DOCUMENT RESUME

ED 368 780 TM 021 291

AUTHOR Budescu, David V.; And Others
TITLE A Revised Modified Parallel Analysis (RMPA) for the Construction of Unidimensional Item Pools.
SPONS AGENCY Office of Naval Research, Arlington, Va.
PUB DATE Apr 94
CONTRACT N00014-91-J-1666; T-4428034
NOTE 42p.; Paper presented at the Annual Meeting of the National Council on Measurement in Education (New Orleans, LA, April 5-7, 1994).
PUB TYPE Reports - Evaluative/Feasibility (142) -- Speeches/Conference Papers (150)

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS Equations (Mathematics); Heuristics; Item Analysis; *Item Banks; *Matrices; Models; *Test Construction; Test Content; Test Items; Test Reliability
IDENTIFIERS Eigenvalues; Modified Parallel Analysis; *Revised Modified Parallel Analysis; *Unidimensionality (Tests)

ABSTRACT

Modified Parallel Analysis (MPA) is a heuristic method for assessing "approximate unidimensionality" of item pools. It compares the second eigenvalue of the observed correlation matrix with the corresponding eigenvalue extracted from a "parallel" matrix generated by a unidimensional and locally independent model. Revised Modified Parallel Analysis (RMPA) generalizes MPA and alleviates some of its technical limitations. An important and useful feature is a new method for eliminating items which violate the test's unidimensionality. This is achieved by eliminating items, one at a time, to determine their contribution to the matrices' eigenvalues. We propose a test for detecting items with larger impact in the observed data set, and eliminating them. The new method was tested in several situations in which unidimensional item pools were "contaminated" by various proportions of items from a secondary pool. The results indicate that RMPA does an excellent job in detecting low (10%) and moderate (25%) levels of contamination, but fails in cases of maximal (50%) contamination. Eleven tables. (Contains 46 references.) (Author)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED 368 780

A REVISED MODIFIED PARALLEL ANALYSIS (RMPA) FOR THE CONSTRUCTION OF UNIDIMENSIONAL ITEM POOLS

David V. Budescu

University of Illinois
Champaign, IL. 61821

Yoav Cohen **Anat Ben-Simon**

National Institute for Testing and Evaluation (NITE)
POBox 26015, Jerusalem 91260, ISRAEL

A paper presented in the annual meeting of the NCME, New Orleans, April, 1994

U S DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☑ This document has been reproduced as received from the person or organization originating it

☐ Minor changes have been made to improve reproduction quality

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY

DAVID V. BUDESCU

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

This research was sponsored by: Manpower, Personnel and Training R & D Program
Cognitive Science program Office of Naval Research (ONR)
Under Contract No. N00014-91-J-1666 & T No. 4428034

BEST COPY AVAILABLE

A REVISED MODIFIED PARALLEL ANALYSIS (RMPA) FOR THE CONSTRUCTION OF UNIDIMENSIONAL ITEM POOLS

ABSTRACT

Modified Parallel Analysis (MPA) is a heuristic method for assessing "approximate unidimensionality" of item pools. It compares the second eigenvalue of the observed correlation matrix with the corresponding eigenvalue extracted from a "parallel" matrix generated by a unidimensional and locally independent model.

Revised Modified Parallel Analysis (RMPA) generalizes MPA and alleviates some of its technical limitations. An important and useful feature is a new method for eliminating items which violate the test's unidimensionality. This is achieved by eliminating items, one at a time, to determine their contribution to the matrices' eigenvalues.

We propose a test for detecting items with larger impact in the observed data set, and eliminating them. The new method was tested in several simulations in which unidimensional item pools were "contaminated" by various proportions of items from a secondary pool. The results indicate that RMPA does an excellent job in detecting low (10%) and moderate (25%) levels of contamination, but fails in cases of maximal (50%) contamination.

A REVISED MODIFIED PARALLEL ANALYSIS (RMPA) FOR THE CONSTRUCTION OF UNIDIMENSIONAL ITEM POOLS

The increasing popularity of Item Response Theory (IRT) (e.g. Hambleton, 1983; Hulin, Drasgow & Parsons 1983; Lord, 1980) in educational, personnel and psychological testing has caused a revolution in this domain. Though multidimensional item response models have been developed (e.g. Reckase, 1985; Sympson, 1978), most readily applicable IRT models used today assume that the test takers responses to all items depend on a single latent trait (ability). Thus, it is crucial to establish that any item used in estimating the examinee's position along this ability continuum measures, in fact, the same trait. In other words, the need to demonstrate that a given item pool is truly unidimensional is a necessary condition for many applications of IRT.

Defining and Assessing Unidimensionality

Consider a test consisting of n items selected from a larger item pool. Let $\vec{U}_i$ be the vector of n binary responses to the test's items (taking values of 1 and 0 for correct and incorrect response, respectively), generated by the *ith* test taker (i=1...N), and let U_{ij} be her response to the *jth* item (j=1...n). Finally, let $\vec{\theta}_i$ be a vector of t latent traits characterizing the examinee's abilities. The strong principle of local independence (McDonald, 1981) states that:

$$P(\vec{U}_i = \vec{u}_i | \vec{\theta}_i) = \prod_{j=1}^{n} P_j(U_{ij} = u_{ij} | \vec{\theta}_i) \qquad (1)$$

This principle asserts that the responses to any pair of items are statistically mutually independent for any individual, or any subpopulation with fixed latent traits. The dimensionality of $\vec{U}$ is, simply, the minimal number of latent traits necessary to produce a (strong) locally independent model for $\vec{U}$. Thus, a pool is unidimensional if responses to all its items can be produced by unidimensional locally independent models.

Although a voluminous literature exists on the issue of unidimensionality of items and tests (see Berger and Knol, 1990; Hattie, 1984 and 1985 for partial reviews), currently there is no single approach which is fully satisfactory and/or universally accepted. Hattie (1984) compiled a list of 87 measures of unidimensionality and classified them into five nonoverlapping classes according to their underlying rationale. He distinguished between indices based on

(i) closeness to specific answer patterns,

(ii) reliability coefficients,

(iii) principal components (PC),

(iv) factor analysis (FA) and

(v) goodness of fit to various IRT models.

Hattie questioned the theoretical rationale of indices based on response patterns and reliability and showed empirically that the measures based on PC, FA and one parameter IRT (the Rasch model) are outperformed by methods quantifying deviation from multi-parameter IRT models.

"Approximate" Unidimensionality

Many researchers have argued, based on theoretical and empirical observations, that purely unidimensional tests, or pools, are quite rare (e.g. Ackerman, 1989; Humphreys, 1985; Reckase, Ackerman & Carlson, 1988; Traub, 1983; Yen, 1984, 1985). If, in fact, unidimensionality is frequently violated it is important to determine the practical implications of such violations. Following Reckase's original work (1979), several researchers (e.g. Drasgow & Parson, 1983; Yen, 1984, 1985) have shown that unidimensional models are quite robust under multidimensionality as long as there is a single "dominant" factor, and item difficulty is not confounded with dimensionality.

These, and other similar, studies suggest that strict unidimensional pools are not necessary for many practical applications of unidimensional IRT models (e.g. CAT). It is, however, important to develop methods that can identify pools which deviate from strict unidimensionality to a degree which does not seriously affect the fit or accuracy of the unidimensional IRT model.

This is the motivation behind recent work by Stout, who developed a test of the *essential unidimensionality* o a data set (Stout, 1987,1990; Nandakumar, 1991). Essential independence is achieved if the mean covariance (conditional on $\vec{\theta}_i$, the test taker's vector of t latent traits) between all n(n-1)/2 pairs of items approaches 0 as the number of items increases to infinity, and the essential dimensionality of a pool is the smallest number of latent traits necessary to satisfy essential independence. Essential independence is a weaker requirement than strong local independence and, in practice, it is obtained whenever there is a single dominant dimension in the data (e.g. Nandakumar, 1991).

In the same spirit Drasgow and Lissak(1983) presented Modified Parallel Analysis (MPA for short) as "a technique that can determine when an item pool is *sufficiently unidimensional* for the use of IRT" (Drasgow and Lissak, 1983, page 365). Modified Parallel Analysis relies on FA, a well understood method which is widely available to users in most statistical packages. Thus, it is (conceptually and computationally) easier to use than Stout's methods. This study will develop a revised and improved version of MPA.

Parallel and Modified Parallel Analysis

Parallel Analysis (PA) was proposed by Horn (1965) as an alternative to traditional factor analytical methods for identifying the number of latent factors. (e.g. Kaiser, 1960, Cattell, 1966, Bartlett, 1950).

The rationale behind PA is intuitively compelling, and its application is simple and straightforward: Random correlation matrices are generated, and their eigenvalues are extracted and averaged. The eigenvalues of the actual correlations are compared to these means and those factors with eigenvalues larger than their counterparts from the randomly generated data are retained. Crawford and Koopman (1973), Humphreys and Montanelli (1975) and Zwick and Velicer (1986), among others, report that PA works well in both Principal Components (PC) and Factor Analysis (FA). Recently Longman, Cota, Holden and Fekken (1989) published regression equations that eliminate the need to actually generate random matrices for each PA (for the PC case).

Parallel Analysis is used to determine the true dimensionality of a given data set, whereas in most applications of IRT models one seeks to determine whether a data set deviates significantly from unidimensionality. Modified Parallel Analysis (Drasgow & Lissak, 1983) provides an ingenious way of answering this question, using the rationale of PA. Its basic stages are:

(1) The intercorrelations (preferably tetrachoric) of the test's items are factor analyzed and the eigenvalues of the unrotated solution are calculated.

(2) A "parallel" unidimensional data set is generated by an IRT model. This data set parallels the observed one along all its attributes: It has an equal number of examinees with identical abilities, and it has the same number of items with identical parameters. Since responses are generated by an unidimensional IRT model satisfying the strong local independence principle the data set is, by definition, unidimensional.

(3) The (tetrachoric) correlations of the parallel data set are factor analyzed, and the eigenvalues of the unrotated solution are calculated.

(4) The dimensionality of the pool is assessed by comparing the magnitude of the second eigenvalues of the two data sets: If the empirical value is "sufficiently close" to the one obtained from the parallel data set, the pool is considered unidimensional.

Drasgow and Lissak (1983) reported five empirical studies providing strong empirical support for the procedure.

Eigenvalue based factor analytical techniques are not always successful in recovering the true dimensionality of binary data and, consequently, can't always distinguish between unidimensional and multidimensional data sets (e.g. Collins, Cliff, McCormick and Zatlin, 1986; Hattie, 1984; Knol and Berger, 1991; Roznowsky, Tucker & Humphreys, 1991; Zwick and Velicer, 1986). Thus it may seem surprising that some of the same measures perform very well in the framework of PA, and MPA. It is important to stress that the key to the success of these methods is their comparative nature. Whatever deficiencies these statistics have, they affect equally the results of the two data sets. Both PA and MPA focus on, and highlight, whatever differences exist between the empirical and parallel data sets above and beyond the systematic biases that the FA based measures may share.

Thus, in Hattie's (1984) typology MPA should not be considered a "factor analytic approach". In fact, it is closer to the "measures of fit to IRT models". MPA is a general method for assessing the similarity, or closeness, between two parallel data sets (one of which is known to be unidimensional). In this case similarity is quantified by some of the statistics usually employed in FA.

A critique of MPA

Modified Parallel Analysis suffers from a few technical limitations. In this section we describe these limitations and the problems they may cause in applying the method:

(i) MPA is a randomized procedure, i.e. its results depend to a certain degree on a random process, namely, the selection of the parallel data set. Thus, with small enough samples, researchers applying exactly the same procedure to the same set of data may reach different conclusions because of the variance between the random data sets generated in their simulations.

(ii) The simulated and the empirical data sets are equated along most important dimensions and any discrepancy between their eigenvalues can ,supposedly, be attributed to the multidimensionality of the empirical matrix. Yet, the communalities are estimated in a purely empirical fashion separately for each data set, introducing another important difference between them. This factor may bias (in an unknown direction and to an unknown degree) the comparative analysis.

(iii) MPA is a heuristic procedure, i.e. it lacks a measure of sampling variability for the formal assessment of the closeness of the critical statistic (the second eigenvalue) obtained from the unidimensional and the empirical solutions.

Other important limitations of MPA are:

(iv) It compares only the second pair of eigenvalues of the two matrices. This choice lacks a solid theoretical or empirical justification, and it may miss differences between the other eigenvalues (especially the third).

(v) MPA is too limited in its scope. The technique provides a global omnibus test of the hypothesis concerning the pool's unidimensionality. It lacks, however, a mechanism to follow up rejections of the hypothesized pattern, by eliminating some items and identify a unidimensional subset of the pool.

A REVISED MODIFIED PARALLEL ANALYSIS (RMPA)

In this section we outline a revised procedure (RMPA) which extends and generalizes the MPA. The revised method offers solutions to the technical problems described above and incorporates them into the existing framework of MPA. RMPA also includes a second stage which allows one to extract unidimensional subsets from larger, potentially multidimensional, pools.

To solve the first problem we replace the random generation of a parallel unidimensional population by the theoretical derivation of the expected correlations under the assumptions of (1) local independence, (2) unidimensionality of the parameter space and (3) the three parameter logistic model (e.g. Lord, 1980). The probability of a correct response for item j by a test taker with (a single) ability θ_i is given by $P(U_{ij} = 1 | \theta_i)$ or, in a shorter notation, P_{ji}.

$$P_{ji} = c_j + \frac{1 - c_j}{1 + \exp\{-1.7a_j(\theta_i - b_j)\}} \qquad (2)$$

where a_j is the item's discrimination parameter, b_j is the item's difficulty and c_j is its pseudo-guessing probability (see Hambleton, 1983 or Lord, 1980 for details). Under these assumptions the expected number of correct answers to any pair of arbitrary items, j and k, in a random sample of N examinees is:

$$f_j = \sum_{i=1}^{N} P_{ji} \text{ and } f_k = \sum_{i=1}^{N} P_{ki} \quad (3)$$

Under the assumption of local independence, the expected number of correct answers to *both* items, j and k, is:

$$f_{jk} = \sum_{i=1}^{N} P_{ji}P_{ki} . \quad (4)$$

Given f_{jk} and the two marginals, f_j and f_k, the expected 2x2 contingency table can be constructed, and the expected tetrachoric correlation can be estimated by standard methods (e.g. by solving a polynomial using the Newton Raphson method, as suggested by Kendall & Stuart 1979, pages 324-327). All expectations are (as in the original MPA) conditional upon the abilities and item parameters shared by the two data sets. The calculation can be further refined when the true distribution of the unidimensional abilities (θ_i) in the population is known. In these cases, the summation is replaced by integration over θ of the probability density function of θ_i.

To solve the second problem we replace the separate estimation of the communalities in the two data sets by the expected tetrachoric correlation between (hypothetical) experimentally independent administrations of any item under the assumptions of (1) local independence, (2) unidimensional ability and (3) a three parameter logistic item curve. This procedure amounts to estimating the items' communalities by their expected test-retest reliabilities. It is well known (e.g. Lord & Novick, 1968; Mulaik, 1972) that a measure's reliability provides an upper bound to its communality. The estimation procedure is just a special case of the technique described above for the calculation of the expected correlation. More specifically, if we let j=k, Equation 4 is reduced to:

$$f_{jj} = \sum_{i=1}^{N} P_{ji}^2 . \quad (4a)$$

The solution of the third problem relies on a data analytic procedure known as "jacknifing" (see Arvesen and Salsburg, 1975, Miller, 1974 or Mosteller & Tukey, 1977 for partial reviews) [1]. Assume that the original nxn correlation matrix between the test's items is strictly unidimensional. By eliminating one item at a time (i.e. deleting a row, and the corresponding column, from the original matrix) we obtain n submatrices of order (n-1)x(n-1) which, by definition, are also unidimensional. It is easy to show that under the "one factor model" (i.e. a matrix of rank one), the average first eigenvalue of these n submatrices, scaled by a factor of n/(n-1), is an unbiased estimate of the first eigenvalue of the original intact matrix.

A useful and important consequence of the "eliminate one item at a time" procedure is that it provides a simple method for assessing the impact, or influence[2], of any single item on the test's eigenvalues. The logic of the MPA procedure predicts that, under unidimensionality, the two matrices will have equal eigenvalues. It is generally accepted, and it was confirmed empirically by Drasgow & Lissak(1983), that the first eigenvalue (λ_1) is approximately equal in the observed and the expected matrices, regardless of the dimensionality of the observed responses. Thus, except for sampling error, the ratio of the two eigenvalues, RL_1, should be:

$$RL_1 = \lambda_1(\text{observed})/\lambda_1(\text{expected}) = 1 \quad . \qquad (5)$$

Furthermore, under unidimensionality, the eigenvalues of the n submatrices of the two data sets will be similar, will have equal variances and will be highly correlated. Finally, the removal of any given item from the pool will affect the observed and the expected data sets in identical fashion and to an equal degree. Thus, equality (5) should also hold in all n submatrices obtained by eliminating one item at a time. Let λ_1^{-i} be the first eigenvalue of the submatrix obtained after the deletion of item i, and let RL_1^{-i} be the ratio of the eigenvalues from the two parallel data sets. Then, for all items (i=1...n), the ratio of the jacknifed eigenvalues should equal the ratio of the original values:

$$RL_1^{-i} = \lambda_1^{-i}(\text{observed})/\lambda_1^{-i}(\text{expected}) = RL_1 \quad . \qquad (6)$$

If the responses are unidimensional, similar results are expected to hold for the second, third, and all subsequent eigenvalues. If, on the other hand, the observed responses violate unidimensionality, the analysis of the two data sets should yield differential results. For example, Drasgow and Lissak(1983) based the original MPA on the prediction that the second

eigenvalue of the observed matrix will be larger than its counterpart from the parallel unidimensional data set:

$$RL_2 = \lambda_2(\text{observed}) / \lambda_2(\text{expected}) > 1 \quad . \qquad (7)$$

If the data are generated by a multidimensional model we expect the mean of the n second eigenvalues extracted from the observed submatrices to be larger, and their variance to be higher, than their counterparts from the expected data set. Depending on the type and degree of deviation from unidimensionality, the correlation between the observed and expected values can be low (or even negative). Furthermore, the eigenvalues of the observed responses will be more sensitive to the removal of the foreign (or "contaminating") items. Since the expected matrix is unidimensional, its eigenvalues should not be affected considerably when any arbitrary item is removed. However, when a contaminating item is removed from a multidimensional test, the data set becomes closer to unidimensionality and its eigenvalues should decrease. For example, in a test of length n=50 with 8 foreign items (8/50=16% contamination), after the removal of such an item, the level of contamination is reduced to (7/49=)14%. Thus, whenever a contaminating item is eliminated the matching eigenvalues should be more similar to each other than in those cases in which a regular (noncontaminating) item is removed. Consequently, the ratio of the eigenvalues should be closer to unity in these instances.

To summarize, for any given data set, the ratio between the first eigenvalues, RL_1, in the two data sets can be used as a benchmark against which one can assess and test the ratios derived from the second and third eigenvalues (RL_2 and RL_3, respectively). At the global (i.e. test or pool) level, this approach is attractive because the behavior of RL_2 and RL_3 is assessed by a data based index which is more sensitive to, and reflects, the peculiarities and idiosyncrasies of the specific test being examined. At the local (i.e. item) level, this procedure provides a natural way of ranking, and scaling, the items according to their deviation from the pattern expected under unidimensionality. These properties can be used to develop a procedure for testing the global dimensionality of the observed responses, and a method of selecting unidimensional pools. In the next section we describe the technical details of such a testing procedure.

The "gap test"

To facilitate the comparison of the two data sets we calculate, for all items (i=1...n) and for the first k eigenvalues (k=1,2,3), the ratio of the two matched eigenvalues:

$$RL_k^{-i} = \lambda_k^{-i}(\text{observed}) / \lambda_k^{-i}(\text{expected}) \quad . \qquad (8)$$

The global ratio RL_1, as well as the individual RL_1^{-i} (i=1..n), are insensitive to the dimensionality of the observed data set. Their empirical distribution will be used to test the hypothesis that the ratios of the second and third eigenvalues behave similarly. Formally, we wish to test that $F\{RL_2^{-i}\} = F\{RL_1^{-i}\}$, and $F\{RL_3^{-i}\} = F\{RL_1^{-i}\}$, where $F\{\cdot\}$ stands for the distribution of the relevant statistic. The alternative hypothesis is that the ratios are distributed differentially.

We are particularly interested in the case where an essentially unidimensional data set is contaminated by a second (sometimes called "nuisance") ability. We speculated earlier, that removal of such contaminating items will affect differentially the two matched eigenvalues. When analyzing the correlations from the observed responses we expect to observe two distinct clusters of eigenvalues --- from the unidimensional and the contaminating pool, respectively --- separated by a substantial "gap". No parallel clustering and separation is expected in the corresponding eigenvalues of the matrix of expected correlations.

To detect such unusual gaps we adopt a procedure described by Wainer and Schacht (1978) to detect unusually large gaps in strings of ordered values. The first step in this procedure is to rank order the values in descending order and to calculate the (n-1) gaps, g_i, by subtracting each observation from the immediately previous (i.e. larger) one. The gaps are then weighted by a set of logistic weights to yield weighted gaps, y_i. These weights were selected to account and compensate for the fact that, typically, observations are more dense (hence should be overweighted) near the center and more sparse (and should be underweighted) in the tails of the distribution. Formally:

$$y_i = \sqrt{i\,(n-i)\,g_i} \quad . \qquad (9)$$

Finally, these values are standardized by division by y_m, the midmean (i.e. the mean of the central 50% values) of the weighted gaps. Thus, the standardized weighted gaps (SWGs for short), z_i can be expressed as:

$$z_i = y_i / y_m . \qquad (10)$$

Zero gaps indicate that two adjacent observations are equal, and unit gaps indicate that the distance between two observations is equal to the gaps' midmean. By definition, all gaps are non-negative but are unbounded from above. Wainer and Schacht (1978) suggest that z_i values greater than 2.25 indicate "unusually" large gaps. The probability of observing gaps this wide by chance is approximately 0.03 under the normal distribution, but this value was

shown by Wainer and Schacht (1978) to work quite well for a variety of symmetric t distributions with tails larger than the normal.

We will use this procedure to detect the location of the gap separating the items from the two pools, on the basis of ratios of the matched eigenvalues, RL_k^{-i} ($k > 1$). Thus, the hypothesis will be tested by comparing $MAX(Z_{ki})$, the largest SWG, with a critical rejection threshold. However, in the absence of precise information regarding the form of the distribution of these ratios, and the multiplicity of tests involved, it is not sufficient to rely on the 2.25 universal rule of thumb proposed by Wainer and Schacht. Instead, we find it necessary to develop more conservative rejection rules .

There are various ways of deriving critical rejection points for this decision: one can estimate the desired percentiles (.01, .05, etc.) from the distribution of RL_1^{-i}, or use a version of Chebyshev inequality (e.g. Stuart and Ord, 1987, page 110). The regular Chebyshev inequality states that the probability of finding a value located more than K standard deviations (SDs) from the population's mean is smaller than $1/K^2$, for any distribution with finite moments; A tighter version, invoking the additional assumptions that the distribution is symmetric and unimodal, yields a lower upper limit ($4/9K^2$), for the probability of the same event [3].

The decision, to reject H_0, will be based on a comparison with a critical threshold, $T(z_1)$. The threshold is derived from the distribution of the ratios of the first eigenvalue, RL_1^{-i}, in the same data set. Specifically, for k=2,3 we will reject H_0 if:

$$MAX(Z_{ki}) > T(z_1) = (M_1 + KS_1)$$

where M_1 and S_1 are the mean and SD, respectively, of the SWGs, z_{1i}, calculated from the ratios of first set of matched eigenvalues, RL_1^{-i}.

If $MAX(z_{2i})$ and/or $MAX(z_{3i}) > T(z_1)$, i.e. there is a significant gap in either distribution of ratios, we can eliminate those items which are located above the significant gap(s)[4]. Let m_1 denote the number of items eliminated ($m_1 > 0$) after this first pass through the data. Repeat the whole process with the reduced $(n-m_1)\times(n-m_1)$ correlation matrices. This second analysis may lead to the elimination of additional (say m_2) items. Repeat the procedure with the remaining items, and stop when the test fails to detect items to be rejected.

AN EMPIRICAL STUDY OF RMPA

Method

In this section we report results of an empirical study designed to test RMPA. Like most other studies in this area we simulated artificial test results by combining real item parameters and a set of reasonable assumptions regarding the distribution of abilities in the population of test takers. For the purpose of this study we contaminated a large unidimensional pool by (various proportions of) responses generated by a second (nuisance) ability correlated (at various levels) with the first. The efficiency of the RMPA was assessed by its ability to identify correctly the contaminating items and, consequently, partition the test into its two basic components.

We expect this procedure to be most efficient in cases of approximate unidimensionality. In other words, it should detect accurately relatively low levels of contamination, but not mixtures of two (equal) abilities. We also predict that the accuracy of the detection will be inversely related to the correlation between the two abilities involved.

Design

We generated 20 distinct "artificial tests". The following characteristics were fixed for all the tests:

n = test length = 80 items;
N = sample size = 2000 examinees;
t = number of abilities = 2.

The following variables were manipulated across tests:

p = proportion of contaminating items = 0%, 10%, 25% or 50% (p=0% is a a strictly, uncontaminated, unidimensional test and the other three cases represent low, medium and high levels of contamination);

r = the correlation between θ_1 and θ_2, the two abilities = 0.0, 0.5, 0.7 (the three values are approximately equally spaced in terms of r^2).

Replications: All combinations of p and r were replicated twice (i.e. with different seeds for the generation of the abilities, and different item parameters). In the sequel the two replications are labeled "B" and "R".

With the exception of the control condition (p=0, r=0), this can be viewed as a 3 x 3 factorial design repeated twice.

Item Parameters

The items for half the tests (replication "R") were randomly selected from the item bank of a test of English as a Foreign Language (EFL). This test was developed and is routinely used by the National Institute for Testing and Evaluation (NITE) as part of the Psychometric Entrance Test (PET) which is administered to all applicants to universities in Israel. The item parameters were estimated under the three parameter logistic model (Equation 2) using responses from approximately 7,000 examinees who took the test in 1988. The estimation was performed using the NITEST parameter estimation program (Cohen & Bodner, 1989). These parameter estimates for the n=80 items will henceforth be referred to as "true parameters".

The items for the other 10 tests (replication "B") were generated artificially, according to some distributional assumptions: The discrimination parameters (a's) were sampled from a normal distribution with a mean of 1.1 and a s.d. of 0.3; The difficulty parameters (b's) were obtained from a normal distribution with a mean of 0 and a s.d. of 0.8; The pseudo-guessing parameters (c's) are taken from a uniform distribution over the range 0.1 - 0.3. The values of the three parameters were sampled, from the respective sources, independently.

Table 1 summarizes the information regarding the two sets of true parameters. The two tests are equally difficult, but vary with respect to other aspects. The discrimination parameters of the real items ("R") have a higher mean and variance (m_a=1.33 and s_a=0.51) than the artificial ones ("B") (m_a=1.12 and s_a=0.25). On the average, it is easier to guess in the artificial test (m_c=0.2 vs. 0.16). Finally, whereas the parameters of the artificial items are uncorrelated (by design), the values of the EFL items parameters are moderately correlated.

Insert Table 1 about here

Abilities

All samples include N=2000 simulated "respondents". First we generated four mutually uncorrelated sets of abilities (T, A_1, A_2 and A_3): We sampled 8000 independent observations from the standard (0,1) normal distribution and randomly assigned them to the four sets. Correlated abilities were generated by calculating:

$$T(r) = r \cdot T + \sqrt{1 - r^2} \cdot A_i \qquad (11)$$

where A_i stand for A_1, A_2 or A_3, and r is the desired correlation (0.0, 0.5, 0.7) between the new set of abilities, T(r), and the reference set, T. Thus T(0), T(.5), T(.7) are sets of N=2000 normally distributed abilities which correlate 0.0. 0.5 and 0.7, respectively, with T.

Responses

Four sets of unidimensional response vectors were generated. Each set was simulated with a different set of abilities {T, T(0), T(.5) or T(.7)}, and all responses were generated with the "true" item parameters. The response vectors were simulated with the NITECAT software package (Cohen, Bodner & Ronen, 1989), which implements the process described by Drasgow and Lissak (1983).

The vectors generated with the T abilities are considered the "original" responses based on the dominant ability. Contaminated responses were obtained by replacing the original responses on p% of the items (randomly selected) with the corresponding responses generated by one of the other samples of abilities. Note that for the case of r=0 this procedure simulates a two-dimensional "noncompensatory" model (e.g. Ackerman, 1989, Sympson, 1978), whereas the other cases ($r > 0$) simulate "compensatory" models (e.g. Ackerman, 1989, Reckase, 1985).

Parameter estimation

In each of the artificial tests the three parameters of the n=80 items were estimated with the NITEST program (Cohen & Bodner, 1989). These are the various sets of "estimated parameters", to be used in the generation of the expected correlations.

Consistent with the massive literature on this topic (e.g. Dorans & Kingston, 1985; Miller & Oshima, 1992; Oshima & Miller, 1992), we found that the estimates of the b's and c's were not affected by the contamination. However, the estimates of the a's (the discrimination parameters) are sensitive to the level of contamination. Consistent with other studies in the literature, the estimates for items loaded on the dominant ability are hardly affected, whereas the discrimination measures of the contaminating items are reduced considerably. The

magnitude of this "shrinkage" is related to the level of contamination and the correlation between the two factors.

Results

Standard MPA

The standard MPA procedure, prescribed by Drasgow and Lissak (1983), was performed. Table 2 summarizes the results of this analysis. The table displays the first three eigenvalues of both correlation matrices, as well as their ratios.

Insert Table 2 about here

There is a clear and consistent pattern in the data which can be summarized by three observations:

(i) The first eigenvalues are, practically, equal in the two matrices and their ratio is, essentially, 1. There are no discernible differences between the 18 contaminated data sets and, in this respect, they are indistinguishable from the two uncontaminated tests.

(ii) In all contaminated tests, the second eigenvalue of the observed matrix is larger than its expected counterpart. Consequently, their ratio is greater than unity, as predicted by Drasgow & Lissak (1983). The ratio is a monotonically increasing function of p, the level of contamination, and a monotonically decreasing function of r, the inter-ability correlation.

(iii) The ratio of the third pair of eigenvalues is also greater than one. In fact, in most cases it is greater than the second ratio. The third ratio is not systematically related to r, the inter-ability correlation. However, it increases monotonically as a function of p, the level of contamination. The sharpest effect is obtained for highly (r=0.7) correlated, and the weakest effect is found for uncorrelated (r=0.0) abilities.

RMPA

We performed an informal RMPA by examining the eigenvalues of the jacknifed parallel matrices. Table 3 displays means, and standard deviations, of the first three eigenvalues extracted from the jacknifed submatrices. All the values in the table are based on n=80 matrices of order (n-1)=79. Note that the mean values are related to the eigenvalues from table 2 through multiplication by a scale factor of n/(n-1)=80/79.

Insert Table 3 about here

Table 4 presents ratios of the means, and the variances, of the three jacknifed eigenvalues of the 20 tests.

Insert Table 4 about here

There is a close correspondence between these mean ratios and the ratios presented in table 2, and the same three basic conclusions apply here, as well. The ratios of the variances follow a similar, but not identical, pattern:

(i) The variances of the first eigenvalues are, on the average, very close to each other and their ratio is close to unity. The only exceptions are the cases {r=0, p=50}, which represent mixtures of two unidimensional half-tests involving uncorrelated abilities.

(ii) In most cases (and on the average) the variance of the second (jacknifed) eigenvalues in the observed matrices is higher than in the expected one. The effect is most pronounced in the case of the independent traits (r=0), and for moderate or high levels of contamination (p=25 and 50, respectively).

(iii) In all 20 tests the variances of the third (jacknifed) eigenvalues are substantially higher in the observed matrices. The effect is much stronger than for the second eigenvalue, but there is no systematic pattern of change across levels and types of contamination.

Table 5 presents the correlations between the matched jacknifed eigenvalues for the 20 tests. Each correlation is based on n=80 observations.

Insert Table 5 about here

The pattern of results is clear and consistent with our expectations:

(i) There is a high (almost perfect) linear correlation for the first eigenvalue in most tests. The single exception is the {Rep=R, r=0, p=50} case, which is a mixture of two uncorrelated (unidimensional) half-tests.

(ii) In all cases of moderate and high contamination (p=25 and 50, respectively) the correlations based on the second and third eigenvalue are low, or negative.

(iii) In most cases of low contamination (p=10) the correlations based on the second eigenvalue are high (almost like for the first eigenvalue), but the correlations based on the third eigenvalue are always low, or negative.

This pattern indicates that, as suggested by Drasgow and Lissak (1983) and others, the first eigenvalues of the two parallel matrices are practically indistinguishable, across all types and levels of contamination. However, contrary to Drasgow and Lissak's speculation, not all the differences between the two data sets can be detected by comparing the second pair of eigenvalues. The means, variances and correlations of the jacknifed values seem to suggest that in some cases of low contamination (p=0.10) departure from unidimensionality can only be detected by examining the third pair of eigenvalues.

Rejection Thresholds

Table 6 presents six rejection thresholds calculated from the distribution of the first ratio in the 20 tests. These six rejection thresholds are obtained by crossing two confidence levels (95% and 99%) with three rules of detection --- an empirical value, a value calculated by the "tight" (i.e. assuming unidimodality and symmetry) Chebyshev inequality, and a value derived from the unconstrained Chebyshev inequality.

Insert Table 6 about here

In all tests, and for both confidence levels, the empirical percentile is more liberal than the corresponding Chebyshev bounds. Thus, the three rules can be ranked, from the most to the least conservative, identically for all tests and for both levels of confidence:

Unconstrained Chebyshev > Constrained Chebyshev > Empirical

One remarkable and reassuring aspect of this table is the relatively low variance of the bounds across the various conditions and replications. This indicates that the ratio of the first pair of jacknifed eigenvalues has a relatively stable distribution across the levels and types of contamination.

The most important issue, from a practical point of view, is to choose the "best" threshold for detection of wide gaps. To address this issue we focus on the performance of the various indices in the uncontaminated (p=0) case. Table 7 displays the proportion of SWGs exceeding the various indices for the three ratios. Since this is a strictly unidimensional test, we expect this proportion to be invariant for all three ratios and not to exceed its nominal confidence level (95% or 99%). Clearly, the empirical percentiles fail the invariance requirement and the 95% constrained Chebyshev bound is too liberal for the third ratio. In light of these results we conclude that is best to identify as "unusually wide gaps" those values that exceed the 95% unconstrained, or the constrained 99% Chebyshev bounds. We will focus primarily on rejections with 99% confidence. However, for completeness sake, we will report in the sequel results according to all the thresholds.

Insert Table 7 about here

Partition of the Tests

Tables 8a - 8c list the maximal SWGs observed in the distributions of the three ratios for each test. The tables also display the pattern of significance achieved by this maximal SWG, and its location. The columns labeled "significance" simply count how many (of the increasingly stringent) thresholds were exceeded in each family of tests. In the 95% and 99% columns, a 1 indicates that the observed value is greater than the empirical percentile lower than both Chebyshev bounds; a value of 2 describes a situation where the actual value is greater than the constrained (but smaller than the unconstrained) Chebyshev bound, and a value of 3 denotes a case where the maximal gap is larger than the most severe rejection rule. Our previous results (see table 7) dictate to interpret as "significant" values of 2 (at 99%), or values of 3 (at 95%).

The location of the gap is described by reporting the number of items above, and below, it. Recall that according to the logic of RMPA the contaminating items should have lower (i.e. closer to unity) ratios. We rank ordered the ratios in ascending order, so these items are expected to cluster "above" the gap. As a rule, we expect the proportion of item above the gap to match, approximately, the proportion of contamination in the specific test.

We reject the null hypothesis of unidimensionality if:

(1) The number of items "above the gap" < n/2 AND
(2) The Maximal SWG of the second AND/OR the third ratio is greater than the designated rejection threshold.

We examine three rejection rules with decreasing levels of conservatism: (1) 99% according to an unconstrained Chebyshev inequality, (2) 99% according to a constrained Chebyshev inequality, and (3) 95% according to the unconstrained Chebyshev inequality.

Insert Tables 8a - 8c about here

As expected, there are no significant gaps in the distribution of the first ratio but, in most tests, the largest SWG in the distribution of the second and/or third ratio is significant. We examine these significant gaps according to the three valid rejection thresholds:

All six cases with low (p=10) contamination are significant at the 99% level (five of them by the most severe criterion). In all six cases the gap separates the top 10% items from the bottom 90%. It appears that the procedure works well for this type of contamination.

Only three of the highly contaminated tests (p=50) are significant at 99%. More important, however, is the fact that in all six tests the widest gap is located at the bottom of the distribution. Although the numbers vary slightly across tests, the proportion of items above the gap is always greater than 80%. Clearly, the gap test does not work well for a mixture of two half tests.

The pattern of results is slightly more complex in the case of moderate (p=25) contamination, and it depends on the level of the inter-ability correlation: For both tests with uncorrelated (r=0) abilities, and one of the tests with moderately correlated (r=0.5) abilities, the significant gap (99%) in the distribution of the second ratio separates the upper 25% items from the rest of the test. In the other test with r=0.5 the gap between the top 25% of the items and the lower 75% is significant at the 95% level. Finally, for the tests involving highly correlated abilities (r=0.7), the maximal gap is located at the lower end of the distribution (69 and 72 items above the gap). In both cases the second largest gap distinguishes between the (most) contaminating items and the original ones. Thus, the gap test operates well only for cases with low inter-ability correlations.

To summarize, RMPA found a significant gap in the distribution of the ratios of matched eigenvalues in all the tests examined. In 14 tests the gap was significant at 99% and in the other six at 95%. A significant gap located in the upper half of the distribution (i.e. with fewer items above the gap than bellow it) is taken as a strong indication of violation of unidimensionality and prescribes elimination of all items above the gap. The ten tests identified by this criterion include all those with low contamination (p=10), as well as the moderately contaminated ones (p=25), with moderate level of inter-ability correlation ($r < 0.7$).

In the sequel we focus only on these 10 shortened tests. Plots of the 10 relevant distributions of standardized weighted gaps (not presented here because of space limitations), clearly show that:

(i) the contaminating items are clustered at one end of the distribution, and
(ii) there is an unusually large gap separating this cluster from the bulk of the items. This gap can be detected in the raw gaps, but it is more pronounced in the standardized weighted form.

The quality of the technique is assessed by its ability to detect the contaminating items and remove them, while retaining the original ones. Table 9 summarizes this analysis for the 10 short tests. For each one we report the hit rate (i.e. contaminating items rejected correctly) and the false alarm rate (i.e. original items rejected incorrectly). The figures are very impressive ---- for all the tests with p=10%, the hit rate is 100% and for the tests with p=25% it is 95%. Both figures are accompanied by false alarm rates close to 0.

Insert Table 9 about here

Re-examination of the shortened tests

Having shortened 10 tests according to the results of the initial RMPA we repeated the procedure. The second iteration verifies the unidimensionality of the shortened tests: If the first stage is successful in removing all sources of contamination, we do not expect to detect any significant gaps in this second round.

Tables 10 and 11 report the results of the MPA and the RMPA of the shortened tests. A quick comparison with tables 2 and 4 (summarizing the same results for the original full tests) reveals that all major sources of multidimensionality were eliminated. The ratios of the second

eigenvalues, and the ratios of their variances, are close to unity (We assume that a heuristic MPA would also declare all these tests unidimensional). The third ratios are somewhat higher but are, considerably, lower than those of the original tests.

Insert Tables 10 and 11 about here

The SWGs of the remaining items were calculated, new rejection thresholds were derived, and the gap test was applied again [5].

The only significant gap was found in the {Rep=B r=0.7 p=10} test. In this case the second iteration of the RMPA prescribes removal of five additional items. All contaminated items were successfully detected by the first iteration so these are five "false alarms". The final test consists of 66 unidimensional items (instead of 72).

SUMMARY

The goal of the current research was to develop a practical, yet theoretically sound and computationally feasible, tool for testing the global dimensionality of large item pools and eliminating items which cause violations of the pool's unidimensionality. Both goals are attained in the unified framework of revised modified parallel analysis (RMPA).

RMPA is an extension and generalization of the modified parallel analysis (MPA) method that was developed by Drasgow and Lissak (1983) as an approximate method for testing the unidimensionality of item pools. MPA replies on a heuristic comparison of a statistic (the second eigenvalue) derived from the matrix of items' intercorrelations and the corresponding value extracted from a "parallel" matrix generated by a unidimensional, and locally independent, model (in our case the three parameter logistic model).

RMPA is based on a similar comparative logic, but improves upon MPA in several ways:

(1) It alleviates some minor technical limitations, through the use of expected inter-item correlations and item communalities. While in MPA the inter-item correlations are generated by a randomized procedure, in RMPA their expected values are derived theoretically under the exact assumptions of the 3-P logistic model.
(2) It implements a formal test for comparing the observed data set with its parallel (and unidimensional) counterpart. The test is based on inspection of the ratio of the k'th eigenvalue of the observed intercorrelation matrix to that of the expected matrix. The ratios for the second and third eigenvalues are compared with the ratio of the first eigenvalue which serves as a benchmark. In case of a unidimensional data set, all three ratios should be of equal magnitude. Violations of unidimensionality are manifested in larger observed/ expected ratios of the second and third eigenvalues relative to that of the first.
(3) Contingent upon the results of this test, it provides a method for identifying and eliminating items which violate the test's unidimensionality.

The testing and elimination procedures are based on the "eliminate one item at a time" principle. This methodology allows one to assess the contribution of each item to the test's eigenvalues. Furthermore, one can determine the variance and distribution of these values and analyze the differential impact of any given item in the observed and parallel matrices. Items which have a "significantly" larger impact in the observed data set violate unidimensionality. The detection of these items relies on a conservative version of Wainer & Schacht's (1978) "gapping" test.

The lagest (first) eigenvalues of the observed and expected matrices are practically identical in all cases, regardless of the level of correlation between the two abilities and the degree of contamination. Therefore, we used the distribution of their ratio to determine rejection thresholds for the ratio of the second and third eigenvalues. These thresholds are based on conservative Chebyshev bounds, and are specifically tailored to each data set.

RMPA was tested in several simulations of unidimensional item pools which were contaminated by various proportions of items loaded on a secondary nuisance ability. The method was highly successful in identifying low (10%) levels of departure from unidimensionality, and in detecting moderate (25%) deviations from unidimensionality when the abilities were not highly ($r < 0.7$) correlated. In these cases over 90% of the contaminating items were identified and less than 1% of the original items were eliminated erroneously. The procedure failed however, in tests which are easily detected as multidimensional by the MPA method, namely tests which are equal mixtures (50%) of items loaded on two abilities.

The two methods, MPA and RMPA, can be regarded as complementing each other. Clear cases of multidimensionality can be easily detected by inspecting the observed/expected ratio of the second eigenvalue but the gapping test is unable to separate between the two (unidimensional) sub-pools. Paradoxically, in cases of multidimensionality which are not easily detected by the magnitude of the ratio, the gapping test not only detects the presence of nuisance ability factor but also identifies the contaminating items. The findings of this study suggest that practical application of RMPA should proceed in two stages. In the first stage, the ratio of the second eigenvalues (observed/expected) should be inspected. In the second stage of the analysis the gapping test is applied. The decision as to the unidimensionality of the data set is determined by a conjunction of the results from both stages. Therefore, the unidimensionality hypothesis should be rejected if the results of either stage indicate multidimensionality. Large ratio (in the first stage) combined with nonsignificant gaps (in the second stage) most probably indicate a case of a multidimensional data set with close to equal proportions of items loaded on each ability factor.

It remains to be proved whether RMPA is as effective in separating out items that load on more than two factors and, more generally, whether the procedure is applicable to other factor structures.

We conclude by pointing out that the logic of MPA and RMPA can be generalized to other statistics of closeness between the two data sets. For example, it might be interesting to apply it to indices derived from non linear factor analysis (e.s. McDonald, 1982).

REFERENCES

Ackerman, T. (1989) Unidimensional IRT calibration of compensatory and noncompensatory multidimensional items. *Applied Psychological Measurement, 13*, 113-127.

Arvesen, J.N. & Salsburg, D.S. (1975) Approximate tests and confidence intervals using the jacknife. In R.M. Elashoff (Ed.) *Perspectives in Biometrics*. New York: Academic Press.

Bartlett, M.S. (1950) Tests of significance in factor analysis. *British Journal of Psychology, 3*, 77-85.

Berger, M.P.F. & Knol, D.K. (1990) On the assessment of dimensionality in multidimensional item response theory models. *Paper presented at the annual AERA Meeting*. Boston, Mass.

Cattell, R.B. (1966) The scree test for the number of factors. *Multivariate Behavioral Research, 1*, 245-276.

Cohen, Y. & Bodner, G. (1989) A manual for NITEST: A program for estimating IRT parameters. *Report No. 94. Jerusalem: NITE.*

Cohen, Y., Bodner, G. & Ronen, T. (1989) A manual for ..ECAT: A software package for research on IRT/CAT. *Report No. 100. Jerusalem: NITE.*

Collins, L.M., Cliff, N., McCormick, D.J. & Zatlin, J.L. (1986) Factor recovery in binary data sets: A simulation. *Multivariate Behavioral Research, 23*, 377-392.

Crawford, C.B. & Koopman, P. (1973) A note on Horn's test for the number of factors in factor analysis. *Multivariate Behavioral Research, 8*, 117-125.

Drasgow, F. & Lissak, R.I. (1983) Modified parallel analysis: A procedure for examining the latent dimensionality of dichotomously scored item responses. *Journal of Applied Psychology, 68*, 363 - 373.

Drasgow, F. & Parsons, C. (1983) Applications of unidimensional item response theory models to multidimensional data. *Applied Psychological Measurement, 7*, 189-199.

Dorans, N.J. & Kingston, N.M. (1985) The effects of violations of unidimensionality on the estimation of item and ability parameters and on item response theory equating to the GRE verbal scale. *Journal of Educational Measurement, 22*, 249-262.

Hambleton, R.K. (1983) *Applications of Item Response Theory*. Vancouver BC: Educational Research Institute of British Columbia.

Hattie, J. (1984) An empirical study of various indices for determining unidimensionality. *Multivariate Behavioral Research, 19*, 49-78.

Hattie, J. (1985) Methodology review: Assessing unidimensionality of tests and items. *Applied Psychological Measurement, 9,* 139-164.

Horn, J.L. (1965) A rationale and test for the number of factors in factor analysis. *Psychometrika, 30,* 179-164.

Hulin, C.L., Drasgow, F. & Parsons, C.K. (1983) *Item Response Theory.* Homewood, Ill: Dow Jones-Irwin Publishing Company.

Humphreys, L.G. (1985) General intelligence: An integration of factor, test and simplex theory. IN B.B. Wolman (Ed.) *Handbook of Intelligence* (pp. 201-224). New York: Wiley.

Humphreys, L.G. & Montanelli, R.G. (1975) An investigation of the parallel analysis criterion for determining the number of common factors. *Multivariate Behavioral Research, 10,* 193-205.

Kaiser, H.F. (1960) The application of electronic computers to factor analysis. *Educational and Psychological Measurement, 20,* 141-151.

Kendall, M. & Stuart, A. (1979) *The Advanced Theory of Statistics (Vol 2),* London: McMillan, Fourth Edition.

Knol, D.L. & Berger, M.P.F. (1991) Empirical comparison between factor analysis and multidimensional item response models. *Multivariate Behavioral Research, 26,* 457-478.

Longman, R.S., Cota, A.A., Holden, R.R. & Fekken, G.C. (1989) A regression equation for the parallel analysis criterion in principal components analysis: Mean and 95th percentile eigenvalues. *Multivariate Behavioral Research, 24,* 59-69.

Lord, F.M. (1980) *Applications of Item Response Theory to Practical Testing Problems.* Hillsdale, NJ: Lawrence Erlbaum Publishers.

Lord, F.M., & Novick, M. R. (1968) *Statistical Theories of Mental Test Scores.* Reading,MA: Addison-Wesley Publishing Company.

McDonald, R.P. (1981) The dimensionality of tests and items. *British Journal of Mathematical and Statistical Psychology, 34,* 100-117.

McDonald, R.P. (1982) Linear vs. nonlinear models in item response theory. *Applied Psychological Measurement, 6,* 379-396.

Miller, R.G. (1974) The jacknife - A review. *Biometrika, 61,* 1-15.

Miller, M.D, & Oshima, T.C. (1992) Effect of sample size, number of biased items, and magnitude of bias on a two-stage item bias estimation method. *Applied Psychological Measurement, 16,* 381-388.

Mosteller, F. & Tukey, J.W. (1977) *Data Analysis and Regression: A Second Course in Statistics*. Reading, Mass: Addison-Wesley Publishing Co.

Mulaik, S.A. (1972) *The Foundations of Factor Analysis*. New York:McGraw Hill Book Company.

Nandakumar, R. (1991) Traditional dimensionality versus essential dimensionality. *Journal of Educational Measurement, 28,* 99-117.

Oshima, T.C. & MIller, M.D. (1992) Multidimensionality and item bias in item response theory. *Applied Psychological Measurement, 16,* 237-248.

Reckase, M.D. (1979) Unifactor latent trait models applied to multifactor tests: results and implications. *Journal of Educational Statistics, 4,* 207-230..

Reckase, M.D. (1985) The difficulty of test items that measure more than one ability. *Applied Psychological Measurement, 9,* 401-412.

Reckase, M.D., Ackerman, T.A., & Carlson, J.E. (1988) Building unidimensional tests using multidimensional items. *Journal of Educational Measurement, 25,* 193-203.

Roznowsky, M., Tucker, L.R., & Humphreys, L.G. (1991) Three approaches to determining the dimensionality of binary items. *Applied Psychological Measurement, 15,* 109-127.

Stout, W.F. (1987) A nonparametric approach for assessing latent trait unidimensionality. *Psychometrika, 52,* 589-617.

Stout, W.F. (1990) A new item response theory modeling approach with applications to unidimensionality assessment and ability estimation. *Psychometrika, 55,* 293-325.

Stuart, A., & Ord, J.K. (1987) *Kendall's Advanced Theory of Statistics (Vol 1),* New York: Oxford University Press, Fifth Edition.

Sympson, J.B. (1978) A model for testing with multidimensional items. In Weiss, D.J. (Ed.) *Proceedings of the 1977 CAT Conference* (pp 82-98). Minneapolis: University of Minnesota, Department of Psychology.

Traub, R.E. (1983) A priori considerations in choosing an item response model. In R.K. Hambleton (Ed.) *Applications of Item Response Theory* (pp. 55-70). Vancouver BC:Educational Research Institute of British Columbia.

Wainer, H. & Schacht, S. (1978) Gapping. *Psychometrika, 43,* 203-212.

Yen, W. M. (1984) Effects of local item dependence on the fit and equating performance of the three parameter logistic model. *Applied Psychological Measurement, 8,* 125-145.

Yen, W. M. (1985) Increasing item complexity: A possible cause of scale shrinkage for unidimensional item response theory. *Psychometrika, 50,* 399-410.

Zwick, W.R. & Velicer, W.F. (1986) Comparison of five rules for determining the number of components to retain. *Psychological Bulletin, 99,* 432-442.

FOOTNOTES

(1) Strictly speaking "jacknifing" refers to an analysis in which observations (i.e. respondents) are eliminated one at a time from the sample. In this case, we eliminate variables (items) in a similar fashion. Several item analysis computer programs use a similar approach in order to identify subscales with maximal reliability.

(2) I_i, the sample influence function (Devlin, Gnanadesikan and Kettenring, 1975; Hampel, 1974) of parameter, T, is given by:

$$I_i = (n-1)(T - T_{-i}),$$

where n is the number of items, and T_{-i} is an estimate of the parameter T obtained after the elimination of item i. Note that I_i is, simply, a linear transformation of T_{-i}.

(3) Strictly speaking, Chebyshev inequality requires knowledge of the parameters (mean and variance) of the population of interest. However, Saw, Yang and Mo (1984) have shown that sample estimates of these parameters can be used, with very little loss of precision, in moderately large samples.

(4) Occasionally a large (and significant) gap will be detected in the lower tail of the distribution, i.e. separating the bulk of the data from a minority of items with unusual low ratios of observed/expected jacknifed eigenvalues. Clearly, these cases are not relevant for our hypothesis.

(5) Since the procedure is data driven, we opt not to use the thresholds values employed in the first stage. Thus, when analyzing a test consisting of $(n - m_1)$ items one should obtain the same results, and reach the same conclusions, whether it is treated as "an original" or "a reduced" test.

Acknowledgements

The authors would like to thank Mr. David Chinn for his assistance in all stages of this work.

Table 1:

Means, standard deviations and correlations of the two sets of item parameters

Rep=B Parameter	n	Mean	Std. Dev.
a	80	1.123	0.245
b	80	0.172	0.873
c	80	0.202	0.057

Correlations	a	b	c
a	1.000	0.090	-0.196
b	0.090	1.000	-0.260
c	-0.196	-0.260	1.000

Rep=R Parameter	n	Mean	Std. Dev.
a	80	1.328	0.511
b	80	-0.026	0.985
c	80	0.161	0.098

Correlations	a	b	c
a	1.000	0.518	0.397
b	0.518	1.000	0.754
c	0.397	0.754	1.000

Table 2:

Modified Parallel Analysis (MPA) of 20 tests: The first three eigenvalues for the observed and expected matrices, and their ratios

			Eigenvalue 1			Eigenvalue 2			Eigenvalue 3		
Rep	r	p	Exp	Obs	Obs/Exp	Exp	Obs	Obs/Exp	Exp	Obs	Obs/Exp
B	-	0	24.34	25.10	1.03	1.79	1.78	0.99	0.17	0.67	4.02
B	0.0	10	22.03	22.71	1.03	1.62	2.55	1.58	0.17	1.66	9.72
B	0.0	25	18.15	18.80	1.04	1.43	5.59	3.92	0.15	1.53	10.17
B	0.0	50	11.93	12.58	1.05	0.61	12.01	19.77	0.08	0.85	10.35
B	0.5	10	22.72	23.55	1.04	1.66	1.78	1.07	0.18	1.66	9.13
B	0.5	25	19.93	20.77	1.04	1.46	3.87	2.65	0.16	1.54	9.75
B	0.5	50	17.76	18.71	1.05	0.92	6.05	6.60	0.07	1.07	14.38
B	0.7	10	23.31	24.16	1.04	1.66	1.73	1.04	0.18	1.26	7.09
B	0.7	25	21.27	22.13	1.04	1.46	2.30	1.58	0.15	1.55	10.51
B	0.7	50	19.91	20.82	1.05	1.18	3.67	3.11	0.09	1.34	14.26
R	-	0	26.20	26.23	1.00	3.47	3.22	0.93	0.36	0.63	1.78
R	0.0	10	23.59	23.84	1.01	2.84	2.86	1.01	0.25	2.57	10.20
R	0.0	25	19.57	19.81	1.01	2.38	6.20	2.60	0.20	2.33	11.91
R	0.0	50	12.23	12.90	1.05	1.45	12.20	8.39	0.15	1.77	11.53
R	0.5	10	23.68	24.14	1.02	2.76	2.76	1.00	0.28	1.90	6.69
R	0.5	25	21.45	21.90	1.02	2.48	4.31	1.74	0.25	2.45	9.70
R	0.5	50	19.30	19.88	1.03	1.79	6.35	3.55	0.13	1.93	14.89
R	0.7	10	24.45	24.86	1.02	2.88	2.87	1.00	0.30	1.18	4.01
R	0.7	25	22.91	23.30	1.02	2.64	3.16	1.20	0.23	2.28	10.09
R	0.7	50	21.78	22.25	1.02	2.32	3.89	1.68	0.17	2.39	14.17

Notes:
All results based on n=80 items and N=2000 respondents.
Exp = Derived from matrix of expected correlations
Obs = Derived from matrix of observed correlations.

Table 3:

Revised Modified Parallel Analysis (RMPA) of 20 tests: Means and standard deviations of the first three eigenvalues of the jacknifed submatrices (observed and expected)

Rep	r	p	Source	Eigenvalue 1 Mean	Eigenvalue 1 SD	Eigenvalue 2 Mean	Eigenvalue 2 SD	Eigenvalue 3 Mean	Eigenvalue 3 SD
B	-	0	Exp	24.036	0.122	1.764	0.029	0.165	0.004
			Obs	24.786	0.118	1.756	0.027	0.664	0.009
B	0.0	10	Exp	21.752	0.145	1.598	0.026	0.168	0.004
			Obs	22.429	0.147	2.515	0.103	1.637	0.027
B	0.0	25	Exp	17.920	0.163	1.409	0.027	0.148	0.004
			Obs	18.563	0.169	5.519	0.143	1.507	0.027
B	0.0	50	Exp	11.778	0.136	0.600	0.013	0.081	0.002
			Obs	12.432	0.158	11.842	0.165	0.842	0.016
B	0.5	10	Exp	22.437	0.133	1.640	0.027	0.179	0.004
			Obs	23.255	0.128	1.769	0.036	1.630	0.038
B	0.5	25	Exp	19.685	0.136	1.443	0.026	0.156	0.004
			Obs	20.507	0.128	3.818	0.084	1.525	0.027
B	0.5	50	Exp	17.536	0.089	0.906	0.014	0.073	0.002
			Obs	18.472	0.085	5.972	0.035	1.053	0.019
B	0.7	10	Exp	23.023	0.125	1.640	0.026	0.175	0.004
			Obs	23.858	0.119	1.713	0.025	1.241	0.043
B	0.7	25	Exp	21.008	0.126	1.437	0.024	0.145	0.003
			Obs	21.850	0.122	2.267	0.048	1.526	0.027
B	0.7	50	Exp	19.660	0.103	1.164	0.019	0.093	0.002
			Obs	20.564	0.101	3.625	0.025	1.326	0.021
R	-	0	Exp	25.874	0.128	3.424	0.040	0.351	0.006
			Obs	25.899	0.132	3.178	0.036	0.628	0.007
R	0.0	10	Exp	23.292	0.155	2.808	0.040	0.249	0.006
			Obs	23.547	0.160	2.829	0.035	2.531	0.113
R	0.0	25	Exp	19.330	0.174	2.354	0.039	0.193	0.005
			Obs	19.561	0.182	6.121	0.149	2.301	0.035
R	0.0	50	Exp	12.078	0.142	1.436	0.031	0.152	0.004
			Obs	12.735	0.187	12.045	0.175	1.751	0.037
R	0.5	10	Exp	23.389	0.141	2.726	0.038	0.281	0.006
			Obs	23.835	0.138	2.726	0.034	1.876	0.078
R	0.5	25	Exp	21.181	0.152	2.449	0.038	0.249	0.006
			Obs	21.629	0.143	4.251	0.084	2.415	0.035
R	0.5	50	Exp	19.063	0.107	1.763	0.023	0.128	0.004
			Obs	19.636	0.101	6.263	0.037	1.906	0.025
R	0.7	10	Exp	24.145	0.136	2.839	0.038	0.291	0.006
			Obs	24.552	0.132	2.8[illegible]5	0.033	1.166	0.043
R	0.7	25	Exp	22.627	0.138	2.605	0.037	0.223	0.005
			Obs	23.006	0.132	3.123	0.040	2.252	0.033
R	0.7	50	Exp	21.504	0.118	2.292	0.030	0.166	0.004
			Obs	21.970	0.116	3.843	0.027	2.357	0.030

Notes:
All results based on n=80 items and N=2000 respondents.
Exp = Derived from matrix of expected correlations
Obs = Derived from matrix of observed correlations.

Table 4:

Revised Modified Parallel Analysis (RMPA) of 20 tests:
Ratio of means and variances of eigenvalues of the jacknifed submatrices
(Ratio = observed / expected)

Rep	r	p	Eigenvalue 1		Eigenvalue 2		Eigenvalue 3	
			Mean	Var	Mean	Var	Mean	Var
B	-	0	1.031	0.933	0.995	0.884	4.030	6.007
B	0.0	10	1.031	1.034	1.574	15.449	9.730	44.102
B	0.0	25	1.036	1.074	3.917	28.726	10.187	52.047
B	0.0	50	1.056	1.355	19.748	151.159	10.350	83.432
B	0.5	10	1.036	0.920	1.078	1.779	9.084	80.253
B	0.5	25	1.042	0.885	2.646	10.454	9.763	44.765
B	0.5	50	1.053	0.909	6.594	5.886	14.379	133.295
B	0.7	10	1.036	0.916	1.045	0.908	7.086	116.778
B	0.7	25	1.040	0.934	1.578	3.836	10.524	62.910
B	0.7	50	1.046	0.955	3.114	1.681	14.278	92.026
R	-	0	1.001	1.054	0.928	0.804	1.790	1.296
R	0.0	10	1.011	1.067	1.008	0.776	10.183	370.061
R	0.0	25	1.012	1.094	2.600	14.554	11.918	53.081
R	0.0	50	1.054	1.746	8.387	31.775	11.527	78.033
R	0.5	10	1.019	0.946	1.000	0.786	6.685	165.024
R	0.5	25	1.021	0.888	1.736	4.773	9.705	36.368
R	0.5	50	1.030	0.892	3.554	2.554	14.892	45.144
R	0.7	10	1.017	0.944	0.999	0.770	4.004	44.388
R	0.7	25	1.017	0.916	1.199	1.205	10.082	42.093
R	0.7	50	1.022	0.966	1.677	0.763	14.166	68.059

Notes:
All results based on n=80 items and N=2000 respondents.

Table 5:

Revised Modified Parallel Analysis (RMPA) of 20 tests:
Correlations of eigenvalues of the observed and the expected jacknifed submatrices

Rep	r	p	Ev 1	Ev 2	Ev 3		Rep	r	p	Ev 1	Ev 2	Ev 3
B	-	0	0.996	0.888	0.650		R	-	0	0.976	0.956	0.195
B	0.0	10	0.995	-0.237	0.603		R	0.0	10	0.995	0.963	-0.147
B	0.0	25	0.996	-0.337	0.655		R	0.0	25	0.996	-0.414	0.306
B	0.0	50	0.981	-0.459	-0.049		R	0.0	50	-0.641	0.537	0.321
B	0.5	10	0.997	-0.256	0.299		R	0.5	10	0.998	0.968	-0.129
B	0.5	25	0.998	-0.320	0.593		R	0.5	25	0.996	-0.320	0.394
B	0.5	50	0.990	-0.180	0.310		R	0.5	50	0.996	-0.007	0.218
B	0.7	10	0.997	0.863	-0.111		R	0.7	10	0.998	0.968	-0.111
B	0.7	25	0.996	-0.311	0.608		R	0.7	25	0.995	0.350	0.140
B	0.7	50	0.997	-0.206	0.508		R	0.7	50	0.996	0.060	0.127

Notes:
All results based on n=80 items and N=2000 respondents.
Ev = Eigenvalue

Table 6:

Revised Modified Parallel Analysis (RMPA) of 20 tests: Six detection thresholds based on the distribution of the standardized Weighted Gaps (SWGs) based on the ratio of the first observed and expected jacknifed eigenvalues

			SWG		Threshold					
					95 %			99 %		
Rep	r	p	Mean	S.D.	Emp	UChe	Cheb	Emp	UChe	Cheb
B	-	0	0.94	0.49	1.89	2.42	3.17	2.38	4.24	5.88
B	0.0	10	0.91	0.58	1.94	2.67	3.54	2.58	4.81	6.76
B	0.0	25	0.97	0.52	1.96	2.54	3.32	2.87	4.45	6.19
B	0.0	50	0.84	0.56	2.07	2.53	3.37	3.09	4.59	6.46
B	0.5	10	0.90	0.42	1.61	2.15	2.77	1.92	3.67	5.06
B	0.5	25	0.96	0.53	2.01	2.56	3.36	2.93	4.52	6.30
B	0.5	50	0.97	0.51	1.84	2.50	3.26	2.26	4.37	6.07
B	0.7	10	0.98	0.50	1.89	2.48	3.23	2.80	4.31	5.97
B	0.7	25	0.95	0.57	2.13	2.65	3.50	2.66	4.73	6.61
B	0.7	50	0.89	0.52	1.92	2.45	3.23	2.13	4.35	6.08
B	Mean		0.93	0.52	1.93	2.49	3.27	2.56	4.40	6.14
R	-	0	0.95	0.44	1.65	2.27	2.93	2.12	3.89	5.36
R	0.0	10	0.99	0.56	1.85	2.66	3.50	3.40	4.72	6.58
R	0.0	25	0.94	0.54	1.90	2.56	3.37	2.65	4.54	6.33
R	0.0	50	0.79	0.51	1.77	2.30	3.06	2.47	4.16	5.84
R	0.5	10	0.99	0.48	1.82	2.42	3.14	1.93	4.18	5.78
R	0.5	25	0.95	0.52	2.03	2.52	3.31	2.24	4.44	6.18
R	0.5	50	0.93	0.53	1.82	2.53	3.33	3.00	4.48	6.25
R	0.7	10	1.07	0.61	2.16	2.90	3.81	2.87	5.12	7.14
R	0.7	25	0.92	0.45	1.95	2.28	2.95	2.23	3.93	5.43
R	0.7	50	0.96	0.47	1.85	2.38	3.09	2.03	4.12	5.70
R	Mean		0.95	0.51	1.88	2.48	3.25	2.49	4.35	6.06
Mean			0.94	0.52	1.90	2.49	3.26	2.53	4.38	6.10

Notes:

All results based on n=80 items and N=2000 respondents.

Emp = Empirical distribution

UChe = Chebyshev bound assuming unimodality

Cheb = Chebyshev bound

Table 7:

Revised Modified Parallel Analysis (RMPA):
Proportion of Standardized Weighted Gaps (SWGs) exceeding each of the six thresholds in the uncontaminated unidimensional test

	Threshold					
	95 %			99 %		
Eigenvalue	Emp	UChe	Cheb	Emp	UChe	Cheb
1	0.051	0.000	0.000	0.013	0.000	0.000
2	0.177	0.063	0.019	0.070	0.006	0.000
3	0.215	0.108	0.038	0.120	0.006	0.000
Mean	0.148	0.057	0.019	0.068	0.004	0.000

Notes:
All results based on n=80 items and N=2000 respondents.
Emp = Empirical distribution
UChe = Chebyshev bound assuming unimodality
Cheb = Chebyshev bound

Table 8a:

Revised Modified Parallel Analysis (RMPA) of 20 tests: Maximal Standardized Weighted Gap (SWG) and significance according to three types of thresholds (First eigenvalue)

					Significance*		No. of items	
Rep	r	p	Gap	Max(SWG)	95%	99%	Below	Above
B	-	0	0.00007800	2.37665	1	1	40	40
B	0.0	10	0.00011536	2.58303	1	1	29	51
B	0.0	25	0.00017863	2.87133	2	1	52	28
B	0.0	50	0.00074733	3.09437	2	1	46	34
B	0.5	10	0.00007026	1.92493	1	1	42	38
B	0.5	25	0.00023166	2.93434	2	1	20	60
B	0.5	50	0.00011191	2.25905	1	1	49	31
B	0.7	10	0.00011347	2.80103	2	1	35	45
B	0.7	25	0.00059257	2.66357	2	1	76	4
B	0.7	50	0.00035381	2.12819	1	1	76	4
R	-	0	0.00013665	2.12450	1	1	30	50
R	0.0	10	0.00083803	3.39892	2	1	4	76
R	0.0	25	0.00019422	2.65083	2	1	22	58
R	0.0	50	0.00513373	2.47295	2	1	40	40
R	0.5	10	0.00004417	1.92817	1	1	33	47
R	0.5	25	0.00010552	2.24198	1	1	38	42
R	0.5	50	0.00017510	2.99795	2	1	28	52
R	0.7	10	0.00010866	2.87002	1	1	16	64
R	0.7	25	0.00009185	2.22529	1	1	33	47
R	0.7	50	0.00017596	2.02883	1	1	7	73

**Note:*

1 --> Max(SWG) > Empirical percentile

2 --> Max(SWG) > Chebyshev + unimodality

3 --> Max(SWG) > Chebyshev

Table 8b:

Revised Modified Parallel Analysis (RMPA) of 20 tests: Maximal Standardized Weighted Gap (SWG) and significance according to three types of thresholds (Second eigenvalue)

					Significance*		No. of items	
Rep	r	p	Gap	Max(SWG)	95%	99%	Below	Above
B	-	0	0.002948	4.6028	3	2	13	67
B	0.0	10	0.154001	10.3842	3	3	72	8*
B	0.0	25	0.058056	5.5693	3	2	61	19*
B	0.0	50	0.420034	2.9129	2	0	6	74
B	0.5	10	0.061734	7.9202	3	3	72	8*
B	0.5	25	0.020733	3.9763	3	1	62	18*
B	0.5	50	0.016864	2.7078	2	1	30	50
B	0.7	10	0.017799	3.6832	3	1	78	2
B	0.7	25	0.017849	4.2707	3	1	64	16
B	0.7	50	0.061186	2.6536	2	1	3	77
R	-	0	0.002074	2.6308	2	1	4	76
R	0.0	10	0.001358	3.1794	2	0	10	70
R	0.0	25	0.041428	5.1730	3	2	60	20*
R	0.0	50	0.063916	7.3824	3	3	15	65
R	0.5	10	0.002657	2.7704	2	1	4	76
R	0.5	25	0.032784	6.7929	3	3	61	19*
R	0.5	50	0.027717	3.0870	2	1	7	73
R	0.7	10	0.000493	2.7074	1	0	23	57
R	0.7	25	0.005675	2.7842	2	1	71	9
R	0.7	50	0.005415	2.7835	2	1	19	61

**Note:*

1 --> **Max(SWG) > Empirical percentile**

2 --> **Max(SWG) > Chebyshev + unimodality**

3 --> **Max(SWG) > Chebyshev**

Table 8c:

Revised Modified Parallel Analysis (RMPA) of 20 tests: Maximal Standardized Weighted Gap (SWG) and significance according to three types of thresholds (Third eigenvalue)

					Significance*		No. of items	
Rep	r	p	Gap	Max(SWG)	95%	99%	Below	Above
B	-	0	0.197155	5.0647	3	2	2	78
B	0.0	10	0.407039	3.5361	2	1	1	79
B	0.0	25	0.087976	4.3483	3	1	9	71
B	0.0	50	0.065372	3.5205	3	1	12	68
B	0.5	10	0.110290	3.4420	3	1	73	7
B	0.5	25	0.138565	4.4914	3	1	6	74
B	0.5	50	0.237589	5.5091	3	2	9	71
B	0.7	10	0.184939	7.1549	3	3	71	9*
B	0.7	25	0.140215	5.1461	3	2	8	72
B	0.7	50	0.116452	3.3573	3	1	7	73
R	-	0	0.013257	3.6044	3	1	9	71
R	0.0	10	0.419092	9.4241	3	3	72	8*
R	0.0	25	0.115029	4.1520	3	1	10	70
R	0.0	50	0.225673	8.3301	3	3	10	70
R	0.5	10	0.287058	10.3919	3	3	72	8*
R	0.5	25	0.118107	5.3490	3	2	12	68
R	0.5	50	0.464085	4.3802	3	1	3	77
R	0.7	10	0.111553	6.4281	3	2	72	8*
R	0.7	25	0.072746	3.5625	3	1	11	69
R	0.7	50	0.272261	4.7821	3	2	7	73

**Note:*

1 --> Max(SWG) > Empirical percentile

2 --> Max(SWG) > Chebyshev + unimodality

3 --> Max(SWG) > Chebyshev

Table 9:

Revised Modified Parallel Analysis (RMPA) of 10 short tests:
Total number of items eliminated and accuracy of the elimination procedure

			Items eliminated			
Rep	r	p	Total	% of "hits"	% of "false alarms"	Significant Eigenvalue
B	0.0	10	8	100	0	2
B	0.5	10	8	100	0	2
B	0.7	10	9	100	1	3
R	0.0	10	8	100	0	3
R	0.5	10	8	100	0	3
R	0.7	10	8	100	0	3
Mean			8.2	100	0.2	-
B	0.0	25	19	95	0	2
*B	0.5	25	18	90	0	2
R	0.0	25	20	100	0	2
R	0.5	25	19	95	0	2
Mean			19	95	0	
Mean			-	98	0.1	

Note:
Tests shortened by 99% criterion
* These tests shortened by a 95% criterion

Table 10:

Modified Parallel Analysis (MPA) of 10 short tests:
The first three eigenvalues for the observed and expected matrices, and their ratios

			Eigenvalue 1			Eigenvalue 2			Eigenvalue 3		
Rep	r	p	Exp	Obs	Obs/Exp	Exp	Obs	Obs/Exp	Exp	Obs	Obs/Exp
B	0.0	10	21.86	22.71	1.04	1.62	1.65	1.02	0.17	0.62	3.64
B	0.0	25	17.89	18.78	1.05	1.42	1.48	1.05	0.14	0.56	3.86
B	0.5	10	21.95	22.71	1.03	1.65	1.65	1.00	0.18	0.62	3.51
B	0.5	25	18.00	18.85	1.05	1.40	1.49	1.06	0.15	0.58	3.89
B	0.7	10	21.49	22.27	1.04	1.58	1.61	1.01	0.17	0.62	3.63
R	0.0	10	23.35	23.84	1.02	2.84	2.85	1.01	0.25	0.54	2.19
R	0.0	25	19.17	19.79	1.03	2.37	2.30	0.98	0.19	0.50	2.63
R	0.5	10	22.79	23.21	1.02	2.70	2.71	1.00	0.28	0.58	2.10
R	0.5	25	19.27	19.73	1.02	2.33	2.45	1.05	0.24	0.53	2.21
R	0.7	10	22.83	23.21	1.02	2.73	2.71	0.99	0.28	0.58	2.04

Notes:
All results based on N=2000 respondents, and various number of items.
Exp = Derived from matrix of expected correlations
Obs = Derived from matrix of observed correlations.

Table 11:

Revised Modified Parallel Analysis (RMPA) of 10 short tests:
Ratio of means and variances of eigenvalues of the jacknifed submatrices
(Ratio = observed / expected)

			Eigenvalue 1		Eigenvalue 2		Eigenvalue 3	
Rep	r	p	Mean	Var	Mean	Var	Mean	Var
B	0.0	10	1.039	0.928	1.023	1.066	3.660	5.473
B	0.0	25	1.049	0.923	1.046	1.032	3.889	4.006
B	0.5	10	1.035	0.922	1.005	0.983	3.526	5.112
B	0.5	25	1.047	0 959	1.064	1.065	3.913	6.538
B	0.7	10	1.036	0.919	1.015	1.072	3.641	5.785
R	0.0	10	1.021	0.975	1.005	0.816	2.202	1.152
R	0.0	25	1.033	0.947	0.976	0.760	2.648	1.670
R	0.5	10	1.018	0.960	1.001	0.785	2.107	1.299
R	0.5	25	1.024	0.997	1.054	0.863	2.227	1.418
R	0.7	10	1.017	0.949	0.991	0.772	2.047	1.102