Derivation and Applicability of Asymptotic Results for Multiple Subtests Person-Fit Statistics

Applied Psychological Measurement
2016, Vol. 40(4) 274–288
© The Author(s) 2016
Reprints and permissions:
sagepub.com/journalsPermissions.nav
DOI: 10.1177/0146621615622832
apm.sagepub.com

Casper J. Albers¹, Rob R. Meijer¹, and Jorge N. Tendeiro¹

Abstract

In high-stakes testing, it is important to check the validity of individual test scores. Although a test may, in general, result in valid test scores for most test takers, for some test takers, test scores may not provide a good description of a test taker's proficiency level. Person-fit statistics have been proposed to check the validity of individual test scores. In this study, the theoretical asymptotic sampling distribution of two person-fit statistics that can be used for tests that consist of multiple subtests is first discussed. Second, simulation study was conducted to investigate the applicability of this asymptotic theory for tests of finite length, in which the correlation between subtests and number of items in the subtests was varied. The authors showed that these distributions provide reasonable approximations, even for tests consisting of subtests of only 10 items each. These results have practical value because researchers do not have to rely on extensive simulation studies to simulate sampling distributions.

Keywords

item response theory, person-model fit, validity test scores

In high-stakes testing, individual test scores are being used to make important decisions for individual test takers. In these circumstances, it is important to check the validity of the individual test scores. Although test scores may be valid for most persons in a particular population, for some test takers, these scores may not reflect their true proficiency level. For example, Meijer and Tendeiro (2014) showed that for some test takers on a high-stakes test, the proficiency scores did not seem to reflect their true proficiency level. Several methods have been proposed to check the validity of individual test scores. In this study, the authors focus on methods that are sensitive to the fit of individual response patterns to an item response theory (IRT) model. The idea behind this approach is that, when an item score pattern is very unexpected given the estimated proficiency level, this estimated proficiency level might not provide a good estimate of their true proficiency level. These methods are often denoted as person-fit methods or person-fit statistics (Meijer & Sijtsma, 2001).

¹University of Groningen, The Netherlands

Corresponding Author:

Casper J. Albers, Department of Psychometrics and Statistics, Faculty of Behavioral and Social Sciences, University of Groningen, Grote Kruisstraat 2/1, 9712 TS Groningen, The Netherlands.
Email: c.j.albers@rug.nl

One of the most popular statistics is the standardized log-likelihood statistic, denoted l_z , proposed by Drasgow, Levine, and Williams (1985). Based on asymptotic arguments, Snijders (2001; see also Magis, Raiche, & Béland, 2012) suggested an improved version of this statistic, denoting it l_z^* . In assessing person fit, the person parameter θ is unknown and needs to be estimated by $\hat{\theta}$. This estimation process biases the (asymptotic) behavior of l_z and Snijders's version accounts for this bias.

Both l_z and l_z^* are developed for unidimensional tests. In practice, however, many tests consist of several (correlated) subtests. For example, the Law School Admission Test consists of four subtests which total scores are combined into one total score. For these types of tests, it would be useful to have a person-fit statistic that combines information from the multiple subtests into one person-fit value. Drasgow, Levine, and McLaughlin (1991) proposed a multiple subtest extension of l_z , which they denoted l_{zm} . Conijn, Emons, and Sijtsma (2014) compared several approaches based on l_z with studying person fit for noncognitive multiple subtests consisting of polytomous items. Because of the advantage of l_z^* over l_z , Tendeiro, Meijer, and Albers (2014) recently studied the performance of Conijn et al.'s approaches applied to l_z^* rather than l_z for multiple subtests settings based on dichotomous items. This study by Tendeiro et al. was performed on the basis of a simulation design. The aim of the current article is to study the distributional properties of multisubtest modifications of l_z^* statistic through statistical (asymptotic) theory.

The outline of this report is as follows. In the next section, the l_z (Drasgow et al., 1985) and l_z^* (Snijders, 2001) statistics are introduced. Next, multiple subtests extensions based on l_z (Conijn et al., 2014; Drasgow et al., 1991) will be discussed. As explained above, the l_z^* statistic is a bias-removing improvement upon the l_z statistic. In the main theoretical section of this article, we study the theoretical null distribution of the multiple subtests person-fit statistics based on l_z^* instead of on l_z . The distributional theory is based on asymptotic arguments. The length of each subtest and the correlation of the latent traits between subtests are manipulated by means of a simulation study. The goal is to study possible effects of these factors on the quality of the asymptotic approximations. It will be shown that the asymptotic approximations are fairly good for subtest lengths as low as 10 items.

The l_z and l_z^* Statistics

This section shortly describes the l_z and l_z^* statistics. For a more extensive discussion of the l_z statistic, see, for example, Armstrong, Stoumbos, Kung, and Shi (2007); Magis et al. (2012); and van Krimpen-Stoop and Meijer (1999).

The l_z Statistic

A test taker with trait level θ is administered a univariate test consisting of n items. The random variable X_i equals 0 or 1, depending on whether item i was answered correctly or incorrectly, respectively. The probability of answering correctly, $P(X_i = 1 \mid \theta)$, is denoted by $p_i(\theta)$. The three-parameter logistic model (3PLM; see Embretson & Reise, 2000), or its constrained versions known as the two- and one-parameter logistic models, is commonly used in IRT to describe the stochastic relationship between θ and X_i . The 3PLM is given by

$$p_i(\theta) = c_i + (1 - c_i) \frac{e^{a_i(\theta - b_i)}}{1 + e^{a_i(\theta - b_i)}}, \quad (1)$$

where a_i , b_i , and c_i denote the discrimination, difficulty, and pseudo-guessing parameters of item i . The two-parameter logistic model (2PLM), which results from constraining c_i to zero in Equation 1, will be used in this simulation study. However, the theory in this article applies to other models as well.

The likelihood function of a response vector $\mathbf{X} = (X_1, \dots, X_n)$ is given by

$$L(\theta) = \prod_{i=1}^n p_i(\theta)^{X_i} (1 - p_i(\theta))^{1-X_i},$$

and the maximum likelihood (ML) estimator $\hat{\theta}_{ML}$ is obtained by maximizing $L(\theta)$ or, equivalently, by maximizing the log likelihood:

$$l_0(\theta) = \log(L(\theta)) = \sum_{i=1}^n \{X_i \log(p_i(\theta)) + (1 - X_i) \log(1 - p_i(\theta))\}. \quad (2)$$

As this function depends on the number of items, it is not directly applicable as a person-fit statistic. To this end, Drasgow et al. (1985) proposed to use the standardized version

$$l_z = \frac{l_0 - E(l_0)}{\sqrt{\text{Var}(l_0)}},$$

as a person-fit statistic. Here, the expectation and variance are given by

$$E(l_0) = \sum_{i=1}^n \{p_i(\theta) \log(p_i(\theta)) + (1 - p_i(\theta)) \log(1 - p_i(\theta))\} \quad (3)$$

and

$$\text{Var}(l_0) = \sum_{i=1}^n \left\{ p_i(\theta)(1 - p_i(\theta)) [\log(p_i(\theta)) - \log(1 - p_i(\theta))]^2 \right\}. \quad (4)$$

The l_z^* Statistic

Snijders (2001) argued that, in practice, it is rarely the case that true trait values θ are known. He showed that the l_z statistic is biased when the true θ is replaced by an estimate $\hat{\theta}$. He proposed a correction, actually applicable to a wider range of estimators than l_z . Snijders studied the class of standardized person-fit statistics described through

$$\frac{W_n(\theta)}{\sqrt{\text{Var}(W_n(\theta))}}, \quad (5)$$

where $W_n(\theta) = \sum_{i=1}^n (X_i - p_i(\theta))w_i(\theta)$, with $w_i(\theta)$ a particular choice of a weight function. For $w_i(\theta) = \log(p_i(\theta)) - \log(1 - p_i(\theta))$, the l_z statistic is obtained.

Snijders (2001) showed that the bias introduced by replacing θ by its estimate $\hat{\theta}$ does not vanish asymptotically, causing, for example, conservative inferences in case of parameter estimation through the 3PLM. A solution to this problem is obtained by modifying the weights $w_i(\theta)$ via

$$\tilde{w}_i(\theta) = w_i(\theta) - c_n(\theta)r_i(\theta),$$

where

$$c_n(\theta) = \frac{\sum_{i=1}^n p'_i(\theta) w_i(\theta)}{\sum_{i=1}^n p'_i(\theta) r_i(\theta)}, \quad (6)$$

$p'_i(\theta)$ is the first-order derivative of $p_i(\theta)$ with respect to θ , and the $r_i(\theta)$ is chosen such that

$$r_0(\hat{\theta}) + \sum_{i=1}^n (X_i - p_i(\hat{\theta})) r_i(\hat{\theta}) = 0. \quad (7)$$

Various choices of $r_i(\hat{\theta})$ satisfy this relation. The most common choice is that of the ML estimates, given by $r_0(\hat{\theta}) = 0$ and $r_i(\hat{\theta}) = \frac{p'_i(\hat{\theta})}{p'_i(\hat{\theta})(1-p'_i(\hat{\theta}))}$ ($i > 0$). Snijders (2001) showed that $W_n(\hat{\theta})$ is asymptotically normally distributed with expected value

$$E(W_n(\hat{\theta})) = -c_n(\hat{\theta}) r_0(\hat{\theta})$$

and variance $\text{Var}(W_n(\hat{\theta})) = n\tau_n^2(\hat{\theta})$, with

$$\tau_n^2(\hat{\theta}) = \frac{1}{n} \sum_{i=1}^n \tilde{w}_i^2(\hat{\theta}) p_i(\hat{\theta}) (1 - p_i(\hat{\theta})). \quad (8)$$

As a consequence, the statistic

$$I_z^* = \frac{W_n(\hat{\theta}) - E(W_n(\hat{\theta}))}{\sqrt{\text{Var}(W_n(\hat{\theta}))}}$$

is asymptotically standard normally distributed.

In case the estimation of θ is done through the method of ML, things become slightly easier. In this case, $r_0(\hat{\theta}) = 0$, which implies that $E(W_n(\hat{\theta})) = 0$. In this article, the authors will only work with ML estimators, but the authors shall continue to use the generalized notation of Snijders (2001).

The $I_{zm}^{(D)}$ and $I_{zm}^{(C)}$ Statistics and Proposed Corrections

The multiple subtests statistic developed by Drasgow et al. (1991), here denoted $I_{zm}^{(D)}$, is based on the sum of the l_0 statistics for each subtest. For each subtest s ($s = 1, \dots, S$), one computes $l_{0(s)}$, $E(l_{0(s)})$, and $\text{Var}(l_{0(s)})$ as described above. Then,

$$I_{zm}^{(D)} = \frac{\left(\sum_{s=1}^S l_{0(s)} \right) - \left(\sum_{s=1}^S E(l_{0(s)}) \right)}{\left[\sum_{s=1}^S \text{Var}(l_{0(s)}) \right]^{1/2}}. \quad (9)$$

The lack of covariance in the denominator is a consequence of the IRT assumption of local independence. Note that this assumption implies that the $l_{0(s)}$ scores across subtests are independent, but it still allows for the latent traits $\theta_{i(s)}$ across subtests to be correlated, which often is

the case in practice. Assuming that the true scores θ are known, $l_{zm}^{(D)}$ is standard normally distributed.

Conijn et al. (2014) suggested a slightly different approach to compute the multiple subtests statistic. Rather than summing the $l_{0(s)}$ statistics over the S subtests and then standardizing the sum, Conijn et al. suggested to first standardize each $l_{0(s)}$ and then to sum the standardized $l_{z(s)}$ statistics:

$$l_{zm}^{(C)} = \sum_{s=1}^S l_{z(s)}.$$

This approach is based on the same assumptions as the approach by Drasgow et al. (1991). In this simulation study, we shall study which method performs better.

Just as the l_z approach is biased when θ is unknown, so are the multiple subtest extensions $l_{zm}^{(D)}$ and $l_{zm}^{(C)}$. The solution by Snijders (2001) is actually directly applicable to $l_{zm}^{(D)}$ and $l_{zm}^{(C)}$. The authors, therefore, propose two new person-fit statistics, which are denoted by $l_{zm}^{*(D)}$ and $l_{zm}^{*(C)}$, respectively. In the following section, the asymptotic null distribution of each of these statistics is derived.

Asymptotic Null Distribution of $l_{zm}^{*(D)}$ and $l_{zm}^{*(C)}$

In this section, the asymptotic distributions of $l_{zm}^{*(D)}$ and $l_{zm}^{*(C)}$ are derived. In the next section, the applicability of this asymptotic theory for tests of finite length is studied.

Null Distribution of $l_{zm}^{(D)}$ and $l_{zm}^{*(D)}$

Statistic $l_{zm}^{(D)}$ is asymptotically normally distributed if the true θ values are used (Drasgow et al., 1991). However, replacing true with estimated θ values introduces a bias, as explained above. We shall now try to correct this bias by applying Snijders's approach.

It is possible to rewrite $l_{zm}^{(D)}$ given by Equation 9 in the form of Equation 5:

$$l_{zm}^{(D)} = \frac{W(\theta)}{\sqrt{\text{Var}(W(\theta))}}, \quad (10)$$

with

$$W(\theta) = \sum_{s=1}^S \sum_{i=1}^{n_s} [X_{i(s)} - p_{i(s)}(\theta_s)] w_{i(s)}(\theta_s),$$

$$w_{i(s)}(\theta_s) = \log(p_{i(s)}(\theta_s)) - \log(1 - p_{i(s)}(\theta_s)). \quad (11)$$

Here, θ denotes the vector $\theta_1, \dots, \theta_S$ of latent trait parameters per subtest. That Equation 9 can be written as Equation 10 can be seen as follows. First, for the numerator (recall Equations 2 and 3),

$$\sum_{s=1}^S l_{0(s)} = \sum_{s=1}^S \sum_{i=1}^{n_s} \{X_{i(s)} \log(p_{i(s)}(\theta_s)) + (1 - X_{i(s)}) \log(1 - p_{i(s)}(\theta_s))\}$$

and

$$\sum_{s=1}^S E(l_{0(s)}) = \sum_{s=1}^S \sum_{i=1}^{n_s} \{p_{i(s)} \log(p_{i(s)}(\theta_s)) + (1 - p_{i(s)}) \log(1 - p_{i(s)}(\theta_s))\},$$

thus, $\sum_{s=1}^S l_{0(s)} - \sum_{s=1}^S E(l_{0(s)}) = W(\theta)$. For the denominator, we have (recall Equation 4)

$$\sum_{s=1}^S \text{Var}(l_{0(s)}) = \sum_{s=1}^S \sum_{i=1}^{n_s} p_{i(s)}(\theta_s)(1 - p_{i(s)}(\theta_s)) [\log(p_{i(s)}(\theta_s)) - \log(1 - p_{i(s)}(\theta_s))]^2$$

and

$$\text{Var}(W(\theta)) = \sum_{s=1}^S \sum_{i=1}^{n_s} p_{i(s)}(\theta_s)(1 - p_{i(s)}(\theta_s)) (w_{i(s)}(\theta_s))^2,$$

and therefore $\sum_{s=1}^S \text{Var}(l_{0(s)}) = \text{Var}(W(\theta))$.

The authors have established that $l_{zm}^{(D)}$ belongs to the family of statistics considered by Snijders. It therefore follows that, for trait estimates $\hat{\theta}_{(s)}$ satisfying Equation 7, $W(\hat{\theta})$ is asymptotically normally distributed with expected value

$$E(W(\hat{\theta})) = \sum_{s=1}^S -c_{n_s}(\hat{\theta}_{(s)}) r_{0(s)}(\hat{\theta}_{(s)})$$

and variance

$$\text{Var}(W(\hat{\theta})) = \sum_{s=1}^S n_s \tau_{n_s}^2(\hat{\theta}_{(s)}),$$

where $c_{n_s}(\hat{\theta}_{(s)})$ and $\tau_{n_s}^2(\hat{\theta}_{(s)})$ are the functions defined by Equations 6 and 8 applied to subtest s .

Thus, in the above, the authors established that, under the assumption of local independence, asymptotically

$$l_{zm}^{*(D)} = \frac{W(\hat{\theta}) - E(W(\hat{\theta}))}{\sqrt{\text{Var}(W(\hat{\theta}))}} \sim N(0, 1).$$

Null Distribution of $l_{zm}^{(C)}$ and $l_{zm}^{*(C)}$

Conijn et al.'s (2014) approach is similar to Drasgow's et al. (1991) approach, but the order of operations is reversed. That is, $l_{z(s)}$ is computed for each subtest s and then all values are added. $l_{z(s)}$ is asymptotically standard normally distributed for true θ values. Furthermore, due to the local independence assumption, the $l_{z(s)}$ statistics are independent. As a consequence, $l_{zm}^{(C)}$ is the sum of S independent standard normally distributed variables and is therefore normally distributed with mean and variance equal to the sums of the means and variances, respectively. Thus, the asymptotic null distribution of $l_{zm}^{(C)}$, assuming known θ and local independence, is given by

$$l_{zm}^{(C)} = \sum_{s=1}^S l_{z(s)} \sim N(0, S).$$

The asymptotic distribution of $I_{zm}^{*(C)}$ can be derived along the same lines. Because $I_{zm}^{*(C)}$ is the sum of the $I_{z(s)}^*$ values, each independent and asymptotically $N(0, 1)$ distributed (Snijders, 2001), we immediately have that

$$I_{zm}^{*(C)} \sim N(0, S)$$

for trait estimates $\hat{\theta}_{(s)}$ satisfying Equation 7.

Design of the Simulation Study

In practice, to assess whether response patterns are unusual, one can either (a) simulate a large number of response patterns under the null distribution of normal behavior and compare the observed with the simulated response patterns; or (b) compute the critical value on the basis of the asymptotic distribution. Obviously, the first approach is time-consuming but has the benefit of not having to rely on asymptotic theory. The second approach is computationally much more efficient but does rely on asymptotic theory.

The main goal of this simulation study was to study the quality of the asymptotic results discussed in the previous sections. In particular, the authors wanted to verify how the asymptotic approximations for person-fit statistics $I_{zm}^{*(D)}$ and $I_{zm}^{*(C)}$ hold for relatively short subtest lengths (say, of 10 items). The goal is to understand whether the asymptotic results are accurate enough for most practical purposes. Discussing the univariate I_z^* statistic, Snijders (2001) expected that $n \geq 15$ would be sufficient for the asymptotic approximations to work well (in case of univariate scales). Subtests might be of shorter length than the 15 items mentioned by Snijders. How much shorter the subtests can be is dependent on the relation between the subtests. If the latent traits for the subtests correlate perfectly (i.e., test taker's θ is the same for each subtest), the test is actually univariate and, according to Snijders, subtests of lengths $n_i \approx 15 / S$ should suffice. When the correlation between subtest traits is smaller than 1, one may expect to need subtests of longer length. Studying how subtest length and trait correlations relate to the quality of the asymptotic approximations is the main goal of this simulation study.

The simulation study was set up as follows. Item scores of 1,000 test takers on four subtests were generated. Four subtest lengths were considered: 10, 25, 50, and 100. The shorter subtest lengths (10, 25) are of most practical interest, whereas the longer subtest lengths (50, 100) are mostly of theoretical interest. All subtests within the same data set had the same length. The 2PLM was used to generate the item scores, with discrimination parameters uniformly distributed between $[0.5, 2.0]$ and difficulty parameters standard normally distributed (bounded between -2.5 and $+2.5$). Moreover, four person θ parameters were generated for each simulated test taker, one per subtest. These parameters were randomly drawn from a multivariate normal distribution. Seven between-subtests correlations of θ were considered: 0.4(0.1)1.0. These item and person parameters resulted in data that were very similar to the empirical data from a number of large-scale high-stakes educational admission tests (see also Rupp, 2013).

The simulation study consisted therefore of a 4 (number of subtest lengths) by 7 (number of between-subtests correlations of θ) completely crossed design, hence 28 experiment conditions in total. One hundred replications were simulated per condition. For each replicated data set, six multiple subtests person-fit statistics were computed. Of these, $I_{zm}^{*(D)}$ and $I_{zm}^{*(C)}$, which the authors proposed and developed in this report, were of most interest. Furthermore, $I_{zm}^{(D)}$ and $I_{zm}^{(C)}$ were computed to compare these uncorrected statistics with their starred versions. The corrected starred statistics were expected to outperform the uncorrected statistics. Finally, I_z and I_z^* were computed by concatenating the four subtests together, that is, by ignoring the multiple

subtests data structure. This approach was expected to work well for large correlation values between the θ s but not so well for lower correlation values between the θ s.

The simulation was coded in R (R Core Team, 2014). The item parameters were estimated by means of the function `est()` in the “irtoys” package (Partchev, 2014). The ML person parameters were estimated by means of the function `mlbme()`, also in the “irtoys” package.

Results of the Simulation Study

Findings are reported in various tables and figures. For I_z^* , $I_{zm}^{*(C)}$, and $I_{zm}^{*(D)}$, Table 1 lists the following values: (a) the mean of the 1,000 statistic values per replication, averaged across the 100 replications; (b) the standard deviation of the 1,000 statistic values per replication, averaged across the 100 replications; (c) the Kolmogorov–Smirnov (KS) distance between the empirical and theoretical (asymptotic) normal cumulative distribution function; and (d) the level of significance when applying critical values from the asymptotic distribution, at $\alpha = .05$. The KS distance (Smirnov, 1948) is a method to assess whether the empirical results lie close to the asymptotic distribution. This metric is a common method for density comparisons and reports the maximum vertical distance between both cumulative distributions. When both distributions completely agree, this value is zero; when they completely disagree, it is one. For I_z , $I_{zm}^{*(C)}$, and $I_{zm}^{*(D)}$, Table 2 lists the means and standard deviations over the replications. (Reporting KS distances and levels of significance for these statistics is undesirable, as the asymptotic distribution only holds if all θ are known.)

First, the results in Table 1 are focused. The values for the means and standard deviations can be directly compared with the means and standard deviations of the asymptotic distribution. With respect to the means, Table 1 shows that (a) the empirical means are structurally larger than zero across all methods and that (b) the means decrease as n_i increases. Furthermore, the value of the means seem unrelated to the value of the subtest correlations ρ , with the exception of I_z^* which seems to have slightly larger mean values for larger ρ (for instance, for $n_i = 100$, the means range from 0.047 ($\rho = .4$) through 0.076 ($\rho = 1$)). With respect to the standard deviations, Table 1 shows that the empirical SDs are very close to their asymptotic values (1 for I_z^* and $I_{zm}^{*(D)}$; 2 for $I_{zm}^{*(C)}$) across all methods.

Table 1 shows that I_z^* is sensitive to ρ : The KS distance increases when ρ decreases, especially for larger subtest lengths. This result is to be expected, as the idea of ignoring the multiple subtests structure is incompatible with increasingly lower values of correlations between the θ s. The two multiple subtest methods, $I_{zm}^{*(C)}$ and $I_{zm}^{*(D)}$, do not show this dependence on ρ . For all methods, it holds that, when n_i increases, the asymptotic approximation lies closer to the empirical distribution. Furthermore, the KS distances decrease when the subtest length increases. This is obvious: faced with more data, better predictions can be made.

The KS distance measures how close the complete empirical distribution is with respect to the asymptotic distribution. This is actually more than what was needed: What happens in the critical region of the distribution (i.e., the lower tail) is what is important when looking for aberrant patterns. We are not (primarily) interested in whether the I_z^* scores of fitting response patterns are estimated without bias; what matters most is that the scores for misfitting response patterns are measured accurately. Figures 1 and 2, based on the $I_{zm}^{*(D)}$ values for the experiment condition defined by $n_i = 25$ and $\rho = .7$, display the empirical and asymptotic density functions (left) and cumulative distribution functions (right), with the 1% and 5% critical values added. The full empirical distribution has a significant misfit compared with the asymptotic standard normal because of its skew to the right. The skewness of this distribution has been noted in practice as well (e.g., Meijer & Tendeiro, 2012, Figure 1). However, the left tail of the empirical distribution is relatively well approximated by the asymptotic distribution, especially at the 1%

Table 1. Results of the Simulation Study for Person-Fit Statistics I_z^* , $I_{zm}^{(D)}$, and $I_{zm}^{(D)}$: Mean, Standard Deviation, Kolmogorov–Smirnov Distances, and Empirical Proportion of Statistic Values Scoring Below the 5% Quantile of the Asymptotic Distribution.

ρ	n_i	I_z^*				$I_{zm}^{(D)}$				$I_{zm}^{(D)}$			
		M	SD	KS	α_{asympt}	M	SD	KS	α_{asympt}	M	SD	KS	α_{asympt}
.4	10	0.083	1.017	0.076	.056	0.642	1.999	0.166	.037	0.336	0.996	0.176	.036
	25	0.053	1.067	0.059	.064	0.415	2.006	0.110	.040	0.212	1.002	0.113	.040
	50	0.045	1.142	0.064	.077	0.312	2.005	0.086	.042	0.155	1.001	0.087	.042
	100	0.047	1.269	0.084	.096	0.264	2.008	0.072	.042	0.129	1.002	0.072	.042
.5	10	0.079	1.009	0.071	.055	0.609	2.016	0.157	.038	0.318	1.003	0.167	.037
	25	0.055	1.054	0.057	.062	0.413	2.012	0.110	.040	0.211	1.005	0.113	.040
	50	0.048	1.108	0.059	.070	0.313	1.999	0.087	.042	0.155	0.998	0.087	.042
	100	0.052	1.220	0.074	.085	0.265	2.003	0.073	.042	0.130	1.000	0.072	.041
.6	10	0.082	1.008	0.072	.055	0.622	2.009	0.160	.037	0.322	1.002	0.169	.038
	25	0.056	1.038	0.056	.058	0.413	2.009	0.111	.040	0.211	1.004	0.114	.041
	50	0.051	1.075	0.053	.063	0.314	2.002	0.086	.041	0.156	0.999	0.086	.042
	100	0.056	1.165	0.065	.076	0.263	1.997	0.071	.042	0.130	0.997	0.071	.041
.7	10	0.083	1.003	0.072	.054	0.620	2.002	0.160	.037	0.326	0.998	0.172	.037
	25	0.058	1.020	0.053	.055	0.411	2.004	0.109	.040	0.210	1.000	0.112	.040
	50	0.054	1.048	0.050	.059	0.315	2.001	0.085	.041	0.156	0.999	0.085	.041
	100	0.061	1.115	0.055	.065	0.262	2.005	0.071	.042	0.130	1.000	0.071	.041
.8	10	0.084	0.999	0.073	.053	0.625	2.010	0.161	.037	0.326	1.000	0.172	.037
	25	0.060	1.010	0.054	.054	0.413	2.008	0.110	.041	0.211	1.003	0.112	.040
	50	0.056	1.030	0.050	.055	0.313	2.014	0.087	.043	0.156	1.006	0.087	.043
	100	0.065	1.074	0.049	.058	0.262	2.008	0.071	.042	0.130	1.003	0.071	.042
.9	10	0.084	0.997	0.073	.053	0.615	2.005	0.160	.038	0.320	0.999	0.168	.037
	25	0.062	0.999	0.055	.052	0.417	2.000	0.111	.040	0.212	1.000	0.113	.040
	50	0.060	1.008	0.048	.050	0.315	2.007	0.086	.042	0.157	1.002	0.086	.042
	100	0.071	1.048	0.046	.051	0.263	2.008	0.073	.042	0.130	1.003	0.073	.043
1.0	10	0.084	1.001	0.075	.054	0.614	2.015	0.161	.038	0.319	1.005	0.169	.038
	25	0.062	1.002	0.054	.053	0.410	2.011	0.110	.041	0.208	1.006	0.111	.042
	50	0.063	1.001	0.048	.049	0.316	2.000	0.085	.041	0.158	0.999	0.085	.041
	100	0.076	1.039	0.047	.048	0.263	2.007	0.072	.042	0.131	1.003	0.072	.042

Note. Results are averaged across replications. KS = Kolmogorov–Smirnov; ρ = correlation between subtest θ values; n_i = subtest length.

Table 2. Results of the Simulation Study for Person-Fit Statistics I_z , $I_{zm}^{(C)}$, and $I_{zm}^{(D)}$: Mean and Standard Deviation.

ρ	n_i	I_z		$I_{zm}^{(C)}$		$I_{zm}^{(D)}$	
		<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
.4	10	0.070	0.912	0.543	1.595	0.267	0.780
	25	0.039	0.956	0.353	1.673	0.173	0.818
	50	0.029	1.028	0.260	1.699	0.129	0.832
	100	0.025	1.146	0.217	1.719	0.108	0.844
.5	10	0.065	0.885	0.508	1.582	0.249	0.770
	25	0.040	0.937	0.351	1.679	0.173	0.822
	50	0.032	0.990	0.261	1.697	0.129	0.833
	100	0.029	1.093	0.218	1.720	0.108	0.846
.6	10	0.069	0.882	0.519	1.584	0.255	0.773
	25	0.043	0.914	0.352	1.674	0.173	0.821
	50	0.036	0.952	0.263	1.702	0.130	0.837
	100	0.033	1.030	0.217	1.712	0.108	0.843
.7	10	0.069	0.874	0.524	1.582	0.258	0.773
	25	0.044	0.890	0.349	1.667	0.172	0.818
	50	0.038	0.918	0.261	1.696	0.130	0.835
	100	0.038	0.972	0.217	1.713	0.108	0.846
.8	10	0.071	0.866	0.526	1.587	0.259	0.776
	25	0.048	0.878	0.352	1.675	0.174	0.825
	50	0.042	0.899	0.263	1.712	0.130	0.846
	100	0.043	0.927	0.217	1.723	0.108	0.853
.9	10	0.072	0.856	0.516	1.583	0.254	0.775
	25	0.052	0.864	0.356	1.672	0.177	0.825
	50	0.046	0.873	0.263	1.706	0.131	0.846
	100	0.048	0.890	0.216	1.722	0.108	0.855
1.0	10	0.074	0.851	0.517	1.587	0.255	0.779
	25	0.053	0.857	0.345	1.670	0.172	0.827
	50	0.050	0.861	0.264	1.700	0.131	0.846
	100	0.054	0.874	0.216	1.719	0.108	0.857

Note. Results are averaged across replications. ρ = correlation between subtest θ values; n_i = subtest length.

level. In Table 1, the α_{asympt} values are reported, which consist of the proportion of empirical data to the left of the 5% quantile of the asymptotic distribution. Thus, α_{asympt} describes for what proportion of the empirical results the null hypothesis of no aberrant behavior would be rejected (a Type I error), if this decision is made based on asymptotic theory and $\alpha = .05$. Values close to .05 are indicative of the adequacy of the asymptotic approximation. Figure 3 presents a visualization of the same results. For comparison, Figure 3 also shows the α_{asympt} values for the statistics without Snijders’s bias correction, where the asymptotic distribution is derived under the additional (and incorrect) assumption of known θ .

From Figure 2, the following can be concluded. (a) One should not rely on asymptotic theory for the nonbias removed statistics (right panels): Even for large subtests ($n_i = 50$) the α_{asympt} values can be less than half the nominal values (around 2%). Thus, the critical values based on the asymptotic approximation are too conservative in the case of noncorrected statistics (i.e., there is lack of power). The problem is also present for the bias-corrected statistics (left panels) but to a lesser extent. (b) I_z^* works very well for very small subtests ($n_i = 10$), but for lower correlations and lengthier subtests, the critical values from the asymptotic distribution yield are too liberal: Too many response patterns are flagged as aberrant. (c) The performance of $I_{zm}^{*(C)}$ and $I_{zm}^{*(D)}$ is

Figure 1. Histogram of the $1,000 \times 100 = 100,000$ computed $I_{zm}^{*(D)}$ values (bars) and the standard normal distribution (curve).

Note. The blue vertical lines correspond to the $\alpha = 1\%$ (left) and $\alpha = 5\%$ (right) critical values. This figure is based on the experiment condition defined by the parameters $n_i = 25$ and $\rho = .7$.

Figure 2. The corresponding cumulative distribution functions (solid: empirical; dashed: standard normal).

Note. It can be seen that the maximal vertical distance in the right-hand side display occurs around $I_{zm}^{*(D)} \approx 0$. This figure is based on the experiment condition defined by the parameters $n_i = 25$ and $\rho = .7$.

comparable, and both methods are unaffected by the value of ρ . (d) Even for subtests of moderate length ($n_i = 25$), the approximation by asymptotic theory provides accurate approximation.

Deviations between the reported α_{asympt} values and the nominal $\alpha = 0.05$ can be due to (a combination of) two reasons: (a) sampling variation (results are based on 100 replications of 1,000 simulated persons) and (b) the approximation is asymptotic and the sample size is clearly finite. However, sampling variation was controlled almost entirely by this experiment design. When sampling $100 \times 1,000 = 100,000$ values from a normal distribution, then in 95% of cases, the α_{asympt} value would be in (.0499, .05001).

Table 3 uses formal regression models to underpin the conclusions of Tables 1 and 2 and Figure 2. For each of the six types of person-fit statistic, first the regression model

Figure 3. The α_{asymp} values of the 100 replications of the 1,000 person-fit statistics.
Note. Left panels, from top to bottom: I_z^* , $I_{zm}^{(C)}$, and $I_{zm}^{(D)}$. Right panels, from top to bottom: I_z , $I_{zm}^{(C)}$, and $I_{zm}^{(D)}$. In the online supplementary material, a more detailed version of this figure is provided.

$M_j = \beta_0 + \beta_1(n_i)_j + \beta_2\rho_j + \beta_3(n_i \times \rho_j) + \varepsilon_j$ is fitted to the 4×7 combinations of subtest length n_i and subtest correlation ρ , and the p values and effect sizes are reported. Next, a similar model, now with α_{asymp} as the dependent variable, is fitted for the three starred methods. The authors decided to include an interaction term because Tables 1 and 2 and Figure 2 indicate that such interaction might be present. It has to be noted that this elementary linear model is not perfect; especially the subtest lengths seem to have a nonlinear relation with the dependent variable. However, the model seems adequate for a rough indication. The results from the table are clear: For every method, the size of the subtest is a significant factor with large effect sizes. The subtest correlation is only significant and relevant for I_z^* and I_z : The multiple subtests approaches indeed are capable of dealing with correlated subtests without distortions in the mean I_z^* or I_z value nor their corresponding α_{asymp} value. For the α_{asymp} values, a clear interaction is present

Table 3. Effect Sizes (η^2) and p Values for the Regression Models Predicting the Mean Values and α_{asympt} Values, With Sample Size and Subtest Correlation as Predictors.

		I_z^*		$I_{zm}^{*(C)}$		$I_{zm}^{*(D)}$		I_z		$I_{zm}^{(C)}$		$I_{zm}^{(D)}$	
		p	η^2	p	η^2	p	η^2	p	η^2	p	η^2	p	η^2
M	n_i	.013	.182	.000	.726	.000	.720	.000	.417	.000	.741	.000	.742
	ρ	.021	.152	.933	.000	.939	.000	.003	.171	.933	.000	.985	.000
	$n_i \times \rho$	.121	.065	.938	.000	.839	.000	.145	.036	.927	.000	.947	.000
α_{asympt}	n_i	.000	.227	.000	.584	.000	.540						
	ρ	.000	.476	.792	.001	.593	.006						
	$n_i \times \rho$	.000	.270	.962	.000	.792	.001						

only for I_z^* ; for the mean values, none of the interactions are significant (at the usual 5% level) nor do they have considerable effect sizes.

Discussion

In this article, the authors investigated the theoretical asymptotic distributions of three person-fit statistics for tests that consist of multiple subtests. In both psychological and educational measurement, these types of tests are often used, but thus far, there were no studies that investigated these asymptotic distributions. A recent study that used the multiple subtest extensions $I_{zm}^{*(C)}$ and $I_{zm}^{*(D)}$ made use of simulation to determine the critical values on the basis of which item score patterns could be classified as normal or aberrant (Tendeiro et al., 2014). A drawback of this approach is that it is time-consuming. In the present study, the authors showed that asymptotic theory can adequately be used for both statistics even for subtest lengths as low as 10 items. And that, at least, at for a 95% confidence interval, Type I errors are in agreement with what is expected.

Type I errors are controlled for by the simple univariate I_z^* statistic when correlations between tests are relatively high (larger than, say .7-.8). This is the case for many large-scale educational tests. Drasgow et al. (1991), for example, reported a correlation $r = .73$ between SAT Verbal and Quantitative tests and $r = .80$ between the enhanced ACT English and Mathematics test. Thus, in these cases, the theoretical asymptotic distribution of both the I_z^* statistic and the multiple subtests extensions $I_{zm}^{*(C)}$ and $I_{zm}^{*(D)}$ can be used. As many studies showed (e.g., Conijn et al., 2014), correlations between test scores for noncognitive instruments are often lower than for cognitive tests. In these cases, the asymptotic distribution of both $I_{zm}^{*(C)}$ or $I_{zm}^{*(D)}$ can be used, at least for $\alpha = .05$. We should note, however, that as shown in Figures 1 and 2, because the empirical distributions are skewed, at an α level of, for example, .10, results may be less optimal. However, almost all person-fit statistics use α levels of .05 or lower, so in practice, the authors conclude that this study showed that researchers can use the discussed asymptotic distributions to classify item score patterns as normal or aberrant for multiple subtests settings.

There can be benefit in applying bootstrap methods (such as those in Tendeiro et al., 2014) rather than resorting to asymptotic theory. These benefits especially hold in small studies, when the use of asymptotic theory is questionable. However, in this article, the authors show that, even for fairly short subtest lengths, asymptotic results already provide decent approximations. Finally, the benefit of not having to use the bootstrap distribution is saving computing time and, to a lesser degree, it is less technical: Understanding the bootstrap is quite hard; flagging all I_z^* scores below -1.65 is extremely simple.

Authors' Note

The opinions and conclusions contained in this report are those of the authors and do not necessarily reflect the position or policy of LSAC.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This study received funding from the Law School Admission Council (LSAC).

Supplementary Material

The R code used to generate the results, figures, and tables, as well as a detailed version of Figure 2, is provided as online supplementary material.

References

- Armstrong, R. D., Stoumbos, Z. G., Kung, M. T., & Shi, M. (2007). On the performance of the I_z person-fit statistic. *Practical Assessment, Research & Evaluation*, 12(16). Retrieved from <http://pareonline.net/getvn.asp?v=12&n=16>
- Conijn, J. M., Emons, W. M., & Sijtsma, K. (2014). Statistic I_z -based person-fit methods for noncognitive multiscale measures. *Applied Psychological Measurement*, 38, 122-136. doi:10.1177/0146621613497568
- Drasgow, F., Levine, M. V., & McLaughlin, M. E. (1991). Appropriateness measurement for some multidimensional test batteries. *Applied Psychological Measurement*, 15, 171-191. doi:10.1177/014662169101500207
- Drasgow, F., Levine, M. V., & Williams, E. A. (1985). Appropriateness measurement with polychotomous item response models and standardized indices. *British Journal of Mathematical and Statistical Psychology*, 38, 67-86. doi:10.1111/j.2044-8317.1985.tb00817.x
- Embretson, S. E., & Reise, S. P. (2000). *Item response theory for psychologists*. Mahwah, NJ: Lawrence Erlbaum.
- Magis, D., Raîche, G., & Béland, S. (2012). A didactic presentation of Snijders's I_z^* index of person fit with emphasis on response model selection and ability estimation. *Journal of Educational and Behavioral Statistics*, 37, 57-81. doi:10.3102/1076998610396894
- Meijer, R. R., & Sijtsma, K. (2001). Methodology review: Evaluating person fit. *Applied Psychological Measurement*, 25, 107-135. doi:10.1177/01466210122031957
- Meijer, R. R., & Tendeiro, J. N. (2012). The use of I_z and I_z^* person-fit statistics and problems derived from model misspecification. *Journal of Educational and Behavioral Statistics*, 37, 758-766.
- Meijer, R. R., & Tendeiro, J. N. (2014). *The use of person-fit scores in high-stakes educational testing: How to use them and what they tell us* (LSAC Research Report 14-03). Newtown, PA: Law School Admission Council.
- Partchev, I. (2014). *irtoys: Simple interface to the estimation and plotting of IRT models* (R package version 0.1.7). Retrieved from <http://CRAN.R-project.org/package=irtoys>
- R Core Team. (2014). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. Available from <http://www.R-project.org/>
- Rupp, A. A. (2013). A systematic review of the methodology for person fit research in Item Response Theory: Lessons about generalizability of inferences from the design of simulation studies. *Psychological Test and Assessment Modeling*, 55, 3-38.

- Smirnov, N. (1948). Table for estimating the goodness of fit of empirical distributions. *Annals of Mathematical Statistics*, 19, 279-281. doi:10.1214/aoms/1177730256
- Snijders, T. B. (2001). Asymptotic null distribution of person fit statistics with estimated person parameter. *Psychometrika*, 66, 331-342. doi:10.1007/BF02294437
- Tendeiro, J. N., Meijer, R. R., & Albers, C. J. (2014). *Detection of invalid test scores on admission tests: A simulation study using person-fit statistics* (LSAC Research Report RR-15-03). Newtown, PA: Law School Admission Council. Available from: <http://www.lsac.org/lisacresources/research/all/rr/rr-15-03>
- van Krimpen-Stoop, E. A., & Meijer, R. R. (1999). The null distribution of person-fit statistics for conventional and adaptive tests. *Applied Psychological Measurement*, 23, 327-345. doi: 10.1177/01466219922031446