sample to the population. Hence, student weights $w_{\mathcal{P},n}$ are used where $w_{\mathcal{P},n}$ is the inverse probability that person n is sampled (Meinck, 2020; Rust et al., 2017). The subscript $\mathcal{P}$ indicates that the weights refer to the population of persons (e.g., students). The inference for a statistic of the ability distribution (e.g., mean achievement) from the sample to the population of students in a country is also referred to as a design-based inference (Lohr, 2010; Särndal et al., 2003).

We illustrate the typical approach for statistical inference in LSA studies for the estimation of a mean μ and a variance σ^2 of an ability distribution (e.g., reading ability for a country in the PISA study). Suppose that there are N sampled students within a country and abilities θ_n would have been directly observed for all $n = 1, \dots, N$. Then, in a design-based (db) approach, sample estimates for the mean μ and the variance σ^2 are given by

$$\hat{\mu}_{\text{db}} = \frac{\sum_{n=1}^N w_{\mathcal{P},n} \theta_n}{\sum_{n=1}^N w_{\mathcal{P},n}} \quad \text{and} \quad \hat{\sigma}_{\text{db}}^2 = \frac{\sum_{n=1}^N w_{\mathcal{P},n} (\theta_n - \hat{\mu}_{\text{db}})^2}{\sum_{n=1}^N w_{\mathcal{P},n}}, \quad (1)$$

where values θ_n are weighted by student weights $w_{\mathcal{P},n}$. There are two obstacles in applying the estimation formulas in Equation (1). First, abilities cannot be directly measured in LSA studies but have to be inferred from a multivariate vector $\mathbf{X}_n$ of discrete item responses. In the following, we only consider dichotomous items for the sake of notational simplicity. A scoring rule $\hat{\theta}_n = f(\mathbf{X}_n)$ is required that relates the item responses to ability values.

Typically, the ability is considered as a latent variable, and estimated abilities $\hat{\theta}_n$ are prone to measurement errors. The extent of measurement errors relies on a specified measurement model (i.e., an item response theory (IRT) model). The probability for item responses are modeled conditional on a latent ability θ_n by posing a local independence assumption:

$$P(\mathbf{X}_n = \mathbf{x} | \theta_n) = \prod_{i=1}^I P(X_{ni} = x_i | \theta_n; \boldsymbol{\gamma}_i), \quad (2)$$

where I is the number of items, X_{ni} is the item response of student n on item i , and $\boldsymbol{\gamma}_i$ denotes a vector of item parameters for item i . Note that error-prone ability estimates result in biased

estimates of distribution parameters, particularly for the standard deviation and quantiles, and biased correlation of abilities with covariates (Lechner et al., 2021; Wu, 2005). The second obstacle in LSA studies like PISA is that not all students receive items in all ability domains (OECD, 2014; see also Frey et al., 2009). Hence, imputation procedures must be used to borrow for each student information from administered ability domains to obtain estimates for non-administered ability domains (Little & Rubin, 2002). The issue of non-administered ability domains is addressed by using a so-called latent background model (LBM; Mislevy, 1991). For a vector of observed covariates $\mathbf{Z}_n$ (e.g., variables such as gender or sociodemographic status), the LBM for a target ability θ_n (e.g., reading) and a vector of additional abilities $\boldsymbol{\eta}_n$ (e.g., mathematics and science) is defined as

$$\begin{pmatrix} \theta_n \\ \boldsymbol{\eta}_n \end{pmatrix} = \mathbf{Z}_n \mathbf{B} + \boldsymbol{\varepsilon}_n, \quad \boldsymbol{\varepsilon}_n \sim \text{MVN}(\mathbf{0}, \mathbf{T}), \quad (3)$$

where MVN denotes the multivariate normal distribution. Note that Equation (3) can be rewritten as a conditional unidimensional normal distribution

$$\theta_n = (\mathbf{Z}_n | \boldsymbol{\eta}_n) \mathbf{B}^* + \varepsilon_n^*, \quad \varepsilon_n^* \sim N(0, \tau^2). \quad (4)$$

It can be seen in Equation (4) that in the LBM, the ability θ_n is inferred from $\mathbf{Z}_n$ and $\boldsymbol{\eta}_n$. If items are administered in the target ability domain for student n , the IRT model in Equation (2) typically provides the major amount of information for the target ability, while for non-administered ability domains, only the LBM delivers information for the ability θ_n . That is, administered ability domains $\boldsymbol{\eta}_n$ and covariates $\mathbf{Z}_n$ are used for imputing the target ability. In the operational practice of LSA studies, the imputations are called plausible values (Mislevy, 1991; von Davier & Sinharay, 2014). Plausible values $(\tilde{\theta}_n, \tilde{\boldsymbol{\eta}}_n)$ are drawn from subject-specific posterior distributions $P(\theta_n, \boldsymbol{\eta}_n | \mathbf{X}_n, \mathbf{Z}_n)$:

$$\text{simulate } \begin{pmatrix} \tilde{\theta}_n \\ \tilde{\boldsymbol{\eta}}_n \end{pmatrix} \text{ from } P(\theta_n, \boldsymbol{\eta}_n | \mathbf{X}_n, \mathbf{Z}_n) \propto P(\mathbf{X}_n | \theta_n, \boldsymbol{\eta}_n; \boldsymbol{\gamma}) P(\theta_n, \boldsymbol{\eta}_n | \mathbf{Z}_n; \mathbf{B}, \mathbf{T}) \quad (5)$$

In the case of a unidimensional ability and normally distributed measurement errors of $\hat{\theta}_n$, plausible values $\tilde{\theta}_n$ can be written as

$$\tilde{\theta}_n = \rho_c \hat{\theta}_n + (1 - \rho_c)(\mathbf{Z}_n | \boldsymbol{\eta}_n)^T \mathbf{B}^* + e_n, \quad e_n \sim N(0, \kappa^2) \quad (6)$$

where the conditional reliability ρ_c and the posterior variance κ^2 are determined by

$$\rho_c = \frac{\tau^2}{\tau^2 + \text{SE}^2(\hat{\theta}_n)} \quad \text{and} \quad \kappa^2 = (1 - \rho_c)(1 - R^2), \quad (7)$$

and where $R^2 = \frac{\tau^2}{\text{Var}(\theta_n)}$ is the proportion of explained variance in Equation (4) (see Mislevy, 1991).

If the IRT model in Equation (2) is misspecified, the likelihood part $P(\mathbf{X}_n|\theta_n, \boldsymbol{\eta}_n; \boldsymbol{\gamma})$ in Equation (5) will be misspecified. As a consequence, the model-implied reliability will be incorrect and plausible values do not correctly reflect the uncertainty associated with the ability variable θ_n . In practice, item parameters $\boldsymbol{\gamma}$ are fixed in Equation (5) when drawing plausible values and the likelihood part can be written as a function of θ_n and $\boldsymbol{\eta}_n$, that is, there is a multidimensional function $h_n(\theta_n, \boldsymbol{\eta}_n) = P(\mathbf{X}_n|\theta_n, \boldsymbol{\eta}_n; \boldsymbol{\gamma})$. The amount of error associated with $(\theta_n, \boldsymbol{\eta}_n)$ is quantified by the peakedness in h_n . The measurement error assumption can be modified by adjusting the function h_n to be steeper (i.e., increase reliability) or more flat (i.e., decrease reliability; see Chandler & Bate, 2007; Mislevy, 1990). In more detail, the unidimensional person-specific likelihood function is approximated with an unnormalized normal density function; that is,

$$h_n(\theta) = P(\mathbf{X}_n|\theta_n; \boldsymbol{\gamma}) = c_n \phi(\theta; \mu_n, \sigma_n) \quad (8)$$

where ϕ is the normal density and c_n is a scaling factor. We set $\mu_n = \hat{\theta}_n$ and $\sigma_n = SE(\hat{\theta}_n)$. Methods that resample items (see Section 2.2) can be used to estimate reliability (or the standard error $SE(\hat{\theta}_n)$) in misspecified IRT models. Hence, the person-specific standard deviation σ_n in Equation (8) can be modified by posing different assumptions about the reliability of the ability scores.

The statistical inference in LSA studies almost exclusively relies on plausible values (von Davier & Sinharay, 2014). It is evident that misspecifications in the LBM vanish with an increasing number of items because individual standard errors $SE^2(\hat{\theta}_n)$ converge to zero (and ρ_c in Equation (7) will be close to 1; see also Marsman et al., 2016). The current approach in LSA studies that relies on plausible values can be described as a model-assisted design-based inference (Binder & Roberts, 2003; Brewer, 2013; Little, 2004; Särndal et al., 2003; Ståhl et al., 2010) because the inference to a concrete population of students in a country is warranted but—at the same time—a model (IRT model and the LBM) is utilized for computation of plausible values that substitute the non-observable ability θ_n . In practice, for reducing the simulation error and enabling the estimation of standard errors with imputed data, several plausible values (e.g., $M = 10$) are generated; that is, for each student n , there are M plausible values $\tilde{\theta}_n^{(m)}$ ($m = 1, \dots, M$). The sample estimates based on the m -th plausible value is given as

$$\hat{\mu}_{\text{db,PV}(m)} = \frac{\sum_{n=1}^N w_{\mathcal{P},n} \tilde{\theta}_n^{(m)}}{\sum_{n=1}^N w_{\mathcal{P},n}} \quad \text{and} \quad \hat{\sigma}_{\text{db,PV}(m)}^2 = \frac{\sum_{n=1}^N w_{\mathcal{P},n} (\tilde{\theta}_n^{(m)} - \hat{\mu}_{\text{db,PV}(m)})^2}{\sum_{n=1}^N w_{\mathcal{P},n}} . \quad (9)$$

The final point estimates $\hat{\mu}_{\text{db,PV}}$ and $\hat{\sigma}_{\text{db,PV}}^2$ are obtained as averages of the estimates $\hat{\mu}_{\text{db,PV},(m)}$ and $\hat{\sigma}_{\text{db,PV},(m)}^2$, respectively (Mislevy, 1991).

Note that the subject-specific posterior distribution $P(\theta_n, \boldsymbol{\eta}_n | \mathbf{X}_n, \mathbf{Z}_n)$ is a continuous function of θ_n and $\boldsymbol{\eta}_n$. Hence, the statistics in Equation (9) relying on plausible values are shortcuts for evaluating person-specific integrals. In more detail, for an infinite number of plausible values, the estimates can be written

$$\hat{\mu}_{\text{db,PV}} = \frac{\sum_{n=1}^N w_{\mathcal{P},n} \int \theta_n P(\theta_n, \boldsymbol{\eta}_n | \mathbf{X}_n, \mathbf{Z}_n) d\theta_n d\boldsymbol{\eta}_n}{\sum_{n=1}^N w_{\mathcal{P},n}} \quad (10)$$

$$\hat{\sigma}_{\text{db,PV}}^2 = \frac{\sum_{n=1}^N w_{\mathcal{P},n} \int (\theta_n - \hat{\mu}_{\text{db,PV}})^2 P(\theta_n, \boldsymbol{\eta}_n | \mathbf{X}_n, \mathbf{Z}_n) d\theta_n d\boldsymbol{\eta}_n}{\sum_{n=1}^N w_{\mathcal{P},n}} \quad (11)$$

Comparing these estimates with the design-based estimates $\hat{\mu}_{\text{db}}$ and $\hat{\sigma}_{\text{db}}^2$ (see Equation (1)) makes evident that $\hat{\mu}_{\text{db,PV}}$ and $\hat{\sigma}_{\text{db,PV}}^2$ depend on both the design (i.e., relying on weights $w_{\mathcal{P},n}$) and model assumptions (i.e., relying on individual posterior distributions $P(\theta_n, \boldsymbol{\eta}_n | \mathbf{X}_n, \mathbf{Z}_n)$). Hence, the choice of a particular IRT model (see Equation (2)) and the specification of the LBM (see Equation (3)) has the potential to change the meaning of θ , and, hence, can affect the meaning of μ and σ^2 and their corresponding estimates.

From Equations (10) and (11), it can be seen that the design-based inference is represented by including student weights $w_{\mathcal{P},n}$, but it is model-assisted because the ability θ_n is represented by the posterior distribution $P(\theta_n, \boldsymbol{\eta}_n | \mathbf{X}_n, \mathbf{Z}_n)$ that relies on the chosen IRT model and the LBM. In a hybrid approach of design-based inference and model-based inference (see Ståhl et al., 2016), subjects can additionally be weighted by including weights $v_{\mathcal{P},n}$ according to their fit to a statistical model. For example, model-based student-specific weights $v_{\mathcal{P},n}$ can be derived according to their fit to the scaling model (person fit; see Hong & Cheng, 2019; Schuster & Yuan, 2011). In such an approach, students whose item responses are atypical with respect to the IRT model (e.g., non-scalable students; see Haertel, 1989) would be downweighted compared to students whose item responses are consistent with the IRT model. Corresponding sample estimates in such a hybrid design-model-based (dmb) are given by

$$\hat{\mu}_{\text{dmb,PV}} = \frac{\sum_{n=1}^N w_{\mathcal{P},n} v_{\mathcal{P},n} \int \theta_n P(\theta_n, \boldsymbol{\eta}_n | \mathbf{X}_n, \mathbf{Z}_n) d\theta_n d\boldsymbol{\eta}_n}{\sum_{n=1}^N w_{\mathcal{P},n} v_{\mathcal{P},n}} \quad (12)$$

$$\hat{\sigma}_{\text{dmb,PV}}^2 = \frac{\sum_{n=1}^N w_{\mathcal{P},n} v_{\mathcal{P},n} \int (\theta_n - \hat{\mu}_{\text{dmb,PV}})^2 P(\theta_n, \boldsymbol{\eta}_n | \mathbf{X}_n, \mathbf{Z}_n) d\theta_n d\boldsymbol{\eta}_n}{\sum_{n=1}^N w_{\mathcal{P},n} v_{\mathcal{P},n}} \quad (13)$$

It might be tempting to identify subgroups of students that do not fit the IRT model as a threat of validity and, subsequently, to eliminate these students from the final analysis by effectively setting $v_{p,n}$ to zero. Clearly, the estimates $\hat{\mu}_{db,pv}$ and $\hat{\mu}_{dmb,pv}$ will turn out to be different in practice and likely target different estimands. There is a danger that estimates in Equations (12) and (13) generalize to a different population of students compared to the model-assisted design-based estimates in Equations (10) and (11). In a hybrid design-model-based inference, the specification of a model allows the target estimand to differ from the estimand in a design-based approach because, in the former, observations are weighted by $w_{p,n}v_{p,n}$, while in the latter observations are weighted by $w_{p,n}$.

Standard errors can be computed by resampling methods (e.g., jackknife or balanced repeated replication methods; Kolenikov, 2010; Rust et al., 2017) in which subgroups of students are resampled. The multi-stage clustered sampling with explicit and implicit stratification can easily be accommodated in such resampling designs (Meinck, 2020).

We tend to argue that a fully design-based inference should be the first option in LSA studies. Obviously, this could only be realized if an infinite (or very large) number of items would be administered in the ability domain of interest so that the variance of the measurement error is negligible. However, the number of administered items in most applications is not so large that measurement errors in abilities can be neglected. Hence, the statistical inference employed in LSA studies (i.e., model-assisted design-based inference) depends on measurement error assumptions in the IRT model and the specified LBM. However, we would argue that misspecifications in the IRT model can be accepted (see Section 3.1) because the choice of the IRT model should be driven by the meaning of the ability variable (e.g., equal weighting of items in the scoring rule) and not by the model fit. In contrast, the degree of misspecification in the LBM should be minimized, even though it can be challenging to adequately treat the high dimensionality of the predictor variables (Grund et al., 2021; von Davier, Khorramdel, et al., 2019). Overall, we believe that the hybrid design-model-based inference poses threats to validity because the fit of each subject in a model can redefine the contribution of subjects by additionally incorporating weights $v_{p,n}$ in the analysis. Thus, a statistical model (and, hence, psychometrics) is allowed to change the target of inference. We would prefer a design-based approach that is less sensitive to specific modeling assumptions when reporting LSA results.

2.2. Design-Based or Model-Based Inference for Items?

In the previous subsection, we discussed the kind of statistical inference for the population of persons. It is not apparent which kind of statistical inference is needed to

represent the process of choosing test items in LSA studies. The test items should cover the ability domain defined by the test framework (test blueprint; see also Pellegrino & Chudowsky, 2003; Reckase, 2017). It might be legitimate to assume that there exists a larger population of test items (henceforth, labeled by $\mathcal{J}$) from which the items are chosen in a particular study and true ability values would be defined as outcomes in a study in which all items from the population would have been chosen (Cronbach & Shavelson, 2004; see also Kane, 1982; Brennan, 2001). Interestingly, it has been argued that classical test theory (CTT) or generalizability theory (GT; Cronbach et al., 1963) treats items in a study as random and, as a consequence, allows the inference to a larger set of items in a population of items (see also Nunnally & Bernstein, 1994; Markus & Borsboom, 2013). In contrast, IRT treats items as fixed (Brennan, 2010) and restricts the statistical inference to the items chosen in a test. This distinction is strongly related to the question of whether the representation of item responses in the ability θ_n follows a design-based (i.e., CTT or GT) or a model-based inference (i.e., IRT). In CTT or GT, items are defined as exchangeable by posing assumptions about the sampling process. Notably, if the selection (or sampling) of items from the domain of test items is appropriately conducted, the inference for the ability from the chosen items to the population of items would be valid. In a design-based perspective, substantive theory (e.g., by test domain experts, item developers) should define the contribution of each chosen item. In more detail, there are a priori defined item-specific weights $w_{j,i}$ that enter the scoring rule for the ability estimate $\hat{\theta}_n$:

$$\hat{\theta}_n = f\left(\sum_{i=1}^I w_{j,i} x_{ni}\right) \quad (14)$$

If the administered test mimics the population of items, all item weights will be set to be equal to each other, that is $w_{j,i} = 1$ for all $i = 1, \dots, I$. If the item selection in a study is adequately made, a subsequent post-hoc elimination of items based on a fit in the IRT model (e.g., item fit statistics for the IRT model in Equation (2)) potentially changes the target of inference (Brennan, 1998). By choosing an IRT model, there are model-based derived item weights $v_{j,i}(\theta)$ (so-called locally optimal item weights) that define a local scoring rule for the ability (see Equation (A4) in Appendix A)

$$\hat{\theta}_n = f\left(\sum_{i=1}^I v_{j,i}(\theta_n) x_{ni}; \theta_n\right) \quad , \quad (15)$$

where the item weights $v_{j,i}(\theta_n)$ are given by (Birnbaum, 1968; Chiu & Camilli, 2013; Yen & Fitzpatrick, 2006)

$$v_{j,i}(\theta) = \frac{P'_i(\theta)}{P_i(\theta)[1 - P_i(\theta)]} \quad . \quad (16)$$

The main consequence of the local scoring rule in Equation (15) is that the choice of the IRT model implicitly defines the contribution of items in the ability, and the model-based approach (see Equation (15)) can deviate from the design-based approach (see Equation (14)) in which weights $w_{j,i}$ are defined by sampling considerations (Camilli, 2018). By posing a particular IRT model, locally optimal item weights $v_{j,i}(\theta_n)$ are determined that provide the best-fitting model in terms of the potentially misspecified maximum likelihood function (White, 1982). Items that are most informative for θ_n in the IRT model receive the largest weights which, in turn, can influence the interpretations of the ability score.

Camilli (2018) defined effective country-specific item weights $v_{j,ic}$ that integrate locally optimal item weights with respect to the country-specific ability density f_c :

$$v_{j,ic} = \int v_{j,i}(\theta) f_c(\theta) d\theta \quad . \quad (17)$$

The quantity $v_{j,ic}$ allows the evaluation of whether the effective contribution of an item in the ability score θ varies across countries.

If an IRT model would be used for scoring, the measurement error in estimated abilities $\hat{\theta}_n$ is mainly driven by the observed information function (Magis, 2015). Hence, the statistical model defines the extent of error associated with ability scores. In contrast, in a design-based approach of CTT or GT, sampling assumptions regarding selecting items from the population of items define the extent of measurement errors. In such a design-based perspective, no assessment of model fit for the set of item responses $\mathbf{X}_n$ is required. For example, the use of Cronbach's alpha (Cronbach, 1951) as a reliability measure does not require that a model with equal item loadings and uncorrelated residual errors have to fit the data of item responses (Cronbach, 1951; Cronbach & Shavelson, 2004; Meyer, 2010; Nunnally & Bernstein, 1994; Tryon, 1957). In the same manner, as for persons, resampling methods for items can be used for determining standard errors in estimated abilities (Liou & Yu, 1991; Wainer & Thissen, 1987; Wainer & Wright, 1980) by resampling items or groups of items for which abilities are reestimated (see also Michaelides & Haertel, 2014). It is also possible to include additional dependence by item stratification (e.g., multiple test components; Cronbach et al., 1965; Meyer, 2010) or item clustering (e.g., due to the arrangement of items in testlets, that is, several items share a common item stimulus; Bradlow et al., 1999) in resampling methods for items.

We tend to favor the scoring rules from a design-based perspective in Equation (14) over the model-based perspective in Equation (15) because, in our view, substantive theory

should define the contribution of items in the ability score for carefully constructed test items. We also want to emphasize that item fit statistics are related to the local fit of single items in an IRT model that treats items as fixed. Notably, the assessment of item fit statistics does not follow the perspective that treats items as random, and removing items from the computation of the ability changes the target of statistical inference. We elaborate on these issues in detail in Section 3.

2.3. A Plea for a Symmetric Role of Persons and Items

In the last two subsections, we discussed statistical inference for the populations of persons and items in LSA studies. For both populations (persons and items), design-based or model-based or hybrid variants of statistical inference can be employed. In most LSA studies, statistical inference for the population of persons is primarily handled under a design-based perspective, while model-based inference is also present for the population of items. Based on the previous arguments, we would argue that persons and items should have symmetric roles in LSA studies. For both facets, we believe that design-based inference should rule out model-based inference. There seems to be consent among researchers that students who do not fit a particular IRT model should not be removed from the analysis. By doing so, the sample of students would no longer be representative of the population of students. We argue that the same perspective should be taken for items: one should not simply remove items from the scoring rule for ability or country comparisons because they do not follow a particular IRT model. In contrast, items should be considered random, and IRT models should be regarded as statistical devices used to achieve the inferential goals of LSA studies. In this sense, these psychometric models merely define estimating equations, and the fit of the chosen model is not of central relevance. The employed likelihood functions in the estimation of abilities in LSA studies are likely to be misspecified. We argue that their sole role is the (implicit) definition of target estimands of interest. Statistical inference should preferably rely on resampling methods for persons and items because these do not rely on a correctly specified statistical model. Also, note that local fit statistics can be computed for each person and item. However, atypical persons or items (with respect to a model) do not invalidate statistical inference in a design-based perspective.

3. Specific Analytical Choices in Scaling Models

In the following, we discuss five topics that are of central relevance in the specification of the scaling modeling in LSA studies: (1) the choice of the functional of the item response function, (2) the role of local dependence and multidimensionality, (3) the treatment of additional information from the test-taking behavior (e.g., response times), and

the role of country DIF in (4) cross-country comparisons and (5) trend estimation. In this discussion, we highlight the consequences of a design-based perspective for the specification of the scaling model.

3.1. Functional Form of Item Response Functions

As argued in Section 2, the choice of the IRT model can affect the meaning of the latent ability variable θ_n . Of particular importance is the specification of the item response function (IRF) that describes the relationship between item responses and ability. In the following, we discuss the most common IRFs and use locally optimal weights (see Section 2.2.) to show how the choice of different IRFs affects item contributions in the scoring rule for the latent ability variable (see Equation (19)).

Probably the most popular IRT model is the one-parameter logistic (1PL) IRT model (also known as the Rasch model; Rasch, 1960) that employs the IRF

$$P(X_{ni} = 1|\theta_n) = \Psi(\theta_n - b_i) , \quad (18)$$

where $\Psi(x) = \exp(x)/[1 + \exp(x)]$ denotes the logistic distribution function and b_i is the item difficulty. For the 1PL model, the sum score is a sufficient statistic; that is, the scoring rule is given by

$$\theta_n = f\left(\sum_{i=1}^I 1 \cdot x_{ni}\right) \quad (19)$$

Hence, all items are equally weighted in the ability variable θ_n and receive the local item score $v_{j,i}(\theta) = 1$. Note that this weight is independent of θ . If the set of selected items in the test adequately represent the population of items (i.e., $w_{j,i} = 1$), it can be argued that the 1PL should be the preferred measurement model because the uniform weighting in the sum score in Equation (19) can be considered as a proxy of an equally weighted sum score for the population of items (see also Stenner et al., 2008, 2009). The 1PL model was used in PISA as a scaling model until PISA 2012 (OECD, 2014).

In the two-parameter logistic (2PL; Birnbaum, 1968) model, items are allowed to have different item discriminations a_i :

$$P(X_{ni} = 1|\theta_n) = \Psi(a_i(\theta_n - b_i)) \quad (20)$$

The sufficient statistic is given by the weighted sum score in which locally optimal item weights $v_{j,i}(\theta)$ are given by a_i that are independent of θ :

$$\theta_n = f\left(\sum_{i=1}^I a_i x_{ni}\right) \quad (21)$$

In most applications, item discriminations a_i are estimated from data and are determined to maximize model fit in terms of the log-likelihood function. However, the empirically determined weights can differ from a priorily specified item weights $w_{j,i}$ in Equation (14) in a design-based inference. In this case, model-based and design-based inference will not provide the same results. However, if a design-based inference and the scoring rule in Equation (14) is desired, the 2PL model can be utilized as a measurement model with fixed item discriminations; that is, $a_i = w_{j,i}$. (see, e.g., Haberkorn et al., 2016). The 2PL model is used in PISA as a scaling model since PISA 2015 (OECD, 2017; see also Jerrim, Parker, et al., 2018).

In the three-parameter logistic (3PL; Birnbaum, 1968) model, an additional guessing parameter g_i is included in the IRF:

$$P(X_{ni} = 1|\theta_n) = g_i + (1 - g_i)\Psi(a_i(\theta_n - b_i)) \quad (22)$$

For the 3PL model, locally optimal item weights indeed depend on the ability θ (Chiu & Camilli, 2013)

$$v_{j,i}(\theta) = \frac{a_i}{1 + g_i \exp[-a_i(\theta - b_i)]} \quad (23)$$

Note that the contribution of item i in the ability value θ increases as a function of θ . In this sense, the model-implied effective item scores for countries depend on the country-specific ability distributions (Camilli, 2018). Another objection against the 3PL model is that g_i is not the probability of guessing for multiple-choice items (Aitkin & Aitkin, 2006; von Davier, 2009). Alternative IRT models have been proposed that circumvent this issue (Aitkin & Aitkin, 2006). Occasionally, arguments against using the 3PL model are made for reasons of a lack (or weak empirical) identification of model parameters of the 3PL model (Maris & Bechger, 2009; San Martin et al., 2015). However, these concerns vanish with sufficiently large samples, distributional assumptions for the ability variable, or weakly informative prior distributions for item parameters. The 3PL model is in operational use in PIRLS (Foy & Yin, 2017) and TIMSS (Foy et al., 2020).

In the psychometric literature, there is recent interest in the four-parameter logistic (4PL; e.g., Culpepper, 2017; Loken & Rulison, 2010) that also allows a slipping parameter s_i in the IRF:

$$P(X_{ni} = 1|\theta_n) = g_i + (1 - s_i - g_i)\Psi(a_i(\theta_n - b_i)) . \quad (24)$$

Students can receive a very large ability θ_n despite the fact that their item response probabilities can be substantially smaller than one due to the presence of slipping parameters.

As a consequence, a failure on some items is not so strongly penalized in the 4PL model because a wrong item response can be attributed to a slipping behavior. Like in the 3PL model, the locally optimal item weights in the 4PL model also depend on the ability (see Magis, 2013). It is unlikely that these θ -dependent item weights in a model-based perspective will coincide with a priorly specified item weights in a design-based perspective. To our knowledge, the 4PL model is not currently in operational practice in any international LSA study.

Alternatively, asymmetric IRFs (Bolt et al., 2014; Goldstein, 1980) can be used that allow item weights to depend on item difficulty (see also Dimitrov, 2016). The most flexible approach would be achieved by a semiparametric or a nonparametric specification of IRFs (Falk & Cai, 2016; Feuerstahler, 2019; Ramsay & Winsberg, 1991). These IRFs imply model-based item weights that might strongly differ from weights that are specified under a design-based perspective and may therefore distort the test composition that is defined in test blueprints (Camilli, 2018).

To sum up, choosing a particular IRF implies different item weights and scoring rules for the ability variable θ . It can be questioned whether IRFs should be chosen for the sole purpose of increasing reliability (and model fit) because different IRFs correspond to different estimation targets. In our view, the choice of an IRF should be mainly a question of validity and cannot be answered by model fit or item fit statistics. However, if the superior model fit would be defined as the main goal of model choice in LSA studies, more complex IRFs (3PL, 4PL, semiparametric IRFs) will almost always outperform simpler IRFs (1PL, 2PL). The switch from the 1PL to the 2PL model in recent PISA studies can, therefore, in our opinion, not be defended by reasons of better model fit, because in our experience (from unpublished research and in line with Culpepper, 2017), the 4PL model outperforms the 1PL, 2PL, and 3PL model in PISA in terms of model fit. However, the crucial question is whether the derived ability from the 4PL model is constituted valid. Following Brennan (1998), we believe that a psychometric model should not prescribe the contribution of items in the ability score. If items in the test represent items in a (hypothetical) larger item domain, the 1PL model can be defended even though it will likely not fit empirical data. It should be emphasized that in a design-based inference, the employed likelihood function in the ability estimation is intentionally misspecified (see Section 2.1). Overall, we tend to favor the operational use of the 1PL model in LSA studies because the ability score primarily reflects an equal contribution of items that appear in the test. Nevertheless, the likelihood part

associated with the misspecified IRT model has to be modified adequately to reflect the reliability (see Section 2.1.).

3.2. Local Dependence and Multidimensionality

In the previous section, we discussed the choice of the IRF in unidimensional IRFs. Typical abilities assessed in LSA studies will be multidimensional, which, in turn, causes a violation of the local independence assumption. In Section 2.1., we argued that a misspecified unidimensional IRT model could be defended from a design-based inference point of view. However, the model-implied reliability obtained from fitting a unidimensional model can be incorrect. If an ability domain is multidimensional (e.g., subdimensions in reading ability), and the multidimensionality is considered construct-relevant (Shealy & Stout, 1993), the model-implied reliability from a fitted unidimensional model will underestimate the true reliability (Zinbarg et al., 2005). Moreover, items are frequently arranged in testlets such that different items share the same item stimulus. This deviation from local independence can introduce additional error (Monseur et al., 2011) if testlet effects are considered construct-irrelevant¹. In this case, the reliability in ability scores decreases (Zinbarg et al., 2005). In empirical LSA data, there will be construct-relevant multidimensionality as well as construct-irrelevant testlet effects. However, the unidimensional IRT model is used as a scoring model because a unidimensional summary is required for country comparisons. Resampling methods for items (see Section 2.2) can be used for determining standard errors associated with estimated abilities $\hat{\theta}_n$. The estimated standard errors can be used to adjust the likelihood part of the measurement model to generate plausible values (see Section 2.1.; Bock et al., 2002; Mislevy, 1990).

Current operational practice in LSA studies ignores deviations from local independence in the scaling of item responses. While we do not think this introduces large bias in country means or individual ability estimates, the estimated uncertainty associated with plausible values might be incorrect. However, the reliability estimate obtained from the misspecified IRT model could be defended on the condition that deviations from unidimensionality cancel out on average. That is residual correlations among items approximately cancel out on average. Continuing this argument, the recent practice of ignoring local dependence could be defended if the underestimation of reliability due to the multidimensionality is compensated by an overestimation due to testlet effects. However, there is no chance to test this property with empirical data. One always has to make

¹ For testlets administered for the ability domain reading, one might argue that testlets effects are (partly) construct-relevant because the common item stimulus (i.e., the whole reading text) have to be processed for answering the items.

assumptions about the true average residual correlation in latent variable models (Westfall et al., 2012). This view on latent variable models corresponds to a design-based perspective in which one defines the interchangeability of items by design assumptions that cannot be guaranteed by any statistical test.

To sum up, we tend to argue that the approximate unidimensionality assumption for ability in the scaling model can be defended in practice due to the following two empirical findings. First, there is frequently only a low amount of multidimensionality found in data (e.g., for subdimensions in reading or mathematics in PISA data; OECD, 2017). Second, when the number of items is tending to infinity (with a bounded number of items in testlets), the local dependence of testlet effects asymptotically vanishes for the estimation of model parameters (Ellis & Van den Wollenberg, 1993; Stout, 1990). For example, with about 100 administered items per domain and at most 5 to 7 items per testlet, introduced biases in scaling models due to local dependence might be small to moderate.

3.3. The Role of Test-Taking Behavior in the Scaling Model

It has frequently been argued that measured student performance in LSA studies is affected by test-taking strategies. For example, in a recent paper that was published in the highly-ranked *Science* journal, Pohl et al. (2021) argued that “current reporting practices, however, confound differences in test-taking behavior (such as working speed and item nonresponse) with differences in competencies (ability). Furthermore, they do so in a different way for different examinees, threatening the fairness of comparisons, such as country rankings.” (Pohl et al., 2021, p. 338). Hence, the reported student performance (or, equivalently, student ability) would be confounded by a “true” ability and test-taking strategies. Importantly, the authors question the validity of country comparisons that are currently reported in LSA studies and argue for an approach that separates test-taking behavior (i.e., item response propensity and working speed) from a purified ability measure. In the following, we clarify that the additional consideration of test-taking behavior has the potential to change the meaning of the measured abilities substantially (within and also between countries). As the proposed approach focuses on the modeling of omitted responses (Pohl et al., 2021), we start with a brief summary of how missing responses are treated in LSA studies.

Missing item responses can be classified into omitted items within the test and not-reached items at the end of the test (see Rose et al., 2017). Until PISA 2012, not reached items were treated as incorrect in ability estimation while they were not scored as incorrect since PISA 2015 (OECD, 2017). In more detail, the proportion of not reached items is used as a

covariate in the LBM since PISA 2015 while recoding all not-reached item responses as not administered in the scaling. We would argue that this treatment of not-reached items has the potential of decreasing the validity of ability scores because countries can easily manipulate the scores on not-reached items by advising test takers to work slowly through the test and only produce missing item responses if there are many items they do not know. Thus, we would not concur with Pohl et al. (2021, p. 339), who conclude that “[...] scoring not-reached items as incorrect—as done in some LSAs—results in scores that differ in their meaning, depending on whether examinees do or do not show missing values. This jeopardizes the comparability of performance scores across examinees and, thus, fairness.”. Unfortunately, the role of not-reached items becomes even more critical in scaling with the implementation of multi-stage testing (see Section 4.3.).

Pohl et al. (2021) propose the speed-accuracy and omission (SA+O) model (Ulitzsch et al., 2020b) that simultaneously models item responses, response indicators that indicate whether students omit items, and response times. Not-reached items are treated as non-administered. In the SA+O model, these observed variables are associated with four latent variables: an ability, a response propensity variable, and two speed variables (for observed and omitted items). In the following, we discuss the potential implications of the use of this model in LSA studies. Let X_{ni} be the item response of person n on item i , and R_{ni} be the response indicator that takes a value of 1 if the item is observed and a value of 0 if it is missing (i.e., omitted). Moreover, let T_{ni} be the logarithmized response time for person n on item i . In the SA+O model, the joint distribution of indicator variables (X_{ni}, R_{ni}, T_{ni}) is modelled as (Ulitzsch et al., 2020b)

$$P(X_{ni} = x_{ni}, R_{ni} = r_{ni}, T_{ni} = t_{ni} | \theta_n, \xi_n, \eta_{0n}, \eta_{1n}) = P(X_{ni} = x_{ni} | \theta_n)^{r_{ni}} P(R_{ni} = r_{ni} | \xi_n) f_1(t_{ni} | \eta_{1n})^{r_{ni}} f_0(t_{ni} | \eta_{0n})^{1-r_{ni}} \quad (25)$$

where ξ_n denotes the response propensity, η_{1n} is the speed variable associated with observed items, η_{0n} is the omission speed, and f_1 and f_0 are normal densities for response times of observed and omitted items, respectively². To further illustrate the meaning of Equation (25), we first consider the decomposition for observed item responses (i.e., $R_{ni} = 1$):

$$P(X_{ni} = x_{ni}, R_{ni} = 1, T_{ni} = t_{ni} | \theta_n, \xi_n, \eta_{0n}, \eta_{1n}) = P(X_{ni} = x_{ni} | \theta_n) P(R_{ni} = 1 | \xi_n) f_1(t_{ni} | \eta_{1n}) \quad (26)$$

For missing item responses (i.e., $R_{ni} = 0$ and $X_{ni} = \text{NA}$), Equation (25) simplifies to

$$P(X_{ni} = \text{NA}, R_{ni} = 0, T_{ni} = t_{ni} | \theta_n, \xi_n, \eta_{0n}, \eta_{1n}) = P(R_{ni} = 0 | \xi_n) f_0(t_{ni} | \eta_{0n}) \quad (27)$$

² It might be possible to include separate distributions for response times of correct and incorrect item responses (see Bolsinova et al., 2017).

In a model-based estimation approach of the SA+O model, effectively, missing item responses X_{ni} have to be imputed based on the latent variables $(\theta_n, \xi_n, \eta_{1n}, \eta_{0n})$ and response time T_{ni} . We now derive how item responses, response indicators, and response times are used for estimating the ability in the SA+O model. For the derivation, it is convenient to reparametrize the vector of latent variables $(\theta_n, \xi_n, \eta_{1n}, \eta_{0n})$ to $(\theta_n, \xi_n^*, \eta_{1n}^*, \eta_{0n}^*)$, where ξ_n^* , η_{1n}^* and η_{0n}^* are residualized latent variable in which the ability θ_n is partialled out:

$$\xi_n = \alpha_{\xi\theta}\theta_n + \xi_n^* \quad \text{and} \quad \eta_{hn} = \alpha_{\eta_h\theta}\theta_n + \eta_{hn}^* \quad \text{for } h = 0, 1 \quad (28)$$

Then, the IRT model in Equation (25) can be equivalently written as

$$\begin{aligned} &P(X_{ni} = x_{ni}, R_{ni} = r_{ni}, T_{ni} = t_{ni} | \theta_n, \xi_n^*, \eta_{0n}^*, \eta_{1n}^*) \\ &= P(X_{ni} = x_{ni} | \theta_n)^{r_{ni}} P(R_{ni} = r_{ni} | \theta_n, \xi_n^*) f_1(t_{ni} | \theta_n, \xi_n^*, \eta_{1n}^*)^{r_{ni}} f_0(t_{ni} | \theta_n, \xi_n^*, \eta_{0n}^*)^{1-r_{ni}} \end{aligned} \quad (29)$$

For item responses X_{ni} , a 2PL model is assumed (see Equation (20)). The probability of responding to an item is assumed to be

$$P(R_{ni} = r_{ni} | \xi_n) = \Psi(\gamma_i(\xi_n - \beta_i)) \quad (30)$$

Logarithmized response times T_{ni} are modelled as conditional normal distributions

$$T_{ni} | \eta_{1n}, \eta_{0n}, R_{ni} = h \sim N(\eta_{hn} - \tau_{ih}, \lambda_{ih}^{-2}) \quad \text{for } h = 0, 1 \quad (31)$$

In Appendix B, local item scores are derived that define local sufficient statistics of indicator variables for θ_n . For observed item responses, the item weight is given by the item discrimination from the 2PL model (i.e., a_i). The local item score for θ_n for a missing item response (see Equation (A12) in Appendix B) is given by

$$a_i P_i(\theta) - \alpha_{\xi\theta} \gamma_i - (\alpha_{\eta_{1\theta}} \lambda_{i1} - \alpha_{\eta_{0\theta}} \lambda_{i0}) t_{ni} \quad (32)$$

From Equation (32), it can be seen that two (in general) positive terms will be subtracted from the item score $a_i P_i(\theta)$. Note that $a_i P_i(\theta)$ would be considered as an appropriate item score if the missingness mechanism is ignorable (i.e., treatment of omitted items as non-administered provides a valid strategy; Pohl & Carstensen, 2013; Pohl et al., 2014). In case of a positive correlation of θ and ξ , the imputed score is adjusted by the value $\alpha_{\xi\theta} \gamma_i$. Furthermore, because omitted items are typically associated with shorter response times, the adjustment term $(\alpha_{\eta_{1\theta}} \lambda_{i1} - \alpha_{\eta_{0\theta}} \lambda_{i0}) t_{ni}$ also plays a role in the scoring rule. Hence, the ability variable θ_n also enters the log-likelihood contributions of response indicators and response times.

Consequently, response indicators and response times contribute to the imputation of omitted item responses and thereby influence the model-based estimation of abilities defined in Equation (25).

Like in our discussion for not-reached items, we would argue that the scoring rule implied by the SA+O model has substantial consequences for the interpretation of ability scores in LSA studies. Study results can be simply manipulated at the country level if students are advised to skip items they do not know or to produce very short response times in such cases. In our opinion, the possibility to influence students' test-taking behavior poses severe threats to the validity and fairness of country comparisons. Furthermore, in our research with LSA data, we found that the conditional independence assumptions of item responses and response indicators in the SA+O model are strongly violated, resulting in a worse model fit of the SA+O model (see Robitzsch, 2020). There is empirical evidence that students who do not know the answer to an item have a high probability of omitting this item even after controlling for latent variables. This seems to be particularly the case for constructed response items. Thus, we believe that the dependence of responding to an item from the true but unknown item response must be considered even after conditioning on latent variables (Robitzsch, 2020). Given these concerns about the less plausible assumptions of the SA+O model and its consequences for the validity of country comparisons, we would argue that it cannot be recommended for operational use in large-scale assessment studies.

It is important to emphasize that the adjustments—and hence the scoring rules for ability—in the SA+O model will differ from country to country because the relationship of ability, response propensity, and speed differ across countries (Sachse et al., 2019; Pohl et al., 2021). In our view, a country comparison that does not employ the same scoring rule for each country cannot be considered characterized as valid (or fair).

Much psychometric work seems to imply that simulation studies demonstrate that missing item responses should never be scored as incorrect (Pohl & Carstensen, 2013; Pohl et al., 2014; Rose et al., 2017). We strongly oppose such a perspective because simulation studies are not helpful in decisions about how to handle missing item responses (Rohwer, 2013; Robitzsch, 2020). One can simulate data that introduce missing item responses only for incorrectly solved items (Rohwer, 2013). In this case, all IRT models that do score missing items as incorrect provide biased model parameter estimates. Moreover, the scoring of items should always be conducted under validity considerations. We think that omitted (constructed response) items should always be scored as incorrect because alternative scoring rules decrease validity.

Moreover, continuing the arguments of Pohl et al. (2021), other test-taking behaviors could be used for purifying ability. For example, response effort such as rapid guessing (Deribo et al., 2021; Ulitzsch et al., 2020a) or performance decline (Debeer & Jansen, 2013;

Jin & Wang, 2014) could be taken into account. Moreover, the ability variable θ could also be redefined in a scaling model in which item responses and response times load on θ , resulting in a purified latent variable for speed (Costa et al., 2021). Furthermore, measurement models could also involve an additional student latent variable α_n that characterizes person fit (Conijn et al., 2011; Ferrando, 2019; Raiche et al., 2012)

$$P(X_{ni} = 1|\theta_n, \alpha_n) = \Psi(a_i \alpha_n (\theta_n - b_i)) \quad (33)$$

Such a model would weigh persons in the log-likelihood by a model-based weight α_n and the model would certainly be justified by reasons of model fit against simpler alternatives. As a consequence, the local scoring rule for abilities also depends on the person fit variable α_n which would further complicate interpretations. We strongly believe that including latent variables that capture test-taking behavior in measurement models should be avoided in the official reporting of LSA results. In our view, the explicit modeling of test-taking behaviors leads to the opposite of fairer country comparisons, and a design-based approach should be preferred for inferences regarding abilities in LSA. In this approach, test-taking behavior is always coupled to a realized test design. Researchers (as well as the public and policy) have to judge whether the assessed abilities —under the given design—are deemed valid.

3.4. Country DIF and Cross-sectional Country Comparisons

For most international LSA studies, the comparability of tests across countries is of crucial importance (Rutkowski & Rutkowski, 2019). There are conceptual and statistical approaches for assessing comparability. Statistical approaches include the assessment of differential item functioning (DIF; Holland & Wainer, 1993; Penfield & Camilli, 2007), which focuses on the heterogeneity in item parameters across countries. In the 2PL model, there is empirical evidence that item difficulties vary from country to country (von Davier, Khorramdel et al., 2019):

$$P(X_{nci} = 1|\theta_{nc}) = \Psi(a_i(\theta_{nc} - b_{ic})) \quad , \quad (34)$$

where the index c denotes country-specific parameters, and discrimination parameters a_i are assumed to be constant across countries (uniform DIF; see OECD, 2017). The presence of DIF with respect to countries is denoted as (cross-sectional) country DIF (see Monseur et al., 2008; Robitzsch & Lüdtke, 2019). If only a few item difficulties b_{ic} are allowed to deviate from a common item difficulty, it is said that partial invariance holds (von Davier, Khorramdel, et al., 2019). Since PISA 2015, the assumption of partial invariance (Oliveri & von Davier, 2011, OECD, 2017; von Davier, Yamamoto, et al., 2019) is employed in the scaling model. Non-invariant item parameters are determined utilizing item fit statistics

(Tijmstra et al., 2020). In the partial invariance approach, the majority of item parameters is assumed to be equal (i.e., invariant) across countries, and there is a low proportion of country-specific item parameters. In contrast, until PISA 2012, the 1PL scaling model with invariant item parameters was assumed, and country DIF was ignored unless it could be attributed to technical issues in item administration (e.g., translation errors; Adams, 2003).

The DIF treatment has fundamentally changed since PISA 2015. The assumption of country-specific item parameters effectively eliminates some items from pairwise country comparisons (Robitzsch & Lüdtke, 2020, 2021). Moreover, the set of effectively used items differs across comparisons (e.g., the comparison between country A and country B could be based on different items than the comparison between country A and country C; see also Zieger et al., 2019). It has been argued that this property poses a threat to validity, and researchers are comparing apples and oranges when pursuing the partial invariance approach (Robitzsch & Lüdtke, 2021). We believe that the decision of whether an item induces bias for country comparisons is not primarily of statistical nature. Camilli (1993) pointed out (see also Penfield & Camilli, 2007) that expert reviews of items showing DIF should accompany DIF detection procedures. Only those items should be excluded from country comparisons for which it is justifiable to argue that construct-irrelevant factors caused DIF (see also El Masri & Andrich, 2020; Zwitser et al., 2017). However, the purely statistical approach since PISA 2015 based on partial invariance disregards that DIF items could be construct-relevant. Also, note that PIRLS and TIMSS do not seem to use country-specific item parameters and rely on a scaling model that assumes full invariance of item parameters across countries (Foy et al., 2020). From a validity perspective, we would prefer the approach that ignores country DIF if the DIF cannot be attributed to test administration issues. This strategy more closely follows a design-based inference perspective for items because the test design and not a psychometric model should guarantee whether the set of items in a test are presentative for a specific item population (Brennan, 1998).

In contrast to the partial invariance approach that assumes that most items do not show DIF (and only a few items possess large country DIF), the assumption of full noninvariance can be made, which assumes that all items show country DIF effects (Fox & Verhagen, 2010). In our experience and in line with other researchers (e.g., Sachse et al., 2016), we find the partial invariance assumption unlikely to hold in empirical data. Instead, it seems that DIF effects are symmetrically distributed and closely follow a normal distribution.

To acknowledge the dependence of country comparisons on the chosen set of items due to country DIF, linking errors (LE; Robitzsch & Lüdtke, 2019; Wu, 2010) have been

proposed to quantify the heterogeneity in the country means due to the selection (or sampling) of items. The inclusion of the item facet for describing the uncertainty in group means has been studied in GT for a long time (Brennan, 2001; Kane & Brennan, 1977). Assume that the 2PL model with country DIF (see Equation (34)) holds and that the country-specific item parameters b_{ic} can be decomposed into a common item difficulty b_i and country-specific deviations e_{ic} :

$$b_{ic} = b_i + e_{ic} \quad , \quad e_{ic} \sim N(0, \tau_{\text{DIF},c}^2) \quad (35)$$

where $\tau_{\text{DIF},c}^2$ is the country-specific DIF variance. For I items, the uncertainty due to the selection of items is quantified in the following cross-sectional linking error

$$\text{LE}_{\text{cs},c} = \sqrt{\frac{\sum_{i=1}^I a_i^2}{(\sum_{i=1}^I a_i)^2}} \tau_{\text{DIF},c} \quad (36)$$

For the 1PL model that assumes equal item discriminations a_i , Equation (36) simplifies to (see Robitzsch & Lüdtke, 2019)

$$\text{LE}_{\text{cs},c} = \frac{1}{\sqrt{I}} \tau_{\text{DIF},c} \quad (37)$$

Instead of establishing a scaling model assuming partial invariance, we prefer the additional error component associated with items for reporting in LSA studies. The total error $\text{TE}_{\text{cs},c}$ for a cross-sectional country mean contains the standard error $\text{SE}_{\text{cs},c}$ due to the sampling of persons as well the linking error $\text{LE}_{\text{cs},c}$ due to the selection of items (Wu, 2010)³

$$\text{TE}_{\text{cs},c} = \sqrt{\text{SE}_{\text{cs},c}^2 + \text{LE}_{\text{cs},c}^2} \quad (38)$$

The uncertainty due to item selection also affects other statistical parameters (e.g., standard deviation, quantile, regression coefficient). Linking errors can be more flexibly obtained by resampling items (Brennan, 2001). It should be emphasized that linking errors also occur if invariant item parameters are assumed in the scaling model. There are still consequences of heterogeneity in item selection for the country means, even if no country-specific item parameters are explicitly modeled in the scaling model.

In this section, we argued against a partial invariance approach that removes items from particular country comparisons (see Robitzsch & Lüdtke, 2021). In empirical data, country DIF effects will almost always occur. There are two options for handling the presence of DIF effects in the scaling models. First, DIF effects can be ignored in a concurrent scaling

³ The estimated linking error $\text{LE}_{\text{cs},c}$ in Equations (36) or (37) also contains a standard error associated with sampling of students. Hence, this variance contribution should be subtracted in Equation (38). In practice, the positive bias in $\text{TE}_{\text{cs},c}$ might be regarded as negligible.

approach in which the incorrect assumption of invariant item parameter is posed. Second, separate scaling can be performed at the country level, and linking methods are used to compare countries (Robitzsch & Lüdtke, 2020, 2021). Notably, concurrent scaling can only be more efficient than separate scaling for correctly specified IRT models, that is, in the absence of country DIF. In the presence of DIF, DIF effects are weighted by a likelihood discrepancy function in concurrent scaling and the estimates of country means can in general be less precise than separate scaling with subsequent linking. In these approaches, the weighing of DIF effects is determined by choosing a linking function. We think that a linking function should be chosen that does not automatically eliminate items with large DIF effects from comparisons (i.e., in robust linking; see He & Cui, 2020; Robitzsch & Lüdtke, 2021). Moreover, the concurrent scaling approach is probably based on a misspecified IRT model that can result in biased estimation of the latent ability distribution parameters (i.e., standard deviation, quantiles). Finally, we believe that the sample sizes in typical LSA studies are large enough to apply a separate scaling approach with subsequent linking for the 1PL or the 2PL model.

3.5. Country DIF and Trend Estimation

One of the primary outcomes in LSA studies is trend estimation that enables monitoring educational systems in abilities. The original trend estimate for two time points is computed by subtracting the cross-sectional country mean of the first time point from the second time point. As an alternative, a marginal trend estimate has been proposed that performs the linking across the two time points only on the link items administered in both studies (Gebhardt & Adams, 2007; Robitzsch & Lüdtke, 2019). Original trends have the advantage that officially reported cross-sectional country means can be utilized for computing the trend estimate (e.g., the difference). However, Robitzsch and Lüdtke (2019) showed analytically and with simulation studies that original trend estimates can be less precise than marginal trend estimates if there is a sufficiently large number of unique items; that is, items that are only administered at one of the two time points (see also Gebhardt & Adams, 2007). The primary reason for the increased precision of marginal trend estimates is that cross-sectional country DIF turns out to be relatively stable across time points. Consequently, unique items introduce additional variability in the country means due to the presence of DIF effects (Robitzsch & Lüdtke, 2019). In PISA, there is a switch from major to minor domains (or the other way around) for two of the three primary domains mathematics, reading, or science. If the number of unique items is large compared to the number of link items, original trend estimates in PISA tended to be much more variable than marginal trend estimates

(Carstensen, 2013; Gebhardt & Adams, 2007; Robitzsch & Lüdtke, 2019; Robitzsch et al., 2020). On the other hand, by relying on the link items, country DIF effects are automatically controlled for in marginal trend estimates because the stable country DIF effects occur to the same extent at both time points and therefore cancel when calculating achievement trends.

We would like to note that marginal trend estimates were originally proposed at the country level, based on separate scaling with subsequent linking for each country (Gebhardt & Adams, 2007). In our experience based on simulation studies, it can be demonstrated that this requirement is not the essential reasoning for the fact that marginal trends can be more efficient than original trends. The linking could be conducted at the (international) level of all countries (i.e., in a joint scaling approach that involves all countries and assumes invariant item parameters across countries). However, the crucial point is that it should only involve the link items and not the unique items.

The variability in trend estimates due to the selection of items is quantified by linking errors (Gebhardt & Adams, 2007; OECD, 2015) in LSA studies. The linking error employed in PISA until PISA 2012 quantifies the uncertainty in trend estimates for the 1PL model based on the variance of item parameter drift (IPD; i.e., a difference of item difficulties across time; OECD, 2015). Notably, this error only assesses variability due to link items, ignoring the variability due to country DIF. Robitzsch and Lüdtke (2019) proposed a linking error for the 1PL model that correctly reflects the variability of original trend estimates due to item selection. Since PISA 2015, the computational approach for the linking error changed by utilizing a recalibration method (OECD, 2017; see also Martin et al., 2012). The motivation for the change in computation was that it should also be applicable for the recently implemented analytic changes (i.e., the 2PL model and the partial invariance approach). In the newly proposed method, data from the first time point are recalibrated using item parameters from the second time point. The linking error is defined as the variance of the average squared difference of original and recalibrated country means (OECD, 2017). If all item parameters were assumed to be invariant, the same item parameters as in the original calibration for the link items would be used. Hence, it can be expected that the newly proposed linking error will turn out to be small if most of the item parameters are assumed to be invariant. We would like to emphasize that there is an essential conceptual change in how linking errors are defined in PISA since PISA 2015. In our opinion, the recalibration method might be helpful to define an effect size of the extent of noninvariance in item parameters in terms of variability in the recalibrated country means. However, we do not believe that the new approach (implemented since PISA 2015) correctly reflects the variability in original trend estimates due to item

selection because variability cannot vanish by just assuming invariant item parameters. While we admit that the computation of linking errors has to be modified in recent PISA cycles due to the use of the 2PL model, we would question that the recently proposed linking error provides a solid basis for statistical inference for trend estimates.

Finally, it can be discussed how several cycles of an LSA study should be optimally analyzed when trend estimates are of primary interest. While previous PISA cycles link subsequent PISA studies to each other in a chain linking (OECD, 2015), other researchers opted for a multiple group IRT concurrent scaling approach that assumes that most item parameters are invariant across time points and countries (von Davier, Yamamoto, et al., 2019). It has been argued that the concurrent scaling approach provides more stable trend estimates (von Davier, Yamamoto, et al., 2019; p. 485) by relying on the assumption of partial invariance (i.e., only a few item parameters are not invariant) because more stable item parameter estimates would be obtained. However, it should be noted that the superiority of the concurrent scaling statement was not confirmed by simulation studies. Moreover, the validity of such a statement would require that the partial invariance assumption holds. Like for cross-sectional LSA data, we suppose that there is no empirical evidence for this assumption. Hence, we would argue that there is lacking support for the higher efficiency of the concurrent scaling approach compared to separate scalings for each time point with subsequent linking.

4. Discussion

In this article, we reflected on several analytical choices in LSA studies. We illustrated that it can be crucial to distinguish between a design-based or model-based perspective on statistical inference. When it comes to official reporting in LSA studies, we argued that a design-based perspective should predominate a model-based perspective. In a part of the methodological LSA literature, there is a tendency to prefer more complex psychometric models with the promise that these complex models produce more stable and less biased estimates of student abilities. However, these claims are primarily made from a model-based perspective, and we clarified that model-based approaches often redefine the meaning of the abilities of interest. For example, using the 2PL model instead of the 1PL model implies a different weighing of items that is primarily defined through optimizing model fit. This contradicts a design-based perspective in which the contribution of items in a score is a priorily defined by a test framework. The reliance on a partial invariance model for scaling is another example of how a model-based perspective can change the meaning of country comparisons. In some sense, it can be argued that the partial invariance approach compares

apples and oranges because the set of effectively used items differs across country comparisons.

In a design-based perspective, the likelihood function that involves the IRT model and the LBM is typically misspecified in LSA studies. It can always be acknowledged that models are only approximately true. However, we even do not believe that the concept of approximate fit makes sense when favoring the 1PL model over the 2PL model. In our view, the 1PL model is preferable because the main goal is to use an equally weighted sum score as a sufficient statistic for θ . The specified likelihood function can be interpreted as a pseudo-likelihood function that is only used to provide an estimating equation for the parameters of interest. As a consequence, we suppose that model fit should not play a (primary) role in choosing psychometric models for LSA studies. Also, note that the likelihood-based inference (i.e., standard errors) obtained from a misspecified model will also be incorrect. We believe that resampling techniques for persons (see Section 2.1) and items (see Section 2.2) allow valid statistical inference, even if the model is misspecified (Berk et al., 2014; White, 1982). In our view, statistical models should be defended from a design-based perspective. Hence, different researchers might opt for different psychometric models for modeling LSA data if the model fit is not considered the primary criterion⁴. The concept of model weighting (or model uncertainty) can be used to quantify the extent of consequences in an uncertain space of models (Young & Holsteen, 2017). For example, it might be beneficial to study the sensitivity of trend estimates for a country for different choices of the linking function. Researchers would be less confident in trend estimates that strongly depend on the chosen estimation method.

4.1. Transparency and Reproducibility

Because LSA data and its methodology are very complex, it is important for researchers and the public that the analytical steps producing the official reports are defensible and transparent. Optimally, external researchers should be able to reproduce officially reported scores based on the published technical report and datasets that are available as public use files. Moreover, researchers should also have easy access to the software used by data producers (i.e., consortia that are contracted, such as ETS or ACER). This does not preclude commercial statistical software. However, if the software is used by data producers that is not open source, careful technical documentation is required for allowing external

⁴ Note that there are typically very different approaches for assessing model fit. Depending on how model fit is defined, the complexity of chosen psychometric models will vary considerably. Hence, there will also be disagreement among psychometricians with respect to model choice in the case in which model fit would serve as the main criterion.

researchers to validate analytic steps by their own implementations. These requirements in LSA methodology are currently not yet fully met. Frey and Hartig (2021, p. 45) note that “[...] there is room for improvement in ILSAs when it comes to meeting open science standards with respect to the transparency and accessibility of methods. Because ILSAs are often used as a reference by researchers working on smaller studies, their methods must be well documented and well-founded.” In our view, we admit that full reproducibility of all officially reported results might be a too strong demand. However, we sometimes found that not all technical details were published for the PISA study. Moreover, empirical or simulation evidence was often not provided for convincing external researchers that the implemented methodological changes improve estimation or increase the validity of the measured abilities. We do not want to argue that there is always a right modeling strategy. However, we feel that contracted consortia in LSA studies should be more open about the main motivations and consequences of their analytical choices. Occasionally, practical consequences of modeling choices appear hidden in a complex psychometric modeling framework. This article tried to argue that these model-based approaches should play a minor role compared to a design-based perspective of statistical inference. In the last part, we provide two examples from PISA that may illustrate the potential for improvement concerning transparency and reproducibility.

4.2. Assessment and Linking of Mode Effects

One of the central administrative changes in LSA studies is the switch from paper-based (PBA) to computer-based administration (CBA) that occurred in PISA (Jerrim, Micklewright, et al., 2018; OECD, 2017; Robitzsch et al., 2020) and TIMSS (Fishbein et al., 2018). It is debated whether trend estimates for countries are reliable under mode changes. In PISA, PBA was used in PISA 2012, and CBA has been implemented (for most countries) since PISA 2015. PISA conducted a field test study in 2014 containing nonrepresentative student samples at the country level, for which PBA and CBA modes were administered. It has been argued that findings from the field test study were used to establish a link from PISA 2012 to PISA 2015 by identifying invariant item parameters (ETS, 2015; OECD, 2017). However, these crucial steps are not reproducible by external researchers because the PISA field test data is not publicly available. Researchers can only rely on statements from the official PISA technical report (OECD, 2017) that the linking approach was successful because most of the item parameters turned out to be invariant across mode effects. It cannot be ruled out from official PISA reports that mode effects are not country-specific. We do not want to argue that the linking approach from PISA 2012 to PISA 2015 is not valid. However, because of the lack of details and publicly available datasets, it cannot be judged whether reported

trend estimates in PISA are reliable and robust to alternative model specifications. In contrast, in TIMSS, a comprehensive bridge study has been conducted that allows researchers to validate the official linking procedure (Fishbein et al., 2018).

4.3. Multi-Stage Testing

In PISA 2018, multi-stage testing (MST) was implemented for the first time. For the reading domain, three item clusters (so-called modules) were administered, and a routing path after the first and second module was defined. Students with higher scores in the first or second module received a more difficult following module (OECD, 2020). The primary motivation for implementing MST seems to be to increase the measurement precision in students' ability, particularly for students in countries with a low average ability (Rutkowski & Rutkowski, 2021). Surprisingly, even researchers who advocate MST only report negligible to small precision gains (Yamamoto, Khorramdel & Shin, 2018; Yamamoto, Shin & Khorramdel, 2018). Notably, the gains can only be realized if the IRT models would be correctly specified and test administration effects do not moderate item parameters (i.e., item difficulties). As a consequence, there is always a bias-variance tradeoff when test administration switches to MST. MST can reduce variance, but it can introduce bias compared to fixed linear testing because every routing path introduces additional context effects. Admittedly, fixed test booklets also introduce context effects. However, in a non-adaptive administration, the allocation of booklets (or routing paths) can be controlled and fixed across (subgroups of) countries. If the main goal is to reduce the standard error for statistics referring to subgroups of students, it would be more promising to sample more students or improve the implemented sampling procedures.

In MST, the proportions of particular routing paths vary from country to country and, hence, the influence of context effects has the potential to vary across countries. We see this property of MST as a potential threat to validity because the test design in MST is not held fixed when comparing countries with each other. In our experience, MST introduces a substantially larger proportion of not reached item responses in PISA that seems to depend on particular routing paths strongly. This also questions the comparability of abilities obtained from different routing paths and, hence, countries.

In the best scenario, MST would provide similar (i.e., point estimates) but more precise (smaller standard errors) results than fixed linear testing. However, the adaptive and non-adaptive testing should be compared in a bridge study in PISA that involves equivalent groups at the country level (like for mode effects in TIMSS; Fishbein et al., 2018). Otherwise,

it is difficult to evaluate whether the implemented MST in PISA provides more precise and not distorted trend estimates.

The issue that ability should be assessed for countries that strongly vary in their average performance has been addressed in the past by administering test booklets of two different average difficulties (OECD, 2014, 2017). This design enables a valid comparison of countries that receive booklets in the same test design. We believe that there are no plausible reasons for switching from this design with a fixed linear test in two difficulty versions to MST.

Appendix

Appendix A: Locally Optimal Item Weights

In this appendix, we derive the locally optimal item weight that is based on the individual log-likelihood function $l_n(\theta) = \sum_{i=1}^I l_{ni}(\theta)$. The log-likelihood contribution $l_{ni}(\theta)$ of item i for person n is given by

$$l_{ni}(\theta) = x_{ni} \log P_i(\theta) + (1 - x_{ni}) \log(1 - P_i(\theta)) \quad (\text{A1})$$

We now derive a Taylor approximation of $l_n(\theta)$ around an ability value θ_0 for deriving the contribution of items in a local sufficient statistic for θ . The derivative l_{ni} with respect to θ is given by

$$\frac{dl_{ni}}{d\theta} = x_{ni} \frac{P'_i(\theta)}{P_i(\theta)} - (1 - x_{ni}) \frac{P'_i(\theta)}{1 - P_i(\theta)} = (x_{ni} - P_i(\theta)) v_i(\theta) \quad (\text{A2})$$

Using (A2) for a Taylor approximation of the log-likelihood function, we obtain

$$l_n(\theta) \simeq l_n(\theta_0) + \left(\sum_{i=1}^I v_i(\theta) x_{ni} \right) (\theta - \theta_0) + \text{const}(\theta) \quad (\text{A3})$$

From Equation (A3) it can be seen that $\sum_{i=1}^I v_i(\theta) x_{ni}$ is a local sufficient statistic for θ . Item weights $v_i(\theta)$ are contributions of items in the weighted sum score and are denoted as locally optimal item weights (Birnbaum, 1968; Chiu & Camilli, 2013), where

$$v_i(\theta) = \frac{P'_i(\theta)}{P_i(\theta)[1 - P_i(\theta)]} \quad (\text{A4})$$

We now derive the locally optimal item weights for the 3PL model (see Equation (22)). It holds that

$$P'_i(\theta) = (1 - g_i) a_i \Psi(a_i(\theta - b_i)) [1 - \Psi(a_i(\theta - b_i))] \quad (\text{A5})$$

We now use the short notation $\psi = \Psi(a_i(\theta - b_i))$. Then, we obtain

$$v_i(\theta) = \frac{P'_i(\theta)}{P_i(\theta)[1 - P_i(\theta)]} = \frac{(1 - g_i) a_i \psi (1 - \psi)}{[g_i + (1 - g_i) \psi] (1 - g_i) (1 - \psi)} = \frac{a_i}{1 + g_i (1 - \psi) \psi^{-1}} \quad (\text{A6})$$

A further simplification of (A6) provides

$$v_i(\theta) = \frac{a_i}{1 + g_i \exp(-a_i(\theta - b_i))} \quad (\text{A7})$$

For the 2PL model, we get $v_i(\theta) = a_i$ because it holds that $g_i = 0$. Furthermore, we get $v_i(\theta) = 1$ in the 1PL model.

Appendix B: Local Item Scores in the SA+O Model

In this appendix, we derive local item contributions for the ability score θ for the SA+O model studied in Pohl et al. (2021). The log-likelihood contribution for student n and item i in the reparametrized SA+O model (see Equations (28) and (29)) is given by

$$l_{ni}(\theta) = \text{const}(\theta) + x_{ni}r_{ni}a_i\theta + r_{ni}\alpha_{\xi\theta}\gamma_i\theta + (1 - r_{ni})\log(1 + \exp(a_i(\theta - b_i))) + r_{ni}t_{ni}\lambda_{i1}\alpha_{\eta_1\theta}\theta + (1 - r_{ni})t_{ni}\lambda_{i0}\alpha_{\eta_0\theta}\theta \quad (\text{A8})$$

Using a Taylor approximation around $\theta = \theta_0$, we obtain

$$\log(1 + \exp(a_i(\theta - b_i))) \simeq \text{const}(\theta_0) + a_i\Psi(a_i(\theta - b_i))(\theta - \theta_0) \quad (\text{A9})$$

Where $\text{const}(\theta_0)$ is a function of θ_0 . Using the approximation in Equation (A9) it can be seen that the multiplication factors of θ in Equation (A8) are given by

$$x_{ni}r_{ni}a_i + r_{ni}\alpha_{\xi\theta}\gamma_i + (1 - r_{ni})a_iP_i(\theta) + r_{ni}t_{ni}\lambda_{i1}\alpha_{\eta_1\theta} + (1 - r_{ni})t_{ni}\lambda_{i0}\alpha_{\eta_0\theta} \quad , \quad (\text{A10})$$

where $P_i(\theta) = \Psi(a_i(\theta - b_i))$. We can extract the local item scores for θ from Equation (A10):

$$\text{observed item responses } (r_{ni} = 1): \quad a_i + \alpha_{\xi\theta}\gamma_i + \alpha_{\eta_1\theta}\lambda_{i1}t_{ni} \quad (\text{A11})$$

$$\text{omitted item responses } (r_{ni} = 0): \quad a_iP_i(\theta) + 0 + \alpha_{\eta_0\theta}\lambda_{i0}t_{ni}$$

These statistics are defined on the logit metric and are unique up to the addition of a constant c ; that is

$$P(X_{ni} = 1) = \frac{\exp(p_1)}{\exp(p_0) + \exp(p_1)} = \frac{\exp(p_1 + c)}{\exp(p_0 + c) + \exp(p_1 + c)} \quad (\text{A12})$$

By defining $c = -(\alpha_{\xi\theta}\gamma_i + \alpha_{\eta_1\theta}\lambda_{i1}t_{ni})$, the local item scores for observed and omitted item responses in Equation (A11) can be equivalently rewritten

$$\text{observed item responses } (r_{ni} = 1): \quad a_i \quad (\text{A13})$$

$$\text{omitted item responses } (r_{ni} = 0): \quad a_iP_i(\theta) - \alpha_{\xi\theta}\gamma_i - (\alpha_{\eta_1\theta}\lambda_{i1} - \alpha_{\eta_0\theta}\lambda_{i0})t_{ni}$$

References

- Adams, R. J. (2003). Response to ‘Cautions on OECD’s recent educational survey (PISA)’. *Oxford Review of Education*, 29(3), 379–389. <https://doi.org/10.1080/03054980307445>

- Aitkin, M. & Aitkin, I. (2006). *Investigation of the identifiability of the 3PL model in the NAEP 1986 math survey*. Technical report. <https://bit.ly/35b79X0>
- Berk, R., Brown, L., Buja, A., George, E., Pitkin, E., Zhang, K., & Zhao, L. (2014). Misspecified mean function regression: Making good use of regression models that are wrong. *Sociological Methods & Research*, 43(3), 422–451. <https://doi.org/10.1177/0049124114526375>
- Birnbaum, A. (1968). Some latent trait models and their use in inferring an examinee's ability. In F. M. Lord & M. R. Novick (Eds.), *Statistical theories of mental test scores* (pp. 397–479). Reading, MA: MIT Press.
- Binder, D. A., & Roberts, G. R. (2003). Design-based and model-based methods for estimating model parameters. In R. L. Chambers, & C. J. Skinner (Eds.), *Analysis of survey data* (pp. 29–48). New York: Wiley. <https://doi.org/10.1002/0470867205.ch3>
- Bock, R. D., Brennan, R. L., & Muraki, E. (2002). The information in multiple ratings. *Applied Psychological Measurement*, 26(4), 364–375. <https://doi.org/10.1177/014662102237794>
- Bolsinova, M., Tijmstra, J., Molenaar, D., & De Boeck, P. (2017). Conditional dependence between response time and accuracy: An overview of its possible sources and directions for distinguishing between them. *Frontiers in Psychology*, 8, 202. <https://doi.org/10.3389/fpsyg.2017.00202>
- Bolt, D. M., Deng, S., & Lee, S. (2014). IRT model misspecification and measurement of growth in vertical scaling. *Journal of Educational Measurement*, 51(2), 141–162. <https://doi.org/10.1111/jedm.12039>
- Bradlow, E. T., Wainer, H., & Wang, X. (1999). A Bayesian random effects model for testlets. *Psychometrika*, 64, 153–168. <https://doi.org/10.1007/BF02294533>
- Brennan, R. L. (1998). Misconceptions at the intersection of measurement theory and practice. *Educational Measurement: Issues and Practice*, 17, 5–9. <https://doi.org/10.1111/j.1745-3992.1998.tb00615.x>
- Brennan, R. L. (2001). *Generalizability theory*. New York: Springer. <https://doi.org/10.1007/978-1-4757-3456-0>
- Brennan, R. L. (2010). Generalizability theory and classical test theory. *Applied Measurement in Education*, 24(1), 1–21. <https://doi.org/10.1080/08957347.2011.532417>
- Brewer, K. (2013). Three controversies in the history of survey sampling. *Survey Methodology*, 39(2), 249–262. <https://bit.ly/3mhYPxx>
- Camilli, G. (1993). The case against item bias detection techniques based on internal criteria: Do item bias procedures obscure test fairness issues? In P. W. Holland & H. Wainer (Eds.), *Differential item functioning: Theory and practice* (pp. 397–417). Hillsdale, NJ: Erlbaum. <https://doi.org/10.4324/9780203357811>
- Camilli, G. (2018). IRT scoring and test blueprint fidelity. *Applied Psychological Measurement*, 42(5), 393–400. <https://doi.org/10.1177/0146621618754897>

- Carstensen, C. H. (2013). Linking PISA competencies over three cycles – Results from Germany. In M. Prenzel, M. Kobarg, K. Schöps, & S. Rönnebeck (Eds.), *Research on PISA* (pp. 199–213). Amsterdam: Springer. https://doi.org/10.1007/978-94-007-4458-5_12
- Chandler, R. E., & Bate, S. (2007). Inference for clustered data using the independence loglikelihood. *Biometrika*, 94(1), 167–183. <https://doi.org/10.1093/biomet/asm015>
- Chiu, T. W. & Camilli, G. (2013). Comment on 3PL IRT adjustment for guessing. *Applied Psychological Measurement*, 37(1), 76–86. <https://doi.org/10.1177/0146621612459369>
- Conijn, J. M., Emons, W. H., van Assen, M. A., & Sijtsma, K. (2011). On the usefulness of a multilevel logistic regression approach to person-fit analysis. *Multivariate Behavioral Research*, 46(2), 365–388. <https://doi.org/10.1080/00273171.2010.546733>
- Costa, D. R., Bolsinova, M., Tijmstra, J., & Andersson, B. (2021). Improving the precision of ability estimates using time-on-task variables: Insights from the PISA 2012 computer-based assessment of mathematics. *Frontiers in Psychology*, 12, 579128. <https://doi.org/10.3389/fpsyg.2021.579128>
- Cronbach, L. J. (1951). Coefficient alpha and the internal structure of tests. *Psychometrika*, 16, 297–334. <https://doi.org/10.1007/BF02310555>
- Cronbach, L. J., Rajaratnam, N., & Gleser, G. C. (1963). Theory of generalizability: a liberalization of reliability theory. *British Journal of Mathematical and Statistical Psychology*, 16, 137–163. <https://doi.org/10.1111/j.2044-8317.1963.tb00206.x>
- Cronbach, L. J., Schoenemann, P., & McKie, D. (1965). Alpha coefficient for stratified-parallel tests. *Educational and Psychological Measurement*, 25, 291–312. <https://doi.org/10.1177/001316446502500201>
- Cronbach, L. J., & Shavelson, R. J. (2004). My current thoughts on coefficient alpha and successor procedures. *Educational and Psychological Measurement*, 64(3), 391–418. <https://doi.org/10.1177/0013164404266386>
- Culpepper, S. A. (2017). The prevalence and implications of slipping on low-stakes, large-scale assessments. *Journal of Educational and Behavioral Statistics*, 42(6), 706–725. <https://doi.org/10.3102/1076998617705653>
- Debeer, D., & Janssen, R. (2013). Modeling item-position effects within an IRT framework. *Journal of Educational Measurement*, 50(2), 164–185. <https://doi.org/10.1111/jedm.12009>
- Deribo, T., Kroehne, U., & Goldhammer, F. (2021). Model-based treatment of rapid guessing. *Journal of Educational Measurement*, 58(2), 281–303. <https://doi.org/10.1111/jedm.12290>
- Dimitrov, D. M. (2016). An approach to scoring and equating tests with binary items: Piloting with large-scale assessments. *Educational Psychological Measurement*, 76(6), 954–975. <https://doi.org/10.1177/0013164416631100>

- Ellis, J. L., & Van den Wollenberg, A. L. (1993). Local homogeneity in latent trait models. A characterization of the homogeneous monotone IRT model. *Psychometrika*, 58(3), 417–429. <https://doi.org/10.1007/BF02294649>
- ETS (2015). *PISA 2015 Field trial analysis report: Outcomes of the cognitive assessment*. Internal Report. Princeton, NJ.
- Falk, C. F., & Cai, L. (2016). Maximum marginal likelihood estimation of a monotonic polynomial generalized partial credit model with applications to multiple group analysis. *Psychometrika*, 81(2), 434–460. <https://doi.org/10.1007/s11336-014-9428-7>
- Ferrando, P. J. (2019). A comprehensive IRT approach for modeling binary, graded, and continuous responses with error in persons and items. *Applied Psychological Measurement*, 43(5), 339–359. <https://doi.org/10.1177/0146621618817779>
- Feuerstahler, L. M. (2019). Metric transformations and the filtered monotonic polynomial item response model. *Psychometrika*, 84(1), 105–123. <https://doi.org/10.1007/s11336-018-9642-9>
- Fox, J.-P., & Verhagen, A. J. (2010). Random item effects modeling for cross-national survey data. In E. Davidov, P. Schmidt, & J. Billiet (Eds.), *Cross-cultural analysis: Methods and applications* (pp. 461–482). Routledge Academic.
- Foy, P., Fishbein, B., M. von Davier, & Yin, L. (2020) Implementing the TIMSS 2019 scaling methodology. In M. O. Martin, M. von Davier, & I. V. Mullis (Eds.), *TIMSS 2019 technical report*. Boston College: IEA.
- Foy, P. & Yin, L. (2017) Scaling the PIRLS 2016 achievement data. In M. O. Martin, I. V. Mullis, & M. Hooper (Eds.), *Methods and procedures in PIRLS 2016*. Boston College: IEA.
- Frey, A., & Hartig, J. (2020). Methodological challenges of international student assessment. In H. Harju-Luukkainen, N. McElvany & J. Stang (Eds.), *Monitoring Student Achievement in the 21st Century* (pp. 39–49). Cham: Springer. https://doi.org/10.1007/978-3-030-38969-7_4
- Frey, A., Hartig, J., & Rupp, A. A. (2009). An NCME instructional module on booklet designs in large-scale assessments of student achievement: Theory and practice. *Educational Measurement: Issues and Practice*, 28(3), 39–53. <https://doi.org/10.1111/j.1745-3992.2009.00154.x>
- Gebhardt, E., & Adams, R. J. (2007). The influence of equating methodology on reported trends in PISA. *Journal of Applied Measurement*, 8(3), 305–322. <https://bit.ly/2UDjWib>
- Goldstein, H. (1980). Dimensionality, bias, independence and measurement scale problems in latent trait test score models. *British Journal of Mathematical and Statistical Psychology*, 33(2), 234–246. <https://doi.org/10.1111/j.2044-8317.1980.tb00610.x>
- Grund, S., Lüdtke, O., & Robitzsch, A. (2021). On the treatment of missing data in background questionnaires in educational large-scale assessments: An evaluation of different procedures. *Journal of Educational and Behavioral Statistics*, 46(4), 430–465. <https://doi.org/10.3102/1076998620959058>

- Guo, J., Hu, X., Marsh, H. W., & Pekrun, R. (2021). Relations of epistemic beliefs with motivation, achievement, and aspirations in science: Generalizability across 72 societies. *Journal of Educational Psychology*. Advance online publication. <https://doi.org/10.1037/edu0000660>
- Haberkorn K., Pohl S., Carstensen C. (2016) Scoring of complex multiple choice items in NEPS competence tests. In H. P. Blossfeld, J. von Maurice, M. Bayer & J. Skopek (Eds.), *Methodological issues of longitudinal surveys*. Springer VS, Wiesbaden. https://doi.org/10.1007/978-3-658-11994-2_29
- Haertel, E. H. (1989). Using restricted latent class models to map the skill structure of achievement items. *Journal of Educational Measurement*, 26(4), 301–321. <https://doi.org/10.1111/j.1745-3984.1989.tb00336.x>
- He, Y., & Cui, Z. (2020). Evaluating robust scale transformation methods with multiple outlying common items under IRT true score equating. *Applied Psychological Measurement*, 44(4), 296–310. <https://doi.org/10.1177/0146621619886050>
- Holland, P. W., & Wainer, H. (Eds.) (1993). *Differential item functioning: Theory and practice*. Hillsdale, NJ: Erlbaum. <https://doi.org/10.4324/9780203357811>
- Hong, M. R., & Cheng, Y. (2019). Robust maximum marginal likelihood (RMML) estimation for item response theory models. *Behavior Research Methods*, 51(2), 573–588. <https://doi.org/10.3758/s13428-018-1150-4>
- Jerrim, J., Micklewright, J., Heine, J.-H., Sälzer, C., & McKeown, C. (2018). PISA 2015: how big is the ‘mode effect’ and what has been done about it? *Oxford Review of Education*, 44(4), 476–493. <https://doi.org/10.1080/03054985.2018.1430025>
- Jerrim, J., Parker, P., Choi, A., Chmielewski, A. K., Sälzer, C., & Shure, N. (2018). How robust are cross-country comparisons of PISA scores to the scaling model used? *Educational Measurement: Issues and Practice*, 37(4), 28–39. <https://doi.org/10.1111/emip.12211>
- Jin, K. Y. and Wang, W. C. (2014). Item response theory models for performance decline during testing. *Journal of Educational Measurement*, 51(2), 178–200. <https://doi.org/10.1111/jedm.12041>
- Kane, M. (1982). A sampling model for validity. *Applied Psychological Measurement*, 6(2), 125–160. <https://doi.org/10.1177/014662168200600201>
- Kane, M. T., & Brennan, R. L. (1977). The generalizability of class means. *Review of Educational Research*, 47(2), 267–292. <https://doi.org/10.3102/00346543047002267>
- Keller, L. K., Preckel, F., Eccles, J. S., & Brunner, M. (2020). Top-performing math students in 82 countries: An integrative data analysis of gender differences in achievement, achievement profiles, and achievement motivation. *PsyArXiv*, 2 May 2020. <https://doi.org/10.31234/osf.io/73wap>
- Kolenikov, S. (2010). Resampling variance estimation for complex survey data. *Stata Journal*, 10(2), 165–199. <https://doi.org/10.1177/1536867X1001000201>

- Lechner, C. M., Bhaktha, N., Groskurth, K., & Bluemke, M. (2021). Why ability point estimates can be pointless: A primer on using skill measures from large-scale assessments in secondary analyses. *Measurement Instruments for the Social Sciences*, 3, 2 <https://doi.org/10.1186/s42409-020-00020-5>
- Liou, M., & Yu, L. C. (1991). Assessing statistical accuracy in ability estimation: A bootstrap approach. *Psychometrika*, 56(1), 55–67. <https://doi.org/10.1007/BF02294585>
- Little, R. J. (2004). To model or not to model? Competing modes of inference for finite population sampling. *Journal of the American Statistical Association*, 99(466), 546–556. <https://doi.org/10.1198/016214504000000467>
- Little, R. J., & Rubin, D. B. (2002). *Statistical analysis with missing data*. Hoboken, NJ: Wiley. <https://doi.org/10.1002/9781119013563>
- Lohr, S. L. (2010). *Sampling: Design and analysis*. Brooks/Cole Cengage Learning.
- Loken, E., & Rulison, K. L. (2010). Estimation of a four-parameter item response theory model. *British Journal of Mathematical and Statistical Psychology*, 63(3), 509–525. <https://doi.org/10.1348/000711009X474502>
- Magis, D. (2013). A note on the item information function of the four-parameter logistic model. *Applied Psychological Measurement*, 37(4), 304–315. <https://doi.org/10.1177/0146621613475471>
- Magis, D. (2015). A note on the equivalence between observed and expected information functions with polytomous IRT models. *Journal of Educational and Behavioral Statistics*, 40(1), 96–105. <https://doi.org/10.3102/1076998614558122>
- Markus, K. A., & Borsboom, D. (2013). *Frontiers of test validity theory: Measurement, causation, and meaning*. New York: Routledge. <https://doi.org/10.4324/9780203501207>
- Maris, G., & Bechger, T. (2009). On interpreting the model parameters for the three parameter logistic model. *Measurement: Interdisciplinary Research and Perspectives*, 7(2), 75–88. <https://doi.org/10.1080/15366360903070385>
- Marsh, H. W. (2016). Cross-cultural generalizability of year in school effects: Negative effects of acceleration and positive effects of retention on academic self-concept. *Journal of Educational Psychology*, 108(2), 256–273. <https://doi.org/10.1037/edu0000059>
- Marsman, M., Maris, G., Bechger, T., & Glas, C. (2016). What can we learn from plausible values?. *Psychometrika*, 81(2), 274–289. <https://doi.org/10.1007/s11336-016-9497-x>
- Martin, M. O., Mullis, I. V., Foy, P., Brossman, B., & Stanco, G. M. (2012). Estimating linking error in PIRLS. *IERI Monograph Series: Issues and Methodologies in Large-Scale Assessments*, 5, 35–47. <https://bit.ly/3yraNrd>
- Meinck, S. (2020). Sampling, weighting, and variance estimation. In H. Wagenmaker (Ed.), *Reliability and validity of international large-scale assessment* (pp.113–129). Cham: Springer. https://doi.org/10.1007/978-3-030-53081-5_7

- Meyer, P. (2010). *Understanding measurement: Reliability*. Cambridge: Oxford University Press.
<https://doi.org/10/grrt>
- Michaelides, M. P., & Haertel, E. H. (2014). Selection of common items as an unrecognized source of variability in test equating: a bootstrap approximation assuming random sampling of common items. *Applied Measurement in Education*, 27, 46–57.
<https://doi.org/10.1080/08957347.2013.853069>
- Mislevy, R. (1990). Scaling procedures. In E. Johnson & R. Zwick (Eds.), *Focusing the new design: The NAEP 1988 technical report* (ETS RR 19-20). Princeton, NJ: Educational Testing Service.
<https://bit.ly/3zuC5OQ>
- Mislevy, R. J. (1991). Randomization-based inference about latent variables from complex samples. *Psychometrika*, 56, 177–196. <https://doi.org/10.1007/BF02294457>
- Monseur, C., Baye, A., Lafontaine, D., & Quittre, V. (2011). PISA test format assessment and the local independence assumption. *IERI Monographs Series. Issues and Methodologies in Large-Scale Assessments*, 4, 131–158. <https://bit.ly/3k6wIyU>
- Monseur, C., Sibberns, H., and Hastedt, D. (2008). Linking errors in trend estimation for international surveys in education. *IERI Monographs Series. Issues and Methodologies in Large-Scale Assessment*, 1, 113–122. <https://bit.ly/38aTVeZ>
- Nunnally, J. C. & Bernstein, I. H. (1994). *Psychometric theory*. New York: McGraw-Hill.
- OECD (2014). *PISA 2012 technical report*. Paris: OECD Publishing.
- OECD (2017). *PISA 2015 technical report*. Paris: OECD Publishing.
- OECD (2020). *PISA 2018 technical report*. Paris: OECD Publishing.
- Oliveri, M. E., & von Davier, M. (2011). Investigation of model fit and score scale comparability in international assessments. *Psychological Test and Assessment Modeling*, 53(3), 315–333.
<https://bit.ly/3mkaRGO>
- Pellegrino, J. W., & Chudowsky, N. (2003). The foundations of assessment. *Measurement: Interdisciplinary Research and Perspectives*, 1(2), 103–148.
https://doi.org/10.1207/S15366359MEA0102_01
- Penfield, R. D., & Camilli, G. (2007). Differential item functioning and item bias. In C. R. Rao & S. Sinharay (Eds.), *Handbook of statistics, Vol. 26: Psychometrics* (pp. 125–167). Amsterdam: Elsevier. [https://doi.org/10.1016/S0169-7161\(06\)26005-X](https://doi.org/10.1016/S0169-7161(06)26005-X)
- Pohl, S., & Carstensen, C. H. (2013). Scaling of competence tests in the national educational panel study - Many questions, some answers, and further challenges. *Journal of Educational Research Online*, 5(2), 189–216.
- Pohl, S., Gräfe, L., & Rose, N. (2014). Dealing with omitted and not-reached items in competence tests: Evaluating approaches accounting for missing responses in item response theory models. *Educational and Psychological Measurement*, 74(3), 423–452.
<https://doi.org/10.1177/0013164413504926>

- Pohl, S., Ulitzsch, E., & von Davier, M. (2021). Reframing rankings in educational assessments. *Science*, 372(6540), 338–340. <https://doi.org/10.1126/science.abd3300>
- Raiche, G., Magis, D., Blais, J. G., & Brochu, P. (2012). Taking atypical response patterns into account: A multidimensional measurement model from item response theory. In M. Simon, K. Ercikan & M. Rousseau (Eds.), *Improving large-scale assessment in education* (pp. 238–259). New York: Routledge. <https://doi.org/10.4324/9780203154519>
- Ramsay, J. O. & Winsberg, S. (1991). Maximum marginal likelihood estimation for semiparametric item analysis. *Psychometrika*, 56(3), 365–379. <https://doi.org/10.1007/BF02294480>
- Rasch, G. (1960). *Probabilistic models for some intelligence and attainment tests*. Copenhagen: Danish Institute for Educational Research.
- Reckase, M. D. (2017). A tale of two models: Sources of confusion in achievement testing. *ETS Research Report*, ETS RR-17-44. <https://doi.org/10.1002/ets2.12171>
- Robitzsch, A. (2020). About still nonignorable consequences of (partially) ignoring missing item responses in large-scale assessment. *OSF Preprints*. 20 October 2020. <https://doi.org/10.31219/osf.io/hmy45>
- Robitzsch, A., & Lüdtke, O. (2019). Linking errors in international large-scale assessments: calculation of standard errors for trend estimation. *Assessment in Education: Principles, Policy & Practice*, 26(4), 444–465. <https://doi.org/10.1080/0969594X.2018.1433633>
- Robitzsch, A., & Lüdtke, O. (2020). A review of different scaling approaches under full invariance, partial invariance, and noninvariance for cross-sectional country comparisons in large-scale assessments. *Psychological Test and Assessment Modeling*, 62(2), 233–279. <https://bit.ly/3kFiXaH>
- Robitzsch, A., & Lüdtke, O. (2021). Mean comparisons of many groups in the presence of DIF: An evaluation of linking and concurrent scaling approaches. *Journal of Educational and Behavioral Statistics*. Advance online publication. <https://doi.org/10.3102/10769986211017479>
- Robitzsch, A., Lüdtke, O., Goldhammer, F., Kroehne, U., & Köller, O. (2020). Reanalysis of the German PISA data: A comparison of different approaches for trend estimation with a particular emphasis on mode effects. *Frontiers in Psychology*, 11, 884. <https://doi.org/10.3389/fpsyg.2020.00884>
- Rohwer, G. (2013). *Making sense of missing answers in competence tests*. NEPS Working Paper No. 30. Bamberg: Otto-Friedrich-Universität, Nationales Bildungspanel. <https://bit.ly/3kzmEPc>
- Rose, N., von Davier, M., & Nagengast, B. (2017). Modeling omitted and not-reached items in IRT models. *Psychometrika*, 82(3), 795–819. <https://doi.org/10.1007/s11336-016-9544-7>
- Rust, K. F., Krawchuk, S., & Monseur, C. (2017). Sample design, weighting, and calculation of sampling variance. In P. Lietz, J. C. Creswell, K. F. Rust, & R. J. Adams (Eds.), *Implementation of large-scale education assessments* (pp.137–167). Chichester: Wiley. <https://doi.org/10.1002/9781118762462.ch5>

- Rutkowski, D., & Rutkowski, L. (2021). Running the wrong race? The case of PISA for development. *Comparative Education Review*, 65(1), 147–165. <https://doi.org/10.1086/712409>
- Rutkowski, L., & Rutkowski, D. (2019). Methodological challenges to measuring heterogeneous populations internationally. In L. E. Suter, E. Smith, & B. D. Denman (Eds.), *The SAGE handbook of comparative studies in education* (pp. 126–140). Thousand Oaks: Sage. <https://doi.org/10.4135/9781526470379>
- Rutkowski, L., von Davier, M., & Rutkowski, D. (Eds.) (2013). *A handbook of international large-scale assessment: Background, technical issues, and methods of data analysis*. London: Chapman Hall/CRC Press. <https://doi.org/10.1201/b16061>
- Sachse, K. A., Mahler, N., & Pohl, S. (2019). When nonresponse mechanisms change: Effects on trends and group comparisons in international large-scale assessments. *Educational and Psychological Measurement*, 79(4), 699–726. <https://doi.org/10.1177/0013164419829196>
- Sachse, K. A., Roppelt, A., & Haag, N. (2016). A comparison of linking methods for estimating national trends in international comparative large-scale assessments in the presence of cross-national DIF. *Journal of Educational Measurement*, 53(2), 152–171. <https://doi.org/10.1111/jedm.12106>
- Särndal, C. E., Swensson, B., & Wretman, J. (2003). *Model assisted survey sampling*. New York: Springer. <https://doi.org/10.1007/978-1-4612-4378-6>
- San Martín, E., González, J., & Tuerlinckx, F. (2015). On the unidentifiability of the fixed-effects 3PL model. *Psychometrika*, 80(2), 450–467. <https://doi.org/10.1007/s11336-014-9404-2>
- Schuster, C., & Yuan, K. H. (2011). Robust estimation of latent ability in item response models. *Journal of Educational and Behavioral Statistics*, 36(6), 720–735. <https://doi.org/10.3102/1076998610396890>
- Shealy, R., & Stout, W. A. (1993). Model-based standardization approach that separates true bias/DIF from group ability differences and detects test bias/DTF as well as item bias/DIF. *Psychometrika*, 58, 159–194. <https://doi.org/10.1007/BF02294572>
- Singer, J. D., & Braun, H. I. (2018). Testing international education assessments. *Science*, 360(6384), 38–40. <http://doi.org/10.1126/science.aar4952>
- Ståhl, G., Saarela, S., Schnell, S., Holm, S., Breidenbach, J., Healey, S.P., Patterson, P.L., Magnussen, S., Næsset, E., McRoberts, R.E., & Gregoire, T.G. (2016). Use of models in large-area forest surveys: Comparing model-assisted, model-based and hybrid estimation. *Forest Ecosystems*, 3, 5. <https://doi.org/10.1186/s40663-016-0064-9>
- Stenner, A. J., Burdick, D. S., & Stone, M. H. (2008). Formative and reflective models: Can a Rasch analysis tell the difference?. *Rasch Measurement Transactions*, 22(1), 1152–1153. <https://www.rasch.org/rmt/rmt221d.htm>
- Stenner, A. J., Stone, M. H., & Burdick, D. S. (2009). Indexing vs. measuring. *Rasch Measurement Transactions*, 22(4), 1176–1177. <https://www.rasch.org/rmt/rmt224b.htm>

- Stout, W. F. (1990). A new item response theory modeling approach with applications to unidimensionality assessment and ability estimation. *Psychometrika*, 55(2), 293–325. <https://doi.org/10.1007/BF02295289>
- Tijmstra, J., Liaw, Y., Bolsinova, M., Rutkowski, L., & Rutkowski, D. (2020). Sensitivity of the RMSD for detecting item-level misfit in low-performing countries. *Journal of Educational Measurement*, 57(4), 566–583. <https://doi.org/10.1111/jedm.12263>
- Tryon, R. C. (1957). Reliability and behavior domain validity: Reformulation and historical critique. *Psychological Bulletin*, 54(3), 229–249. <https://doi.org/10.1037/h0047980>
- Ulitzsch, E., von Davier, M., & Pohl, S. (2020a). A hierarchical latent response model for inferences about examinee engagement in terms of guessing and item-level non-response. *British Journal of Mathematical and Statistical Psychology*, 73(S1), 83–112. <https://doi.org/10.1111/bmsp.12188>
- Ulitzsch, E., von Davier, M., & Pohl, S. (2020b). Using response times for joint modeling of response and omission behavior. *Multivariate Behavioral Research*, 55(3), 425–453. <https://doi.org/10.1080/00273171.2019.1643699>
- von Davier, M. (2009). Is there need for the 3PL model? Guess what?. *Measurement: Interdisciplinary Research and Perspectives*, 7(2), 110–114. <https://doi.org/10.1080/15366360903117079>
- von Davier, M., & Sinharay, S. (2014). Analytics in international large-scale assessments: Item response theory and population models. In L. Rutkowski, M. von Davier & D. Rutkowski (Eds.), *Handbook of international large-scale assessment* (pp. 155–174). Boca Raton: CRC Press. <https://doi.org/10.1201/b16061>
- von Davier, M., Khorramdel, L., He, Q., Shin, H. J., & Chen, H. (2019). Developments in psychometric population models for technology-based large-scale assessments: An overview of challenges and opportunities. *Journal of Educational and Behavioral Statistics*, 44(6), 671–705. <https://doi.org/10.3102/1076998619881789>
- von Davier, M., Yamamoto, K., Shin, H. J., Chen, H., Khorramdel, L., Weeks, J., Davis, S., Kong, N., & Kandathil, M. (2019). Evaluating item response theory linking and model fit for data from PISA 2000–2012. *Assessment in Education: Principles, Policy & Practice*, 26(4), 466–488. <https://doi.org/10.1080/0969594X.2019.1586642>
- Yamamoto, K., Khorramdel, L., & Shin, H. J. (2018). Introducing multistage adaptive testing into international large-scale assessments designs using the example of PIAAC. *Psychological Test and Assessment Modeling*, 60(3), 347–368.
- Yamamoto, K., Shin, H. J., & Khorramdel, L. (2018). Multistage adaptive testing design in international large-scale assessments. *Educational Measurement: Issues and Practice*, 37(4), 16–27. <https://doi.org/10.1111/emip.12226>
- Yen, W. M., & Fitzpatrick, A. R. (2006). Item response theory. In R. L. Brennan, R. L. (Ed.), *Educational measurement* (pp. 111–154), Praeger Publishers.

- Young, C., & Holsteen, K. (2017). Model uncertainty and robustness: A computational framework for multimodel analysis. *Sociological Methods & Research*, 46(1), 3–40.
<https://doi.org/10.1177/0049124115610347>
- Wainer, H., & Thissen, D. (1987). Estimating ability with the wrong model. *Journal of Educational Statistics*, 12(4), 339–368. <https://doi.org/10.3102/10769986012004339>
- Wainer, H., & Wright, B. D. (1980). Robust estimation of ability in the Rasch model. *Psychometrika*, 45(3), 373–391. <https://doi.org/10.1007/BF02293910>
- Westfall, P. H., Henning, K. S., & Howell, R. D. (2012). The effect of error correlation on interfactor correlation in psychometric measurement. *Structural Equation Modeling*, 19(1), 99–117.
<https://doi.org/10.1080/10705511.2012.634726>
- White, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica*, 50(1), 1–25. doi: 10.2307/1912526
- Wu, M. (2005). The role of plausible values in large-scale surveys. *Studies in Educational Evaluation*, 31(2-3), 114–128. <https://doi.org/10.1016/j.stueduc.2005.05.005>
- Wu, M. (2010). Measurement, sampling, and equating errors in large-scale assessments. *Educational Measurement: Issues and Practice*, 29, 15–27. <https://doi.org/10.1111/j.1745-3992.2010.00190.x>
- Zieger, L., Sims, S., & Jerrim, J. (2019). Comparing teachers' job satisfaction across countries: A multiple-pairwise measurement invariance approach. *Educational Measurement: Issues and Practice*, 38(3), 75–85. <https://doi.org/10.1111/emip.12254>
- Zinbarg, R. E., Revelle, W., Yovel, I., & Li, W. (2005). Cronbach's α , Revelle's β , and McDonald's ω_H : Their relations with each other and two alternative conceptualizations of reliability. *Psychometrika*, 70(1), 123–133. <https://doi.org/10.1007/s11336-003-0974-7>
- Zwitser, R. J., Glaser, S. S. F., & Maris, G. (2017). Monitoring countries in a changing world: A new look at DIF in international surveys. *Psychometrika*, 82(1), 210–232.
<https://doi.org/10.1007/s11336-016-9543-8>