IRT Models for Ability-Based Guessing

Ernesto San Martín and Guido del Pino, Pontificia Universidad Católica de Chile, Chile
Paul De Boeck, K. U. Leuven, Belgium

An ability-based guessing model is formulated and applied to several data sets regarding educational tests in language and in mathematics. The formulation of the model is such that the probability of a correct guess does not only depend on the item but also on the ability of the individual, weighted with a general discrimination parameter. By so doing, the possibility that the individual uses his or her ability to some extent for differentiating among responses while guessing is also considered. Some important properties of the

model are described and compared with analogous properties of related models. After simulation studies, the model is applied to different data sets of the Chilean Sistema de Medición de la Calidad de la Educación (SIMCE) tests of mathematics and language. The conclusion of this analysis seems relevant—namely, that the examinees use their ability to guess in the language test but not in the mathematics test. *Index terms:* guessing parameter, partial knowledge, propensity to guess, Rasch model

For multiple-choice tests, it is reasonable to assume that the respondents guess when they believe that they do not know the correct response. The solution for this problem is to include a so-called guessing parameter in the model that would have been used without any guessing, which is commonly the one-parameter logistic (1PL) or the two-parameter logistic (2PL) model. Three possibilities are found in the literature: (a) a fixed value $\frac{1}{L}$, with L being the number of response categories; (b) an overall guessing parameter to be estimated from the data, with the same value for all items; and (c) an item-specific guessing parameter. The third possibility is the one used in the three-parameter logistic (3PL) model, the extension of the 2PL with an item-specific guessing parameter. The parameter reflects the probability of a correct guess. The 3PL is introduced by Birnbaum (1968) and is discussed in most handbooks of item response theory (IRT; Embretson & Reise, 2000; Hambleton, Swaminathan, & Rogers, 1991; McDonald, 1999; Thissen & Wainer, 2001; van der Linden & Hambleton, 1997), and the option is available in various computer programs (e.g., BILOG, LOGIST, MIRTE, MULILOG, PARSCALE, and RASCAL). The 3PL is a popular model, but it restricts the guessing parameter to be item dependent instead of allowing this parameter to be also person dependent.

The aim of this study is to investigate a model with also individual differences in the guessing parameter and to apply this model to test whether the probability of a correct guess is ability related. Section 1 briefly summarizes some of the results regarding the 3PL and the formulation of alternative models for multiple-choice data because they may be considered also relevant for the model considered in this study. In Section 2, a model with individual differences in guessing will be formulated, and its properties will be described, followed by a simulation study as described in Section 3. In Section 4, the application of this model to data of two educational tests will be reported.

The 3PL Model

Formulation of the Model

The 3PL model is a model for binary data, $Y_{ij} \in \{0, 1\}$, with index $i = 1, \dots, I$ for persons and index $j = 1, \dots, J$ for items. It is an extension of the 2PL model with an item-specific guessing parameter. In the 2PL model, $P(Y_{ij} = 1 | \theta_i) = \text{expit}(\alpha_j \theta_i - \beta_j)$, with $\text{expit}(x) = \exp(x) / [1 + \exp(x)]$ and $\theta_i \sim \mathcal{N}(0, \sigma^2)$, with α_j , β_j , and θ_i denoting the discrimination parameter, the difficulty parameter, and the person ability, respectively. When the degrees of discrimination are all equal, so that the discrimination parameter can be omitted, the 1PL model is obtained. Whether or not the 2PL or the 1PL model is used, the resulting probability is denoted with p_{ij} . The 3PL model can now be formulated as

$$P(Y_{ij} = 1 | \theta_i) = p_{ij} + g_j - g_j p_{ij}. \quad (1)$$

Interpretations of the Model

Several interpretations can be found for the 3PL model (Hutchinson, 1991). One of the popular interpretations is that the 3PL results from two processes, a p -process, and a g -process. The p -process consists of working on the item to find the correct response. The result of this process is the realization of a binary variable $V_{ij} \sim \text{Bernoulli}(p_{ij})$. Here, $V_{ij} = 1$ means that a person knows the correct response. The g -process consists of guessing. Also the result of this second process is the realization of a binary variable, $W_{ij} \sim \text{Bernoulli}(g_j)$, with the same distribution for all encounters of respondents with item j . Here, $W_{ij} = 1$ means that a person answers the item correctly by guessing.

There are three possible arrangements for the execution of the two processes. The first is that the p -process comes first and that, depending on the result, the g -process follows. This would mean that the respondent first works on the item, with a success probability p_{ij} . When the correct response is found, that response is given; when the correct response is not found, the respondent makes a guess, with a success probability g_j . This interpretation corresponds to the left panel of Figure 1 and is directly reflected in equation (2):

$$P(Y_{ij} = 1 | \theta_i) = p_{ij} + (1 - p_{ij}) g_j. \quad (2)$$

The second arrangement is that the g -process comes first and that, depending on the result, the p -process follows. In principle, there are two interpretations of the second order, but neither of the two is appealing. Either the respondent first makes a guess (with a success probability g_j) and, when the guess is not correct, starts working on the item (with a success probability p_{ij}), or the respondent first decides to make a guess (with probability g_j) and makes a correct guess with probability 1, and when the respondent does not make a guess, the regular solving process is started (with a success probability p_{ij}). The second order corresponds to the right panel of Figure 1 and is directly reflected in a variant of the model formulation:

$$P(Y_{ij} = 1 | \theta_i) = g_j + (1 - g_j) p_{ij}. \quad (3)$$

The third arrangement is that both processes are executed and that a disjunctive mapping applies so that if $V_{ij} = 1$ or $W_{ij} = 1$ or both, then $Y_{ij} = 1$, whereas $Y_{ij} = 0$ otherwise. The probability that both $V_{ij} = 0$ and $W_{ij} = 0$ is $(1 - p_{ij})(1 - g_j)$, so that the probability that $V_{ij} = 1$ or $W_{ij} = 1$ or both equals $1 - (1 - p_{ij})(1 - g_j)$, which also leads to equation (1). It is difficult to find an interpretation implying that both processes are executed in all cases. If one knows the response, guessing is very unlikely, and when one decides to guess, it is very unlikely that one also tries to solve the item. However, when the g -process would no longer be interpreted as guessing, two parallel or sequential strategies can make sense, with a disjunctive combination of the result of both, and it would perhaps no longer be evident to use an ability interpretation.

Figure 1
Interpretations of the Three-Parameter Logistic Model

Of course, the disjunctive mapping is formally equivalent with two drop-off mapping processes for latent responses (E. Maris, 1995). These two drop-off rules correspond with the first and second arrangements of the two processes and with the left and right panels of Figure 1, respectively.

Robustness and Accuracy of Estimation

Different aspects about the inclusion of guessing parameters are touched upon in the literature. Pelton (2002) makes an empirical study of the accuracy of estimates from the 1PL, 2PL, and 3PL models. For person abilities, he found that accuracy is fairly comparable across models, whereas for item difficulties, the 3PL or the 2PL produces the most accurate results, depending on the amount of guessing present in the data. Estimation of the guessing parameters is the most unstable. These results are consistent with previous work (DeMars, 2001; Divgi, 1986). The 3PL model can be recommended only for large samples, unless the guessing parameters are made equal to a known or unknown constant. One argument is that alternatives are guessed at random for a person of very low ability, which implies that a probability of guessing right is equal to $\frac{1}{L}$. However, this does not seem to be a reasonable assumption independent of the set of distractors.

Roughly speaking, as the number of model parameters increases, less data per parameter are available, which may cause serious instability. This may become manifest if it is possible to compare estimates computed for different samples of the same population. van der Linden and Hambleton (1997) state that for the 3PL, small changes in the values of the lower asymptote can be compensated for by small changes in the slope of the curve. This suggests a possible trade-off between the two kinds of parameters and therefore also some instability. Samejima (1973) and Yen, Burket, and Sykes (1991) discuss numerical problems for computing the maximum likelihood estimation (MLE), particularly the existence of multiple modes in the likelihood function.

Pelton (2002) mentions another source of problems. The calibration of the guessing parameters based on capable and weak samples of students may produce substantially different item parameter sets. In fact, most information about the lower asymptote in the item characteristic curve (ICC) is obtained for relatively easy items, but the discrepancy between capable and less capable persons may also stem from the probability of a correct guess being dependent on ability.

Apart from the well-known identification problem caused by the translational invariance, which occurs also in the 1PL and 2PL, some other identification problems concerning the 3PL model are mentioned in the literature. Cressie and Holland (1983) have presented a parameterization of the Rasch model so that the difficulty parameter is a function of the lower asymptote of the ICC. To realize this parameterization, one of the transformations is that the θ -range $(-\infty, \infty)$ is mapped into the range $(0, \infty)$ by setting the new ability θ^* equal to $\exp(\theta)$; thus, θ^* is used as a logarithmic parameter. A similar parameterization is possible for the 2PL. In a similar vein, G. Maris (2002) has pointed out that the Rasch model with a guessing parameter, and consequently also the 3PL, is not identified. The parameterization again implies that the θ^* -scale (the new θ -scale) is mapped into the domain of positive real numbers. These identification problems can be circumvented in a model with a normal distribution for θ with a fixed mean of zero because the θ -values are then distributed symmetrically around zero, from the positive range into the negative range, so that the mapping of θ into the range of positive reals is impossible.

Alternative Models for Multiple-Choice Data

The 3PL is a model for binary data (correct/incorrect), whereas the initial observations are multicategorical. Therefore, it is not surprising also that models are used that take the multicategorical nature of the data into account. For example, Bock's (1972) nominal model is an evident choice. Because this model leaves no room for guessing, Thissen and Steinberg (1984) have formulated a response model for multiple-choice items that does include a guessing component. The basis of the model is Bock's nominal model. When a respondent cannot make a choice between the regular responses, which corresponds to a latent "don't know" response, the guessing component determines what the response will be. Also, Bechger, Maris, and Verstralen (2005) have developed a model with a guessing component for multicategorical data from multiple-choice items. In their model, called the Nedelsky (1954) model, each response is first judged on its acceptability, and then a random guess is made among the nonrejected responses.

Note that in both models, ability matters also when the correct response is not known, whereas this is not the case in the 3PL. In the Thissen and Steinberg (1984) model, an incorrect response may follow from the first process, which is ability based. In the Nedelsky (1954) model, the success probability of a guess depends on the number of rejected responses, and therefore it also depends on ability.

Both models are of interest when one wants to analyze the process of responding to multiple-choice items in a finer way, but then one needs to use the original responses as data. Although this is an appealing way to go, this article stays within the 3PL framework, with a model for binary data. The model as formulated in equations (1) to (3) will be extended to accommodate individual differences in the success probability of guessing.

As a consequence, one may obtain information from the guessing process for the estimation of ability. This approach may seem similar but is in fact different from Weitzman's (1996) proposal to use the total score corrected for guessing in the maximum likelihood equation for estimating ability based on the Rasch model. The proposal by Weitzman implies a discrepancy between the model specification and the estimation because the model specification does not include a guessing parameter. Another set of multiple-choice models is described and investigated regarding their theoretical and potentially practical merits by Hutchinson (1991) in a rather extensive analysis. An important category of models that are considered by Hutchinson is based on what he calls "mismatch theory," implying that respondents evaluate each response separately, an idea that is also found in Bechger et al.'s (2005) Nedelsky model. The models Hutchinson discusses either require alternative response formats (such as confidence ratings, answer-until-correct, etc.), or they require that the respondents are offered the possibility of nonresponse. For the various modes Hutchinson discusses, the estimation of the models remains to be investigated.

The 1PL-AG Model

As mentioned earlier, a model with ability-based guessing (AG) is investigated. This is the 1PL-AG model, with the chances of a correct guess being dependent on ability and with a difficulty but not a discrimination parameter.

The motivation for the 1PL-AG is twofold. First, there are good reasons to believe that in some cases, the success of guessing is related to ability. The first reason is that the guessing parameter depends on the item and is often not equal to $\frac{1}{L}$. This means that the guessing is not just a kind of random guessing. Several authors have pointed out that the reason for $g_j < \frac{1}{L}$ is the attractiveness of the incorrect responses (Hambleton et al., 1991; Hutchinson, 1991; Lord, 1983; McDonald, 1989). If this is the case, it seems not unreasonable that some respondents would be more seducible than others and that those with a higher sensitivity to the attraction of incorrect responses are also less able. Second, when $g_j > \frac{1}{L}$, a plausible explanation is that respondents eliminate one or more of the responses and then guess one among the noneliminated responses. Also, this phenomenon may be related to ability. When it is indeed the case that the guessing parameter is related to the ability of the person, then partial knowledge may be reflected in this parameter.

Second, as mentioned earlier, even when the model is identified, the simultaneous inclusion of a guessing parameter and a discrimination parameter may lead to some trade-off between the estimates of both, and the estimates of these two parameters are found to be less robust than those of the difficulty and the ability parameters. A lower asymptote with a value that is clearly higher than zero is also visually not always easy to differentiate from a low degree of discrimination, especially in the low-ability range. Because the intended study is concentrated on the guessing parameter and because the model proposed in this article creates more flexibility with respect to this parameter, a discrimination parameter will not be used to avoid interference and trade-offs in the estimation.

The two choices also each have a possible drawback. The introduction of individual differences in the guessing parameter can create an interpretation problem because both the p -process and the g -process would then depend on ability, given that an ability interpretation is used. One may wonder which one is the regular ability-based solving process and which is the guessing process and whether any of them may be interpreted as based on ability. As a way out of this first problem, and assuming an ability interpretation (as explored in this article), the size of the ability weight in the guessing parameter will be looked at. It may be expected that ability is of much less importance in guessing than in the regular solving process; see the formulation of the 1PL-AG model in equation (4) for how to assess the weight of ability. To interpret the g -process as a guessing component, the weight of ability should be much smaller than in the p -process, a fact observed in the application of the 1PL-AG model to educational data sets (see the results reported in Section 4).

The omission of a discrimination parameter may affect the goodness of fit of the model and makes it less general. As a way out of this second problem, a model with item discrimination parameters also will be estimated; both the absolute and the relative goodness of fit of the model without such parameters will consequently be tested.

Formulation of the 1PL-AG Model

How important ability is for the probability of a correct guess is an issue considered in the 1PL-AG model by inserting ability with a general discrimination parameter α in the g -process of the model. The probability of a correct response of examinee i to item j is therefore given by

$$P(Y_{ij} = 1 | \theta_i) = \text{expit}(\theta_i - \beta_j) + [1 - \text{expit}(\theta_i - \beta_j)] \text{expit}(\alpha\theta_i + \gamma_j), \quad (4)$$

where $\theta_i \sim \mathcal{N}(0, \sigma^2)$ is a latent variable representing the ability of the examinees; β_j corresponds to the difficulty of item j ; γ_j is the guessing parameter of item j on the logistic scale, corresponding to a person with average ability (if set equal to zero); and α is the weight of the ability in the guessing component. The normality assumption for θ_i is a common one for the 2PL and 3PL models because a marginal maximum likelihood (MML) estimation instead of a conditional maximum likelihood (CML) estimation is required. In principle, other specifications are possible, as well as a semiparametric one with a general probability distribution G . The actual choice is motivated by the use of normal distributions for random effects, which SAS-NLMIXED requires. Furthermore, as in 1PL, 2PL, and 3PL models, the 1PL-AG model is not identified without further restrictions. This is why the means of the random effect distribution are restricted to be 0.

The research issue that the model in equation (4) deals with is whether there is ability in making a correct guess or, more precisely, whether the ability of the p -process also plays a role in the g -process and whether $\alpha \neq 0$ correspondingly. It is expected that $\alpha > 0$, assuming that the ability contributes in a positive way to a correct guess. It seems rather plausible that respondents with higher abilities have a more favorable distribution of the probabilities over the possible responses when they guess, raising in this way the probability of a correct guess. This may lead to the elimination of incorrect responses so that one can guess among the remaining ones. For lower abilities, it may also lead to a higher attractiveness of incorrect responses because they seem correct for the wrong reasons.

When $\alpha = 0$, 1PL-AG reduces to a model as in equation (2)—namely,

$$P(Y_{ij} = 1 | \theta_i) = \text{expit}(\theta_i - \beta_j) + [1 - \text{expit}(\theta_i - \beta_j)] \text{expit}(\gamma_j). \quad (5)$$

To emphasize that the guessing part of this model does not incorporate the person's ability, model (5) is called 1PL-G. Thus, 1PL-AG and 1PL-G are nested models, so that a likelihood ratio can be used to test the hypothesis $H_0 : \alpha = 0$ versus $H_1 : \alpha \neq 0$. If λ denotes the log-likelihood ratio, under H_0 , $-2 \log \lambda$ is asymptotically distributed according to $\chi^2(1)$. In other words, the 1PL-AG model is a tool to test whether guessing correctly is ability related ($\alpha > 0$) or not ($\alpha = 0$)¹ (again conditionally upon an ability interpretation).

Estimation Procedure

As pointed out above, MML estimations are required. Thus, when the interest is focused on the 1PL-G model, the parameters to be estimated are $(\sigma^2, \beta, \gamma)$, whereas if the interest is focused on the 1PL-AG model, the parameters to be estimated are $(\alpha, \sigma^2, \beta, \gamma)$. The mean of θ is set equal to 0. Here, $\beta = (\beta_1, \dots, \beta_J)$ and $\gamma = (\gamma_1, \dots, \gamma_J)$. In both the simulation study and the application to a set of educational testing data, these estimates were obtained with the NLMIXED procedure from SAS with the following specification: a nonadaptative numerical integration method with 15 Gauss-Hermite quadrature points and the Newton-Raphson optimization technique (SAS Institute, 1999). An alternative procedure, which will not be used here, is a Markov chain Monte Carlo (MCMC) approach using, for example, WINBUGS.

Properties of the 1PL-AG Model

Some properties of the 1PL-AG model are discussed to differentiate it from the 1PL model and the 1PL-G model (and, by implication, also from the 3PL). The properties all relate to the ICC. It is assumed for all properties that the same parameter values apply as far as they are relevant: β s (difficulties) for the 1PL and the 1PL-G, as well as γ s (logistic guessing parameters) for the 1PL-G.

For the illustration of the properties, two figures are used, Figures 2 and 3. Each panel in the figures shows an ICC for an item with $\beta = 0$ and $\gamma = -1.09861$ (corresponding to an asymptote of

0.25 in the 1PL-G, chosen to reflect $1/L$, with $L = 4$). Two of the values for α ($\alpha = 0.036$ and $\alpha = 0.228$) are chosen on the basis of the results obtained in the application (see Section 4), and they may therefore be considered realistic. The other values are chosen for theoretical reasons to illustrate the properties of the 1PL-AG. For each model, $\text{ICC}(\theta)$ will be denoted by $\text{model-name}(\theta)$.

Property 1. For $\alpha > 0$, the probability of success is always higher for the 1PL-AG than for the 1PL, independent of the value of θ . More formally,

$$1\text{PL-AG}(\theta) > 1\text{PL}(\theta) \quad \text{for all } \theta \in \mathbb{R}, \alpha > 0.$$

It is most prominent in the two upper panels of Figure 2. The property follows from the fact that $[1 - \text{expit}(\theta - \beta_j)] \text{expit}(\alpha\theta + \gamma_j) > 0$ for all $\theta \in \mathbb{R}$ and for all $\alpha > 0$; in particular, it also implies that $1\text{PL-G}(\theta) > 1\text{PL}(\theta)$ for all $\theta \in \mathbb{R}$.

Property 2. The ICCs of the 1PL-AG and the 1PL-G cross at the point where $\theta = 0$. More formally,

$$\begin{aligned} 1\text{PL-AG}(\theta) &< 1\text{PL-G}(\theta) \quad \text{for all } \theta < 0, \\ 1\text{PL-AG}(\theta) &> 1\text{PL-G}(\theta) \quad \text{for all } \theta > 0, \text{ and} \\ 1\text{PL-AG}(\theta) \text{ and } 1\text{PL-G}(\theta) &\text{ cross at } \theta = 0. \end{aligned}$$

The property follows from the fact that if $\alpha > 0$, then

$$\alpha\theta + \gamma_j < \gamma_j \iff \theta < 0, \quad \alpha\theta + \gamma_j > \gamma_j \iff \theta > 0,$$

and that $1\text{PL-AG}(0) = 1\text{PL-G}(0)$. It implies that, depending on the item difficulty and the item guessing parameter and for $\alpha > 0$, the success rate for the higher ability range is higher for the 1PL-AG than for the 1PL-G, whereas the reverse is true for the lower ability range. One can clearly see the crossing from the second panel in Figure 2. This property may explain why, if the true model is 1PL-AG and $\alpha > 0$, the guessing parameter is estimated to be smaller than $\frac{1}{L}$ by the 1PL-G model and also by the 3PL.

Property 3. For a positive value of α , the ICC of the 1PL-AG converges to the ICC of the 1PL (converges to 0) when θ keeps decreasing, whereas this is not the case for the 1PL-G. More formally,

$$\lim_{\theta \rightarrow -\infty} 1\text{PL-AG}(\theta) = \lim_{\theta \rightarrow -\infty} 1\text{PL}(\theta) < \lim_{\theta \rightarrow -\infty} 1\text{PL-G}(\theta) \equiv \text{expit}(\gamma_j).$$

This property implies that the overall ICC shape of the 1PL-AG is not very different from the corresponding ICC shape of the 1PL, especially when α is sufficiently large. This makes it important to compare the goodness of fit of the 1PL-AG with that of the 1PL and not just with that of the 1PL-G (its more natural competitor). The two lower panels of Figure 2 clearly illustrate this point.

This third property may be considered a problem because it implies that persons with extremely low ability would make guesses with extremely low chances of being correct and certainly much lower than $1/L$. Two possible interpretations for this property of the model can be offered. First, perhaps the model should be understood as an approximate model, as a model that one intends to be precise for a broad range of θ but only approximate (and, strictly speaking, incorrect) for extremely low values of θ . This is not necessarily a problem because there would be not much information in the data for such low values anyhow. Second, the fact that the probability of a correct guess converges to zero when θ keeps decreasing could mean that the sensitivity to attractive features of the distractors keeps increasing with decreasing values of θ . A low ability would therefore also mean that one is highly seducible by the distractors.

Property 4. For $0 < \alpha \leq \frac{1}{\sqrt{2}}$,

$$\text{Inflexion Point}_{1\text{PL-AG}} < \text{Inflexion Point}_{1\text{PL}} \quad \text{for all } \gamma \in \mathbb{R}.$$

Figure 2

ICC for the 1PL-AG Model (Continuous Line), for the 1PL-G Model (Dotted Line),
and for the 1PL Model (Dashed Line) ($\beta = 0$ and $\gamma = -1.09861$)

Note. 1PL-AG = one-parameter logistic model with ability-based guessing; 1PL-G = one-parameter logistic model with guessing.

The proof of this property is given in the appendix. It means that for the range of α values likely to appear in applications, $0 \leq \alpha < 0.707$, the maximum information point of 1PL-AG items is located to the left of the corresponding point for 1PL items. Therefore, the items have their maximum information value for lower abilities than under the 1PL. As will be pointed out in the proof, for values of α higher than 0.707, it depends on β and γ whether the inflexion point is located to the right or to the left of that for the 1PL.

Property 5. The fifth property is that the contribution of the second term on the right in equation (4) to the success probability is a single-peaked function of θ , whereas for the 1PL-G, it is a monotonic decreasing function. This second term is what is added in comparison with the 1PL model:

$$[1 - \text{expit}(\theta - \beta)] \text{expit}(\alpha\theta + \gamma).$$

This property is illustrated in Figure 3, with low values of α , as found in the application, and with higher values of α to illustrate the effect. In the first panel, the contribution (dotted curve) of this term is shown for the case of $\alpha = 0$ and thus for the 1PL-G. From the third of the upper panels on, it is clearly visible that the contribution has a maximum. The implication of the property is that the maximum contribution of the second term to the success probability is to be found somewhere

Figure 3

Contribution of the One-Parameter Logistic Term (Dashed Line) and the Guessing Term (Dotted Line) to the ICC of the 1PL-AG Model (Continuous Line)

Note. 1PL-AG = one-parameter logistic model with ability-based guessing; ICC = item characteristic curve.

between the two extremes of the ability scale. For very able persons, it is almost exclusively the 1PL term that contributes to the success probability, and for persons with very weak ability, the second term does not contribute much to success.

Simulation Study

Rationale of the Simulation Study

A simulation study consisting of three parts was performed. The first part corresponds to a standard simulation study—namely, the following:

- (a) data generated from the 1PL-AG analyzed under the same model, to check the recovery of the true parameter values.

The second and third parts are motivated by the fact that the 1PL-AG model is a tool to test whether the ability of the nonguessing part plays ($\alpha > 0$) or does not play ($\alpha = 0$) a role in the guessing part. In practice, this leads to a likelihood ratio test for choosing between the 1PL-AG and the 1PL-G models. Therefore, two types of analysis were planned—namely, the following:

- (b) data generated from the 1PL-G analyzed under both the 1PL-G and the 1PL-AG models, to check whether the analysis would indicate the 1PL-G as the best model;
- (c) data generated from the 1PL-AG analyzed under both the 1PL-G and the 1PL-AG models, to study the effect of omitting ability from the guessing component.

Design of the Simulation Study

Each generated data set consists of 2,000 examinees taking a test with 43 items. The “true parameters,” the β s and γ s employed in the simulations, were obtained from a pilot sample of the Sistema de Medición de la Calidad de la Educación (SIMCE) test used for the applications (see the Section 4). The average parameter values are $\bar{\beta} = 0.85$ and $\bar{\gamma} = -1.23$. The β s range from -1.63 to 3.16 , and the γ s range from -3.16 to 0.90 . These values were used to generate 12 data sets with either the 1PL-G or the 1PL-AG models. More specifically, in part (a), four databases (denoted as Sample 1a, Sample 2a, Sample 3a, and Sample 4a) were simulated according to 1PL-AG: the first two with $\alpha = 0.1$ and the second two with $\alpha = 0.2$. In part (b), four databases (denoted as Sample 1b, Sample 2b, Sample 3b, and Sample 4b) were also simulated according to the 1PL-G model. In part (c), again four databases (denoted as Sample 1c, Sample 2c, Sample 3c, and Sample 4c) were generated according to the 1PL-AG model: the first two with $\alpha = 0.1$ and the second two with $\alpha = 0.2$. The random effect used to generate the 12 data sets was a standard normal distribution. The α values used in parts (a) and (c) are chosen because of Property 4 and on the basis of the results obtained in the application—namely, in the neighborhood of 0.2 or lower. The estimation procedure that was followed is the one that is described for the estimation procedure in the previous section (using SAS NL MIXED).

It should be remarked that the number of replications per cell is limited: four replications in part (b) and two replications for each α value in parts (a) and (c). This is due to the cost of computing time in fitting simulated databases: the time varied between 15 and 40 hours (on a Pentium® 4 CPU, 2.80 GHz, 1 GB of RAM). Because the number of replications is small, no standard errors can be derived based on replications. However, it will be checked in an explicit way how much different the estimates are for the replications.

Results of the Simulation Study

For part (a), namely for parameter recovery, Table 1 summarizes the results for the estimated α and σ^2 ; the recovery is good. Figure 4 shows a representative case for the recovery of the β s and the γ s using the 1PL-AG model. More informative is Figure 5, in which for each item of the same data set, the true and estimated values are plotted along with the confidence band bounded by dotted lines

Table 1
 Parameter Recovery for the 1PL-AG for Simulated Samples

	$\alpha = 0.1$		$\alpha = 0.2$	
	Sample 1a	Sample 2a	Sample 3a	Sample 4a
$\hat{\alpha}$	0.101 (0.062)	0.103 (0.056)	0.224 (0.071)	0.217 (0.077)
$\hat{\sigma}^2$	0.968 (0.071)	1.043 (0.096)	0.964 (0.071)	0.972 (0.069)

Note. 1PL-AG = one-parameter logistic model with ability-based guessing.

Figure 4

Comparison Between True and Estimated Values for the 1PL-AG—True Alpha, $\alpha = 0.1$

Note. 1PL-AG = one-parameter logistic model with ability-based guessing.

corresponding to an interval of the type [est. value $-2 \cdot$ stand. error, est. value $+2 \cdot$ stand. error]. From these figures, it may be concluded that the recovery is quite good.

In part (b), both 1PL-AG and 1PL-G models were estimated for each of the four databases generated with 1PL-G. The deviance ($-2l$) was never more than one unit smaller for the 1PL-AG than for the 1PL-G (103,704 vs. 103,704 for Sample 1b, 104,310 vs. 104,311 for Sample 2b, 103,944 vs. 103,945 for Sample 3b, and 103,702 vs. 103,702 for Sample 4b). Given these small differences in deviance, it is evident that the likelihood ratio test is far from significant, so that, as expected, the 1PL-G (the true model) must be considered the best model. Given the minor differences in deviance, of course also the Akaike information criterion (AIC) and the Bayesian information criterion (BIC) indicate the 1PL-G as the best model. Let us also mention that when the four databases were analyzed with the 1PL-AG model, the α estimates were around 0.05 with standard errors around 0.05.

In part (c), again model 1PL-G and model 1PL-AG were fitted to each of the databases. For Sample 1c and Sample 2c (i.e., the databases generated with $\alpha = 0.1$), the values of the likelihood ratio test are 3 and 1 ($p > .05$), respectively; for Sample 3c and Sample 4c (i.e., generated with $\alpha = 0.2$), the corresponding values are 10 and 5 ($p < .05$), respectively. The Wald tests for the

Figure 5

Comparison Between True Value ($\circ$), Estimated Value ($\times$), Estimated Values Plus Two Standard Errors, and Estimated Values Minus Two Standard Errors (Dotted Lines) for the 1PL-AG—True Alpha, $\alpha = 0.1$

Note. 1PL-AG = one-parameter logistic model with ability-based guessing.

estimates of α yield results with $p > .05$ for $\alpha = 0.1$ and $p < .01$ for $\alpha = 0.2$. The estimates have values close to the true ones.

The results of the simulation study are very consistent, so that even with a small number of replications, one can trust that the recovery is good and that the model can be differentiated from its most evident competitor (the 1PL-G model). This is an important result, given our interest in determining whether ability contributes to the chances of a correct guess.

Application to Educational Data Sets

The main purpose of this section is to report the results obtained after analyzing different data sets of the Chilean SIMCE test in mathematics and language. The analyses were performed using the 1PL-G and the 1PL-AG models. The conclusion of the analyses seems relevant—namely, that the examinees use their ability to guess in language but not in mathematics, as will be explained. Before

reporting the results, in the first subsection, the design of the SIMCE test is described and general information about the responses is provided. The second subsection gives details about the design of the analysis and the estimation procedure that was used. The results are reported in the following two subsections. Also, the robustness of the estimations is discussed, and some additional analyses with the 1PL and the 3PL-AG models are reported.

The SIMCE test

The SIMCE project in Chile has developed mandatory tests to assess regularly the educational progress in three levels: 4th, 8th, and 10th graders. All students in the grade level in the country are expected to take the tests when they are scheduled (every 3 or 4 years). For the application, data from two tests applied in 2001 to 10th graders (second year of the secondary school, corresponding to an age of about 16) are used: a language test (37 items with four alternatives) and a mathematics test (48 items with four alternatives), both given to examinees from public schools and mixed schools in Chile. Mixed schools are schools that receive private as well as state financial support. For the language test, the total numbers of examinees are 23,495 and 28,801 for the public schools and the mixed schools, respectively. For the mathematics test, the corresponding figures are 36,118 and 25,310. Because SIMCE is not a high-stakes test, there are many reasons for not having responded to an item that are not associated with the ability of the respondent. For this reason, it was more meaningful to use only the data from respondents who have responded to all items of the test: 92.22% and 92.17% of the respondents for the language test and 79.56% and 79.44% for the mathematics test. In any case, conclusions were checked by fitting the 1PL-AG to a sample for both tests where respondents with incomplete data were not omitted and with a nonresponse considered as an incorrect response, and the results were in full agreement.

The language test focuses on reading comprehension. There are no questions on topics such as grammar, synonyms, spelling, or punctuation. The test consists of different types of texts, followed by a number of questions about reading comprehension and making inferences based on the text. For instance, the student may be asked to analyze how changes in the title of an ad influences the interpretation by the reader. The mathematics test has a variety of questions ranging from problem formulation, functions, simple algebra, geometry, and probability—for instance, simplifying $\frac{4}{x^2} / \frac{2}{x}$ or computing 30% of \$2,000 in the context of an applied problem.

For the 2001 application of the SIMCE test, the averages of the percentages of correct responses over the 37 language items were 53.0% and 53.9%, with a standard deviation of 15.2% and 15.1%, for mixed and public schools, respectively. For the mathematics test, the corresponding average percentages over the 48 items were 53.1% and 54.5%, with standard deviations of 17.2% and 16.8%, for public and mixed schools, respectively.

Design of the Analysis and Estimation Procedure

From each of the data sets, corresponding to the 2×2 cross-classification (public vs. mixed; language vs. mathematics), two samples of 2,000 respondents are drawn for the analysis (hence, eight samples in all, denoted as Samples 1L to 4L for language and Samples 1M to 4M for mathematics). The analysis proceeds in three steps:

- (a) For all eight samples, the 1PL-AG and the 1PL-G models were estimated to find out whether the success probability is related to ability ($\alpha > 0$), using a likelihood ratio to compare these two nested models. Because the same analysis was used for two different samples from the same cell in the 2×2 design, a check on the robustness of the result is available.

- (b) For two samples, one for language (Sample 1L) and one for mathematics (Sample 3M), also the 1PL and the 3PL-AG models were estimated. The 1PL is estimated to check also whether this model can explain the data, based on a concern that stems from Property 3 that was discussed earlier. The 3PL-AG is a model that is analogous to the 1PL-AG but with a discrimination parameter for each item. The reason for estimating this more complex model, at the risk of some instability of the parameter estimates, is to check whether the complexity pays off in comparison to the 1PL-AG.
- (c) Finally, also for two unselected samples (independent of whether all items were responded to), the 1PL-AG model was estimated.

As pointed out earlier, all models were estimated with SAS NLMIXED using nonadaptive Gaussian quadrature with 15 nodes and with the Newton-Raphson technique for the optimization. The distribution of the abilities always was a normal with 0 mean and variance σ^2 . For the estimation of the 1PL-AG, the initial values of the β s and γ s were taken from the corresponding estimation results of the 1PL-G; the computing time varied between 10 and 16 hours (on a Pentium® 4 CPU, 2.80 GHz, 1 GB of RAM), which was less than that of the simulation study (see the Design of the Simulation Study subsection). Let us mention that an adaptive procedure would require much more time: Such a procedure was followed for one sample, obtaining similar results as with the nonadaptive procedure.

Main Results

Tables 2 and 3 summarize the goodness of fit of the 1PL-AG and the 1PL-G for the eight samples, in terms of log-likelihood, and the AIC and BIC values. It is clear from Table 2 that the goodness of fit of the 1PL-AG is superior for the four language data sets. From Table 3, it can be concluded that 1PL-G is preferred to the 1PL-AG for the mathematics test, except maybe for mixed schools (Sample 3M), if the AIC is considered. These conclusions are essentially the same if both the 1PL-G and 1PL-AG are fitted with zero as initial values for the β s and γ s.

It should be mentioned that for three of the $37 + 48 = 85$ items, poor estimations (i.e., with high standard errors) were obtained:

1. Language, public schools, Sample 1L, 1PL-AG: The estimate of γ_8 has a standard error equal to 104.57; the estimates are $\hat{\beta}_8 = -2.40$ and $\hat{\gamma}_8 = -10.26$ for Sample 1L and $\hat{\beta}_8 = -1.27$ and $\hat{\gamma}_8 = 0.11$ for Sample 2L.

Table 2
1PL-G Versus 1PL-AG: Goodness of Fit for the Language Test Data

	Language, Public Schools				Language, Mixed Schools			
	Sample 1L		Sample 2L		Sample 3L		Sample 4L	
	1PL-G	1PL-AG	1PL-G	1PL-AG	1PL-G	1PL-AG	1PL-G	1PL-AG
-2l	89,805	89,786	89,276	89,256	89,690	89,673	89,918	89,880
AIC	89,955	89,938	89,426	89,408	89,840	89,825	90,068	90,032
BIC	90,375	90,364	89,846	89,833	90,260	90,250	90,488	90,458

Note. 1PL-AG = one-parameter logistic model with ability-based guessing; 1PL-G = one-parameter logistic model with guessing; AIC = Akaike information criterion; BIC = Bayesian information criterion.

Table 3
1PL-G Versus 1PL-AG: Goodness of Fit for the Mathematics Test Data

	Mathematics, Public Schools				Mathematics, Mixed Schools			
	Sample 1M		Sample 2M		Sample 3M		Sample 4M	
	1PL-G	1PL-AG	1PL-G	1PL-AG	1PL-G	1PL-AG	1PL-G	1PL-AG
-2l	110,454	110,452	110,251	110,249	109,212	109,206	108,891	108,890
AIC	110,648	110,648	110,445	110,445	109,406	109,402	109,085	109,086
BIC	111,192	111,197	110,988	110,994	109,949	109,951	109,629	109,635

Note. 1PL-AG = one-parameter logistic model with ability-based guessing; 1PL-G = one-parameter logistic model with guessing; AIC = Akaike information criterion; BIC = Bayesian information criterion.

2. Language, mixed schools, Samples 3L and 4L, 1PL-AG: The estimates of β_{12} have standard errors equal to 67.20 and 125.09, respectively; the estimates are $\hat{\beta}_{12} = 15.31$ and $\hat{\gamma}_{12} = -0.46$ for Sample 3L and $\hat{\beta}_{12} = 14.53$ and $\hat{\gamma}_{12} = -0.47$ for Sample 4L.
3. Mathematics, public schools, Sample 2M, 1PL-G: The estimate of γ_{38} has a standard error equal to 81.58; the estimates are $\hat{\beta}_{38} = -1.73$ and $\hat{\gamma}_{38} = -2.78$ for Sample 1M and $\hat{\beta}_{38} = -1.91$ and $\hat{\gamma}_{38} = -9.81$ for Sample 2M.

Figure 6 shows the corresponding ICCs for the three items and for the two corresponding samples.

For Item 8 of the language test, the problem might have been that it was too easy to obtain a reliable estimation of γ , which may also explain the discrepancy between the ICCs in the lower region. For Item 12 of the language test, the apparent problem is that the discrimination is too low. This was confirmed by the 3PL-AG analysis. For Item 38 of the mathematics test, the two ICCs are very similar, but again this item is perhaps too easy for a reliable estimation of β and γ . Because estimation problems for only three of the 85 items were encountered and because the problems have a reasonable explanation, the results of the analyses should be trusted, especially for the overall parameter α , the focus of interest. When the analyses are repeated without the three problematic items, essentially the same results were obtained, with only very small differences in the parameter estimates.

To evaluate the absolute goodness of fit, a bootstrap approach was used: 100 samples were generated from the estimated 1PL-AG model, with the estimates obtained for β s, γ s, σ^2 , and α and with a uniform distribution between -4 and 4 for the abilities θ s. This last choice was made to obtain reliable estimates of the ICC also in the upper and lower ranges of ability. It should be remarked that 100 samples were also generated with a normal distribution with 0 mean and variance equal to 1.56, corresponding to the estimates obtained with the 1PL-AG model for Sample 1L, but the number of cases in the upper and lower ranges of ability was low. Based on the sum scores obtained from the generated data, 100 generated ICCs were obtained, so that the empirical ICCs based on the sum scores of the real data can be compared with these. An acceptable goodness of fit means that the real data ICC is contained in the confidence band based on the 100 samples. This test was performed for Sample 1L (language test, public schools) and for Sample 4L (language test, mixed schools); the two items with unreliable estimates were omitted from this test. Because frequencies of sum scores in the real data lower than 4 and higher than 33 (maximum is 35) were very low (< 20 on a total

Figure 6

Item Characteristic Curve for Items With Estimation Problems:
Sample 1L/3L/1M (Full Line) and Sample 2L/4L/2M (Dotted Line)

larger than 21,000), the empirical ICCs were plotted for the range from 4 to 33. The empirical ICC of the real data is completely contained in the confidence band for all the items. The four panels of Figure 7 illustrate this result for four items corresponding to the language test, public school.

Table 4 presents the estimates of α in the eight samples. In line with the goodness-of-fit indices, the value of α is clearly different from 0 for the language test in both data sets. Given the values of the standard errors, the α estimates are clearly significant. Although clearly significant, the value of α is not very high, only about one fifth of the overall discrimination in the IPL term of the model (where it is 1). This means that there is some but not very much ability involved in the g -process. That α is positive but clearly smaller than 1 also means that there can be no doubt which of the two processes must be considered the guessing process (if any of the two is, as assumed earlier).

As can be seen in Table 4, the results are different for the mathematics test. Only in one of the four samples a significant α estimate is found, and this significant value is very low (.087 for Sample 3M from the mixed-schools sample) and clearly lower than the corresponding values for the language test. This means that for the mathematics test, ability does not play much of a role in the chances of making a correct guess.

Figure 7
 Empirical Item Characteristic Curves of the Real Data (Continuous Line)
 Plotted Against Bootstrap Results—Language, Public School

Table 4
 α Estimates Using the 1PL-AG for the Language and Mathematics Tests

	Language				Mathematics			
	Public Schools		Mixed Schools		Public Schools		Mixed Schools	
	Sample 1L	Sample 2L	Sample 3L	Sample 4L	Sample 1M	Sample 2M	Sample 3M	Sample 4M
$\hat{\alpha}$	0.212	0.228	0.191	0.246	0.052	0.036	0.087	0.036
Standard error	0.033	0.036	0.031	0.026	0.030	0.030	0.030	0.033

Note. 1PL-AG = one-parameter logistic model with ability-based guessing.

Additional Results

The robustness of the 1PL-AG estimates can be evaluated based on the two samples that were used in each cell of the design, for example, by correlating the estimates of β and γ between the two samples from each data set. Always high correlations were obtained: from 0.931 to 0.985 for the β s and from 0.904 to 0.987 for the γ s. Note that the three items with unreliable estimates were excluded from the correlations. These results mean that the estimation should be trusted and that there is a good basis for the inference regarding α as well. Not just the α estimates themselves but also the estimates of the model it is embedded in seem to have some stability from one sample to the other.

From the estimation of the 1PL in two different samples (Sample 1L for language, public schools; Sample 3M for mathematics, mixed schools), it should be concluded that the goodness of fit is clearly inferior to that of the 1PL-AG and also to that of the 1PL-G. This means that the examinees have guessed also while taking the mathematics test. For language, the goodness-of-fit values for the first sample are 90,461, 90,537, and 90,750 for $-2l$ (deviance), the AIC, and the BIC, respectively, and for the second sample, the corresponding values are 89,956, 90,031, and 90,245 (compare to Table 2). For mathematics, the goodness-of-fit values for the first sample are 110,639, 110,737, and 111,012 for $-2l$ (deviance), the AIC, and the BIC, respectively, and for the second sample, the corresponding values are 110,585, 110,683, and 110,957 (compare to Table 3). This is in all cases worse than for the 1PL-AG and 1PL-G.

As mentioned earlier, for two samples (language, public schools, Sample 1L; mathematics, mixed schools, Sample 3M), also the 3PL-AG was estimated. Two important observations were made. First, the values of the standard errors of some parameter estimates were out of range, which is an indication of an unstable solution and trade-offs between parameter estimates. However, equally or even more important is that when the BIC was used as a criterion, the 1PL-AG performs better for the language sample (90,364 vs. 90,495) and almost as good for the mathematics sample (109,635 vs. 109,608). This means that adding a discrimination parameter is not really necessary (unless perhaps for an item such as Item 12 of the language test, as discussed earlier).

Finally, for the two samples that were drawn independent of nonresponse (with nonresponse coded as 0), also the same results regarding α were obtained: for language, $\hat{\alpha} = 0.298$ (standard error = 0.050); for mathematics, $\hat{\alpha} = -0.007$ (standard error = 0.055).

Discussion

From the results of the simulation study, it should be concluded that the recovery of the parameter values of the 1PL-AG model is reasonably good and that the model can be differentiated from the 1PL-G model (with $\alpha = 0$), so that one can make inferences regarding the role of ability in the probability of making a correct guess. These results were obtained for a sample size comparable to the one in the application.

From the application, it should be concluded that the 1PL-AG model has a reasonably good fit, better than the fit of the 1PL and 1PL-AG models, and that the difficulty and guessing parameters are quite stable from one sample to the other. When the absolute goodness of fit was tested with a bootstrap approach, the results were again satisfactory. This means that the findings regarding α can be trusted, as well as the probability of a correct guess being related to ability (conditionally upon an ability interpretation, as discussed earlier).

From the results regarding α , and assuming the earlier interpretation of guessing for the second process, it should be concluded that guessing has occurred in both kinds of tests but that guesses

for the mathematics test are not (or, at most, to a very minor degree) based on ability. For the language test, however, ability seems to make a difference when a guess is made. This makes sense given the kind of test. Mathematics is perhaps more an all-or-none matter, so that guessing based on partial knowledge can play less of a role. Responses to language problems are perhaps more gradual with respect to their correctness because correctness is a matter of judgment, so that partial knowledge can change the pattern of guessing probabilities. This difference between the two tests is perhaps also reflected in the higher nonresponse rate for the mathematics test, suggesting that the examinees were less inclined to guess when they did not know the correct response. For language, it may have paid off better to guess because educated guesses are possible.

The differential and replicated result for the two tests shows that the 1PL-AG can be used as a diagnostic tool to find out whether guessing is related to ability. The differential result for the two tests eliminates the possibility that the result follows from a general kind of artifact related to the kind of model or method of estimation. The differentiation is a rather robust result, in that it is replicated in a second sample, and obtained also in samples with nonresponse. For all these reasons, the 1PL-AG is a useful model as a diagnostic tool to find out whether there is ability in guessing (provided that one assumes the interpretation given earlier to the two processes).

The model is in fact a compromise between using a multicategorical model in which partial knowledge can evidently play a role and a model for binary (recoded) data. It shares with the models for binary data that one does not run into the complexities of the various categories with possibly different parameters, depending on the item. It shares with the models for multicategorical data that there is room for partial knowledge through educated guesses. These features also apply to the 3PL-AG, but this model is of course more complex, and as one can infer from the literature and from this study's results, it is also more vulnerable to instabilities in the estimation. Alternative models may be worth being investigated, especially when they would belong to the exponential family and approach the model used here. Such models may solve estimation problems as the present estimation procedure is really slow, and they would have an advantage with respect to statistical testing.

Appendix

Proof of Property 4

Denote by $y(\theta)$ and $z(\theta)$ the item characteristic functions of the 1PL and the 1PL-AG model, respectively. It is well known that $y''(\beta) = 0$ (i.e., β is the inflexion point of the ICC for the 1PL). Given the asymptotes of $z(\theta)$, it also follows that $z''(\theta)$ changes its sign from positive to negative at a single point. Therefore, if $z''(\beta) < 0$ (respectively, > 0), then the inflexion point of z is smaller (respectively, larger) than that of y . The following result is helpful in computing $z''(\theta)$:

$$\text{If } g(x) = \frac{e^x}{1 + e^x}, \text{ then } g'(x) = g(x)[1 - g(x)], g''(x) = g(x)[1 - g(x)][1 - 2g(x)]. \quad (6)$$

Letting $u(\theta) = g(\theta - \beta)$ and $w(\theta) = g(\alpha\theta + \gamma)$, $z(\theta) = u(\theta) + w(\theta) - u(\theta)w(\theta) = 1 - [1 - u(\theta)][1 - w(\theta)]$, and therefore $z''(\theta) = u''(\theta) + w''(\theta) - 2u'(\theta)w'(\theta)$. Applying (6) and evaluating at $\theta = \beta$, it follows that $u(0) = \frac{1}{2}$, $u'(0) = u(0)[1 - u(0)] = \frac{1}{4}$, $u''(0) = u(0)[1 - u(0)][1 - 2u(0)] = 0$, and therefore

$$z''(\beta) = \alpha w(\beta)[1 - w(\beta)] \left\{ \alpha^2 [1 - 2w(\beta)] - \frac{1}{2} \right\}.$$

Because $\alpha > 0$, the sign of $z''(\beta)$ is determined by the sign of $(\alpha^2[1 - 2w(\beta)] - \frac{1}{2})$. Although $w(\beta)$ depends on γ , $w(\beta) > 0$ implies $1 - 2w(\beta) < 1$, and so the sign of $z''(\beta)$ is determined by the sign of $(\alpha^2[1 - 2w(\beta)] - \frac{1}{2}) < \alpha^2 - \frac{1}{2}$. Therefore, $\alpha < \frac{1}{\sqrt{2}} \approx 0.707$ implies $z''(\beta) < 0$, and so the inflexion point of 1PL-AG is to the left of that of 1PL. $\square$

Note

1. In principle, it is possible also that $\alpha < 0$. The result of a negative α is that the ICC increases again for lower abilities. The more a negative α approaches zero, the further the turning point of the ICC moves to the lower end.

References

- Bechger, T. M., Maris, G., & Verstralen, H. H. M. (2005). The Nedelsky model for multiple choice items. In A. van der Ark, M. Croon, & K. Sijtsma (Eds.), *New developments in categorical data analysis for the social and behavioral sciences*. Mahwah, NJ: Lawrence Erlbaum.
- Bock, R. D. (1972). Estimating item parameters and latent ability when responses are scored in two or more latent categories. *Psychometrika*, 37, 29-51.
- Birnbaum, A. (1968). *Some latent trait models and their use in inferring an examinee's ability*. In F. M. Lord & M. R. Novick (Eds.), *Statistical theories of mental test scores*. London: Addison Wesley.
- Cressie, N., & Holland, P. W. (1983). Characterizing the manifest probabilities of latent trait models. *Psychometrika*, 48, 129-141.
- DeMars, C. (2001). Group differences based on IRT scores: Does the model matter? *Educational and Psychological Measurement*, 61, 60-70.
- Divgi, D. R. (1986). Applications of unidimensional item response theory models to multidimensional data. *Applied Psychological Measurement*, 7, 189-199.
- Embretson, S. E., & Reise, S. P. (2000). *Item response theory for psychologists*. Mahwah, NJ: Lawrence Erlbaum.
- Hambleton, R. K., Swaminathan, H., & Rogers, H. J. (1991). *Fundamentals of item response theory*. Newbury Park, CA: Sage.
- Hutchinson, T. P. (1991). *Ability, partial information and guessing: Statistical modelling applied to multiple-choice tests*. Rundle Mall, South Australia: Rumsby Scientific Publishing.
- Lord, F. M. (1983). Maximum likelihood estimation of item response parameters when some responses are omitted. *Psychometrika*, 48, 477-482.
- Maris, E. (1995). Psychometric latent response models. *Psychometrika*, 60, 523-547.
- Maris, G. (2002). *Concerning the identification of the 3PL model* (Measurement and Research Department Reports 2003-3). Arnhem, The Netherlands: CITO National Institute for Educational Measurement.
- McDonald, R. P. (1989). Future directions for item response theory. *International Journal of Educational Research*, 13, 205-220.
- McDonald, R. P. (1999). *Test theory: A unified treatment*. Mahwah, NJ: Lawrence Erlbaum.
- Nedelsky, L. (1954). Absolute grading standards for objective tests. *Educational and Psychological Measurement*, 16, 159-176.
- Pelton, T. W. (2002). *The accuracy of unidimensional measurement models in the presence of deviations for the underlying assumptions*. Unpublished doctoral dissertation, Brigham Young University, Department of Instructional Psychology and Technology.
- Samejima, F. (1973). A comment on Birbaum's three parameter logistic model in the latent trait theory. *Psychometrika*, 38, 221-233.
- SAS Institute. (1999). SAS online doc (Version 8) [Software manual on CD-ROM]. Cary, NC: Author.
- Thissen, D., & Steinberg, L. (1984). A response model for multiple choice items. *Psychometrika*, 49, 501-519.
- Thissen, D., & Wainer, H. (Eds.). (2001). *Item response models for items scored in two categories*. Mahwah, NJ: Lawrence Erlbaum.
- van der Linden, W. J., & Hambleton, R. K. (1997). *Handbook of modern item response theory*. New York: Springer Verlag.

- Weitzman, R. A. (1996). The Rasch model plus guessing. *Educational and Psychological Measurement*, 56, 779-790.
- Yen, W. M., Burket, G. R., & Sykes, R. C. (1991). Nonunique solutions to the likelihood equation for the three parameter logistic model. *Psychometrika*, 56, 39-54.

Acknowledgment

This study is financially supported by different grants: the FONDECYT-1030801 grant from Fondo Nacional de Ciencia y Tecnología de Chile; the BIL01/1 grant to Paul De Boeck and Geert Molenberghs (Flanders) for a collaboration with Pilar Iglesias, Guido del Pino, and Ernesto San Martín (Chile); and

a grant to the first author from the COIMBRA Group for a visit at the K. U. Leuven. The SIMCE Office from the Chilean government kindly allowed us access to the databases used in this study. The first author also acknowledges A. Jara (Biostatistics Center, K. U. Leuven) for his valuable help with SAS. Finally, the authors are grateful to the editor and the reviewers for the valuable and helpful comments they have received.

Author's Address

Address correspondence to E. San Martín, Department of Statistics, Pontificia Universidad Católica de Chile, Casilla 306, Santiago 22, Chile; e-mail: esanmart@mat.puc.cl.