Robustness of Marginal Maximum Likelihood Estimation in the Rasch Model

Aeilko H. Zwinderman and Arnold L. van den Wollenberg
University of Nijmegen

Simulation studies examined the effect of misspecification of the latent ability (θ) distribution on the accuracy and efficiency of marginal maximum likelihood (MML) item parameter estimates and on MML statistics to test sufficiency and conditional independence. Results were compared to the conditional maximum likelihood (CML) approach. Results showed that if θ is assumed to be normally distributed when its distribution is actually skewed, MML estimators lose accuracy and efficiency when compared to CML estimators. The effects are not large, though they increase as the skewness of the number-correct score distribution increases. However, statistics to test the sufficiency and conditional independence assumptions of the Rasch model in the MML approach are very sensitive to misspecification of the θ distribution. *Index terms: ability distribution, conditional likelihood, efficiency, goodness of fit, marginal likelihood, Rasch model, robustness.*

Estimation of the item parameters of the Rasch model and of the parameters of the latent ability (θ) distribution by means of marginal maximum likelihood (MML; see Bock & Aitkin, 1981; Rigdon & Tsutakawa, 1983) is considered to be equivalent to estimation of the item parameters by means of conditional maximum likelihood (CML) and estimation of the parameters of the θ distribution according to methods described by Andersen and Madsen (1977; see Thissen, 1982).

In simulation studies, CML estimates of the item parameters were found to be equal to MML estimates for all practical purposes (see Thissen, 1982). However, this is only true when the Rasch model holds and when the assumed θ distribution is valid.

In the conditional approach, the item parameters and the parameters of the θ distribution are estimated separately. The item parameters are estimated by means of the conditional likelihood of the item responses given number-correct scores. These conditional estimators are consistent and asymptotically normally distributed, with an asymptotic covariance matrix that can be estimated from the information matrix. The parameters of the θ distribution are estimated by means of the population likelihood as defined by Andersen and Madsen (1977), given the conditionally estimated item parameters and given some assumed density function of the θ distribution.

In contrast to the conditional approach, the item parameters and the parameters of the θ distribution are estimated simultaneously in the marginal approach. In the marginal approach, the person parameters are considered to be independent identically distributed random variables with density $\phi(\theta)$. The item

APPLIED PSYCHOLOGICAL MEASUREMENT
Vol. 14, No. 1, March 1990, pp. 73–81

0146-6216/90/010073-09$1.70

parameters and the parameters of $\phi(\theta)$ are estimated by means of the marginal likelihood (see Rigdon & Tsutakawa, 1983). Although a rigorous proof is unknown to the present authors, the marginal estimates are consistent and are asymptotically normally distributed (see Bock & Aitkin, 1981).

The marginal approach has some advantages. It can be used to estimate the parameters of the two- and three-parameter models (among others), whereas the conditional approach can only be used within the Rasch model. Furthermore, the marginal approach can easily be extended to estimate the parameters in designs with structurally empty cells. Moreover, the estimation of the person parameters is handled elegantly. Instead of (unstable) point estimates, posterior distributions can be obtained.

The marginal approach is, in principle, more efficient than the conditional procedure. The standard errors (SEs) of MML estimators of the item parameters are smaller than the SEs of CML estimators. This is because the information about the item parameters contained in the density of the conditioning statistic is not used in the conditional approach. The number-correct scores are not ancillary (see Rao, 1975) to the item parameters (Andersen, 1970). In contrast to CML estimates, MML estimates are fully efficient. Of course, this is only true when the model holds.

A serious disadvantage of the marginal approach is its dependence on the assumed θ distribution. Although the distribution can be of any type, provided it has a finite number of parameters, normal densities are ordinarily used (see Bock & Aitkin, 1981; Bock & Lieberman, 1970; Mislevy, 1984; Rigdon & Tsutakawa, 1983; Thissen, 1982). If such an assumed θ distribution is invalid, it will affect both the accuracy and the efficiency of the estimates. Moreover, it will affect the statistics used to test the fit of the Rasch model.

As the item parameters and the parameters of the θ distribution are estimated separately in the conditional approach, the assumptions pertaining to the Rasch model can also be tested independently from testing the assumption pertaining to the distribution of θ. In the conditional approach, the estimation of the item parameters is independent of the actual values of the person parameters, and hence of the distribution of θ. The statistics used within the conditional framework are also independent of the person parameters, as they are functions of the estimated item parameters only; hence they can be used to test the specific assumptions of the Rasch model independently of any assumption concerning the θ distribution.

Several statistics have been developed to test the assumptions of the Rasch model within the conditional approach (Andersen, 1973; Glas, 1988; Martin-Löf, 1973; van den Wollenberg, 1982). The T statistic developed by Martin-Löf (1973) is a Pearson chi-square statistic comparing the observed and expected number of persons responding correctly to an item given the number-correct score:

$$T = \sum_{r=1}^{k-1} \mathbf{d}_r' \mathbf{V}_r^{-1} \mathbf{d}_r \quad , \tag{1}$$

where k is the number of items,
r is the index of number-correct score groups,
$\mathbf{d}_r$ is a vector of length k with elements $n_{ri} - n_r\hat{\pi}_{i.r}$,
n_{ri} is the observed number of persons with score r who respond correctly to item i,
n_r is the observed number of persons with score r, and
$\hat{\pi}_{i.r}$ is the estimated conditional probability of responding correctly to item i conditional on score r.

$\mathbf{V}_r$ is a matrix of order $k \times k$ with elements

$$\mathbf{V}_r(i,j) = \begin{cases} n_r\hat{\pi}_{i.r} & (i = j;\ i,j = 1, \ldots, k) \\ n_r\hat{\pi}_{ij.r} & (i \neq j) \end{cases} \tag{2}$$

with

$$\hat{\pi}_{i.r} = \hat{\varepsilon}_i \frac{\gamma_{r-1}^{(i)}(\hat{\boldsymbol{\varepsilon}})}{\gamma_r(\hat{\boldsymbol{\varepsilon}})} \tag{3}$$

and

$$\hat{\pi}_{ij.r} = \hat{\varepsilon}_i \hat{\varepsilon}_j \frac{\gamma_{r-2}^{(i,j)}(\hat{\boldsymbol{\varepsilon}})}{\gamma_r(\hat{\boldsymbol{\varepsilon}})} \quad , \tag{4}$$

where $\hat{\boldsymbol{\varepsilon}}$ is the vector of item easiness parameters,
γ_r is the basic symmetric function of $\hat{\boldsymbol{\varepsilon}}$ of order r, and
$\gamma_{r-1}^{(i)}$ and $\gamma_{r-2}^{(i,j)}$ are the first-order and second-order partial derivatives of γ_r with respect to ε_i and ε_j, respectively. (For a detailed discussion of the basic symmetric functions, see Verhelst, Glas, & van der Sluis, 1984.)

T is asymptotically chi-square distributed with $(k - 1)(k - 2)$ degrees of freedom if the number of persons with score r is large for all r. If the latter is not the case, adjacent score groups can be grouped together. T is especially sensitive to violation of the sufficiency assumption in the Rasch model (see van den Wollenberg, 1982).

In the marginal approach, the item parameters and the parameters of the θ distribution are estimated simultaneously. This means that the assumptions pertaining to the Rasch model cannot be tested independently of testing the assumption pertaining to the θ distribution. Glas (1988) developed two statistics for the marginal approach for testing the assumptions pertaining to the Rasch model (i.e., the sufficiency assumption and the assumption of conditional independence). His first statistic, R_{1m}, is comparable to the statistic developed by Martin-Löf and is meant to test the sufficiency assumption of the Rasch model:

$$R_{1m} = \frac{d_0^2}{\mathrm{E}(N_0)} + \sum_{r=1}^{k-1} \mathbf{d}_r' \mathbf{V}_r^{-1} \mathbf{d}_r + \frac{d_k^2}{\mathrm{E}(N_k)} \quad , \tag{5}$$

where $\mathbf{d}_r$ is a vector of length k with elements $n_{ri} - N\hat{\pi}_{ri}$,
N is the total number of persons in the sample,
$\hat{\pi}_{ri}$ is the estimated marginal probability of score r and a correct response on item i,
$\mathrm{E}(N_0)$ is the expected number of persons with zero number-correct score,
d_0 is the deviation between the observed and the expected number of persons with zero score,
$\mathrm{E}(N_k)$ is the expected number of persons with a perfect score, and
d_k is the corresponding deviation between the observed and the expected number of persons with a perfect score.

$\mathbf{V}_r$ is a matrix of order $k \times k$ with elements defined as for T evaluated with the estimated marginal probabilities:

$$\mathbf{V}_r(i,j) = \begin{cases} N\hat{\pi}_{ri} & (i = j;\, i,j = 1, \ldots, k) \\ N\hat{\pi}_{rij} & (i \neq j) \end{cases} \tag{6}$$

with

$$\hat{\pi}_{ri} = \int_{-\infty}^{+\infty} \frac{\hat{\varepsilon}_i \gamma_{r-1}^{(i)}(\hat{\boldsymbol{\varepsilon}}) \exp(r\theta)}{\prod_{i=1}^{k} 1 + \exp(\theta - \hat{\alpha}_i)} \phi(\theta|\hat{\mu},\hat{\sigma}^2)\, d\theta \tag{7}$$

and

$$\hat{\pi}_{rij} = \int_{-\infty}^{+\infty} \frac{\hat{\varepsilon}_i \hat{\varepsilon}_j \gamma_{r-2}^{(i,j)}(\hat{\varepsilon}) \exp(r\theta)}{\prod_{i=1}^{k} 1 + \exp(\theta - \hat{\alpha}_i)} \phi(\theta|\hat{\mu},\hat{\sigma}^2)\, d\theta \quad , \tag{8}$$

where $\phi(\theta|\hat{\mu},\hat{\sigma}^2)$ is the density function of θ with estimated mean $\hat{\mu}$ and variance $\hat{\sigma}^2$,
$\hat{\alpha}_i$ is the estimated item difficulty, and
$\varepsilon_i = \exp(\alpha_i)$.

R_{1m} is asymptotically chi-square distributed with $k(k - 2)$ degrees of freedom.

In the marginal approach, statistics to test the Rasch model assumptions depend on the assumed θ distribution. The θ distribution can be handled in a number of ways. In general, the integral involved in the marginal likelihood can be evaluated to any degree of accuracy for any reasonable θ distribution. One approach is to use Simpson's rule, which requires the specification of a number of suitably chosen quadrature points. In practice, normal θ distributions are usually assumed, in which case Gauss-Hermite quadrature points and weights are used to evaluate the integral involved in the marginal likelihood. However, in general the θ distribution is completely unknown in practice.

In the Gauss-Hermite approach, the points and weights are fixed. Bock and Aitkin (1981) introduced an approach in which the quadrature weights are estimated by means of maximum likelihood. The weights represent the height of the density curve at the corresponding quadrature points and are updated in each iteration. This represents a nonparametric approach to estimating the θ distribution. However, such a procedure does not necessarily lead to convergence (see Mislevy & Bock, 1984), and in the present authors' experience convergence frequently does not occur.

In most cases, application of the marginal estimation approach will be restricted to assuming a normal θ distribution. Usually, a normal distribution seems to be a reasonable approximation. However, if a test is used to select a small group of low- or high-θ persons from some population, items will be used that have easiness parameters in the vicinity of the cutoff point. In such cases the distribution, given the items, will be skewed.

This study addressed the question of how much the accuracy and the efficiency of MML item parameter estimates are affected by assuming a normal θ distribution when the actual θ distribution is skewed. The effect of assuming an invalid θ distribution on the R_{1m} statistic was also examined. The results are compared to the conditional approach.

Method

Data Generation

A number of data matrices were generated containing dichotomous responses of 4,000 simulees on a varying number of items. In the generation program, item parameters were input as fixed parameters (input parameters). By means of IMSL (1982) routines GGUBS and GGEXN, simulee parameters were sampled from the standard normal or exponential distributions. The exponential distribution is skewed depending on its mean (m). The probability density function is $f(x) = 1/m \exp(x/m)$, with $0 \leq x < \infty$ and $0 < m < \infty$, and 0 otherwise.

By means of these simulee and item parameters, response probabilities were obtained according to the basic formula of the Rasch model. These response probabilities were converted into binary manifest responses by means of random numbers sampled from a uniform distribution with domain (0,1). For each simulee-item combination, such a random number was obtained by the IMSL routine GGUBS. A manifest response was considered positive when the probability exceeded that random number and negative otherwise.

CML and MML estimates of the item parameters were first compared with respect to accuracy and efficiency for the case in which θ parameters were sampled from the standard normal distribution [N(0,1)] and θ was assumed to be normally distributed [$N(\mu,\sigma^2)$]. The number of items varied (5, 10, or 15), and the item parameters were spaced equidistantly over the ranges (-3,3) or (-1,1) with mean item difficulty 0. In this case, neither the assumptions pertaining to the Rasch model nor the assumption pertaining to the θ distribution were violated.

Both estimation methods were then compared for the case in which simulee parameters were sampled from exponential distributions, while in the MML approach a normal distribution [$N(\mu,\sigma^2)$] was assumed. Depending on the item parameters, sampling from exponential distributions will result in skewed number-correct score distributions, with a larger skewness corresponding to a larger mean of the exponential sampling distribution. Samples were drawn from three exponential distributions with mean equal to 1, 5, or 10 respectively. The resulting number-correct score distributions were lightly, moderately, and extremely skewed respectively. Again, the number of items varied (5, 10, or 15) and the item difficulties were equidistantly distributed over the ranges (-3,3) or (-1,1) with mean item difficulty 0. In this case, the assumptions pertaining to the Rasch model were not violated but the assumptions pertaining to the θ distribution were violated.

The range (-3,3) over which the item parameters were distributed represents a rather extreme case. Usually, the parameters will be distributed over a smaller range. The simulations were thus repeated for item parameters distributed over the range (-1,1).

Analysis

For each data matrix, conditional and marginal estimates of the item parameters were obtained and the SEs were computed. The Martin-Löf (1973) T test and the Glas (1988) R_{1m} test were also computed. The CML item parameters were obtained using the computer program PML (Gustafsson, 1979); the MML item parameter estimates were obtained simultaneously with the estimates of the mean and the variance of the θ distribution using the computer program RIDA (Glas, 1989).

The mean of the input item parameters was 0. In order to express the estimates of the item parameters on the same scale as the input item parameters, the estimates of the item parameters were normed such that their means also equaled 0. SEs were computed as the square root of the inverse of the estimated information matrix. The simulations were replicated 50 times. The accuracy of the item parameter estimates was determined by means of the square root of the mean squared difference (RMSD) between the input parameters and the estimates averaged over the replications. The RMSDs can be interpreted as the true SEs of the item parameters:

$$\mathrm{RMSD} = \left[\sum_s (\alpha_i - \hat{\alpha}_{is})^2 \right]^{1/2} , \tag{9}$$

where α_i is the input or true parameter of item i, and $\hat{\alpha}_{is}$ is the corresponding estimate in replication s ($s = 1, \ldots, 50$). The RMSDs were averaged over all items in the matrix.

Possible bias in the estimates cannot be determined directly from the RMSD. Bias of the estimates is contaminated with the efficiency of the estimates. If the estimates are biased, the index will be large, which is also true if the variance of the estimates is large. However, if the RMSDs are compared with the

SEs of the item parameter estimates (estimated from the information matrix), then bias can be inferred from this comparison.

Results

Normal Ability Distribution

Table 1 gives the RMSDs between the input parameters and the CML and the MML estimates of the item parameters, and the estimated SEs of the item parameters averaged over replications and items. The RMSDs between the CML and MML estimates and the input parameters vary from .035 to .05. Although the RMSDs of the MML estimates were smaller than the RMSDs of the CML estimates in all cases, the differences are negligible. The RMSDs seem to be large, but the estimated SEs of the estimates are of the same order. In fact, the RMSDs between the estimates and the input parameters are almost totally explained by the estimated variances of the estimates. Hence the estimates seem to be unbiased.

Both the RMSDs and the estimated SEs of MML estimates are smaller than the RMSDs and the estimated SEs of CML estimates. This reflects the fact that in the conditional approach the information with respect to the item parameters, contained in the density of the conditioning statistic, is not used.

The RMSDs between CML and MML estimates are very small. Mean RMSDs are less than .01. Therefore, when simulee parameters are sampled from the standard normal distribution and a normal θ distribution is assumed in the MML method, MML and CML estimates of the item parameters are virtually equal. This was also true of Martin-Löf's (1973) T and Glas' (1988) R_{1m}. Based on these test statistics, two to three times out of the 50 replications, the model was rejected. In each case, both statistics rejected the model.

Skewed Ability Distribution

CML and MML were compared for simulees sampled from extremely, moderately and lightly skewed distributions, while a normal θ distribution was assumed in the marginal approach. Sampling from the exponential distribution (with varying mean parameter) yielded number-correct score distributions with (averaged) skewness of -2.01, $-.80$, and $-.15$ respectively. Figure 1 shows these three score distributions for a 15-item test graphed together with the score distribution resulting from sampling from the standard normal distribution.

Table 2 presents the RMSDs between the input parameters and the CML and MML estimates for item parameters distributed over the range $(-3,3)$, as well as the averaged SEs. These RMSDs vary considerably more than those found for the data with simulees sampled from a standard normal distribution. The RMSDs between MML estimates and input item parameters are as large as .20 to .30 for the extremely skewed

Table 1
RMSD Between CML and MML Item Parameter Estimates and the Input Parameters, Average Estimated Standard Errors of the Item Parameter Estimates, and RMSD Between CML and MML Estimates for Item Parameter Ranges of $\alpha = -1$ to 1 and -3 to 3, for Tests of $k = 5$, 10, and 15 Items

	CML				MML				RMSD(CML,MML)	
	$\alpha = -1,1$		$\alpha = -3,3$		$\alpha = -1,1$		$\alpha = -3,3$			
k	RMSD	SE	RMSD	SE	RMSD	SE	RMSD	SE	$\alpha=-1,1$	$\alpha=-3,3$
5	.036	.039	.050	.045	.036	.031	.048	.039	.0001	.0138
10	.036	.038	.046	.043	.036	.034	.045	.040	.0001	.0045
15	.035	.038	.042	.042	.035	.035	.042	.040	.0001	.0032

Figure 1
Score Distributions for a 15-Item Test

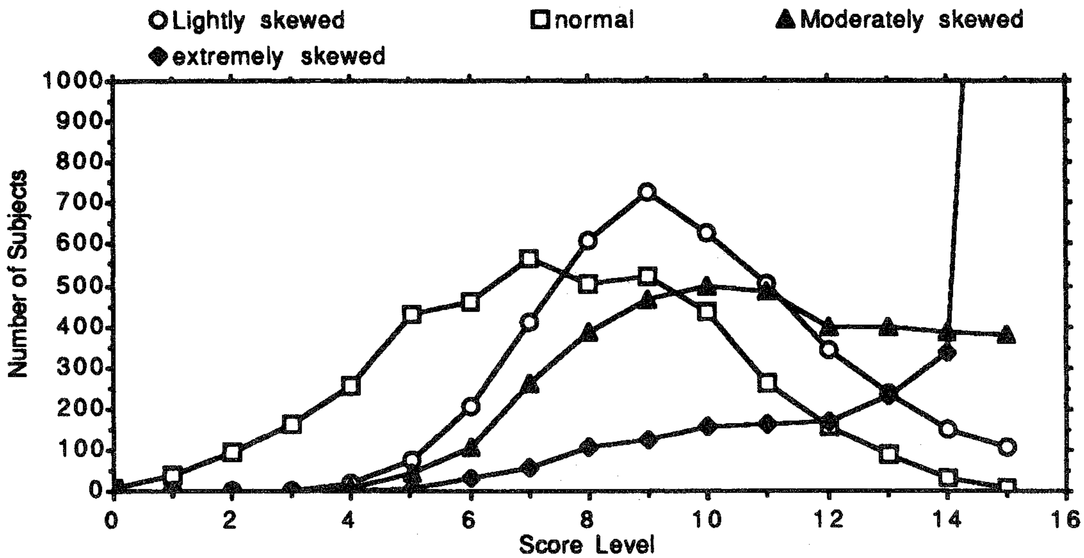

distribution. For the moderately and lightly skewed distributions, they are much smaller. The RMSDs between CML and MML estimates are of the same order of magnitude as the RMSDs between the MML estimates and the input item parameters.

Note that the RMSDs between the CML estimates and the input item parameters are rather large as well. This can be explained by the loss of information due to the large number of simulees giving a correct response to all items. This number varied from 75% of the total number of simulees for a five-item test with an extremely skewed distribution to only a few percent for a 15-item test with a lightly skewed distribution. This loss of information is reflected in the large estimated SEs of the estimates.

The large RMSDs for the MML estimates can be explained only partially by the large estimated SEs of the MML estimates. The RMSDs for the extremely skewed distributions are two to three times as large as the corresponding estimated SEs. Moreover, the estimated SEs of the MML estimates are larger than those for the CML estimates for the extremely skewed distributions. Hence the use of an invalid θ distribution leads to a loss of efficiency for the marginal estimates. Furthermore, the marginal estimates are biased whenever the discrepancy between the assumed and the actual θ distribution is large.

Based on the T statistic in the conditional approach, the model was rejected in two to three cases out of 50 replications for all simulations. However, based on the R_{1m} statistic in the marginal approach, the model was rejected in 50 cases out of 50 replications for all k and for all three skewed distributions. Obviously, R_{1m} is very sensitive to violation of the assumed θ distribution. Even when the skewness of the number-correct score distribution was only .15 (lightly skewed), R_{1m} showed a rejection of the model in 100% of all cases.

Narrow Range of Item Difficulty

The corresponding RMSDs and averaged SEs for this set of simulations are also given in Table 2. The RMSDs are much lower than obtained for the $(-3,3)$ range for cases with very easy and difficult

Table 2
RMSDs Between CML and MML Item Parameter Estimates and the Input Parameters, Average Estimated Standard Errors of the Item Parameter Estimates, and RMSD Between CML and MML Estimates for Extremely, Moderately, and Lightly Skewed Score Distributions, for Tests of $k = 5$, 10, and 15 Items, and for Item Parameter Distributions of $\alpha = -3$ to 3 and -1 to 1

	$\alpha = -3,3$						$\alpha = -1,1$					
	Extremely		Moderately		Lightly		Extremely		Moderately		Lightly	
k	RMSD	SE	RMSD	SE	RMSD	SE	RMSD	SE	RMSD	SE	RMSD	SE
CML												
5	.151	.108	.079	.058	.060	.052	.066	.088	.043	.048	.033	.042
10	.111	.100	.058	.055	.050	.048	.080	.085	.042	.046	.039	.040
15	.102	.097	.057	.054	.056	.047	.075	.083	.047	.046	.039	.040
MML												
5	.278	.122	.120	.056	.084	.050	.087	.076	.048	.037	.035	.034
10	.288	.110	.097	.054	.068	.047	.183	.084	.046	.041	.040	.037
15	.308	.110	.074	.052	.064	.046	.212	.087	.049	.043	.040	.038
RMSD(CML,MML)												
5	.0416		.0200		.0030		.0074		.0009		.0001	
10	.0705		.0101		.0015		.0228		.0008		.0003	
15	.0989		.0053		.0012		.0378		.0007		.0001	

items, although the MML estimates still show larger RMSDs than CML estimates. Moreover, the RMSDs between MML estimates and input parameters remain very large for the extremely skewed distributions.

The estimated SEs indicate that MML estimates still seem biased for this situation, although the bias is smaller. The RMSDs between CML estimates and input parameters are of the same order as their estimated SEs, whereas the RMSDs between MML estimates and input parameters are larger, especially for the extremely skewed distributions. Note that the estimated SEs of the MML estimates are smaller than those of the CML estimates.

In these simulations as well, the model was rejected by the T statistic in the conditional approach in two to three cases out of 50 replications. Based on the R_{1m} statistic in the marginal approach, the model was rejected in 100% of the replications for the extremely and moderately skewed distributions. For the lightly skewed distributions, the model was rejected in 80% of all replications.

Conclusions

The MML estimation method in the Rasch model has some attractive advantages over conditional estimation. The method can be applied in a variety of models and the estimates, in principle, are more efficient. However, these advantages are obtained at the cost of the additional assumption regarding the distribution of θ in the population. Violation of this assumption leads to a loss of efficiency and bias in the estimates. The bias and the loss of precision, relative to the conditional estimates, is not large when the violation is not too large.

More seriously, however, violation of the distribution assumption prohibits the testing of the assumptions of the Rasch model. Statistics developed for the marginal approach to test the sufficiency and conditional independence assumptions of the Rasch model depend very much on the validity of the assumed θ distribution. Even when the actual distribution is only lightly skewed, assuming a normal θ distribution will lead to rejection of the model based on these statistics.

Therefore, the use of R_{1m} and similar statistics to test the sufficiency and conditional independence assumptions of the Rasch model in the marginal framework requires that the assumed θ distribution has already been tested carefully. However, this requirement may be impossible to fulfill. Just as testing the assumptions of the Rasch model is not independent of testing the assumption on the θ distribution in the marginal framework, the reverse is also true. The assumption on the θ distribution is intertwined with the assumptions of the Rasch model in the marginal framework.

This is not so in the conditional framework. The assumptions of the Rasch model can be tested independently from any assumption on the distribution of θ. Any assumption on the θ distribution can be tested using methods suggested by Andersen and Madsen (1977). In their approach, the parameters of an assumed θ distribution are estimated given the conditionally estimated item parameters. The fit of the assumed θ distribution can be tested using the likelihood ratio of the observed numbers of examinees with score r ($r = 0, \ldots, k$) against the expected numbers under the assumed θ distribution (see Andersen & Madsen, 1977).

The Rasch model imposes strong restrictions on the data, and these deserve careful testing. Only the conditional approach enables testing of the assumptions of the Rasch model independently of the assumption of the θ distribution.

References

Andersen, E. B. (1970). Asymptotic properties of conditional maximum likelihood estimates. *Journal of the Royal Statistical Society, Series B, 32,* 283–301.

Andersen, E. B. (1973). A goodness of fit test for the Rasch Model. *Psychometrika, 38,* 123–140.

Andersen, E. B., & Madsen, M. (1977). Estimating the parameters of the latent population distribution. *Psychometrika, 42,* 357–374.

Bock, R. D., & Aitkin, M. (1981). Marginal maximum likelihood estimation of item parameters: Application of an EM algorithm. *Psychometrika, 46,* 443–459.

Bock, R. D., & Lieberman, M. (1970). Fitting a response model for *n* dichotomously scored items. *Psychometrika, 35,* 189–197.

Glas, C. A. W. (1988). The derivation of some tests for the Rasch model from the multinomial distribution. *Psychometrika, 53,* 525–546.

Glas, C. A. W. (1989). *RIDA, computer program for Rasch incomplete design analysis.* Arnhem, The Netherlands: National Institute for Educational Measurement (CITO).

Gustafsson, J. E. (1979). *PML: A computer program for the Rasch model* (Report No. 85). Gothenburg: University of Gothenburg.

IMSL, Inc. (1982). *IMSL reference manual.* Houston: Author.

Martin-Löf, P. (1973). *Statistical models* [College notes edited by R. Sundberg]. Stockholm: Institute of Mathematics and Mathematical Statistics, University of Stockholm. (In Swedish.)

Mislevy, R. J. (1984). Estimating latent distributions. *Psychometrika, 49,* 359–381.

Mislevy, R. J., & Bock, R. D. (1984). *BILOG II user's guide.* Mooresville IN: Scientific Software Inc.

Rao, C. R. (1975). *Linear statistical inference and its applications* (2nd ed.). New York: Wiley.

Rigdon, S. E., & Tsutakawa, R. K. (1983). Parameter estimation in latent trait models. *Psychometrika, 48,* 567–574.

Thissen, D. (1982). Marginal maximum likelihood estimation for the one parameter logistic model. *Psychometrika, 47,* 175–186.

van den Wollenberg, A. L. (1982). Two new test statistics for the Rasch model. *Psychometrika, 47,* 123–140.

Verhelst, N. D., Glas, C. A. W., & van der Sluis, A. (1984). Estimation problems in the Rasch model: The basic symmetric functions. *Computational Statistics Quarterly, 1,* 245–262.

Acknowledgments

This research was supported by Grant No. 56-267-004 from the Dutch Science Foundation (ZWO).

Author's Address

Send requests for reprints or further information to Aeilko H. Zwinderman, University of Nijmegen, Department of Psychology, P.O. 9104, 6500 HE Nijmegen, The Netherlands.