Estimation of the MIRID: A program and a SAS-based approach

DIRK J. M. SMITS and PAUL DE BOECK
Katholieke Universiteit Leuven, Leuven, Belgium
and

NORMAN D. VERHELST
National Institute of Educational Measurement (Cito), Arnhem, The Netherlands

The MIRID CML program is a program for the estimation of the parameter values of two different componential IRT models: the Rasch-MIRID and the OPLM-MIRID (Butter, 1994; Butter, De Boeck, & Verhelst, 1998). To estimate the parameters of both models, the program uses a CML approach. The model parameters can also be estimated with a MML approach that can be implemented in PROC NLMIXED of SAS Version 8. Both the MIRID CML program and the MML SAS approach are explained and compared in a simulation study. The results showed that they did about equally well in estimating the values of the item parameters but that there were some differences in the estimation of the person parameters, as could be expected from the differential assumptions regarding the distribution of the persons. The SAS MML approach is much slower than the MIRID CML program, but it is more flexible.

Butter, De Boeck, and Verhelst (1998) described the model with internal restrictions on item difficulties (MIRID) as a componential model for binary data. In the MIRID, parameters of some items are defined to be a linear combination of the parameters of other items. The model requires that two sets of items be defined: *component items* and *composite items*. A composite item is an item that measures a concept that can be decomposed into components. A component item is an item that measures one of these components. The item parameters of the composite items are decomposed into parts attributed to the component items (the item parameters of the component items). For example, $10 * (5 + 3)$ as a composite item has two component items: $5 + 3$ and $10 * 8$. The first component is of the addition type; the second is of the multiplication type. One can formulate many different items along the same line, using different basic numbers—for example, the composite item $7 * (6 + 8)$ and its component items $6 + 8$, $7 * 14$. The componential approach applies to the affective domain as well. For example, feeling guilty in a given situation may stem from feeling that a norm is violated, from a tendency to brood about what one did, and from a tendency toward restitution for what one did wrong, each component being related to the same given situation (Smits & De Boeck, 2003). The question “Do you feel that you have violated a moral, an ethical, a religious, and/or a personal code in situation A?” is a component item of the

norm violation type, and it is associated with the composite item “Do you feel guilty in situation A?” For the approach to work, for each composite item associated with a particular cognitive task or with an affective situation, a number of component items has to be formulated with respect to the same cognitive task or situation.

In general, J item families ($j = 1, \dots, J$) are defined so that, within each family, there is one composite item, to be conceived of as a dependent variable, and K component items, one for each of the K component types ($k = 1, \dots, K$), to be conceived of as the independent variables. For the composite items, k is set to zero. The total number of items is equal to $J * (K + 1)$.

Suppose we have a questionnaire with five item families and three types of components; then the total number of items equals $5 * [3 \text{ (number of components)} + 1 \text{ (composite item)}] = 20$. The structure of such a questionnaire is given in Table 1.

The crucial assumption of the MIRID is that the item parameter of a composite item is a linear function of the item parameters of the associated component items (Butter et al., 1998). The MIRID models the composite item parameters as a linear function of the component item parameters:

$$\beta_{j0} = \sum_{k=1}^K \sigma_k \beta_{jk} + \tau, \quad (1)$$

where β_{j0} is the item parameter of the composite item from item family j , β_{jk} is the item parameter of the component item of type k from item family j , σ_k is the weight of the component item parameters of type k in determining the composite item parameters, and τ is a normalization constant.

The research was financially supported by a GOA 2000/2 grant from the Katholieke Universiteit Leuven: “Psychometric models for the study of personality.” Correspondence concerning this article should be addressed to D. J. M. Smits, K.U. Leuven, Department of Psychology (H.C.I.V.), Tiensestraat 102, B-3000 Leuven, Belgium (e-mail: dirk.smits@psy.kuleuven.ac.be).

(see, e.g., Baker, 1992; Fischer & Molenaar, 1995; Verhelst, 1993) of the Rasch–MIRID and the OPLM–MIRID. Apart from this, SAS Version 8 also can be used for the estimation of the model parameters of these models, assuming a normal distribution for the person parameter (Wolfinger, 1999) and following a marginal maximum likelihood (MML) approach (see, e.g., Baker, 1992; Fischer & Molenaar, 1995; Verhelst, 1993). Both the MIRID CML program and the SAS procedure for MML will be explained next. In a final section, both approaches will be compared in a small simulation study.

THE MIRID CML PROGRAM

The MIRID CML program (Smits, De Boeck, Verhelst, & Butter, 2001) is a Windows-based program. It is written in Borland Delphi 5.0 and tested under Windows 95, 98, 2000, and NT 4.0. About 8 MB of free disk space is needed to install the program.

Model Estimation

Four models can be estimated with the MIRID CML program: the Rasch model, the Rasch–MIRID, the OPLM, and the OPLM–MIRID. The item parameter values and their standard errors are estimated using a CML approach and the Davidon–Fletcher–Powell technique (a quasi Newton–Raphson optimization technique; Bunday, 1984) or the Newton–Raphson optimization technique (Bunday, 1984; Gill, Murray, & Wright, 1981). Using a CML approach, the item parameters are estimated by conditioning upon the examinees' sufficient statistics (sum of a priori degrees of discrimination for succeeded items). Once the item pool is calibrated, the person parameter estimate corresponding to each sufficient statistic can easily be obtained. A weighted maximum likelihood estimation procedure was implemented for the estimation of the person parameters and their standard errors (Warm, 1989).

Input

Before starting the estimation, one needs to specify the name of the data file, the number of persons and the number of items in the data file, the discrimination values of the items, and the name of the output file. The data files need to be in plain text format. The requested structure of the data files is such that the rows are formed by the persons and the columns are formed by the items. All responses are typed one next to the other, without any spacing between them. The order of the items (columns) is the following: first, all component items of Component Type 1, ordered according to the item family they belong to; next, all component items of Component Type 2 in the same order, and so on; and finally, all composite items, again in the same order.

Since the data sets need to satisfy a rather rigid structure, a module is included in the program to rearrange data files with a different ordering.

Output

During the estimation process, a screen with the current value of the log-likelihood function of the model is shown, so that the state of convergence can be followed. After the estimation procedure has reached the convergence criterion, the output is automatically displayed in the built-in text editor. First, the estimated parameter values are shown: the item parameters of the component items (the β_{jk}) and the linear coefficients (the σ_k and τ), all with their standard errors. Person parameter estimates are optional, and when requested, they are followed by their standard errors in a separate section after the item parameter estimates. As was mentioned earlier, the program provides Warm (1989) estimates. Warm estimates of the person parameters can also be obtained, for example, by using another CML IRT program, such as LPCM–WIN (Fischer, Ponocny-Seliger, Ponocny, & Parzer, 1998) or OPLM (Verhelst et al., 1994).

Second, information is given about the fit of the estimated model. If the estimated model is a Rasch–MIRID or an OPLM–MIRID, the fit of this model is compared with the fit of the corresponding basic model (the Rasch model and the OPLM, respectively), using a likelihood-ratio test, since because of the MIRID principle, the MIRID variants are nested within the corresponding original model. A more specific test will be presented in the section on the SAS MML approach.

Simulation Module

The program also contains a module to simulate data. In addition, error can be added to an existing data set, as is explained in detail in the manual (Smits et al., 2001).

Availability

Two versions of the program are available: one for computers running Windows 95 or Windows NT 4.0 and one for computers running Windows 98 or Windows 2000. Except for some animations, the two versions are equivalent. The program can be obtained by e-mailing the author (miridprogram@hotmail.com) or by sending two 3.5-in. high-density diskettes and a self-addressed stamped diskette mailer to Dirk Smits, Department of Psychology, Tiensestraat 102, B-3000 Leuven, Belgium.

The MIRID CML program comes with a manual in a PDF file.

THE SAS MML APPROACH

The parameters of the Rasch model, the Rasch–MIRID, the OPLM, and the OPLM–MIRID can also be estimated using SAS Version 8. For a discussion on how to use SAS for IRT models, see Rijmen, Tuerlinckx, De Boeck, and Kuppens (2003). The SAS software package includes a procedure, called PROC NL MIXED, to fit nonlinear mixed models. Nonlinear mixed models are regression models that are nonlinear in the predictors—for example, because of a logit link—and that have regression weights

that are of a mixed nature, depending on the predictor: fixed effects or random effects. When the nonlinearity is due to the link function, as in the Rasch model and the OPLM, the models are generalized linear models. In the Rasch–MIRID and the OPLM–MIRID, there is a second type of nonlinearity, because of the product of the parameters σ_k and β_{jk} (see Equations 1 and 4), so that they are not part of the family of generalized linear models. All item parameters can be considered fixed effects, and the person parameter can be regarded as a random intercept, normally distributed over persons.

For all four models, PROC NLMIXED estimates the item parameters and the parameters of the person parameter distribution (and their standard errors) by using an approximation of the likelihood function based on a normally distributed random intercept. This means that PROC NLMIXED uses an MML approach (see, e.g., Baker, 1992; Verhelst, 1993) to estimate all these parameters. The item parameters are estimated by integrating the likelihood function over a prespecified person parameter distribution—here, the normal distribution. PROC NLMIXED also estimates the mean of the person distribution (if not fixed for identification reasons) and either the standard deviation or the variance and their standard errors. Also, individual person parameter estimates can be obtained by requesting empirical Bayes estimates. Various integral approximations, optimization techniques, and approximations for the first and second derivatives of the likelihood function are available in PROC NLMIXED, some of which will be discussed below.

Information about the fit of the estimated model is given by the maximized value of the log-likelihood function (transformed into a deviance), as well as by the information criteria of Akaike (AIC; Akaike, 1977) and Schwartz (BIC; Schwartz, 1978). These statistics can be used to compare the fit of different models (SAS OnlineDoc, 1999). A more specific test that requires the estimation of several model variants will be presented later.

In the remainder of this section, the structure of the data set and the SAS statements will be explained briefly.

Input

The structure of the input file needed for an analysis is the following.

1. There is a separate row for each observation (for each person by item combination).
2. The first column contains a label for the person.
3. The second column contains the observations for the person by item combination in question. Although not required by the program, we will use a fixed order for the items within each person, the same order as that for the MIRID CML program.
4. Finally, there is a column containing the discrimination value of the item that is involved in the observation (discrimination vector $\mathbf{a}$).

The remaining columns of the input file contains the design matrix $\mathbf{X}$. For the models considered here, the design matrix is identical for all persons, and therefore, it is repeated for each person. For the Rasch model and for the OPLM, the design matrix is defined as follows (see also the section on the Rasch model).

1. There is one row for each item and as many columns as there are item parameters.
2. An element of a row equals 1 if the corresponding item parameter is needed for the item corresponding the row in question and 0 otherwise (see item indicator vector $\mathbf{x}'_{jk}$).

Since in the Rasch model, there is one item parameter per item, for each person this results in an identity matrix with the same number of columns as the number of items. An example with 20 items and 284 persons is presented in Figure 1. The 20 items are organized in five item families and three types of components. The additional column with discrimination values is omitted since, for the Rasch model, these values are all equal.

For the Rasch–MIRID and the OPLM–MIRID, the design matrix is defined as follows.

1. There is one row for each item and as many columns as there are component item parameters plus one.
2. An element of a row equals 1 if the corresponding component item parameter is needed for the item corresponding the row in question and 0 otherwise (see item indicator vector $\mathbf{x}'_{jk}$). Again, this part of the design matrix is an identity matrix, but only for the component items (item indicator vector $\mathbf{x}'_{jk}$). Note that since we cannot include the weights σ_k in the SAS design matrix for the composite items, they are replaced with ones. In the SAS code representing the likelihood formula, the weights will be explicitly added, so that for the composite items, this modified item indicator vector containing only ones and zeros will be multiplied by σ_k times the component item parameters β_{jk} .
3. An additional column is needed to denote the kind of item. The elements of this column equal 0 if the item is a component item and 1 if the item is a composite item (last element of $\mathbf{x}'_{jk}$, denoted as X0 in SAS code).

For a data set containing 20 items and five item families and with three types of components, the input file for a Rasch–MIRID with the responses and the design matrix looks like the example in Figure 2. The additional column with discrimination values is omitted since, for the Rasch–MIRID, these values are all equal.

The structures of the input files for the OPLM and the OPLM–MIRID are the same as those for the Rasch model and the Rasch–MIRID, respectively, except for the additional column with a priori discrimination values.

SAS Statements (see SAS OnlineDoc, 1999)

First, the DATA procedure has to be called to read the data file. In this procedure, the directory and name of the data file and the names of variables (columns) it contains

are to be specified. For the Rasch model, the SAS code can be found in Listing 1, and for the Rasch–MIRID, the code can be found in Listing 2. For the OPLM and for the OPLM–MIRID, the DATA procedure is identical to that for the Rasch model or the Rasch–MIRID, respectively, except for one additional column: the column with discrimination values. The discrimination value must be mentioned in the INPUT statement. The code can be found in Listings 3 and 4, respectively.

Subsequent to the DATA procedure, PROC NLMIXED should be called. To construct the SAS code for the Rasch model and the OPLM, Equations 3 and 5 should be used. SAS needs the formula for the probability of giving a correct answer to an item. The dummy variables X correspond to the vectors $\mathbf{x}_{jk}'$ and are used to select the parameters needed for the current item. Since the SAS code does not use vector (or matrix) notations or a summation, the vector multiplication of the item indicator vector $\mathbf{x}_{jk}'$ with the item parameter vector $\boldsymbol{\beta}$ needs to be spelled out completely.

The PROC NLMIXED will now be explained, and for the model-specific part, we will first use the Rasch model for the example with 20 items, three components, and five item families. Later, the model specifications for the other three models will also be presented. Since it would be too difficult to specify the β s with two indices in the SAS code, we will replace the indices j and k with one index r , $r = 1, \dots, R$. If we order the component item parameters similarly to the order of the component items in the data set, it is easy to link the β_r to the original β_{jk} .

The statements for the Rasch model are mentioned in Listing 5. We will now go through the statements and explain them statement by statement.

In applying PROC NLMIXED, some choices have to be made about the estimation procedure. These choices can be discussed independently of the model to be estimated.

General options. The first general option is METHOD=. It is used to specify the method for integral approximation. We have opted for the Gauss–Hermite quadrature approximation (GAUSS), as described in Pinheiro and Bates (1995), in combination with the NOAD option, so that the quadrature points are centered at zero for each random effect and so that the current random-effects variance matrix is used as the scale matrix. This is also the default integration method.

With the general option QPOINTS= the number of quadrature points used during the evaluation of the integral can be specified. In combination with the Gauss–Hermite quadrature approximation of the integral, this number equals the number of points used in each dimension of the random effects (we have only one, for the intercept). We chose to set this option equal to 20, to obtain a reasonable precision in describing the distribution of the random effects, and so that the estimation time is not increased too much.

The general option TECHNIQUE= in combination with UPDATE= can be used to determine the optimization technique. Eight different techniques are available, among which are conjugate gradient (CONGRA), Newton–

Raphson optimization (NEWRAP), Newton–Raphson optimization with ridging (NRRIDG), and quasi Newton–Raphson (QUANew, which is the default option). For the quasi Newton–Raphson, the UPDATE= option is needed in addition. Eight different possibilities are available, but not all update methods can be combined with all optimizers. Here and in the simulation study, a quasi Newton–Raphson approach, together with the Davidon–Fletcher–Powell update of the inverse Hessian matrix, has been used. This approach is also implemented in the MIRID CML program and, in contrast to the original Newton–Raphson optimization method that involves the calculation of the second derivatives of the log-likelihood function, which is very time consuming, for the quasi Newton–Raphson optimization techniques only the first derivatives need to be calculated. The default value for the UPDATE option is the Dual Broyden, Fletcher, Goldfarb, and Shanno (DBFGS) update of the inverse Hessian matrix. For the other available options and alternatives, see SAS *onlineDoc* (1999).

Model-specific statements. To specify a model, three statements are needed: PARMS, MODEL, and RANDOM. The PARMS statement identifies all model parameters and their starting values. Between the PARMS and the MODEL statement, the model equation is given [$ex = \dots, p = ex/(1 + ex)$]. The MODEL statement defines the dependent variable and how it depends on the result of the model equation. In our case, the response variables are Bernoulli variables with a probability as described in Equation 2, which is the p from the SAS code: $y \sim \text{binary}(p)$. In the RANDOM statement, the distribution of the random effect is specified [$\text{theta} \sim \text{normal}(0, \text{VarTheta})$]. Only a normal distribution is supported by SAS, and the variance can be either specified a priori or defined as a parameter to be estimated. In our case, we constrained the mean of the person distribution to be zero, to render the model identifiable, and we defined the variance as a parameter to be estimated. Also the standard deviation (and its standard error) can be obtained as follows [$\text{theta} \sim \text{normal}(0, \text{StdTheta}^2)$]. The SUBJECT option within the RANDOM statement is needed to specify when the random effect obtains new realizations. As in the Rasch model, each person has its own person parameter; the variable “Person” defines the realizations of the random effect (SUBJECT = Person).

In order to fit a Rasch–MIRID to the data set mentioned above, only the statements referring to the specific model (PARMS and the model equation) differ. The SAS statements are given in Listing 6.

For the OPLM and OPLM–MIRID, the PARMS statement is equal to the one for the Rasch model or the Rasch–MIRID, respectively. In the INPUT statement, the degrees of discrimination need to be added. The SAS statements representing the model equation of the OPLM are given in Listing 7.

To demonstrate the flexibility of PROC NLMIXED, we will introduce a second way by which to test the MIRID structure besides the comparison with the basic model. One can test the MIRID structure with PROC

NLMIXED by freeing one of the composite item parameters at a time and reestimating the new model. These new models should not have a better fit than the original MIRID. The relaxation can be made for more than one item family at the time. A likelihood-ratio test can be used to test the difference in fit. If we free the first composite item parameter, the previous code for the Rasch-MIRID, for example, should be modified into Listing 8.

As a result, the MIRID restrictions do not apply for the composite item of the first item family. This procedure of leaving one or more out is a way of testing whether the weights σ_k are equally valid for all composite items. That the weights would not be equal for all composite items is the most likely source of misspecification of the MIRID.

Output

The output of PROC NLMIXED contains the estimates of all parameters (the item parameters, the weights, the normalization constant, and the variance of the person parameter distribution), the corresponding standard errors, a Wald test for testing the significance of the parameter estimates, and the value of the first derivative for the current parameter after the final iteration. In addition, four relative fit statistics are given: the deviance, defined as $-2 * \log\text{-likelihood value}$, the AIC value, the AICC value (a finite-sample corrected version of AIC; Burnham & Anderson, 1998), and the BIC value (Schwartz's information criterion; Schwartz, 1978).

Note that the Wald test in PROC NLMIXED for the variance estimates does not give the correct p value. The reference distribution that is used for the null hypothesis is a normal distribution, whereas a variance cannot be smaller than zero. An appropriate way of testing whether there are individual differences (random vs. fixed intercept) is described by Verbeke and Molenberghs (2000) in terms of a likelihood ratio test. The reference distribution of this likelihood ratio test is a mixture of two χ^2 distributions, one with zero degrees of freedom and one with one degree of freedom, leading to p values which are half the size of the p values obtained under the classical χ^2_1 approximation to the null distribution (Verbeke & Molenberghs, 2003). Since the Wald test asymptotically equals the previously mentioned likelihood ratio test, a similar result applies for the PROC NLMIXED output. As a consequence, the correct p value is half the size of the one mentioned in the PROC NLMIXED output.

Since PROC NLMIXED provides the deviance for each model and the Rasch-MIRID is a restriction of the Rasch model and the OPLM-MIRID of the OPLM, one can test the fit of the MIRID against the more general model, using a likelihood-ratio test.

THE MIRID CML PROGRAM AND THE SAS MML APPROACH COMPARED

To compare the MIRID CML program with the SAS MML approach implemented in PROC NLMIXED, 140

data sets were simulated under the Rasch-MIRID, in which two features were varied: the number of persons and the kind of distribution from which the person parameters were sampled. Eighty data sets contained 200 persons, and 60 contained 100 persons. The person parameters were sampled from three different distributions: a normal distribution with a mean of zero and a standard deviation of one (40 data sets with 200 persons and 20 with 100 persons), a truncated normal distribution (20 data sets with 200 persons and 20 with 100 persons), with the negative half omitted from the previous distribution, and a bimodal distribution, obtained by sampling half of the values for the person parameters from a normal distribution with a mean of 0 and the other half from a normal distribution with a mean of 4 (the standard deviations of both were equal to 1; 20 data sets with 200 persons and 20 with 100 persons). All 140 data sets contained 40 items, 10 item families, and three types of components. The component item parameters (β), the weights (σ), and the normalization constant (τ) were sampled from a normal distribution with a mean equal to zero and a standard deviation equal to one. Note that since we will use the group of data sets containing 200 persons and stemming from a normal distribution as reference condition, 40 data sets were included. In this small simulation study, we concentrated on the Rasch-MIRID, since the MIRID is our primary point of interest and the OPLM-MIRID is very similar to the Rasch-MIRID.

To differentiate among the different conditions, the following names will be used: (1) the *normal group*, denoting the data sets containing 200 persons, with the person parameters sampled from a normal distribution (40 data sets), (2) the *truncated group*, denoting the data sets containing 200 persons, with the person parameters sampled from a truncated normal distribution (20 data sets), and (3) the *bimodal group*, denoting the data sets containing 200 persons, with the person parameters sampled from a bimodal distribution (20 data sets). (4) The remaining three groups are named similarly, but the number of persons (100) is added as a suffix resulting in the *normal 100 group* (20 data sets), the *truncated 100 group* (20 data sets), and the *bimodal 100 group* (20 data sets).

All the data sets were analyzed with the MIRID CML program and PROC NLMIXED, using the previously mentioned options. In both the MIRID CML program and PROC NLMIXED, we used a quasi Newton-Raphson optimization technique together with the Davidon-Fletcher-Powell update method for the inverse Hessian matrix. In SAS, we chose the starting values for the component item parameters, the weights, and the normalization constant to be one, whereas in the MIRID CML program, first, a Rasch model was fitted, and the values for item parameters obtained under the Rasch model were the basis for the starting values for all the parameters of the Rasch-MIRID (based on a regression of the composite item parameters on the component item param-

eters). In this way, we obtained estimates for all item parameters and their standard errors. An estimate for the variance of the person parameter distribution was obtained only by PROC NLMIXED (starting value = 1). Subsequently, estimates of the individual person parameters were calculated. The MIRID CML program uses a weighted likelihood approach (Warm, 1989) to estimate the values for the person parameters, often called *Warm estimates*. This results in one value for each possible sum score from the complete questionnaire. In PROC NLMIXED from SAS, Version 8, empirical Bayes estimates can be obtained for the individual realizations of a random effect—here, the person parameter.

We expected the MIRID CML program to be superior with respect to the goodness of recovery for the data sets generated from a nonnormal distribution: the data sets stemming from a bimodal distribution and the data sets stemming from a truncated normal distribution for the person parameter. The misspecification of the distribution should affect primarily the estimates of the person parameters. The MIRID CML and the Warm estimation method used in the MIRID CML program for the estimation of the person parameters make no assumptions about the distribution of the person parameters, whereas PROC NLMIXED in SAS imposes a normal prior distribution for these parameters, which does not correspond with the generating distribution. This effect should especially be visible in data sets with a smaller number of persons. More specifically, we expected PROC NLMIXED to result in an underestimation of the variance of the person parameter distribution for the bimodal distribution and the truncated normal distribution. For the data generated under a normal distribution, an equal goodness of recovery was expected, but there could be an underestimation of variance of the random effect distribution as produced by PROC NLMIXED. As to the Warm estimates, we did not expect them to perform poorly in any condition, since Hoijtink and Boomsma (1996) found, for example, that Warm estimates perform reasonably well for sets of at least 15 items.

The fits of the models were examined with five different statistics, related to the different kinds of parameters (the component item parameters, the weights, and the person parameters). First, since the component item parameters were defined up to an additive constant, the estimated values for the component item parameters

were correlated with the generating values. The higher these correlations, the better was the goodness of recovery for these parameters. Second, the ratio between the variances of the estimated versus the generating values for the component item parameters was calculated. The ratio of the variances was needed to detect differences in variance, which cannot be detected by a correlation. The closer this ratio is to one, the better is the goodness of recovery. Third, we calculated the mean squared differences between the original and the estimated values for the weights of the component item parameters only, because these parameters remain invariant under scale transformations (Butter et al., 1998). The higher the mean squared differences, the worse is the goodness of recovery of the model. Fourth, since the person parameters were defined up to an additive constant, the estimated values (Warm estimates or empirical Bayes estimates) for the person parameters were correlated with the generating values. The higher these correlations, the better is the goodness of recovery for these parameters. Fifth and finally, the variance of the generating person parameters was compared directly with the variance of the Warm estimates, the variance of the empirical Bayes estimates, and the variance of the random effect distribution as estimated by PROC NLMIXED. The latter estimate is direct, whereas the former requires an estimation of the model first.

A more extensive and more extensively documented simulation study on estimating the parameters of the Rasch-MIRID can be found in the article of Butter et al. (1998).

Results

In Table 2, the mean correlation between the estimated and the generating item parameters values over all data sets of the same kind are given, together with their standard deviations. Two-tailed Fischer Z transformations were made before testing the differences between the mean correlations. The standard deviations mentioned in Table 2 are the standard deviations of the correlations before the Fischer Z transformations.

To test the differences between the different conditions, we performed an analysis of variance (ANOVA) for split plot designs on the Fischer Z transformed correlations, with the kind of generating distribution and the number of persons as between-subjects factors and the estimation

Table 2
Mean Correlations Between Generating and Estimated Parameter Values of the Component Item Parameters Over All Data Sets of the Same Kind

Data Set	MIRID CML Program		PROC NLMIXED		N
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	
Normal	.986	.005	.986	.005	40
Truncated	.987	.004	.984	.014	20
Bimodal	.974	.009	.974	.009	20
Normal 100 persons	.974	.008	.974	.007	20
Truncated 100 persons	.972	.008	.972	.008	20
Bimodal 100 persons	.961	.013	.961	.013	20

method (PROC NL MIXED vs. MIRID CML program) as a within-subjects factor. Only the main effects of the two between-subjects factors turned out to be significant. From post hoc *t* tests, we can conclude that the difference between the normal and the truncated normal groups was not significant, whereas the bimodal group did significantly worse. As to the number of persons, the goodness of recovery was significantly worse when the number of persons decreased from 200 to 100. The one within-subjects factor did not yield a significant difference. Both approaches (PROC NL MIXED vs. MIRID CML Program) did about equally well.

In Table 3, the mean values of the ratio of the variance of the component item parameters as estimated by both approaches, as compared with the variance of the generating item parameters over all data sets of the same kind, are displayed. To test the differences between the different conditions, we performed an ANOVA (split plot design) on the variance ratios, with the same design as that for the previous ANOVA. Again, only the main effects of the two between-subjects factors turned out to be significant. Both approaches performed somewhat less well if the person parameter distribution deviated from the normal distribution and if the number of persons decreased. In addition, according to the *F* tests per single data set, the ratios never differed significantly from 1 (all *p* values are even larger than .24).

In Table 4, the mean values for the mean-squared differences between the original and the estimated weights of the component item parameters over the different data sets of one kind are shown. With PROC NL MIXED, the estimates for one data set deviated strongly from the generating values. Excluding this one data set, the mean of the mean-squared difference and its standard deviation dropped to the values displayed in Table 4. An ANOVA with the same design as that for the previous ANOVAs revealed no significant effects.

To summarize, we were not able to find differences in goodness of recovery between the MIRID CML program and the PROC NL MIXED MML approach for the component item parameters and the weights. Therefore, it can be concluded that both approaches perform equally well for the estimation of the item parameters and the weights.

As for the person parameters, the estimated values were correlated with the generating values. The means of

the correlations and the corresponding standard deviations are displayed in Table 5. An ANOVA on the Fischer *Z* transformed correlations, with the same design as that for the previous ANOVAs, revealed a significant main effect for the kind of distribution and a significant interaction between the kind of distribution and the kind of estimation method. On the basis of post hoc tests, the correlations were lower for the truncated distributions. For these distributions also, a small difference was found between the two estimation methods (.841 vs. .832 and .852 vs. .845 for 200 and 100 persons, respectively), although it was not significant. Note that in one of the data sets of the bimodal group, no Warm estimates could be computed, due to computational problems.

Finally, the variance of the originally simulated person parameters was compared with the variance of the Warm estimates, with the variance of the empirical Bayes estimates, and with the variance of the person parameter distribution as estimated by PROC NL MIXED. In Table 6, the mean difference between the estimated variance and the variance of the simulated person parameters is shown. Each mean difference was also tested against zero. A positive value reflects an overestimation of the variance, and a negative value reflects an underestimation.

An ANOVA with the same design as that for the previous ANOVAs revealed two significant main effects (kind of distribution and estimation method) and three significant interactions (all interactions with the estimation method). A post hoc analysis revealed that all row-wise differences, except for one, in Table 6 were significantly different. The variance of the Warm (1989) estimates overestimated the variance of the generating values in the normal group and in the truncated normal group. In the bimodal group, there was a slight underestimation. The variance, as estimated from the empirical Bayes estimates, was always underestimated, especially for the smaller number of persons. A similar underestimation, but much smaller, was found for the estimate of the random effect variance from PROC NL MIXED. However, in three of the six rows, the difference with zero was not significant.

We also investigated the absolute deviations from the variance of the originally simulated parameters. An ANOVA with the same design as that for the previous ANOVAs revealed two significant main effects (kind of

Table 3
Mean Ratios of Variances
Between Generating and Estimated Parameter Values of the
Component Item Parameters Over All Data Sets of the Same Kind

Data Set	MIRID CML Program		PROC NL MIXED		<i>N</i>
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	
Normal	1.034	0.075	1.034	0.075	40
Truncated	1.076	0.086	1.073	0.081	20
Bimodal	1.061	0.092	1.055	0.093	20
Normal 100 persons	1.048	0.164	1.048	0.164	20
Truncated 100 persons	1.139	0.189	1.141	0.189	20
Bimodal 100 persons	1.158	0.216	1.158	0.230	20

Table 4
Means of the Mean Squared Differences Between
Generating and Estimated Weights of the Component Item Parameters
Over All Data Sets of the Same Kind

Data Set	MIRID CML Program		PROC NL MIXED		<i>N</i>
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	
Normal	.033	.049	.027	.031	40
Truncated	.031	.029	.031	.028	19
Bimodal	.055	.060	.053	.060	20
Normal 100 persons	.052	.100	.050	.095	20
Truncated 100 persons	.098	.237	.099	.244	20
Bimodal 100 persons	.126	.314	.115	.269	20

distribution and estimation method) and one significant interaction (between the kind of distribution and the estimation method). The PROC NL MIXED estimate of the random effect variance was always the closest to the expected variance, except for the bimodal 100 group, where the variance of the Warm estimates was the closest to the expected variance. The bias of the empirical Bayes estimates was always the largest, except in the normal group, where it was equal to the bias in the Warm estimates.

Discussion

First, it was expected that the CML approach would do better in recovering the generating parameter values when the data were generated from a nonnormal person parameter distribution, but there was actually no effect for the item parameters and the weights. The difference we found concerns the person parameters. It was very small in terms of correlations, even restricted to the data sets with a truncated distribution. The difference was larger for the estimated variance of the person parameters.

In general, the direct estimation of the variance of the person parameter distribution from PROC NL MIXED gave the best results. The variance of the empirical Bayes estimates underestimated the variance of the generating parameters in all kinds of data sets, even the variance of those generated from a normal distribution. As was expected, the estimated variance was smaller for the data sets generated from the two nonnormal distributions than that for the data sets generated from the normal distribution, except for the variance as estimated by PROC NL MIXED in the bimodal 100 group.

In contrast, the Warm (1989) estimates overestimated the variance in the normal and in the truncated groups and underestimated the variance in the bimodal groups. Both had not been predicted. The underestimation in the bimodal groups can be related to the fact that Warm estimates are negatively biased for large, positive θ values (Warm, 1989). Since there are more such values expected for the data generated with the bimodal distribution, the effect of the negative bias would be expected to be relatively large, which explains the underestimation of the variance. The overestimation of the variance in the normal sample and the truncated sample is similar to the results obtained by Hoijsink and Boomsma (1996) with a normal generating distribution. We did not find any explanation in the literature for this overestimation.

Despite the differences found, we cannot conclude that one approach is better, in general, than the other. Nevertheless, we can conclude that for our data, MIRID CML, supplemented with the Warm (1989) estimates, should be preferred if estimates for the person parameters are requested. Warm estimates showed equal (for the normal and the truncated groups) or less (for the bimodal group) bias in terms of overestimation or underestimation of the variance of the person parameter distribution than did the empirical Bayes estimates. If one does not need individual estimates, both approaches are inferior to the estimate obtained by from PROC NL MIXED for the variance of the random effect.

A remarkable difference between the MIRID CML program and PROC NL MIXED was the time needed for the estimation of the models: The MIRID CML program, which first fits the Rasch model and only then the

Table 5
Mean Correlations Between Generating and Estimated Person Parameters
Over All Data Sets of the Same Kind

Data Set	MIRID CML Program		PROC NL MIXED		<i>N</i>
	W.E.	<i>SD</i>	E.B.E.	<i>SD</i>	
Normal	.931	.009	.933	.009	40
Truncated	.841	.021	.832	.020	20
Bimodal	.963	.007	.967	.007	19
Normal 100 persons	.929	.015	.931	.015	20
Truncated 100 persons	.852	.025	.845	.025	20
Bimodal 100 persons	.961	.008	.964	.007	20

Note—W.E., Warm estimates; E.B.E., empirical Bayes estimates.

Table 6
Means of the Differences Between Variance of the Estimated Person Parameters and the
Variance of Generating Person Parameters Over All Data Sets of the Same Kind

Data Set	MIRID CML Program		PROC NLMIXED				N
	W.E.	SD	E.B.E.	SD	E.V.	SD	
Normal	.065*	.177	-.059*	.141	.020	.140	40
Truncated	.111**	.037	-.129**	.043	-.030**	.039	20
Bimodal	-.065**	.101	-.167**	.077	-.071**	.095	19
Normal 100 persons	.066**	.054	-.090**	.052	-.011	.062	20
Truncated 100 persons	.109**	.047	-.133**	.059	-.034*	.054	20
Bimodal 100 persons	-.084**	.105	-.187**	.118	.078	.216	20

Note—W.E., Warm estimates; E.B.E., empirical Bayes estimates; E.V., random effect variance as estimated by PROC NLMIXED. *Mean difference is significantly different from 0 at .05 level. **Mean difference is significantly different from 0 at .01 level.

Rasch–MIRID, takes 2–3 min for a single simulated data set. With PROC NLMIXED, we fitted only the Rasch–MIRID, and this took 15–30 min for a single simulated data set, not including the empirical Bayes estimates for the random effect.

A major advantage of the SAS approach is that PROC NLMIXED is a very broad procedure that can be used for fitting many other generalized linear and nonlinear models with fixed and random effects (see, e.g., Rijmen et al., 2003). One can, for example, test the MIRID structure with PROC NLMIXED by freeing one of the composite item parameters at a time and reestimating the new models, as was explained earlier. The price to pay for this generality is computing time.

CONCLUSIONS

Both approaches are useful for fitting MIRIDs and do (about equally) well according to the goodness of recovery statistics for the item parameters and the weights. When the person parameters are taken into account, small differences between both approaches are found: The CML approach, supplemented with Warm (1989) estimates, can be preferred when individual estimates of the person parameter are requested, and PROC NLMIXED can be preferred when an estimate of the variance of the person parameter distribution suffices.

A major advantage of PROC NLMIXED is that it is very flexible because of the many different options and the many different model variants it can fit. On the other hand, PROC NLMIXED is rather time consuming. The MIRID CML program is less flexible, since it has fewer options and can fit only MIRIDs, but it is faster.

REFERENCES

- AKAIKE, H. (1977). On entropy maximization principle. In P. R. Krishnaiah (Ed.), *Applications of statistics* (pp. 27-41). Amsterdam: North-Holland.
- BAKER, F. B. (1992). *Item response theory: Parameter estimation techniques*. New York: Marcel Dekker.
- BUNDAY, B. D. (1984). *Basic optimisation methods*. London: Arnold.
- BURNHAM, K., & ANDERSON, D. (1998). *Model selection and inference: A practical information-theoretic approach*. New York: Springer-Verlag.
- BUTTER, R. (1994). *Item response models with internal restrictions on item difficulty*. Unpublished doctoral dissertation, Katholieke Universiteit Leuven.
- BUTTER, R., DE BOECK, P., & VERHELST, N. D. (1998). An item response model with internal restrictions on item difficulty. *Psychometrika*, **63**, 1-17.
- FISCHER, G. H., & MOLENAAR, I. W. (Eds.) (1995). *Rasch models: Foundations, recent developments and applications*. New York: Springer-Verlag.
- FISCHER, G. H., PONOCNY-SELIGER, E., PONOCNY, I., & PARZER, P. (1998). *LPCM-WIN* [Computer program]. St. Paul, MN: Assessment Systems Corp.
- GILL, P. E., MURRAY, W., & WRIGHT, M. H. (1981). *Practical optimization*. London: Academic Press.
- HOUTINK, H., & BOOMSMA, A. (1996). Statistical inference based on latent ability estimates. *Psychometrika*, **61**, 313-330.
- PINHEIRO, J. C., & BATES, D. M. (1995). Approximations to the log-likelihood function in the nonlinear mixed-effects models. *Journal of Computational & Graphical Statistics*, **4**, 12-35.
- RASCH, G. (1960). *Probabilistic models for some intelligence and attainment tests*. Copenhagen: Danish Institute for Educational Research.
- RIJMEN, F., TUERLINCKX, F., DE BOECK, P., & KUPPENS, P. (2003). A nonlinear mixed model framework for item response theory. *Psychological Methods*, **8**, 185-205.
- SAS ONLINEDOC, VERSION 8 (1999). Cary, NC: SAS Institute.
- SAS SYSTEM V8 FOR WINDOWS [Computer program] (1999). Cary, NC: SAS Institute.
- SCHWARTZ, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, **6**, 461-464.
- SMITS, D. J. M., & DE BOECK, P. (2003). A componential IRT model for guilt. *Multivariate Behavioral Research*, **38**, 161-188.
- SMITS, D. J. M., DE BOECK, P., VERHELST, N., & BUTTER, R. (2001). *The MIRID program (version 1.0)* [Computer program and manual]. Leuven: Katholieke Universiteit Leuven.
- VERBEKE, G., & MOLENBERGHS, G. (2000). *Linear mixed models for longitudinal data*. New York: Springer-Verlag.
- VERBEKE, G., & MOLENBERGHS, G. (2003). The use of the score test for inference on variance components. *Biometrics*, **59**, 254-262.
- VERHELST, N. D. (1993). Itemresponse theorie [Item response theory]. In T. J. H. M. Eggen & P. F. Sanders (Eds.), *Psychometrie in de praktijk* (pp. 83-176). Arnhem, The Netherlands: Citogroep.
- VERHELST, N. D., & GLAS, C. A. W. (1995). One parameter logistic model. In G. H. Fischer & I. W. Molenaar (Eds.), *Rasch models: Foundations, recent developments and applications* (pp. 215-238). New York: Springer-Verlag.
- VERHELST, N. D., GLAS, C. A. W., & VERSTRALEN, H. H. F. M. (1994). *One parameter logistic model* [Computer program and manual]. Arnhem, The Netherlands: Citogroep.
- WARM, T. (1989). Weighted likelihood estimation of ability in item response theory. *Psychometrika*, **54**, 427-450.
- WOLFINGER, R. (1999). *Fitting nonlinear models with the new NLMIXED procedure* (SUGI 24 Conference proceedings, Paper 287). Cary, NC: SAS Institute.

LISTINGS

Listing 1: SAS Statements for Reading Data for the Rasch Model

```

Comments are written between /*...*/:
DATA Rasch; /*name of the data set within the SAS environment*/
INFILE 'c:\data\Rasch.dat'; /*name and location of datafile*/
INPUT Person $ y X1-X20;
/*Variables: Person: person label (followed by $ because
person is a string (character)), y: responses, X1-X20 are the
dummy variables that form the columns of the design matrix*/
RUN;

```

Listing 2: SAS Statements for Reading Data for the Rasch-MIRID

```

Comments are written between /*...*/:
DATA RaschMirid; /*name for SAS data set*/
INFILE 'c:\data\raschmirid.dat'; /*name and location of data
file*/
INPUT Person $ y X1-X15 X0;
/*Variables: Person: person label, y: responses, X1-X15:
dummy variables for component item parameters, X0: dummy
variable denoting the composite item*/
RUN;

```

Listing 3: SAS Statements for Reading Data for the OPLM

```

Comments are written between /*...*/:
DATA Oplm; /*name of the data set within the SAS environment*/
INFILE 'c:\data\oplm.dat'; /*name and location of datafile*/
INPUT Person $ y X1-X20 A;
/*Variables: Person: person label (followed by $ because
person is a string (character)), y: responses, X1-X20 are the
dummy variables that form the columns of the design matrix, A:
discrimination values*/
RUN;

```

Listing 4: SAS Statements for Reading Data for the OPLM-MIRID

```

Comments are written between /*...*/:
DATA OplmMirid; /*name for SAS data set*/
INFILE 'c:\data\oplmmirid.dat'; /*name and location of data
file*/
INPUT Person $ y X1-X15 X0 A;
/*Variables: Person: person label, y: responses, X1-X15:
dummy variables for component item parameters, X0: dummy
variable denoting the composite item, A: discrimination values*/
RUN;

```

Listing 5: SAS Statements for Estimating the Rasch Model

```

Comments are written between /*...*/:
PROC NLMIXED DATA=Rasch METHOD=gauss NOAD QPOINTS=20
TECHNIQUE=QuaNew UPDATE=dfp;
/*Specification of data and estimation procedure*/
PARMS Beta1-Beta20=1 VarTheta=1; /*Parameters and their starting
values*/
ex=exp(theta-X1*Beta1-X2*Beta2-X3*Beta3-X4*Beta4-X5*Beta5-X6*Beta6
-X7*Beta7- . . . -X18*Beta18-X19*Beta19-X20*Beta20);
p=ex/(1+ex); /*Formula of Rasch model, see Equation 3*/
MODEL y ~ binary(p); /*the Rasch model is a model for binary
data*/
RANDOM theta ~ normal(0,VarTheta) SUBJECT=Person;
/*specification of distribution of the random intercept  $\theta$ :
The persons are normally distributed with mean zero and variance
equal to VarTheta. The subject option specifies over which
variable the random effects are distributed. If the option
OUT=SAS-data set is specified, empirical Bayes estimates for the
realizations of the person parameter are calculated and stored
in the specified SAS data set*/
RUN;

```

LISTINGS (Continued)**Listing 6: SAS Statements for Estimating the Rasch–MIRID**

```

Comments are written between /*...*/:
PROC NLMIXED DATA=RaschMIRID METHOD=gauss NOAD QPOINTS=20
  TECHNIQUE=QuaNew UPDATE=dfp;
/*Specification of data and estimation procedure*/
PARMS Beta1-Beta15=1 Sigma1-Sigma3=1 Tau=1 VarTheta=1;
/*Specification of the parameters of the Rasch-MIRID and their
starting values; Beta1-Beta15 are the component item parameters,
Sigma1-Sigma3 are the weights of the three types of components
and Tau is the normalization constant*/
ex=exp(theta+(1-X0)*(-X1*Beta1-X2*Beta2-X3*Beta3-X4*Beta4-X5*Beta5
- . . . -X14*Beta14-X15*Beta15))
/*part specific to component items*/
+X0*(-X1*Beta1*Sigma1-X2*Beta2*Sigma1-X3*Beta3*Sigma1
-X4*Beta4*Sigma1-X5*Beta5*Sigma1-X6*Beta6*Sigma2-X7*Beta7*Sigma2
-X8*Beta8*Sigma2-X9*Beta9*Sigma2-X10*Beta10*Sigma2
-X11*Beta11*Sigma3-X12*Beta12*Sigma3-X13*Beta13*Sigma3
-X14*Beta14*Sigma3-X15*Beta15*Sigma3-Tau));
/*part specific to composite items*/
p=ex/(1+ex); /*inverse logit transformation*/
MODEL y ~ binary(p); /*the Rasch-MIRID is a model for binary
data*/
RANDOM theta ~ normal(0,VarTheta) SUBJECT=Person;
RUN;

```

Listing 7: SAS Statements for Model Equation of the OPLM

```

Comments are written between /*...*/:
ex=exp(A*(theta-X1*Beta1-X2*Beta2-X3*Beta3-X4*Beta4-X5*Beta5-X6*Beta6
-X7*Beta7- . . . -X18*Beta18-X19*Beta19-X20*Beta20));
p=ex/(1+ex);

```

Listing 8: SAS Statements for Estimating a Rasch–MIRID in Which the First Composite Item Parameter is Freed

```

Comments are written between /*...*/:
PROC NLMIXED DATA=RaschMIRID METHOD=gauss NOAD QPOINTS=20
  TECHNIQUE=QuaNew UPDATE=dfp;
/*Specification of data and estimation procedure.*/
PARMS Beta1-Beta15=1 Sigma1-Sigma3=1 Tau=1 VarTheta=1
CompBeta1=1;
/*CompBeta1 is the freed composite item parameter of the first
item family*/
ex=exp(theta+(1-X0)*(-X1*Beta1-X2*Beta2-X3*Beta3-X4*Beta4-X5*Beta5
- . . . -X14*Beta14-X15*Beta15))
/*part specific to component items, nothing changes*/
+X0*(-X1*CompBeta1-X2*Beta2*Sigma1-X3*Beta3*Sigma1-X4*Beta4*Sigma1
-X5*Beta5*Sigma1-X7*Beta7*Sigma2-X8*Beta8*Sigma2-X9*Beta9*Sigma2
-X10*Beta10*Sigma2-X12*Beta12*Sigma3-X13*Beta13*Sigma3
-X14*Beta14*Sigma3-X15*Beta15*Sigma3-Tau));
/*part specific to composite items: X1*Beta1*Sigma1-X6*Beta6*Sigma2
-X11*Beta11*Sigma3 is omitted and replaced by X1*CompBeta1*/
p=ex/(1+ex); /*inverse logit transformation*/
MODEL y ~ binary(p);
RANDOM theta ~ normal(0,VarTheta) SUBJECT=Person;
RUN;

```