Measurement Invariance Violation Indices (MIVIs):**Effect Sizes for (Partial) Non-Invariance of Items and Item Sets**

Katharina Groskurth^{1,2*}, Matthias Bluemke¹, and Clemens M. Lechner¹

¹*GESIS – Leibniz Institute for the Social Sciences, Mannheim, Germany*; ²*Graduate School of Economic and Social Sciences, University of Mannheim, Germany*

* Katharina Groskurth, e-mail: katharina.groskurth@gmx.de

Draft version 3.0, 08-31-2022. This paper has not been peer reviewed.

Suggested citation:

Groskurth, K., Bluemke, M., & Lechner, C. M. (2021). *Measurement Invariance Violation Indices (MIVIs): Effect sizes for (partial) non-invariance of items and item sets*.

<https://doi.org/10.31234/osf.io/fbshu>

24 **Author Note**

25 **Funding details**

26 not applicable

27 **Disclosure statement**

28 The authors report there are no competing interests to declare.

29 **Data availability statement**

30 The code to reproduce the simulation, analysis, and results of this article is included in
31 Additional Files 2, 3, and 4 of the Supplementary Materials, which makes the study eligible
32 for the Open Materials Badge.

33 **Authors' Contributions**

34 **Katharina Groskurth:** Methodology, Software, Validation, Formal Analysis, Investigation,
35 Data Curation, Writing – Original Draft, Writing – Review & Editing, Visualization.

36 **Matthias Bluemke:** Conceptualization, Methodology, Writing – Original Draft, Writing –
37 Review & Editing, Supervision, Project Administration.

38 **Clemens M. Lechner:** Resources, Writing – Review & Editing, Supervision, Project
39 Administration.

40 **Acknowledgments**

41 We thank Tenko Raykov for the initial idea and fruitful comments on earlier versions of the
42 manuscript. This work was supported by the University of Mannheim's Graduate School of
43 Economic and Social Sciences.

44

45 **Measurement Invariance Violation Indices (MIVIs): Effect Sizes for** 46 **(Partial) Non-Invariance of Items and Item Sets**

47 In many applications, measurement invariance does not hold. When a model with a certain
48 level of invariance is rejected, the amount of non-invariance bias may either be consequential
49 or practically irrelevant. So far, few attempts have been made to quantify the extent of bias
50 due to the lack of measurement invariance. We derived new effect sizes from first principles
51 called Measurement Invariance Violation Indices (MIVIs) for items and item sets. MIVIs
52 assume that one can compare the basic measurement model across groups (i.e., configural
53 invariance holds) but cannot compare some factor loadings, intercepts, and/or unique
54 variances. Assuming partial invariance for a set of items, MIVIs quantify non-invariant factor
55 loading, intercept, or uniqueness differences in relation to the pooled latent standard deviation
56 (either per item or as an average for item sets). Thus, parameter differences can be interpreted
57 in standard deviation units (of the pooled latent variable). One can further inspect the
58 compensatory cancelation and non-compensatory aggregation effects of non-invariance bias
59 when maintaining the directional information (signed MIVIs). MIVIs support the group-fair
60 item selection, help to evaluate the questionnaire quality, and allow for assessing the amount
61 of non-invariance bias in index scores (e.g., mean scores).

62 **Keywords:** effect size; factor analysis; metric invariance; scalar invariance; strict invariance

Introduction

Statistical comparisons are ubiquitous in the social sciences. Researchers may, for instance, compare numeracy scores of an exam across two school classes or the learning motivation across two entire school districts. Such comparisons are only valid if the test captures the numeracy scores, alternatively learning motivation, without bias across the two groups. The latent construct (e.g., numeracy or learning motivation) must be placed on a comparable scale. Otherwise, researchers compare apples with oranges (Chen, 2008). Examining the validity of numerical group comparisons is called measurement invariance testing in linear latent variable modeling (and differential item functioning in item response theory frameworks; Osterlind & Everson, 2009).

Crucially, some statistical comparisons (i.e., including variances or means) are only valid if specific levels of measurement invariance hold. If metric invariance does not hold, any differences (or lack of differences) in latent (co)variances might be spurious (Steenkamp & Baumgartner, 1998). If scalar invariance does not hold, any differences (or lack of differences) in latent or observed means might be spurious (Steenkamp & Baumgartner, 1998; Steinmetz, 2013; cf. DeShon, 2004; Wu et al., 2007). The uniqueness invariance level answers the question of whether any differences (or lack of differences) in observed (co)variances might be spurious (Steenkamp & Baumgartner, 1998; Millsap & Olivera-Aguilar, 2012).

Despite the necessity of establishing measurement invariance, applied researchers often compare group-level statistics (such as regression weights or sum scores) without testing for measurement invariance (e.g., Boer et al., 2018). When researchers do report measurement invariance tests, they frequently find that invariance does not hold; especially uniqueness or scalar invariance is often hard to reach (e.g., Davidov et al., 2012, 2014; Dong & Dumas, 2020). A practical question thus arises: How consequential are any violations of

invariance? In other words: How much non-invariance bias is present when comparing relevant group-level statistics (e.g., regression coefficients or mean scores)?

This work proposes an intuitive way to empirically quantify non-invariance bias in factor loadings, intercepts, and/or unique variances. Whereas researchers should refrain from comparing statistics (such as mean scores) across groups when non-invariance bias is large (e.g., the intercepts differ strongly due to the different use of response styles), small non-invariance bias may not be detrimental. Several effect sizes currently compete for acceptance among psychometricians and more widespread recognition by applied researchers (for a review, see Gunn et al., 2020). Our proposed effect sizes, termed Measurement Invariance Violation Indices (MIVIs), quantify the differences in measurement model parameters (i.e., factor loadings, intercepts, or unique standard deviations) in units of latent standard deviations pooled across groups. Like Cohen's d , MIVIs provide an intuitive metric for bias in items and item sets. The indices are readily implemented in standard structural equation modeling (SEM) software for confirmatory factor analysis (CFA). We provide examples for the few code lines which need to be added in Mplus and R to estimate MIVIs.

Testing for Measurement Invariance

Basics of Measurement Invariance Testing

Before defining MIVIs as effect size measures, we outline the basics of measurement invariance testing (Davidov et al., 2014; Millsap, 2011; Millsap & Olivera-Aguilar, 2012; Steenkamp & Baumgartner, 1998). Assuming a common-factor model holds, one can test measurement invariance across multiple groups. In the following, we consider a multi-group application in a unidimensional measurement scenario for individual i : Let $\mathbf{x}$ be a $p \times 1$ vector of observed measures, such as a set of approximately continuous rating scores across p items obtained from a social survey or psychological test. Let w_i be the individual score on the common factor forming the latent dimension underlying the correlated p measures to achieve

model identification (if necessary, by additional constraints). We measure w_i by a set of items $\mathbf{x}_{ik}$ for an individual who belongs to a specific group, such as language, nationality, gender, school class, or age group. We can then express the congeneric factor model for the k th subgroup ($k = 1, \dots, K$), assuming p items and one common factor, as

$$\mathbf{x}_{ik} = \boldsymbol{\tau}_k + \boldsymbol{\lambda}_k w_{ik} + \mathbf{u}_{ik}. \quad (1)$$

Here, $\boldsymbol{\tau}_k$ is a $p \times 1$ vector of regression intercept parameters, $\boldsymbol{\lambda}_k$ is a $p \times 1$ vector of factor loading parameters, and $\mathbf{u}_{ik}$ is a $p \times 1$ vector of unique scores. In line with classical test theory, unique scores are assumed to be statistically independent (i.e., uncorrelated) and do not correlate with common scores of the latent variable. The variance-covariance matrix, $\boldsymbol{\Sigma}$, of the factor model in the k th subgroup is

$$\boldsymbol{\Sigma}_k = \boldsymbol{\lambda}_k \Phi_k \boldsymbol{\lambda}_k' + \boldsymbol{\Theta}_k. \quad (2)$$

Here, Φ_k is the variance of the latent scores w_k and $\boldsymbol{\Theta}_k$ is the variance matrix of the unique scores $\mathbf{u}_k$. Then, the observed item means $E(\mathbf{x}_k)$ are related to the latent factor mean κ_k in the k th subgroup as

$$E(\mathbf{x}_k) = \boldsymbol{\tau}_k + \boldsymbol{\lambda}_k \kappa_k. \quad (3)$$

Because the expected means of uniqueness terms are zero, they drop from Equation 3 (Millsap, 2011; Millsap & Olivera-Aguilar, 2012; Steenkamp & Baumgartner, 1998).¹

Measurement invariance holds if a multi-group CFA model fits after constraining cross-group parameters to equality: For a model to be invariant, the same item-to-factor structure (i.e., configural invariance), factor loadings (i.e., metric invariance; $\boldsymbol{\lambda}_1 = \dots = \boldsymbol{\lambda}_K$)

1. DeShon (2004) and Wu et al. (2007) argue that the means of uniqueness terms are rarely zero in reality. Oftentimes the factor model is not correctly specified and unmodeled latent variables lead to means of uniqueness terms that are different from zero.

and intercepts (i.e., scalar invariance, $\tau_1 = \dots = \tau_K$) must be equivalent across the K groups. For strict invariance to hold, even the unique item variances must be equal (i.e., uniqueness invariance, $\theta_1 = \dots = \theta_K$). Readers can revisit the logic of invariance testing in Additional File 1 of the Supplementary Materials. Additional File 1 not only contains a detailed description of the traditional, sequential approach to measurement invariance testing (e.g., Meredith, 1993; Millsap, 2011; Steenkamp & Baumgartner, 1998) but also the description of a less well-known but more promising approach to measurement invariance by Raykov et al. (2013). Unlike the traditional approach, Raykov et al. (2013) suggested a top-down procedure for measurement invariance to counter the threat of false discoveries.

The Problem of Establishing Measurement Invariance

Measurement invariance is indispensable for specific statistical comparisons to be valid but often hard to reach in empirical settings (Davidov et al., 2012; 2014). Lack of invariance hampers researchers' ability to perform group comparisons. These comparisons, all too often, motivated the research in the first place (e.g., rankings of countries, comparability of validity coefficients across gender, or mean-level comparisons across age groups).

However, the more restrictive the invariance level, the harder it is to reach. In Dong and Dumas's (2020) meta-analysis of personality constructs (across culture, gender, and age groups), only 8 out of 67 group comparisons (11.9%) reached uniqueness invariance and 22.4% reached scalar invariance. Most group comparisons reached metric invariance in their meta-analysis of personality constructs (35.8%). Similarly, Zercher and colleagues (2015) conducted a measurement invariance test of items measuring Schwartz's human value of Universalism across 90 groups. They found that metric invariance held for all groups, but (partial) scalar invariance held for only 37 out of 90 groups (41.1%). In Dong and Dumas's (2020) meta-analysis of personality constructs, about a quarter reached merely configural invariance (23.9%). Some group comparisons did not even attain configural invariance

(6.0%). Chen (2008), who conducted a meta-analysis of cross-cultural studies, supported the finding and found that only 9 out of 130 group comparisons lacked configural invariance (6.9%).

If uniqueness, scalar, or metric invariance is unattainable, there is always the option to improve model fit by testing for partial invariance (Byrne et al., 1989). Equality constraints on the factor loadings, intercepts, or uniquenesses are relaxed for some items, whereas the rest remain constrained (i.e., invariant). For partial invariance to hold, a pure CFA measurement model requires at least two equivalent items across groups (in terms of factor loadings, intercepts, or uniquenesses) with at least one group-specific (i.e., freely estimated) item (Byrne et al., 1989; Steenkamp & Baumgartner, 1998; Steinmetz, 2013).

Partial invariance has an important implication for model identification (if done via fixed values of group statistics; Putnick & Bornstein, 2016): If either scalar or even metric invariance does not hold, latent group statistics (i.e., latent variances in the configural model; latent means in the configural/metric model) must be fixed to a certain value (commonly 1 for latent variances and 0 for latent means) for model identification in all groups. However, partial invariance implies that latent group statistics (i.e., latent variances in the partial metric/scalar model; latent means in the partial scalar model) must only be fixed in one group for identification but can be freely estimated in all others. Thus, partial invariance enables researchers to compare latent group statistics; this does not require full invariance. Differently, comparing observed group statistics requires full invariance (e.g., Steinmetz, 2013).

Underutilization of Effect Sizes in Invariance Testing

Typically, invariance testing stops after reaching either full or partial invariance in a binary accept-or-reject logic (see Additional File 1 of the Supplementary Materials for an overview). The binary accept-or-reject logic of measurement invariance testing discards essential

information: When a certain invariance level does not (fully) hold, it remains unknown how large and serious the resulting bias in observed statistics (e.g., group differences in means) would be. Effect sizes can complement traditional model comparisons by enabling the researcher to quantify the size of non-invariant model parameter differences. In the development process of item sets, one might consider discarding items with large bias due to non-invariance but keeping those with only negligible bias. In fixed item sets, the size of non-invariance bias helps researchers evaluate (next to other information: McNeish & Wolf, 2020; Widaman & Revelle, 2022) whether simplified aggregates of the manifest item set (such as loading-weighted factor scores or unit-weighted mean scores) can be valid ways of approximating the latent variable representing the levels of the construct. If effect sizes indicate that non-invariance bias (a) adds up at the level of simplified aggregates or (b) is large for specific items or the whole item set, one might better refrain from using simplified statistics. If effect sizes indicate that (a) positive and negative parameter differences cancel out at the level of aggregates or (b) non-invariant differences are so tiny that bias is practically inconsequential for any observed scores, one might consider simple statistics for comparing groups.

Millsap and Olivera-Aguilar (2012) provided such effect size measures for factor loadings, intercepts, and unique variances, which we call difference measures here.² To define a difference measure for factor loadings, Millsap and Olivera-Aguilar (2012, see also Pornprasertmanit, 2022) assumed a unidimensional CFA model that shows partial metric invariance across two groups ($K = 2$). Item j ($j = 1, \dots, p$) has non-invariant loadings across

2. Other approaches for quantifying measurement non-invariance, which we do not review here, do exist (e.g., Nye & Drasgow, 2011; Oberski, 2014). Yet, other effect size approaches are not as easy to apply as the effect size measures of Millsap and Olivera-Aguilar (2012), or lack a clear interpretation.

groups. That is, the loading λ_{1j} of item j in group 1 differs from the loading λ_{2j} of item j in group 2, which are, thus, freely estimated across groups. Millsap and Olivera-Aguilar are not explicit about the intercepts; however, to quantify only loading differences, the intercept τ_{1j} of item j in group 1 must be equal to the intercept τ_{2j} of item j in group 2. Otherwise, the loading difference measure does not only quantify non-invariant loading differences but also intercept differences. At any common latent score w , Millsap and Olivera-Aguilar expressed the expected difference in observed item scores x_j between the two groups as

$$\begin{aligned} \text{Loading Difference}_j &= |E(x_{2j} - x_{1j}|w)| \\ &= |(\tau_{2j} - \tau_{1j}) + (\lambda_{2j} - \lambda_{1j})w|. \end{aligned} \quad (4)$$

Equation 4 represents the expected absolute difference in raw metrics between two people's item scores (both from different groups) that arise due to loading invariance despite identical values on the latent variables. Millsap and Olivera-Aguilar proposed to quantify the impact of the loading differences (and/or intercept difference to be correct) on the expected difference in observed item scores in three steps:

- (1) Define a meaningful difference on the observed scale, which might be $x_d = x_{2j} - x_{1j}$. Millsap and Olivera-Aguilar did not propose guidelines on defining a meaningful difference in observed item scores but left it open to the applied researcher.
- (2) Calculate the range of common latent scores $\mathbf{w}$ for which $|E(x_{2j} - x_{1j}|\mathbf{w})| \leq x_d$.
- (3) Evaluate how many individuals are in the region $|E(x_{2j} - x_{1j}|\mathbf{w})| > x_d$. Means and variances of $\mathbf{w}$ estimated in each group (i.e., the latent variance Φ_k and the latent mean κ_k) should help this evaluation.

If the model has no loading differences (i.e., if metric invariance holds), the loading difference measure (Equation 4) is not useful anymore because the expected observed score

differences will always equal the raw intercept differences: $E(x_{2j} - x_{1j}|w) = x_{2j} - x_{1j}$.

Thus, Millsap and Olivera-Aguilar developed a specific intercept difference measure.

To define their intercept difference measure, Millsap and Olivera-Aguilar (2012) assumed a unidimensional CFA model showing partial scalar invariance across two groups ($K = 2$). Item j ($j = 1, \dots, p$) has non-invariant intercepts across groups. That is, the intercept τ_{1j} of item j in group 1 differs from the intercept τ_{2j} of item j in group 2, which are, thus, freely estimated across groups. Millsap and Olivera-Aguilar's intercept difference measure represents the proportion of the cross-group difference in item means attributable to the intercept difference:

$$\text{Intercept Difference}_j = \frac{|\tau_{2j} - \tau_{1j}|}{|E(x_{2j}) - E(x_{1j})|}. \quad (5)$$

Here, $E(x_{1j})$ and $E(x_{2j})$ are approximated by the observed means of item j for groups 1 and 2 (alternatively by Equation 3). Thus, the numerator denotes the intercept difference of item j ; the denominator reflects the mean difference of item j . If there were no intercept differences between groups (i.e., the numerator equals zero), the intercept difference would be zero (unless the item mean difference is zero, in which case the ratio is undefined).

Further, Millsap and Olivera-Aguilar (2012) developed a specific uniqueness difference measure to quantify non-invariance bias due to differing unique variances. For the uniqueness difference measure, Millsap and Olivera-Aguilar assumed a unidimensional CFA model showing partial uniqueness invariance across two groups ($K = 2$). Item j ($j = 1, \dots, p$) has non-invariant unique variances across groups. That is, the unique variance θ_{1j} of item j in group 1 differs from the unique variance θ_{2j} of item j in group 2, which are, thus, freely estimated across groups. Millsap and Olivera-Aguilar's uniqueness difference measure represents the proportion of the cross-group difference in item variances attributable to the difference in unique variances:

$$251 \quad \text{Uniqueness Difference}_j = \frac{|\theta_{2j} - \theta_{1j}|}{|\sigma_{2j}^2 - \sigma_{1j}^2|}. \quad (6)$$

252 The item variance σ_{kj}^2 can be approximated by the observed item variance of item j in group k
 253 (alternatively by $\sigma_{kj}^2 = \lambda_j' \Phi_k \lambda_j + \theta_{kj}$ including loading λ_j of item j and latent variance Φ_k in
 254 group k). The numerator reflects the difference in unique variances of item j ; the denominator
 255 reflects the difference in total variances of item j . If there were no differences in unique
 256 variances between groups (i.e., the numerator equals zero), the uniqueness difference
 257 measure would be zero (unless the item variance difference is zero, in which case the ratio is
 258 undefined).

259 Millsap and Olivera-Aguilar's (2012) difference measures help quantify non-
 260 invariance. They are easy to implement in standard software such as Mplus (Muthén &
 261 Muthén, 1998-2017) or R (using the lavaan package; Rosseel, 2012). Still, they have three
 262 crucial problems: First, they do not come with a unified rationale spanning different model
 263 parameters (i.e., factor loadings, intercepts, and unique variances). Although one should not
 264 compare effect sizes for different measurement parameters, an intuitive logic that applies to
 265 all parameters will surely aid the understanding and, in turn, the applicability of a set of effect
 266 size measures. Second, Millsap and Olivera-Aguilar only vaguely defined what constitutes a
 267 meaningful raw-metric difference in observed item scores, given a specific score on the latent
 268 variable, which hampers at least the application of the loading difference measure. Third, the
 269 denominators of Millsap and Olivera-Aguilar's intercept and uniqueness difference measures
 270 (i.e., the difference in item means or item variances) change their units with each item.
 271 Accordingly, researchers can only evaluate the bias item-by-item but not compare it across
 272 items or estimate their combined effects within an aggregated set of items. Effect size
 273 measures that apply to single items and item sets alike would be favorable.

274

Objective

We develop new effect size measures to quantify the degree of non-invariance, which we term Measurement Invariance Violation Indices (MIVIs). MIVIs are derived for non-invariance in three different model parameters (i.e., factor loadings, intercepts, and unique variances) and for each model parameter in four different versions (item/item set $\times$ absolute/signed). Thus, MIVIs build on the difference measures proposed by Millsap and Olivera-Aguilar (2012) but improve on them in three regards: First, MIVIs quantify non-invariance bias with the same rationale (i.e., denominator) for all model parameter differences (in the numerator, i.e., factor loadings, intercepts, and unique variances). This makes MIVIs easy to grasp. Second, MIVIs have identical points of reference (i.e., denominators) comparable across multiple items. Consequently, they are informative about systematic—compensatory or non-compensatory—non-invariance bias. As a third asset, MIVIs consider sampling error by providing bootstrap confidence intervals around the point estimate (Cumming & Finch, 2001). MIVIs are also readily implemented in standard statistical programs such as Mplus or R.

Depending on the specific research stage (e.g., scale development vs. application of an extant and well-defined item set for measuring a construct comparably across groups), MIVIs have various advantages for quantifying the amount of non-invariance bias. In scale development, MIVIs can support the group-fair item selection (i.e., the decision on keeping non-invariant items or dropping them from the item set). In settings with fixed item sets, MIVIs can help evaluate the quality of a questionnaire in a new context (e.g., an international large-scale assessment, say, with new countries onboarding a survey program). MIVIs are

also helpful in evaluating the amount of bias in index scores when using simple proxies such as sum or mean scores.³ We highlighted all advantages of MIVIs in Table 1.

[Table 1 near here]

Measurement Invariance Violation Indices (MIVIs)

Preconsiderations and Prerequisites

Like Millsap and Olivera-Aguilar's (2012) difference measures, MIVIs are estimated based on a tenable (i.e., well-fitting) partial invariance model with at least one freely estimated parameter of interest (i.e., factor loadings, intercepts, or unique variances). We assume models to be unidimensional. In multidimensional models (assuming no cross-loadings), MIVIs should be separately estimated for each common factor (or latent variable). Before defining MIVIs, we discuss the key idea of the new effect size measures.

The key idea of MIVIs is to use the pooled (across groups) standard deviation of the latent variable as a denominator and the unstandardized parameter differences of interest as a numerator. The pooled standard deviation $SD_{LV, \text{pooled}}$ of the latent variable (Cohen, 1988) is defined as

$$SD_{LV, \text{pooled}} = \sqrt{\frac{(n_2-1)\Phi_2 + (n_1-1)\Phi_1}{n_1+n_2-2}}. \quad (7)$$

3. Notably, sum or mean scores require additional assumptions such as tau-equivalent or parallel measurement. However, methodologists still debate about the necessity to which these assumptions must hold before one can use sum or mean scores (McNeish & Wolf, 2020; Widaman & Revelle, 2022).

One may ask why we chose the pooled standard deviation of the latent variable as a common denominator for all MIVIs, instead of the (observed) pooled standard deviation of an item or the item set (e.g., Pornprasertmanit, 2022). The pooled standard deviation of the latent variable as a common denominator is favorable in at least three regards: (1) It is comparable across items, (2) it is independent of the number of items in an item set, and (3) it solely consists of true score variance.

First, a problem with an item's (observed) pooled standard deviation is that it differs across items of an item set. Consequently, not only differences in parameters of interest (i.e., factor loadings, intercepts, or unique variances) but also differences in the pooled standard deviation of the item would impact the effect size. The effect size would not be comparable across items of an item set. By using the pooled standard deviation of the latent variable as a common denominator, we make MIVIs comparable across items of an item set.

Second, a problem with an item set's pooled standard deviation is that it changes dramatically with the number of items in it. (The pooled standard deviation of the item set is estimated by the sum of item variances plus two times all item covariances, e.g., Nye & Drasgow, 2011). Thus, the effect size would drastically depend on the number of items in an item set. Differently, the pooled standard deviation of the latent variable is relatively independent of the number of items in an item set.

Third, another problem with both the pooled standard deviation of the item and the item set is that they compound true score variation with fluctuating error variation. The compound makes it more challenging to establish comparable quantities and heuristics for evaluating the effect size. Differently, the pooled standard deviation of the latent variable is the true score variation.

However, taking the pooled standard deviation of the latent variable as a common denominator for an effect size also has its limitations: As raw group estimates of the latent

variable change with the identification method, the effect size also does. Typically, researchers use the (pooled) standard deviation per item to standardize model parameters (i.e., factor loadings, intercepts, or unique variances; Muthén, 1998-2004; Pornprasertmanit, 2022). Parameter estimates standardized with the (pooled) standard deviation per item are independent of the identification method (e.g., Putnick & Bornstein, 2016). That is, estimates are the same, independent of whether the model is identified by fixing the latent variance and the latent mean or by fixing an anchor item's loading or intercept to a certain value. Because MIVIs include unstandardized (i.e., raw) parameters in the numerator, standardized by the (pooled) standard deviation of the latent variable (not the item), MIVIs change with the identification method. To obtain common ground for MIVIs, we propose identifying the final model by fixing the reference group's latent variance to 1 (and its latent mean to 0). As MIVIs are based on partial invariance models, latent variances in all groups other than the reference one can be freely estimated.

Effect Size Measures for Items

Absolute Effect Size Measures

To define MIVIs for items, we start on the assumption of a tenable partial invariance model with at least one freely estimated factor loading, intercept, or unique standard deviation.⁴ We

4. We chose to include the differences in unique standard deviations instead of the unique variances in the numerator to be consistent with the entity of the denominator (which is also a standard deviation). If preferred to retain the metric of unique variances (as it is in the measurement model), one might use the difference in unique variances divided by the pooled variance of the latent variable as an alternative version of $MIVI - U$.

define MIVIs for an (unstandardized) loading λ_{kj} , intercept τ_{kj} , or unique standard deviation $\sqrt{\theta_{kj}}$ of the non-invariant item j in group k as

$$MIVI-L_{Item\ j|absolute} = \frac{|\lambda_{2j} - \lambda_{1j}|}{SD_{LV, pooled}}, \quad (8a)$$

$$MIVI-I_{Item\ j|absolute} = \frac{|\tau_{2j} - \tau_{1j}|}{SD_{LV, pooled}}, \quad (8b)$$

$$MIVI-U_{Item\ j|absolute} = \frac{|\sqrt{\theta_{2j}} - \sqrt{\theta_{1j}}|}{SD_{LV, pooled}}. \quad (8c)$$

These MIVIs pertain to the item level, are naturally bounded at zero on one end, and reflect an absolute loading, intercept, or uniqueness difference relative to the pooled standard deviation of the latent variable. The logic of MIVI is similar to Cohen's d (1988, 1992) in (pooled) standard deviation units of the latent variable.

To adequately interpret and understand MIVI, one needs to revisit the logic of Equation 3: In a common factor model, common scores on the latent variable and observed item scores relate to each other through linear regression of the observed item scores on the latent scores (e.g., Wu et al., 2007, for a visualization). Whereas loadings represent the steepness (or slope/weight) of regression lines, intercepts represent their origin. Unique scores describe the differences between the scores of the latent variable on the regression line and the observed scores. For a factor model to be invariant across groups, the regression lines that relate the observed units to the latent property should be identical across groups. Further, the variance of the unique scores should be identical across groups (i.e., observed scores have an equal precision across groups; see DeShon, 2004). MIVIs quantify the differences in those regression lines across groups: $MIVI-L_{Item\ j|absolute}$ quantifies the differences in regression weights, $MIVI-I_{Item\ j|absolute}$ the differences in regression origins, and $MIVI-U_{Item\ j|absolute}$ the differences in the standard deviations of unique scores. Ideally for group comparisons, all differences would be zero.

The different meanings of the three parameters (i.e., factor loadings, intercepts, and unique variances represent different features of a regression) have two important implications. First, an intercept and a unique variance should only be fixed across groups if the corresponding loading can be fixed without inducing misfit. If a loading is non-invariant across groups, the item's intercept (as well as its unique variance) should not undergo invariance tests. They, too, should be estimated freely across groups. Thus, metric non-invariance of an item implies scalar and uniqueness non-invariance of that item. If the units (i.e., loadings) differ across groups, unit shifts (i.e., intercepts) and unique variances are barely comparable (theoretically and, often, empirically; e.g., Millsap & Olivera-Aguilar, 2012). Then, one stops the quantification of item non-invariance at the level of loadings (i.e., $MIVI-L$)—although it might be helpful to obtain a rough estimate of intercept or uniqueness non-invariance too (i.e., $MIVI-I$ or $MIVI-U$) in some cases. Because the expected mean of the unique scores is zero (see Equation 3), scalar invariance is not a prerequisite for testing uniqueness invariance (Steinmetz, 2013; cf. DeShon, 2004; Wu et al., 2007).

Second, one cannot compare the three types of MIVI (i.e., $MIVI-L$, $MIVI-I$, $MIVI-U$) to each other and conclude, for instance, that $MIVI-L_{Item\ j|absolute}$ were larger than $MIVI-I_{Item\ j|absolute}$. Loadings, intercepts, and uniquenesses under multivariate normality are entirely different parameters, and so are the associated effect sizes. However (and more importantly), the different types of MIVI can be compared across items. Thus, one can conclude, for instance, that $MIVI-L_{Item\ j|absolute}$ were larger than $MIVI-L_{Item\ j+1|absolute}$.

Signed Effect Size Measures

Violations of measurement invariance in loadings, intercepts, or uniquenesses across groups manifest as interaction effects of the respective parameters and group membership (e.g., Bauer, 2017), quantified through the MIVIs. Such an interaction term can favor one or the

other group (i.e., group 1 or group 2 has larger non-invariant parameters). So far, $MIVI_{Item\ j|absolute}$ always has a positive sign; it cannot disentangle scenarios where group 1 has a larger parameter than group 2 from those where group 2 has a larger parameter than group 1. Without loss of generality, we can define a signed variant, $MIVI_{Item\ j|signed}$, that traces the mathematical signs of the non-invariant loading, intercept, or unique standard deviation differences back to the groups,

$$MIVI-L_{Item\ j|signed} = \frac{\lambda_{2j} - \lambda_{1j}}{SD_{LV, pooled}}, \quad (9a)$$

$$MIVI-I_{Item\ j|signed} = \frac{\tau_{2j} - \tau_{1j}}{SD_{LV, pooled}}, \quad (9b)$$

$$MIVI-U_{Item\ j|signed} = \frac{\sqrt{\theta_{2j}} - \sqrt{\theta_{1j}}}{SD_{LV, pooled}}. \quad (9c)$$

Different from $MIVI_{Item\ j|absolute}$, the $MIVI_{Item\ j|signed}$ estimates are unbounded.

Item-level MIVIs allow quantifying the size of parameter differences for a single item. Knowing the size of loading, intercept, or uniqueness differences for a single item is highly informative in the scale development process when researchers face the question of retaining or discarding items in an item set. However, in applications with fixed item sets, one usually wants to know whether non-invariance substantially impacts group comparisons, especially when the goal is to use summary statistics such as weighted or unweighted mean scores. Next, we present MIVI variants that consider multiple items in tandem and can answer this question.

Effect Size Measures for Item Sets

Absolute Effect Size Measures

Let us extend the logic of MIVIs for single items to the aggregate level, that is, calculating average MIVIs by considering *all* items of an item set simultaneously. Evaluating the bias by average MIVIs across all items is useful when the goal is to use summative indices such as

manifest mean scores, or when the goal is to use MIVIs as a rough, overall indicator of non-invariance bias at the level of item sets. As with item-level MIVIs, we define absolute and signed versions.

Suppose we have a partial metric invariance model if we want to quantify loading differences, a partial scalar invariance model if we want to quantify intercept differences, and a partial uniqueness invariance model if we want to quantify differences in unique standard deviations. For absolute MIVIs at the level of item sets, the numerators summarize all biases of single items; once again, they do so in terms of absolute parameter differences (i.e., invariant parameters do not contribute to aggregate MIVIs). Again, the denominators contain the pooled standard deviation of the latent variable. Dividing the fraction by the number of items serves as the average amount of bias introduced per item:

$$MIVI-L_{Item\ set|absolute} = \frac{|\lambda_{21} - \lambda_{11}| + |\lambda_{22} - \lambda_{12}| + \dots + |\lambda_{2p} - \lambda_{1p}|}{SD_{LV, \text{ pooled}}} / p = \frac{\sum_{j=1}^p |\lambda_{2j} - \lambda_{1j}|}{SD_{LV, \text{ pooled}}} / p, \quad (10a)$$

$$MIVI-I_{Item\ set|absolute} = \frac{|\tau_{21} - \tau_{11}| + |\tau_{22} - \tau_{12}| + \dots + |\tau_{2p} - \tau_{1p}|}{SD_{LV, \text{ pooled}}} / p = \frac{\sum_{j=1}^p |\tau_{2j} - \tau_{1j}|}{SD_{LV, \text{ pooled}}} / p, \quad (10b)$$

$$MIVI-U_{Item\ set|absolute} = \frac{|\sqrt{\theta_{21}} - \sqrt{\theta_{11}}| + |\sqrt{\theta_{22}} - \sqrt{\theta_{12}}| + \dots + |\sqrt{\theta_{2p}} - \sqrt{\theta_{1p}}|}{SD_{LV, \text{ pooled}}} / p = \frac{\sum_{j=1}^p |\sqrt{\theta_{2j}} - \sqrt{\theta_{1j}}|}{SD_{LV, \text{ pooled}}} / p. \quad (10c)$$

$MIVI_{Item\ set|absolute}$ indicates the average between-group differences in factor loadings, intercepts, or unique standard deviations for a set of items in units of the pooled standard deviation of the latent variable. MIVIs quantify the average amount of bias that arises from all non-comparable factor loadings, intercepts, or unique standard deviations in the latent representation of the item set, and it does so by using the same Cohen's d -like metric as all MIVI versions do.

Given the absolute values of constituents in both MIVI-numerators, $MIVI_{Item\ set|absolute}$ is restricted to values ≥ 0 . $MIVI_{Item\ set|absolute}$ gives an overall, absolute impression and quantifies *how biased* an item set is on average. Crucially, by

dividing the fraction by the number of items, $MIVI_{Item\ set}$ takes the length of an item set into account. Thereby, it considers that, for instance, two non-invariant items are more problematic in an item set of four rather than eight items.

From an applied perspective, bias due to non-invariant parameters may cancel out (at least in part). Thus, despite biased items in an item set, the overall between-group differences in the total variance or total mean score of the item set may remain unbiased (for a formula-based definition, see Nye & Drasgow, 2011). This is reminiscent of Chen (2007; see also Robitzsch & Lüdtke, 2020, for the context of item response theory), who distinguished between two patterns of non-invariance: uniform (i.e., one group has higher parameter values, such as intercepts, on all non-invariant items) and mixed (i.e., one group has higher parameter values on some non-invariant items but lower parameter values on other non-invariant items).

Crucially, any such aggregation or compensation effects may be evaluated for intercepts (Equation 3) or unique variances (Equation 2). Factor loadings work as amplifiers; they result in multiplicative effects (Equations 2 and 3). They multiply with the latent variances/means when producing manifest variances/means of the items (or item sets). Thus, factor loadings ultimately connect to the variances and means of the latent variable, which can differ across groups in partial invariance models. Compensation across groups is, thus, not immediately conceivable for loading differences.

Signed Effect Size Measures

A variant of $MIVI_{Item\ set|absolute}$ with each summand keeping its direction (i.e., sign), $MIVI_{Item\ set|signed}$, allows for complete (or incomplete) compensation of parameter differences of intercepts and unique standard deviations,

$$MIVI-L_{Item\ set|signed} = not\ applicable, \quad (11a)$$

$$MIVI-I_{Item\ set|signed} = \frac{(\tau_{21} - \tau_{11}) + (\tau_{22} - \tau_{12}) + \dots + (\tau_{2p} - \tau_{1p})}{SD_{LV, pooled}} / p = \frac{\sum_{j=1}^p (\tau_{2j} - \tau_{1j})}{SD_{LV, pooled}} / p, \quad (11b)$$

$$MIVI-U_{Item\ set|signed} = \frac{\sqrt{\theta_{21}} - \sqrt{\theta_{11}} + \sqrt{\theta_{22}} - \sqrt{\theta_{12}} + \dots + \sqrt{\theta_{2p}} - \sqrt{\theta_{1p}}}{SD_{LV, pooled}} / p = \frac{\sum_{j=1}^p \sqrt{\theta_{2j}} - \sqrt{\theta_{1j}}}{SD_{LV, pooled}} / p. \quad (11c)$$

Given signed differences as constituents in the numerator, estimates of $MIVI_{Item\ set|signed}$ are unbounded.

Although MIVIs might show that non-invariance bias is negligible or cancels out completely, one should always be aware that the invariance assumption is violated to a certain degree in the item set. MIVIs are no legitimation to ignore non-invariance bias but rather a tool to assess its size and impact, depending on the research focus.

How Much of a Non-Invariant Parameter Difference is Too Large?

We have not yet outlined which parameter differences in the pooled standard deviation units can be considered either negligible or substantial (especially when using simplified statistics such as mean scores). Although there can only be rough rules about acceptable MIVI magnitudes, some initial guidance for applied settings might be helpful. Our goal is not to prescribe a single fixed rule. Many components, such as the substantive research question, type of analysis, empirical context, and, especially, the parameter of interest (i.e., factor loading, intercept, or uniqueness), may influence what is considered a critical value (e.g., Steinmetz, 2013).

Although Cohen's (1988, 1992) effect size guidelines have recently come under scrutiny for not being applicable in all research contexts (e.g., Schäfer & Schwarz, 2019), researchers are well-acquainted with his cutoffs and their meanings. For his d , a difference measure between independent means relative to the pooled standard deviation, he suggested that $|d| = 0.2$ may indicate a substantial but small, $|d| = 0.5$ a medium, and $|d| = 0.8$ a large effect. Arguably, until more evidence has been gathered, these values may serve as initial thresholds for MIVI (especially for $MIVI - I$ quantifying intercept differences), which, when surpassed, alert researchers to become cautious about biased group comparisons in their

substantive analyses. We reiterate that this guideline is rough and empirical data must back it up in future scenarios for the different parameters of interest. Importantly, setting up a guideline must be subject to substantive considerations.

Applications

Example Setup

We simulated data to exemplify the applicability of the MIVI versions. The advantage of simulations is that we can specify beforehand all model, data, and analysis characteristics, such as the magnitude of the model parameters and, accordingly, the magnitude of bias due to non-invariance. We conducted all analyses with Mplus (version 8.4; Muthén & Muthén, 1998-2017). We estimate bootstrap confidence intervals around the point estimate of all MIVIs. We supply the Mplus code for the Simulated Examples 1 and 2 in the Supplementary Materials (see Additional Files 2 and 3). Additional File 4 includes R code for a MIVI application similar to Example 1 (i.e., it uses the same population models for data generation). Both examples include non-invariant intercepts; all factor loadings are invariant. Thus, full metric and partial scalar invariance hold. In empirical reality, this is not always the case. As we also derived MIVIs to quantify loading non-invariance, MIVIs can also be applied when only partial metric invariance holds. Further, MIVIs can also be applied when partial uniqueness invariance holds, as we also derived MIVIs for uniqueness invariance. The logic of applying MIVIs in different empirical settings is analogous to the exemplary one.

Simulated Example 1

To demonstrate the utility of the MIVI versions, we simulated data for a hypothetical four-item set that conforms to a partial scalar invariance model across two groups. We simulated data for four continuous indicator variables underlying a single latent factor. We set the number of observations to 1,000 per group. The intercepts of Items 1 and 2 in Group A

equaled those in Group B. The intercepts of Items 3 and 4 differed across groups. We fit the partial scalar invariance model with the maximum likelihood (ML) estimator. Figure 1 shows the predetermined coefficients from the population model. Further, it shows the empirical coefficients resulting from the partial scalar analysis model. Apart from sampling variation, the empirical coefficients mirrored the predetermined ones. We identified the partial scalar analysis model by fixing the latent variance in Group A at 1 and the latent mean in Group A at 0. The latent variance and latent mean in Group B are freely estimated. Table 2 includes the resulting estimates of the MIVI versions (item- or item-set-level $\times$ absolute or signed) and accompanying bootstrap confidence intervals.

[Figure 1 near here]

[Table 2 near here]

We first inspected the (relative) size of the bias due to non-invariant intercepts by looking at the absolute MIVIs on the item level. As intended, bias was more substantial in Item 3 than in Item 4. Whereas Item 3 had a non-invariant intercept difference of 0.54 pooled standard deviations of the latent variable, Item 4 had a non-invariant intercept difference of 0.46 pooled standard deviations. Both correspond to a substantial medium effect, following Cohen's guidelines. The signed MIVIs on the item level were equal to the absolute ones, which indicates that intercept bias is larger in Group B than in Group A for both items.

As a next step, we took a closer look at the size and shape of the total non-invariance bias of the complete item set. The absolute bias due to all non-invariant intercept differences adding up in the item set was non-negligible according to $MIVI-I_{Item\ set|absolute}$: The average intercept difference was as large as 0.25 pooled standard deviations per item of the

item set. The signed version of $MIVI-I_{Item\ set}$ equaled the absolute one and, thus, showed that intercept differences accumulated at the item set rather than compensating each other.

Considering all MIVI values, we would not consider the bias due to non-invariant intercepts negligible. Especially the absolute and signed versions of $MIVI-I_{Item\ set}$ clearly demonstrated that non-invariance bias added up and substantially contaminated the item set in total (and accordingly statistics of the item set such as mean scores).

Simulated Example 2

The following example illustrates the applicability of MIVIs in scenarios with compensatory non-invariance bias. We simulated data in the same way as in Example 1. The crucial difference was the opposite direction of intercept differences across Items 3 and 4. Figure 2 shows the predetermined coefficients from the population model and the estimated coefficients from the partial scalar analysis model. Table 3 includes the resulting estimates of the MIVI versions and accompanying bootstrap confidence intervals.

[Figure 2 near here]

[Table 3 near here]

Again, the absolute non-invariance bias of each item was relatively substantial, with the intercept difference of Item 3 ($MIVI-I_{Item\ 3|absolute} = 0.58$) being larger than that of Item 4 ($MIVI-I_{Item\ 4|absolute} = 0.46$). However, $MIVI_{Item\ 3|signed}$ had a negative sign for intercept non-invariance bias ($MIVI-I_{Item\ 3|signed} = -0.58$), whereas $MIVI_{Item\ 4|signed}$ had a positive sign ($MIVI-I_{Item\ 4|signed} = 0.46$). The non-invariance bias of Items 3 and 4 operated differently: Item 3 has larger intercepts in Group B, Item 4 has larger intercepts in Group A.

To further investigate the mutual compensation of intercept differences, we considered the average non-invariance bias of the whole item set. The absolute biasing effect was non-negligible for intercepts ($MIVI - I_{Item\ set|absolute}$) and resulted in an average difference of 0.26 pooled standard deviations per item. Given this non-negligible effect size at first inspection, one might refrain from treating the item set as invariant. However, calculating $MIVI_{Item\ set|signed}$ provides further information on possible compensatory differences (i.e., non-invariant intercepts canceling each other out), which are relevant for valid cross-group comparisons of total mean scores (when evaluating non-invariance of intercepts). $MIVI_{Item\ set|signed}$ showed that the two intercept differences accounted for an average non-invariance bias per item of $|0.03|$ latent standard deviations ($MIVI - I_{Item\ set|signed} = -0.03$). The confidence interval was compatible even with the notion of the total absence of any biasing effect ($CI_{95\%} = [-0.07; 0.01]$). The non-negligible amount of (absolute) bias present in all items was negligible when considering the direction of item-specific bias.

The relatively small values of the signed MIVI at the item set level suggested that non-invariant intercept differences compensated for each other. Some researchers may conclude that the item set is practically invariant in this example to justify the comparison of group statistics, especially if one favors treating the items as a fixed item set without the possibility to further select items before running substantive analyses. In different words, applied researchers might conclude from the analysis of all MIVIs (and especially $MIVI_{Item\ set|signed}$) that the non-invariant intercepts do not substantially impact the difference in observed mean scores of the item set.

Discussion and Conclusion

The Usefulness of Measurement Invariance Violation Indices

Research questions involving numerical between-group comparisons of constructs are ubiquitous in psychological research and beyond. However, for group comparisons to be valid, researchers must first establish the measurement invariance of the constructs in question. Unfortunately, as any experienced analyst can testify, measurement non-invariance is a frequently encountered problem. The question then becomes: How wrong are assumptions of invariance? Stated differently: How much non-invariance bias is present?

This paper provided a novel set of tools to answer these essential questions. Specifically, we proposed an effect size approach called MIVI, which quantifies the non-invariance of factor loadings, intercepts, and uniquenesses in a unified and principled way. MIVIs are handy effect size measures that overcome the limitations of existing approaches. Specifically, MIVIs provide a common metric for different parameters (i.e., factor loadings, intercepts, and uniquenesses) that is comparable across all items of an item set. At the item level, MIVIs relate each item's cross-group difference in factor loadings, intercepts, or uniquenesses to the pooled standard deviation of the latent variable. At the item set level, MIVIs reflect the average cross-group differences in factor loadings, intercepts, or uniquenesses in units of the pooled latent standard deviation across all items. Deriving absolute and directional MIVI versions, we can inspect the cancelation of non-invariance bias. Confidence intervals display the accuracy of the MIVI estimates. MIVI estimates and confidence intervals can easily be computed by adding a few code lines for standard SEM software such as Mplus or R.

MIVIs are especially helpful for scale development, as well as in settings with fixed item sets: In scale development, MIVIs can support the decision on keeping non-invariant items or dropping them from the item set—depending on the size of the specific parameter

non-invariance MIVIs identify on the level of single items and item sets. By quantifying the non-invariance bias in a fixed item set, MIVIs can help evaluate the quality of a questionnaire, for instance, in a new context. Further, MIVIs allow to evaluate the amount of bias in index scores when using simple proxies such as sum or mean scores.

Preconditions of Measurement Invariance Violation Indices

Without loss of generality, we have sketched MIVIs as useful for quantifying non-invariance bias when comparing two groups. Researchers can legitimately apply MIVI variants to scenarios with more than two groups: They may pick a reference group and compare all non-invariant parameters against this reference group, yielding multiple reference-group comparisons.

MIVIs also rest on other assumptions or preconditions that must be fulfilled. The analysis model should be unidimensional (without cross-loadings) to employ MIVIs correctly. If analysis models are multidimensional (assuming no cross-loadings), separate MIVIs should be used for each dimension. Alternatively, one might estimate how proximate the multidimensional model is to a unidimensional one, indicating whether the multidimensional model can be treated as essentially unidimensional (Raykov & Bluemke, 2020).

Further, the multi-group measurement model, which is the basis for all MIVI versions, must be correctly identified. Multi-group measurement models are commonly identified by constraining the latent variance or a factor loading to a fixed value (commonly to 1). Additionally, to identify the mean structure, either the latent mean or an intercept must be constrained to a fixed value (commonly to 0). The crucial assumption for measurement invariance tests to be valid is that those fixed parameters are equal (or invariant) across groups (Putnick & Bornstein, 2016). If the fixed parameters are not invariant, measurement invariance tests might be misleading. Whereas one can only assume that latent variances and

latent means are equal across groups, several approaches guide identifying so-called reference indicators or anchor items (e.g., Kopf et al., 2015, Schulze & Pohl, 2021; Thompson et al., 2021). Raykov et al.'s (2013) approach to measurement invariance, which we outlined in Additional File 1 of the Supplementary Information, does not require identifying a reference item (as it constrains multiple items at one time). It is specifically suitable for detecting non-invariant intercepts (Thompson et al., 2021).

To properly quantify bias due to non-invariant parameters (i.e., factor loadings, intercepts, or uniquenesses) and, thus, to estimate MIVIs, one needs a common, congeneric factor model (i.e., configural invariance) as the starting point. Otherwise, one cannot compare parameters at all. Furthermore, partial metric invariance must hold. Otherwise, one cannot compare and pool latent variances for the computation of the MIVI denominators.

For MIVIs to be interpretable for intercepts and unique variances, the non-invariant items should also have equal loadings. If the units (i.e., loadings) differ across groups, unit shifts (i.e., intercepts) or variances of unique scores are barely comparable (theoretically and, often, empirically; e.g., Millsap & Olivera-Aguilar, 2012). Then, MIVI should only be estimated for loadings—although a rough estimate of intercept non-invariance (i.e., $MIVI-I$) or uniqueness non-invariance (i.e., $MIVI-U$) might be helpful in some cases.

Further, MIVIs require unstandardized parameters (i.e., factor loadings, intercepts, and uniquenesses). Standardized parameters cannot quantify parameter-specific non-invariance because those parameters incorporate unique variances (which commonly differ across items, e.g., Dong & Dumas, 2020). MIVIs quantifying parameter-specific non-invariance are only informative (or pure) with unstandardized parameters.

As MIVIs rely on unstandardized parameters, the method of model identification influences the MIVI estimates. Either constraining the latent variance or a factor loading to a fixed value, commonly to 1, identifies the (co)variance structure of the model. Either the

latent mean or an intercept constrained to a fixed value, commonly to 0, identifies the mean structure of the model. The way to identify the model influences the estimates of model parameters (such as cross-group estimates of latent variances, factor loadings, or intercepts) and, in turn, MIVI estimates. To have a common ground for MIVIs, we propose identifying the final partial invariance model by fixing the latent variance to 1 and the latent mean to 0. As MIVIs are estimated from partial invariance models, the latent variance (and the latent mean if at least partial scalar invariance holds) must only be fixed in a reference group. Still, it can be freely estimated in the other group(s).

Limitations and Future Directions

Although MIVIs might show that non-invariance is non-substantial or cancels out at the level of the item set, one should always be aware that non-invariance is present. One should not use MIVI as a legitimization to completely ignore non-invariance but rather as a tool to quantify non-invariance or assess its impact on simple statistics (e.g., mean scores). Theoretical considerations and empirical models, such as multiple-indicators multiple-cause models (or even more flexible variations of it, e.g., Bauer, 2017), can help explain the non-invariance and should still be reflected in addition to applying MIVIs.

Specific heuristics help evaluate whether an effect is small, medium, or large (e.g., Cohen, 1988, 1992). In this paper, we tentatively recommended Cohen's effect size guidelines as initial thresholds to evaluate MIVI values. However, Cohen's guidelines are no more than starting points for MIVI guidelines. Future research should evaluate MIVI values under different conditions to derive guidelines on what constitutes small, medium, or large effects for the different MIVIs of interest. Unless enough empirical evidence is gathered for a meta-analysis of MIVI values in empirical data (e.g., Gignac & Szodorai, 2016), simulation studies evaluating the size of MIVIs at different strengths of non-invariance may aid in deriving guidelines (e.g., Nye et al., 2019).

Conclusion

MIVIs are easily applicable tools for quantifying measurement non-invariance in a comparable and standardized manner. They are available by adding only a few lines of code in standard statistical programs. MIVIs quantify non-invariance bias in partial metric, scalar, or uniqueness invariance. They help overcome the practice of merely accepting or rejecting an invariance model in a binary accept-or-reject logic. Their advantages are manifold: MIVIs can support the group-fair item selection by quantifying the item non-invariance during the scale development process. In settings with fixed item sets, MIVIs can help evaluate the quality of a questionnaire in a new context or evaluate the amount of non-invariance bias in simple index scores. Taken together, MIVIs advance researchers' focus and alert to the practical consequences of measurement non-invariance.

References

- Bauer, D. J. (2017). A more general model for testing measurement invariance and differential item functioning. *Psychological Methods*, 22(3), 507–526.
<http://doi.org/10.1037/met0000077>
- Boer, D., Hanke, K., & He, J. (2018). On detecting systematic measurement error in cross-cultural research: A review and critical reflection on equivalence and invariance tests. *Journal of Cross-Cultural Psychology*, 49(5), 713–734.
<https://doi.org/10.1177/0022022117749042>
- Byrne, B. M., Shavelson, R. J., & Muthén, B. (1989). Testing for the equivalence of factor covariance and mean structures: The issue of partial measurement invariance. *Psychological Bulletin*, 105(3), 456–466. <https://doi.org/10.1037/0033-2909.105.3.456>
- Chen, F. F. (2007). Sensitivity of goodness of fit indexes to lack of measurement invariance. *Structural Equation Modeling*, 14(3), 464–504.
<https://doi.org/10.1080/10705510701301834>

- 722 Chen, F. F. (2008). What happens if we compare chopsticks with forks? The impact of
723 making inappropriate comparisons in cross-cultural research. *Journal of Personality*
724 *and Social Psychology*, 95(5), 1005–1018. <https://doi.org/10.1037/a0013193>
- 725 Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). L. Erlbaum
726 Associates.
- 727 Cohen, J. (1992). A power primer. *Psychological Bulletin*, 112(1), 155–159.
728 <https://doi.org/10.1037/0033-2909.112.1.155>
- 729 Cumming, G., & Finch, S. (2001). A primer on the understanding, use, and calculation of
730 confidence intervals that are based on central and noncentral distributions. *Educational*
731 *and Psychological Measurement*, 61(4), 532–574.
732 <https://doi.org/10.1177/0013164401614002>
- 733 Davidov, E., Dülmer, H., Schlüter, E., Schmidt, P., & Meuleman, B. (2012). Using a
734 multilevel structural equation modeling approach to explain cross-cultural measurement
735 noninvariance. *Journal of Cross-Cultural Psychology*, 43(4), 558–575.
736 <https://doi.org/10.1177/0022022112438397>
- 737 Davidov, E., Meuleman, B., Cieciuch, J., Schmidt, P., & Billiet, J. (2014). Measurement
738 equivalence in cross-national research. *Annual Review of Sociology*, 40(1), 55–75.
739 <https://doi.org/10.1146/annurev-soc-071913-043137>
- 740 DeShon, R. P. (2004). Measures are not invariant across groups without error variance
741 homogeneity. *Psychology Science*, 46(1), 137–149.
- 742 Dong, Y., & Dumas, D. (2020). Are personality measures valid for different populations? A
743 systematic review of measurement invariance across cultures, gender, and age.
744 *Personality and Individual Differences*, 160, Article 109956.
745 <https://doi.org/10.1016/j.paid.2020.109956>

- 746 Gignac, G. E., & Szodorai, E. T. (2016). Effect size guidelines for individual differences
747 researchers. *Personality and Individual Differences*, 102, 74–78.
748 <https://doi.org/10.1016/j.paid.2016.06.069>
- 749 Gunn, H. J., Grimm, K. J., & Edwards, M. C. (2020). Evaluation of six effect sizes of
750 measurement non-invariance for continuous outcomes. *Structural Equation Modeling:
751 A Multidisciplinary Journal*, 27(4), 503–514.
752 <https://doi.org/10.1080/10705511.2019.1689507>
- 753 Kopf, J., Zeileis, A., & Strobl, C. (2015). Anchor selection strategies for DIF analysis:
754 Review, assessment, and new approaches. *Educational and Psychological
755 Measurement*, 75(1), 22–56. <https://doi.org/10.1177/0013164414529792>
- 756 McNeish, D., & Wolf, M. G. (2020). Thinking twice about sum scores. *Behavior Research
757 Methods*, 52(6), 2287–2305. <https://doi.org/10.3758/s13428-020-01398-0>
- 758 Meredith, W. (1993). Measurement invariance, factor analysis and factorial
759 invariance. *Psychometrika*, 58(4), 525–543. <https://doi.org/10.1007/BF02294825>
- 760 Millsap, R. E. (2011). *Statistical approaches to measurement invariance*. Routledge.
- 761 Millsap, R. E., & Olivera-Aguilar, M. (2012). Investigating measurement invariance using
762 confirmatory factor analysis. In R. Hoyle, D. Kaplan, G. A. Marcoulides, & S. West
763 (Eds.), *Handbook of Structural Equation Modeling* (pp. 380–392). Guilford Press.
- 764 Muthén, B.O. (1998-2014). *Mplus technical appendices*. Muthén & Muthén.
- 765 Muthén, L.K., & Muthén, B.O. (1998-2017). *Mplus User's Guide* (8th ed.). Muthén &
766 Muthén.
- 767 Nye, C. D., & Drasgow, F. (2011). Effect size indices for analyses of measurement
768 equivalence: Understanding the practical importance of differences between
769 groups. *Journal of Applied Psychology*, 96(5), 966–980.
770 <https://doi.org/10.1037/a0022955>

- 771 Nye, C. D., Bradburn, J., Olenick, J., Bialko, C., & Drasgow, F. (2019). How big are my
772 effects? Examining the magnitude of effect sizes in studies of measurement
773 equivalence. *Organizational Research Methods*, 22(3), 678–709.
774 <https://doi.org/10.1177/1094428118761122>
- 775 Oberski, D. L. (2014). Evaluating sensitivity of parameters of interest to measurement
776 invariance in latent variable models. *Political Analysis*, 22(1), 45–60.
777 <https://doi.org/10.1093/pan/mpt014>
- 778 Osterlind, S. J., & Everson, H. T. (2009). *Differential item functioning* (2nd ed.). Sage
779 Publications.
- 780 Pornprasertmanit, S. (2022). *A note on effect size for measurement invariance*. [https://cran.r-](https://cran.r-project.org/web/packages/semTools/vignettes/partialInvariance.pdf)
781 [project.org/web/packages/semTools/vignettes/partialInvariance.pdf](https://cran.r-project.org/web/packages/semTools/vignettes/partialInvariance.pdf)
- 782 Putnick, D. L., & Bornstein, M. H. (2016). Measurement invariance conventions and
783 reporting: The state of the art and future directions for psychological
784 research. *Developmental Review*, 41, 71–90. <https://doi.org/10.1016/j.dr.2016.06.004>
- 785 R Core Team (2021). *R: A language and environment for statistical computing*. R Foundation
786 for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>
- 787 Raykov, T., & Bluemke, M. (2020). Examining multidimensional measuring instruments for
788 proximity to unidimensional structure using latent variable modeling. *Educational and*
789 *Psychological Measurement*, 81(2), 319–339.
790 <https://doi.org/10.1177/0013164420940764>
- 791 Raykov, T., Marcoulides, G. A., & Millsap, R. E. (2013). Factorial invariance in multiple
792 populations: A multiple testing procedure. *Educational and Psychological*
793 *Measurement*, 73(4), 713–727. <https://doi.org/10.1177/0013164412451978>
- 794 Robitzsch, A., & Lüdtke, O. (2020). A review of different scaling approaches under full
795 invariance, partial invariance, and noninvariance for cross-sectional country

- 796 comparisons in large-scale assessments. *Psychological Test and Assessment*
797 *Modeling*, 62(2), 233–279.
- 798 Rosseel, Y. (2012). lavaan: An R Package for Structural Equation Modeling. *Journal of*
799 *Statistical Software*, 48(2), 1–36. <http://www.jstatsoft.org/v48/i02/>
- 800 Schäfer, T., & Schwarz, M. A. (2019). The meaningfulness of effect sizes in psychological
801 research: Differences between sub-disciplines and the impact of potential biases.
802 *Frontiers in Psychology*, 10, Article 813. <https://doi.org/10.3389/fpsyg.2019.00813>
- 803 Schulze, D., & Pohl, S. (2021). Finding clusters of measurement invariant items for
804 continuous covariates. *Structural Equation Modeling*, 28(2), 219–228.
805 <https://doi.org/10.1080/10705511.2020.1771186>
- 806 Steenkamp, J. B. E., & Baumgartner, H. (1998). Assessing measurement invariance in cross-
807 national consumer research. *Journal of Consumer Research*, 25(1), 78–90.
808 <https://doi.org/10.1086/209528>
- 809 Steinmetz, H. (2013). Analyzing observed composite differences across groups: Is partial
810 measurement invariance enough? *Methodology*, 9(1), 1–12.
811 <https://doi.org/10.1027/1614-2241/a000049>
- 812 Thompson, Y. T., Song, H., Shi, D., & Liu, Z. (2021). It matters: Reference indicator
813 selection in measurement invariance tests. *Educational and Psychological*
814 *Measurement*, 81(1), 5–38. <https://doi.org/10.1177/0013164420926565>
- 815 Widaman, K.F., & Revelle, W. (2022). Thinking thrice about sum scores, and then some
816 more about measurement and analysis. *Behavior Research Methods*. Advance online
817 publication. <https://doi.org/10.3758/s13428-022-01849-w>
- 818 Wu, A. D., Li, Z., & Zumbo, B. D. (2007). Decoding the meaning of factorial invariance and
819 updating the practice of multi-group confirmatory factor analysis: A demonstration
820 with TIMSS data. *Practical Assessment, Research, and Evaluation*, 12, Article 3.

- 821 Zercher, F., Schmidt, P., Cieciuch, J., & Davidov, E. (2015). The comparability of the
822 universalism value over time and across countries in the European Social Survey: Exact
823 vs. approximate measurement invariance. *Frontiers in Psychology*, 6, Article 733.
824 <https://doi.org/10.3389/fpsyg.2015.00733>

825 **Supplementary Materials**

- 826 Additional File 1: How to Test for Measurement Invariance (PDF)
827 Additional File 2: Mplus Code for the Simulated Example 1 (PDF)
828 Additional File 3: Mplus Code for the Simulated Example 2 (PDF)
829 Additional File 4: R Code for the Simulated Example 1 (PDF)

830 **Tables**

831 Table 1. Advantages of MIVIs.

		Absolute	Signed
Item level		Scale development: + MIVIs support group-fair item selection (decision on keeping or dropping non-invariant items)	
	Item set level	Settings with fixed item set: + MIVIs evaluate the quality of a questionnaire	Settings with fixed item set: + MIVIs evaluate the amount of bias in index scores when using simple proxies such as sum or mean scores

832

833 Table 2. MIVIs for Item and Item Set Levels (Example 1).

		Absolute	Signed
Item level	Item 3	$MIVI - I_{Item\ 3 absolute}$ (Eq. 8b) $= \frac{ 0.95 - 0.47 }{\sqrt{(0.60 + 1.00) / 2}}$ $= 0.54 \quad CI_{95\%}[0.45; 0.63]$	$MIVI - I_{Item\ 3 signed}$ (Eq. 9b) $= \frac{0.95 - 0.47}{\sqrt{(0.60 + 1.00) / 2}}$ $= 0.54 \quad CI_{95\%}[0.45; 0.63]$
	Item 4	$MIVI - I_{Item\ 4 absolute}$ (Eq. 8b) $= \frac{ 0.87 - 0.47 }{\sqrt{(0.60 + 1.00) / 2}}$ $= 0.46 \quad CI_{95\%}[0.37; 0.54]$	$MIVI - I_{Item\ 4 signed}$ (Eq. 9b) $= \frac{0.87 - 0.47}{\sqrt{(0.60 + 1.00) / 2}}$ $= 0.46 \quad CI_{95\%}[0.37; 0.54]$
Item set level		$MIVI - I_{Item\ set absolute}$ (Eq. 10b) $= \frac{ 0.95 - 0.47 + 0.87 - 0.47 }{\sqrt{(0.60 + 1.00) / 2}} / 4$ $= 0.25 \quad CI_{95\%}[0.21; 0.29]$	$MIVI - I_{Item\ set signed}$ (Eq. 11b) $= \frac{0.95 - 0.47 + 0.87 - 0.47}{\sqrt{(0.60 + 1.00) / 2}} / 4$ $= 0.25 \quad CI_{95\%}[0.21; 0.29]$

834

835 Table 3. MIVIs for Item and Item Set Levels (Example 2)

		Absolute	Signed
Item level	Item 3	$MIVI - I_{Item\ 3 absolute}$ (Eq. 8b) $= \frac{ 0.45 - 0.97 }{\sqrt{(0.60 + 1.00) / 2}}$ $= 0.58 \quad CI_{95\%}[0.49; 0.67]$	$MIVI - I_{Item\ 3 signed}$ (Eq. 9b) $= \frac{0.45 - 0.97}{\sqrt{(0.60 + 1.00) / 2}}$ $= -0.58 \quad CI_{95\%}[-0.67; -0.49]$
	Item 4	$MIVI - I_{Item\ 4 absolute}$ (Eq. 8b) $= \frac{ 0.88 - 0.47 }{\sqrt{(0.60 + 1.00) / 2}}$ $= 0.46 \quad CI_{95\%}[0.37; 0.54]$	$MIVI - I_{Item\ 4 signed}$ (Eq. 9b) $= \frac{0.88 - 0.47}{\sqrt{(0.60 + 1.00) / 2}}$ $= 0.46 \quad CI_{95\%}[0.37; 0.54]$
Item set level		$MIVI - I_{Item\ set absolute}$ (Eq. 10b) $= \frac{ 0.45 - 0.97 + 0.88 - 0.47 }{\sqrt{(0.60 + 1.00) / 2}} / 4$ $= 0.26 \quad CI_{95\%}[0.24; 0.28]$	$MIVI - I_{Item\ set signed}$ (Eq. 11b) $= \frac{0.45 - 0.97 + 0.88 - 0.47}{\sqrt{(0.60 + 1.00) / 2}} / 4$ $= -0.03 \quad CI_{95\%}[-0.07; 0.01]$

836

837

Figures

838 Figure 1. Population and Analysis Models (Example 1).

839

840 Note. Non-invariant parameters are printed in red, invariant parameters in green.

841

842

843 Figure 2. Population and Analysis Models (Example 2).

845 Note. Non-invariant parameters are printed in red, invariant parameters in green.