



Published in final edited form as:

Educ Psychol Meas. 2012 April 1; 72(2): 264–290. doi:10.1177/0013164411410056.

Characterizing Sources of Uncertainty in IRT Scale Scores

Ji Seung Yang, MA, Mark Hansen, MPH, and Li Cai, PhD

University of California, Los Angeles

Abstract

Traditional estimators of item response theory (IRT) scale scores ignore uncertainty carried over from the item calibration process, which can lead to incorrect estimates of standard errors of measurement (SEM). Here, we review a variety of approaches that have been applied to this problem and compare them on the basis of their statistical methods and goals. We then elaborate on the particular flexibility and usefulness of a Multiple Imputation (MI) based approach, which can be easily applied to tests with mixed item types and multiple underlying dimensions. This proposed method obtains corrected estimates of individual scale scores, as well as their SEM. Furthermore, this approach enables a more complete characterization of the impact of parameter uncertainty by generating confidence envelopes (intervals) for item tracelines, test information functions, conditional SEM curves, and the marginal reliability coefficient. The MI based approach is illustrated through the analysis of an artificial data set, then applied to data from a large educational assessment. A simulation study was also conducted to examine the relative contribution of item parameter uncertainty to the variability in score estimates under various conditions. We found that the impact of item parameter uncertainty is generally quite small, though there are some conditions under which the uncertainty carried over from item calibration contributes substantially to variability in the scores. This may be the case when the calibration sample is small relative to the number of item parameters to be estimated, or when the IRT model fit to the data is multidimensional.

Keywords

IRT; IRT scoring; calibration; standard error of measurement; multiple imputation; confidence envelope

1 Introduction

In applications of item response theory (IRT), it is common to score individuals using maximum marginal likelihood (MML) estimates of item parameters, which are ideally obtained using a large and independent calibration sample. In this process, the standard errors of measurement are estimated as either the reciprocal of the square root of the Fisher information function evaluated at the posterior mode for maximum *a posteriori* (MAP) scoring (or maximum likelihood scoring if there is no prior), or the standard deviation of the posterior under expected *a posteriori* (EAP) scoring. The use of MML estimates in conjunction with MAP or EAP scoring is a form of Empirical Bayes (EB; Carlin & Louis, 2000).

However, both the scale score and its standard error—the latter in particular—are affected by the uncertainty in the item parameter estimates, which is carried over from the calibration

process (Tsutakawa & Johnson, 1990). For frequentists, the item parameter estimates differ from true parameter values due to sampling error, which is represented in the standard errors of the parameter estimates. From a likelihood perspective, the uncertainty is reflected by the curvature of the log-likelihood, often represented as the information matrix of the item parameters. For Bayesians, point estimates of item parameters alone do not convey information about the width of the posterior distributions, which is non-negligible unless the calibration sample size tends to infinity so that the posteriors become peaked. From any statistical point of view, then, uncertainty in the item parameter estimates is concerning.

Unfortunately, traditional scoring approaches (e.g., MAP or EAP) fail to acknowledge this uncertainty. Chief among the problems is an incorrect statement of measurement error. The scale score's standard error of measurement (SEM) is often underestimated and the scale's marginal reliability overestimated, which means that the scores may be taken as having greater precision than is justified. As Cheng and Yuan (2010) pointed out, an understated SEM can also lead to premature termination of a computerized adaptive test. The problem is most pronounced in situations where a small calibration sample is used, or when item parameters are estimated using the sample to be scored (Tsutakawa & Johnson, 1990).

Researchers have proposed a number of approaches to address this problem and related issues (e.g., Lewis, 1985; Tsutakawa & Soltys, 1988; Tsutakawa & Johnson, 1990; Thissen & Wainer, 1990; Mislevy & Yan, 1991; Mislevy, Wingersky, & Sheehan, 1993; Patz & Junker, 1999; Embretson, 1999; Lewis, 2001; Hoshino & Shigemasu, 2008; Cheng & Yuan, 2010; Zhang, Xie, Song, & Lu, 2011). In the research reported here, we first provide a brief review of existing approaches, comparing them with respect to both their underlying methods and intended objectives. In broad terms, three types of methods have been applied to this problem: analytic approximations, fully Bayesian sampling based approaches, and multiple imputation (MI) based approaches. These approaches have been utilized to accomplish two related, but distinct goals: 1) obtain corrected standard errors of measurement that take uncertainty in the item parameters into account, and 2) characterize the nature and impact of item parameter uncertainty on subsequent estimation and inference.

After reviewing these approaches, we provide a formal justification for the MI based strategy and illustrate its use with a three-item artificial data set. Extending Mislevy et al.'s (1993) results, we use MI to obtain corrected estimates of individual scale scores and their standard errors of measurement. Using MI, we also characterize the specific ways in which uncertainty in the item parameters can affect scoring. Thissen and Wainer (1990) proposed that it may be helpful to visualize the uncertainty by generating confidence envelopes for item characteristic curves. We carry this idea further by constructing confidence envelopes for test information functions and conditional SEM curves. Such depictions make it possible to observe the ways in which the effects of parameter uncertainty vary across the values of the latent trait. The MI based approach may also be used to obtain confidence intervals for the marginal reliability coefficient. After illustrating the approach with the artificial data set, we analyze data from a large educational assessment. Lastly, through simulation studies, we examine the ways in which model complexity, calibration sample size, and test length contribute to these effects.

Our approach improves upon the existing alternatives in several important ways:

- First, in contrast to analytical approximation methods (e.g., Cheng & Yuan, 2010; Zhang et al., 2011) that explicitly require the calculation of a number of nonstandard derivative matrices that are model-specific, the MI based approach is easily applied to any IRT model (e.g., uni- or multi-dimensional, dichotomous or polytomous) and scoring method (e.g., MAP, EAP, or even summed score EAP), provided that the asymptotic covariance matrix of the item parameter estimates is

available. This flexibility is an important feature, as we note that previous studies have not considered item parameter uncertainty in the context of polytomous or multidimensional IRT models. Here, we illustrate the proposed approach for tests of mixed item types (including 2- and 3-parameter logistic models for dichotomous data and a logistic graded response model for ordered responses) analyzed using either a unidimensional or a multidimensional model, specifically the item bifactor model (e.g., Gibbons & Hedeker, 1992 and Cai, Yang, & Hansen, in press)

- Second, in contrast to fully Bayesian methods (e.g., Patz & Junker, 1999) that must rely on Markov chain Monte Carlo (MCMC) to sample the intractable posterior distributions, the MI approach is computationally simple and efficient. It relies on information that is routinely printed in the output of most standard IRT software programs, in conjunction with the easily accomplished task of random sampling from the multivariate normal distribution.
- Third, in contrast to previous MI based methods, in which analyses were either limited to single items or the between-item parameter error covariances were not estimated (and, thus, treated as zero), our approach makes use of a modern estimation algorithm (Supplemented EM; Cai, 2008) for computing the asymptotic covariance matrix of the item parameters that is applicable to any IRT model. This covariance matrix is not an automatic byproduct of the current gold-standard Bock and Aitkin (1981) EM algorithm of item parameter estimation.
- Finally, this approach systematizes a number of seemingly disparate methods under a single framework, including Thissen and Wainer's (1990) confidence envelopes and Lewis's (1985) expected response functions. Even approximation methods based on pseudo maximum likelihood (e.g., Hoshino & Shigemasu, 2008) can be reinterpreted in light of the multiple imputation framework.

2 A Brief Review of Existing Approaches

2.1 Two- vs. Single-Stage

With the exception of fully Bayesian sampling based methods, nearly all existing proposals assume a two-stage process for estimating the IRT scale scores. The items are first calibrated, preferably using MML or Bayesian methods. In this stage, the item parameters are estimated, as well as their error covariance matrix. In the second stage, the item parameters are used to produce the IRT scale scores. Corrections are generally made in the second stage, utilizing a) the point estimates, b) the error covariance matrix, and c) either additional derivatives and linearization arguments for analytic approximations or, in the case of MI methods, random sampling from an approximation to the posterior distributions of the item parameters. Two-stage methods are popular not only because they generally lead to sound estimates, but also because they are consistent with the standard operating procedures in applied educational and psychological testing situations.

Fully Bayesian sampling based methods (e.g., Patz & Junker, 1999), on the other hand, involve a single stage. They rely on MCMC to produce random draws from a Markov chain having the full joint posterior of the item parameters and the individual latent variables as its invariant distribution. Under the ergodicity of the Markov chain, dependent samples from the chain can be used to approximate the full posterior. When inferences regarding individual latent traits are desired, one simply “marginalizes” over the other dimensions of the posterior, i.e., by integrating out the item parameters, and the IRT scale scores thus obtained (e.g., as posterior means) would already have taken the uncertainty in item parameters into account. Given the MCMC output, marginalization amounts to ignoring that part of the MCMC output related to the item parameters. The single-stage approach is appealing conceptually. However, a significant barrier to its wide-spread adoption is its

complexity, computational intensiveness, and in some cases, inflexibility. Even as MCMC gains popularity, its proper usage still requires considerably more effort and statistical expertise when compared to more deterministic and better understood methods such as the EM algorithm. As Edwards (2005) noted, a number of competing MCMC samplers have been proposed for IRT, and the question of their relative algorithmic efficiency has not been entirely settled in the methodological literature. Furthermore, we point out that in contrast to two-stage IRT scoring methods, it is difficult to conceive how the single-stage approach can easily accommodate certain non-standard but nevertheless popular and useful IRT scoring algorithms such as summed-score to EAP translations (Thissen & Wainer, 2001).

2.2 Two Related Goals: Correction vs. Characterization

Most existing approaches have sought solely to obtain corrected estimates of SEM for IRT scale scores that take into account the uncertainty in the item parameters. These include analytic approximations based on Bayesian calculations (Tsutakawa & Soltys, 1988; Tsutakawa & Johnson, 1990), analytic approximations based on pseudo maximum likelihood (Cheng & Yuan, 2010; Hoshino & Shigemasu, 2008), as well as MI-based expected response functions (Lewis, 1985; Lewis, 2001; Mislevy et al., 1993). MCMC methods (e.g., Patz & Junker, 1999) may also be regarded as belonging to this category.

For analytic approximations (Bayesian or likelihood), the asymptotic argument is typically based on Taylor series expansions of the nonlinear estimating equations of the IRT scale scores. From the first (sometimes second) order approximation, corrected standard error formulae are obtained. In contrast, the expected response functions are motivated by an explicit analogy to multiple imputation for treating missing data (Rubin, 1987). Using the point estimates of item parameters to represent the item response functions amounts to a single round of imputation, replacing the unknown (hence “missing”) item parameter values by the modes of their marginal log-likelihood. In contrast, a multiple imputation approach uses more than one random imputations to recover the uncertainty due to not knowing the item parameters exactly. Averaged response functions under multiple random imputations are expected response functions.

Thissen and Wainer (1990) demonstrated an approach with an entirely different goal. Rather than seeking corrected estimates, they used confidence envelopes around the item response functions to reflect the uncertainty in the item parameters. This approach can be viewed as complimentary to the expected response function approach of Mislevy et al. (1993). Unlike Mislevy et al. (1993), however, the focus of Thissen and Wainer (1990) is not on obtaining corrected SEM for the scale scores. The resulting graphical displays (referred to as the M-line plots) are simply used to characterize how errors in the item parameter estimates impact the plausible shapes of item characteristic curves. The confluence of Thissen and Wainer's (1990) confidence envelopes and Mislevy et al.'s (1993) expected response functions leads to the dominating insight of this research. That is, multiple imputation provides a natural framework to integrate the two goals, simultaneously providing corrected standard errors of measurement and visual representations of the uncertainty.

3 The Proposed Approach

3.1 Some Notation

Without loss of generality, let \mathbf{y}_i be an $n \times 1$ vector of observed item responses for an individual randomly sampled from an ability distribution with density $f(\theta)$, where θ is the latent trait. Suppose the item parameters for the n items are contained in γ . The IRT model postulates a conditional distribution for \mathbf{y} on θ and γ : $f(\mathbf{y}|\theta, \gamma)$. We assume that θ and γ are *a priori* independent. Integrating out the incidental parameter θ , we have $f(\mathbf{y}|\gamma) = \int f(\mathbf{y}|\theta,$

$\gamma) f(\theta)d\theta$. For a sample of N independent respondents, the distribution of response patterns is $f(\mathbf{Y}|\gamma) = \prod_{i=1}^N f(y_i|\gamma)$, where \mathbf{Y} is an $N \times n$ matrix of item responses. Given prior distribution $\pi(\gamma)$, the posterior of the item parameters is

$$f(\gamma|\mathbf{Y}) = \frac{f(\mathbf{Y}|\gamma)\pi(\gamma)}{\int f(\mathbf{Y}|\gamma)\pi(\gamma)d\gamma}.$$

If the prior is taken to be uniform, then the posterior is proportional to the marginal log-likelihood $\log L(\gamma|\mathbf{Y})$ based on an observed matrix of item responses \mathbf{Y} (the calibration sample). Numerical optimization of the log-likelihood leads to the MML estimator $\widehat{\gamma}_N$, and Bock and Aitkin's (1981) EM algorithm is often used. The log-likelihood curvature around the mode is usually presented as the asymptotic covariance matrix \mathcal{I}_N^{-1} , where \mathcal{I}_N is the information matrix based on a calibration sample of size N . Recently Cai (2008) proposed the use of a Supplemented EM algorithm (Meng & Rubin, 1991) to compute \mathcal{I}_N for IRT models.

Due to the Bernstein-von Mises phenomenon (see e.g., van der Vaart, 1998), it is well-known that a multivariate normal with mean $\widehat{\gamma}_N$ and a covariance matrix equal to \mathcal{I}_N^{-1} provides a reasonable approximation to the posterior $f(\gamma|\mathbf{Y})$. With an abuse of notation, we may write

$$f(\gamma|\mathbf{Y}) \stackrel{a}{=} \phi(\gamma; \widehat{\gamma}_N, \mathcal{I}_N^{-1}), \quad (1)$$

where $\phi(\cdot; \widehat{\gamma}_N, \mathcal{I}_N^{-1})$ represents a normal distribution with mean $\widehat{\gamma}_N$ and covariance matrix \mathcal{I}_N^{-1} and $\stackrel{a}{=}$ indicates asymptotic equivalence (in N). In other words, the (analytically intractable) true posterior $f(\gamma|\mathbf{Y})$ becomes the gold standard here, and $\phi(\cdot; \widehat{\gamma}_N, \mathcal{I}_N^{-1})$ provides a convenient approximation. If $\phi(\cdot; \widehat{\gamma}_N, \mathcal{I}_N^{-1})$ provides a good enough approximation to $f(\gamma|\mathbf{Y})$, estimands (e.g. IRT scores, information, and standard errors of measurement) that are based on $f(\gamma|\mathbf{Y})$ should also be approximated well.

3.2 Illustration

To illustrate the proposed MI based procedure, we created a simple data set consisting of 500 simulated responses to three items. Each of the items was of a different type: two-parameter logistic (2PL), three-parameter logistic (3PL), and three-category graded response (GRM3). Although this illustration is somewhat unrealistic, the very short test length offers some advantages. Most importantly, it allows us to easily present the generating and estimated values of all eight item parameters, their error covariance matrix, and examples of the randomly imputed parameter sets. In addition, the numbers of possible response patterns (twelve) and summed scores (five) for this three-item test are small, allowing us to demonstrate the MI based scoring for each possibility. Finally, despite the simplicity of this example, it nonetheless features two kinds of complexity not previously considered in studies addressing parameter uncertainty: polytomous items (as represented by item 3) and tests of mixed item types. These features will also be present in the real data set considered later.

The generating item slopes and intercepts for this illustration were randomly drawn from distributions chosen to resemble values commonly observed in educational or psychological

testing. Both data generation and parameter estimation were conducted using IRTPRO (Cai, du Toit, & Thissen, forthcoming).

Table 1 shows the generating values, along with the MML estimates and their error covariance matrix. The standard errors of the item parameter estimates are the square roots of the diagonal elements of this matrix. The off-diagonal elements of the matrix represent covariances between the parameters. The fact that some of these elements are non-zero is notable, as prior methods seeking to account for parameter uncertainty have ignored these covariances. In the case of the 3PL model (used for item 2), the guessing parameter is expressed as the logit of the guessing probability g ; the generating value of -1.39 corresponds to a probability of about 0.2, as might be expected for a multiple choice item with five options. The MML parameter estimates and the error covariance matrix will be used in the MI based approach, as we describe in the following sections.

3.3 Multiple Imputation Inference for EAP Scores

3.3.1 Preliminaries—We now consider EAP scoring under item parameter uncertainty. The logic developed in this section, however, is quite general. Though we will assume a unidimensional θ for simplicity of notation, we note that the methods apply directly to multidimensional IRT models where θ is a vector. In the ideal case where item parameters are known, inference for θ for an individual with $n \times 1$ response pattern \mathbf{x} should be based on the following posterior

$$f(\theta|\mathbf{x}, \gamma) = \frac{f(\mathbf{x}|\theta, \gamma) f(\theta)}{\int f(\mathbf{x}|\theta, \gamma) f(\theta) d\theta} = \frac{f(\mathbf{x}|\theta, \gamma) f(\theta)}{f(\mathbf{x}|\gamma)}. \quad (2)$$

The EAP estimator with given γ is an expectation over the posterior distribution in Equation (2)

$$\widehat{\theta}_\gamma = \int \theta f(\theta|\mathbf{x}, \gamma) d\theta = \frac{1}{f(\mathbf{x}|\gamma)} \int \theta f(\mathbf{x}|\theta, \gamma) f(\theta) d\theta, \quad (3)$$

with SEM given by the square root of posterior variance

$$V(\widehat{\theta}_\gamma) = \int (\theta - \widehat{\theta}_\gamma)^2 f(\theta|\mathbf{x}, \gamma) d\theta = \frac{1}{f(\mathbf{x}|\gamma)} \int \theta^2 f(\mathbf{x}|\theta, \gamma) f(\theta) d\theta - \widehat{\theta}_\gamma^2. \quad (4)$$

If the item parameters are unknown, standard practice is to use the “plug-in” estimator, in which the MML estimates $\widehat{\gamma}_N$ are used in place of γ . Notice that unless N tends to infinity, the traditional estimator

$$\widehat{\theta}_{\widehat{\gamma}_N} = \int \theta f(\theta|\mathbf{x}, \widehat{\gamma}_N) d\theta$$

ignores the inherent variability of $\widehat{\gamma}_N$ as reflected by \mathcal{I}_N^{-1} . The traditional SEM estimator $V(\widehat{\theta}_{\widehat{\gamma}_N})$ also ignores this variability.

3.3.2 A Formal Justification for the Proposed Multiple Imputation Procedure

Tsutakawa and Johnson (1990) demonstrated that to properly account for the uncertainty in $\widehat{\gamma}_N$, one must base the inference for θ on the posterior distribution of θ given \mathbf{x} and \mathbf{Y} , which they represented as (their Equation 14, presented here with slight notational change)

$$f(\theta|\mathbf{x}, \mathbf{Y}) = \frac{f(\theta)}{f(\mathbf{x}|\mathbf{Y})} \int f(\mathbf{x}|\theta, \gamma) f(\gamma|\mathbf{Y}) d\gamma. \quad (5)$$

Note that a critical feature of Equation (5) is that the posterior of the item parameters $f(\gamma|\mathbf{Y})$ from calibration now serves as the prior. It can be shown (see Appendix A) that the EAP estimator of θ can be represented as

$$\widehat{\theta} = \int \widehat{\theta}_\gamma f(\gamma|\mathbf{x}, \mathbf{Y}) d\gamma. \quad (6)$$

This estimator does not depend on any particular values of $\widehat{\gamma}$ because it averaged over all plausible values of γ . The posterior variance (see Appendix A)

$$V(\widehat{\theta}) = \int V(\widehat{\theta}_\gamma) f(\gamma|\mathbf{x}, \mathbf{Y}) d\gamma + \int (\widehat{\theta}_\gamma - \widehat{\theta})^2 f(\gamma|\mathbf{x}, \mathbf{Y}) d\gamma \quad (7)$$

also automatically takes the uncertainty in γ into account. Equation (6) shows that $\widehat{\theta}$ is an expectation of $\widehat{\theta}_\gamma$ with respect to $f(\gamma|\mathbf{x}, \mathbf{Y})$. Equation (7) reveals a familiar variance decomposition. The posterior variance is equal to the sum of two components, the expectation of a variance and the variance of an expectation. The first component may be conceived of as the “within” variance, and the second is the “between” variance.

It turns out that $f(\gamma|\mathbf{x}, \mathbf{Y})$ can be treated as $f(\gamma|\mathbf{Y})$ for a number of reasons. If \mathbf{x} is actually a component of \mathbf{Y} (i.e., calibration and scoring for the same sample), then the appropriate notation should replace \mathbf{Y} by $\mathbf{Y}_{(\mathbf{x})}$, where $\mathbf{Y}_{(\mathbf{x})}$ denotes the response pattern matrix without observation \mathbf{x} . However, $f(\gamma|\mathbf{x}, \mathbf{Y}_{(\mathbf{x})})$ coincides with $f(\gamma|\mathbf{Y})$. If \mathbf{x} is an independent observation (i.e., the scoring sample), IRT scoring requires that $f(\gamma|\mathbf{x}, \mathbf{Y})$ be the same as $f(\gamma|\mathbf{Y})$, as we do not wish to change the scoring criterion once the items have been calibrated. Furthermore, due to Equation (1), we can use a normal approximation of $f(\gamma|\mathbf{Y})$ as $\phi(\gamma; \widehat{\gamma}_N, \mathcal{I}_N^{-1})$. Consequently, we may replace Equations (6) and (7) with the following approximations:

$$\widehat{\theta} \stackrel{a}{=} \int \widehat{\theta}_\gamma \phi(\gamma; \widehat{\gamma}_N, \mathcal{I}_N^{-1}) d\gamma \quad (8)$$

$$V(\widehat{\theta}) \stackrel{a}{=} \int V(\widehat{\theta}_\gamma) \phi(\gamma; \widehat{\gamma}_N, \mathcal{I}_N^{-1}) d\gamma + \int (\widehat{\theta}_\gamma - \widehat{\theta})^2 \phi(\gamma; \widehat{\gamma}_N, \mathcal{I}_N^{-1}) d\gamma. \quad (9)$$

Now, we have a multiple imputation algorithm to approximate $\widehat{\theta}$ and $V(\widehat{\theta})$:

1. Draw $M > 1$ sets of values from a multivariate normal distribution with mean $\widehat{\gamma}_N$ and covariance matrix \mathcal{I}_N^{-1} . Denote them as γ_j for $j = 1, \dots, M$.
2. Plug each γ_j into Equations (3) and (4) and compute $\widehat{\theta}_{\gamma_j}$ as well as $V(\widehat{\theta}_{\gamma_j})$. Write $\widehat{\theta}_j$ as shorthand for $\widehat{\theta}_{\gamma_j}$ and V_j for $V(\widehat{\theta}_{\gamma_j})$.
3. The multiple imputation EAP approximation to $\widehat{\theta}$ is the empirical average,

$$\bar{\theta} \approx \frac{1}{M} \sum_{j=1}^M \widehat{\theta}_j, \quad (10)$$

which approximates the right hand side expression of Equation (8).

4. The multiple imputation variance approximation is

$$V(\bar{\theta}) \approx \bar{V} + (1+M^{-1})B, \quad (11)$$

where $\bar{V} = M^{-1} \sum_{j=1}^M V_j$ is an estimate of the “within” imputation variance (expectation of a variance) and $B = (M-1)^{-1} \sum_{j=1}^M (\hat{\theta}_j - \bar{\theta})^2$ is an estimate of the “between” imputation variance (variance of an expectation). These correspond to the two parts on the right hand side of Equation (9).

The ratio

$$r = \frac{(1+M^{-1})B}{\bar{V}} \quad (12)$$

is known as *relative increase in variance* due to missing data (Schafer, 1997). It can be taken as a crude measure of the impact of item parameter uncertainty. We will use this measure in the simulation studies to be described in Section 5.

As previously mentioned, the MI based approach may be applied to any number of other scoring methods. Thus, in our analyses, we present results not only for pattern EAP scoring, but also for summed score to EAP translations (Lord & Wingersky, 1984). As will be demonstrated, the impact of item parameter uncertainty may differ according to scoring method.

3.3.3 Illustration for the Proposed Procedure—The MI based procedure was applied to the three-item data set, following the steps described above.

1. We drew $M = 20$ sets of plausible item parameter values from a multivariate normal distribution with mean $\hat{\gamma}_N$ and covariance matrix \mathcal{S}_N^{-1} , both given in Table 1. We denote the plausible parameter values as γ_j for $j = 1, \dots, 20$. The imputed sets of γ_j are presented in Table 2.
2. We now plug each γ_j into Equations (3) and (4). We obtain a different posterior distribution for each vector γ_j . This is illustrated in Figure 1, in which the prior distribution, item tracelines, and posteriors distributions are shown for three response patterns. The posteriors have different means and standard deviations.
3. The MI based approximation to $\bar{\theta}$ is calculated by taking the average of $\hat{\theta}_{\gamma_j}$ across all imputations.
4. The MI based variance approximation $V(\bar{\theta})$ is calculated using Equation (11), and the relative increase in variance r is calculated using Equation (12). Table 3 shows \bar{V} , B , $V(\bar{\theta})$, and r for both full pattern EAP scoring and the summed score EAP translations. For the full pattern EAP, the relative increase in variance r ranged from 0.2% to 10.8%. For the summed score EAP translations, the observed values of r had a much narrower distribution, between 1.4% and 1.6%.

The scores obtained using the MML parameter estimates are compared with MI based scores in Figure 2. These two sets of scores are almost perfectly correlated, for both the response pattern EAPs and summed score EAPs. This result is consistent with previous studies; while

item parameter uncertainty may result in overestimates of score precision, it does not necessarily result in biased scores.

In Figure 3 the SEM are plotted against the score estimates. There are differences in the SEM. As expected, the SEMs are generally larger for the MI scores. The magnitude of the correction, however, is substantially larger for the full pattern EAP scores than for the summed score EAPs (consistent with the values of r reported in Table 4).

3.4 Multiple Imputation Confidence Envelopes

3.4.1 Confidence Envelopes for Item Characteristic Curves—Thissen and Wainer (1990) considered confidence envelopes for the item characteristic curves. Let $T_k(\theta|\gamma)$ denote a generic item category response curve for category k . For example, for the 2PL IRT model, the characteristic curve for the “correct” response is

$$T_1(\theta|\gamma) = \frac{1}{1 + \exp[-(c + a\theta)]}$$

where c is the item intercept, a is the item slope, and both are components of γ . When γ is known without error, the curves are simply mathematical functions of the item parameters. When estimates of item parameters are used, $T_k(\theta|\widehat{\gamma}_N)$ contains uncertainty. Thissen and Wainer (1990) reasoned that since the posterior distribution of γ is reasonably well approximated by a normal with mean $\widehat{\gamma}_N$ and covariance matrix \mathcal{J}_N^{-1} (i.e., Equation 1), if one can produce multiple random samples of γ from this approximate posterior, approximate confidence limits of the item characteristic curves can be found by plotting the randomly varying item characteristic curves over repeated imputations. These so-called “M-line” plots are generated by Thissen and Wainer (1990) to illustrate the shape of confidence envelopes for a variety of dichotomous IRT models. However, Thissen and Wainer's (1990) method has the same limitations as Mislevy et al.'s (1993) expected response function method. The random sampling of the item parameters are done in an item by item manner because they do not have access to the full error covariance matrix of the item parameter estimates.

We present here a slight variation of Thissen and Wainer's (1990) basic idea:

1. Draw $M > 1$ sets of values from a multivariate normal distribution with mean $\widehat{\gamma}_N$ and covariance matrix \mathcal{J}_N^{-1} . Denote them as γ_j for $j = 1, \dots, M$.
2. Choose a reasonably fine grid to numerically represent θ , e.g., from -3 to $+3$ in step sizes of .01.
3. Plug each γ_j into $T_k(\theta|\gamma)$ and at a chosen θ level, empirically locate the upper and lower $1 - \alpha/2$ quantile from the M values of $T_k(\theta|\gamma_j)$.
4. Repeat the last step for all θ levels to find a $(1 - \alpha) \times 100\%$ confidence envelope.

To ensure that the boundaries of the confidence envelopes are well-characterized, a large M is necessary. We generally use $M = 1000$ random draws.

Figure 4 presents confidence envelopes for the three items in our illustration. (see. The upper, middle, and lower panels correspond to item 1 (2PL), 2(3PL) and 3 (GRM) in order. The solid curve is the usual item characteristic curve based on the MML item parameter estimates. The dotted curves represent the upper and lower 95% confidence limits, based on $M = 1000$ random samples drawn from the multivariate normal approximation to the posterior distribution of the item parameters. Incidentally, we also computed the expected

response functions (dashed curves), obtained by averaging the response functions across all the imputed parameter sets. As Mislevy et al. (1993) noted, the expected response functions are not logistic and tend to have slightly lower slopes than the item response function based on the MML item parameter estimates.

3.4.2 Confidence Envelopes for Information and SEM Curves—For general multiple-categorical IRT models, item i 's Fisher information function is given by the following expression

$$F_i(\theta|\gamma) = - \sum_{k=0}^{K-1} T_k(\theta|\gamma) \frac{\partial^2}{\partial \theta^2} \log T_k(\theta|\gamma),$$

where K is the number of categories (see Baker & Kim, 2004). The Fisher information functions are additive. For n items, the test information function is the sum of item information functions

$$F(\theta|\gamma) = \sum_{i=1}^n F_i(\theta|\gamma). \quad (13)$$

When a prior for the ability distribution is used in scoring (e.g., MAP scoring), test information must also include the contribution from the prior. The standard error of measurement curve is found as a one-to-one transformation of the test information curve

$$\text{sem}(\theta|\gamma) = \frac{1}{\sqrt{F(\theta|\gamma)}}. \quad (14)$$

Recognizing that $F(\theta|\gamma)$ is a nonlinear transformation of γ , we follow essentially the same strategy used in the previous sections to create confidence envelopes for the test information function:

1. Draw $M > 1$ sets of values from a multivariate normal distribution with mean $\widehat{\gamma}_N$ and covariance matrix \mathcal{J}_N^{-1} . Denote them as γ_j for $j = 1, \dots, M$.
2. Choose a reasonably fine grid to numerically represent θ , e.g., from -3 to $+3$ in step sizes of .01.
3. Plug each γ_j into $F(\theta|\gamma)$ and at a chosen θ level, empirically locate the upper and lower $1 - \alpha/2$ quantile from the M values of $F(\theta|\gamma_j)$.
4. Repeat the last step for all θ levels to find a $(1 - \alpha) \times 100\%$ confidence envelope.

Because $F(\theta|\gamma)$ and $\text{sem}(\theta|\gamma)$ enjoy one-to-one relation, the confidence limits for the SEM curve are found by transforming the confidence limits of the test information function.

As an illustration, consider the confidence envelopes for the test information function and the conditional SEM for the three-item data set, presented in Figure 5.

These curves highlight the extent to which our characterizations of test information or the conditional SEM can be impacted by item parameter uncertainty. This may have implications for test assembly, in which particular combinations of items may be selected in order to produce a particular test information or conditional SEM curve. The width of the confidence envelopes we observe for these functions suggests that assembling tests to

closely match target information of SEM curves may be misguided. The confidence envelopes generated through the MI based approach allow the imprecision due to uncertainty in item parameters to be visualized.

3.4.3 Confidence Intervals for Marginal Reliability—As a final application, we consider confidence intervals for the marginal reliability coefficient, which is an important IRT-based measure of overall scale reliability. Let the prior $f(\theta)$ for the ability distribution have variance σ^2 . Marginal reliability is defined as

$$\bar{\rho}(\gamma) = \frac{\sigma^2 - \int [\text{sem}(\theta|\gamma)]^2 f(\theta) d\theta}{\sigma^2}, \quad (15)$$

(see e.g., Thissen & Wainer, 2001), where the integral in the numerator returns the average error variance, and $\text{sem}(\theta|\gamma)$ is as defined in Equation (14). Of course, $\bar{\rho}$ is also a nonlinear transformation of γ , suggesting the following algorithm:

1. Draw $M > 1$ sets of values from a multivariate normal distribution with mean $\widehat{\gamma}_N$ and covariance matrix \mathcal{S}_N^{-1} . Denote them as γ_j for $j = 1, \dots, M$.
2. Plug each γ_j into $\bar{\rho}(\gamma)$ and empirically locate the upper and lower $1 - \alpha/2$ quantile from the M values of $\bar{\rho}(\gamma_j)$. The result is a $(1 - \alpha) \times 100\%$ confidence interval for marginal reliability.

For the three-item data set, the MML item parameter estimates yield an estimated marginal reliability of .29 (which is low, as expected, given the very short test length). However, the MI based approach offers a more complete characterization—that the marginal reliability has a 95% confidence interval between .17 and .43 based on $M = 1000$ imputations.

3.5 Summation

The proposed multiple imputation approach not only provides corrected scale scores and standard errors of measurement that take parameter uncertainty into account, but also confidence envelopes or intervals for other important quantities such as the item characteristic curve, the test information function, the conditional SEM curve, and the marginal reliability coefficient. As demonstrated with the three-item data set, the approach is flexible enough to be applied to not only full pattern EAP scores, but also other scoring methods (e.g. summed score EAP) and various IRT models (e.g., 2PL, 3PL, or GRM).

4 Application to empirical data

4.1 Data and methods

As an empirical demonstration of the proposed multiple imputation approach, we analyzed data from the 2000 Program for International Student Assessment (PISA; Adams & Wu, 2002). We extracted a random sample of 1500 students from the United States with complete responses for Math Booklet 8. For our analyses, we used 500 students as a calibration sample and 1000 students as a scoring sample. Math Booklet 8 consists of 15 items—8 free response (FR), 5 multiple choice (MC), and 2 complex multiple choice (CMC). Logistic graded response models were fit to the FR and CMC items with the number of ordered categories determined by the number of different scores assigned (either 2 or 3). A 3PL model was used for the MC items. In estimating item parameters, a log-normal prior was placed on the guessing parameter. The mean of this prior was based on the expected probability of a correct response, assuming blind guessing and given the number of response

options. Item parameter estimates and the item parameter asymptotic covariance matrix were obtained using IRTPRO (Cai et al., forthcoming).

Because PISA items are nested within testlets or passages, we obtained estimates for both a unidimensional (ignoring the testlet structure but consistent with how the test data are scored in practice) and a bifactor item response model with specific factors modeling testlet effects. Parameter estimation and scoring for the bifactor model has been described elsewhere (Cai et al., in press). Here, we focus on the effects of item parameter uncertainty on the precision of scores for the general dimension and show the application of the MI based method in both the unidimensional and multidimensional IRT contexts.

As outlined in previous sections, the MML item parameter estimates and the item parameter error covariance matrix were used to generate multiple sets of plausible parameter values. For the unidimensional model, 39 parameters were estimated. The bifactor model required estimation of additional 10 parameters (specific factor slopes for the 10 items loading on 3 specific factors).

As before, 20 sets of parameters were imputed to obtain the MI based score estimates based on the full response patterns and summed scores for individuals in the scoring sample. Confidence envelopes for test information and the conditional SEM and a confidence interval of the marginal reliability coefficient were based on $M = 1000$ imputations.

4.2 Results

As shown in Figure 6, the scale score estimates are almost perfectly correlated for both unidimensional model and bifactor model, regardless of the scoring method (full response pattern or summed score EAP)

Figure 7 shows the scale score and SEM estimates and their corrections based on the MI approach. For most individuals in the scoring sample, the SEM increases under the MI based approach. Across the scoring sample, the average relative increase in variance (r) based on Equation (12) was 3.5% for the unidimensional model and 7.6% for the bifactor model under full response pattern scoring. In contrast, the average relative increase in variance for the summed score EAP estimates was less than 1% for either scoring method. In addition to having larger average increases, the impact of item parameter uncertainty on the full pattern scores appears to be more variable than for the summed score translations. Specifically, there are some response patterns with vary large increases in the SEM (up to 32.9% for the unidimensional model and 67.1% for the bifactor model). The increases for the summed score a much more uniform across all response patterns.

The 95% confidence envelopes for test information and the conditional SEM for the unidimensional model are shown in Figure 8. Graphical representations and proper interpretation of test information in the multidimensional case, though potentially quite interesting, are beyond the scope of this study.

For this 15-item test, the MML parameter estimates produce a marginal reliability of .82. From the MI based approach, its 95% confidence interval was found to be .78 to .84.

5 A Simulation Study

Previous studies have shown that the magnitude of bias in the standard errors of measurement can vary depending on conditions such as item response model complexity, calibration sample size, and test length all of which influence the test information. For example, Tsutakawa and Johnson (1990) observed that relatively simple Rasch models tend to provide decent estimation of the latent ability levels and their standard errors of

measurement, even with small calibration samples. In contrast, a three-parameter logistic model with a calibration sample size of 400 individuals produced biases the posterior means of the latent trait and underestimation the posterior standard deviations by more than 40 percent, on average.

As a complement to the simple illustration and the empirical analysis of PISA data, a simulation study was conducted. The goal was further characterize the conditions in which the uncertainty carried over from item calibration leads to uncertainty in the scoring process. Based on the findings of previous studies, we manipulated the following: model complexity, calibration sample size, and test length. For this study, the number of imputations was fixed to $M = 20$, as previous research has already demonstrated that 20 imputations should provide reasonably good correction of the standard errors of measurement (Mislevy et al., 1993). The combination of the number of items ($J = 5, 10, 20, 40$), the size of calibration sample ($N = 500, 1000, 2000, 5000$), and item type (2PL, 3PL, and logistic GRM with 5 categories) yielded a total of 48 conditions.

For each condition, 500 calibration samples were generated; in addition, one independent scoring sample of 10,000 cases was produced. The item parameters used for data generation were chosen to resemble estimates obtained from real educational and psychological data sets. Following data generation, we calibrated the items using MML estimation and obtained the parameter error covariance matrix. Twenty item parameter sets were randomly drawn from the multivariate normal approximation to the posterior distribution of the item parameters, as described in previous sections. These parameter sets were then used to obtain the MI based scores and SEM for the scoring sample. The relative increase in variance (r) was also calculated for each case in the scoring sample. The values of r were averaged across the many calibrations (500 replications were used for most conditions; however, for conditions). The mean values of r were then averaged within deciles (1000 individuals) of the “true” level of the latent trait and across the entire scoring sample for each condition. The results are reported in Table 5.

For most conditions, the average relative increase in variance is rather small, ranging from 1.8% to 14.7%. However, some trends are evident. For a given test (i.e, for a fixed number of items and item type), the average relative increase in variance decreases as the size of the calibration sample increases. This is to be expected, as larger calibration samples yield more precise estimates of the item parameters. The patterns of relative increase across tests of different length and across item types are more complex. Test information increases with the number of items. Thus, for longer tests, we expect smaller “within” imputation error variance. At the same time, longer tests require estimation of more item parameters, resulting in greater “between” variance. In cases where the calibration sample is small, those parameter estimates may not be very precise. Consequently, the highest relative increase in variance was observed for the 40-item tests with the smallest calibration sample examined. Even for these conditions, however, the increases appear rather modest.

Figure 9 presents a more detailed view of the relative increase in variance for the 5-category graded items across the simulation conditions of varied test length and calibration sample size. Here, the average relative increase in variance is obtained within each decile of the true scores on the latent trait (which are known for these simulated data but, of course, are never known in practice). As in Table 5, it is clear that the relative increase in variance diminishes with increasing sample size. It is also apparent that the percentages are not uniform across the range of the latent trait. For those conditions where there is substantial relative increase in variance (e.g, $n = 40, N = 500$), the percentages are greatest at the highest and lowest deciles of the true θ 's. This is again consistent with existing results obtained by analytic approximation (e.g., Cheng & Yuan, 2010) for simpler IRT models.

6 Discussion

In this research, we have proposed a multiple imputation based approach for characterizing the ways in which uncertainty about item parameters affects item response functions, test information functions, standard errors of measurement, and marginal reliability. We argue that the approach presents several advantages over the existing alternatives. First, it can be applied to a variety of IRT models and scoring methods. Second, it is computationally simple and utilizes information that is readily available in the output of standard IRT software programs. Third, the approach makes use of the asymptotic covariance matrix of the item parameters, obtained through an application of the Supplemented EM algorithm (Cai, 2008). This allows us to conduct analyses of entire tests, whereas past efforts have mostly focused on parameter uncertainty with respect to individual items. Finally, our proposed approach connects a variety of seemingly disparate methods that have been used to handle item parameter uncertainty. In so doing, our approach integrates the related goals of providing corrected standard errors of measurement and the (highly visual) characterization of the effects of uncertainty.

To demonstrate the relevance of the proposed approach to the problem of uncertainty in item parameters, we derived approximations for EAP scores and their SEM (Equations 8 and 9, respectively) that utilize the MML estimates of the item parameters and their covariance matrix. The MI based EAP approximation is an average of the EAP scores obtained with $M > 1$ sets of plausible values for the item parameters. The error variance for this estimate is a combination of the within and between imputation variances. The square root of the total error variance provides a corrected SEM.

After illustrating the proposed approach with a artificial, three-item data set, we examined data from the 2000 PISA math test. Crucial to the MI based method is the random sampling of plausible item parameter values from a multivariate normal approximation to the item parameter posteriors. The MI based score estimates and SEM are obtained by scoring with these imputed parameter sets and combining the results. In addition to these corrections, we constructed confidence envelopes for item characteristic curves, building on the work of Thissen and Wainer (1990). A depiction such as this helps to convey the extent to which error in the item parameters leads to uncertainty about the shape of the response functions. Importantly, the amount of uncertainty may vary across levels of the latent trait. A similar approach was taken to generate confidence envelopes for the test information function and the conditional SEM curve. Finally, we used MI to calculate confidence intervals for marginal reliability. This allows us to convey the uncertainty in the marginal reliability that is due to error in the estimation of the item parameters.

We also conducted simulations, which allowed us to investigate how various factors can influence the uncertainty in scores. In these simulations, we manipulated the size of the calibration sample, the length of the test, and the complexity of the response model. The effect of parameter uncertainty was quantified as the relative increase in error variance. This analysis demonstrated that, on the whole, parameter uncertainty contributes little to total error variance. However, in situations where the calibration sample is small and the number of items is large (and especially in the case of a more complex response model), the error carried over from item calibration may occasionally be non-negligible.

It is important to know the extent to which the latent trait estimates are uncertain because a number of critical decisions are based on the SEM. In variable-length computerized adaptive testing algorithms, for example, the SEM is often used as a termination criterion. In such cases, underestimation of SEM can result in premature termination of the test. In addition, such underestimation could result in flawed inferences concerning individuals' standing

relative to a certain performance standard or to one another. Specifically, scores might be assumed to have greater precision than is warranted, given the known uncertainty in the item parameters. Standard errors have also been incorporated into statistical models such as hierarchical linear models with latent variables (Raudenbush & Bryk, 2002). For such applications, improved SEM estimates will enhance estimation of regression parameters and their associated standard errors. The MI based approach presented in this paper provides a simple and flexible means of obtaining these improved estimates.

Acknowledgments

We acknowledge financial support from the following sources: Institute of Education Sciences (R305B080016 and R305D100039), and National Institute on Drug Abuse (R01DA026943 and R01DA030466).

The views expressed in this paper belong to the authors and do not reflect the views/policies of the funding agencies or grantees.

We thank Dr. David Thissen for comments on an earlier draft.

Appendix A

The EAP estimator of θ is a posterior expectation, i.e.,

$$\widehat{\theta} = \int \theta f(\theta | \mathbf{x}, \mathbf{Y}) d\theta.$$

From Equation (5), the equation above can be rewritten as

$$\widehat{\theta} = \frac{1}{f(\mathbf{x} | \mathbf{Y})} \int \left[\int f(\mathbf{x} | \theta, \gamma) f(\gamma | \mathbf{Y}) d\gamma \right] \theta f(\theta) d\theta.$$

Interchanging the order of integration, we see that

$$\begin{aligned} \widehat{\theta} &= \frac{1}{f(\mathbf{x} | \mathbf{Y})} \int \left[\int \theta f(\mathbf{x} | \theta, \gamma) f(\theta) d\theta \right] f(\gamma | \mathbf{Y}) d\gamma \\ &= \frac{1}{f(\mathbf{x} | \mathbf{Y})} \int \left[\int \theta f(\theta | \mathbf{x}, \gamma) d\theta \right] f(\mathbf{x} | \gamma) f(\gamma | \mathbf{Y}) d\gamma \\ &= \int \widehat{\theta}_\gamma f(\gamma | \mathbf{x}, \mathbf{Y}) d\gamma, \end{aligned}$$

where the last line requires the conditional independence of \mathbf{x} and \mathbf{Y} given γ . By the same token, the posterior variance can be written as

$$\begin{aligned} V(\widehat{\theta}) &= \int \theta^2 f(\theta | \mathbf{x}, \mathbf{Y}) d\theta - \widehat{\theta}^2 \\ &= \int \left[V(\widehat{\theta}_\gamma) + \widehat{\theta}_\gamma^2 \right] f(\gamma | \mathbf{x}, \mathbf{Y}) d\gamma - \widehat{\theta}^2 \\ &= \int V(\widehat{\theta}_\gamma) f(\gamma | \mathbf{x}, \mathbf{Y}) d\gamma + \int (\widehat{\theta}_\gamma - \widehat{\theta})^2 f(\gamma | \mathbf{x}, \mathbf{Y}) d\gamma. \end{aligned}$$

References

Adams, R.; Wu, M. PISA 2000 technical report. Organization for Economic Cooperation and Development; Paris, France: 2002.

Baker, FB.; Kim, S-H. Item response theory: Parameter estimation techniques. Marcel Dekker; New York: 2004.

- Bock RD, Aitkin M. Marginal maximum likelihood estimation of item parameters: Application of an EM algorithm. *Psychometrika*. 1981; 46:443–459.
- Cai L. SEM of another flavour: Two new applications of the supplemented EM algorithm. *British Journal of Mathematical and Statistical Psychology*. 2008; 61:309–329. [PubMed: 17971266]
- Cai, L.; du Toit, SHC.; Thissen, D. IRTPRO: Flexible, multidimensional, multiple categorical IRT modeling [Computer software]. SSI International; Chicago, IL: forthcoming
- Cai L, Yang JS, Hansen MP. Generalized full-information item bifactor analysis. *Psychological Methods*. in press.
- Carlin, BP.; Louis, TA. Bayes and empirical Bayes methods for data analysis. 2nd ed.. Chapman & Hall/CRC; Boca Raton, FL: 2000.
- Cheng Y, Yuan K-H. The impact of fallible item parameter estimates on latent trait recovery. *Psychometrika*. 2010; 75:280–291. [PubMed: 21076651]
- Edwards, MC. Unpublished doctoral dissertation. Department of Psychology, University of North Carolina at Chapel Hill; 2005. A Markov chain Monte Carlo approach to confirmatory item factor analysis.
- Embretson SE. Generating items during testing: Psychometric issues and models. *Psychometrika*. 1999; 64:407–433.
- Gibbons RD, Hedeker D. Full-information item bifactor analysis. *Psychometrika*. 1992; 57:423–436.
- Hoshino T, Shigemasu K. Standard errors of estimated latent variable scores with estimated structural parameters. *Applied Psychological Measurement*. 2008; 32:181–189.
- Lewis, C. Estimating individual abilities with imperfectly known response functions. Annual Meeting of the Psychometric Society; Nashville, TN. 1985.
- Lewis, C. Expected response functions. In: Boomsma, A.; van Duijn, MAJ.; Snijders, TAB., editors. *Essays on item response theory*. Springer-Verlag; New York, NY: 2001.
- Lord FM, Wingersky MS. Comparison of IRT true-score and equipercentile observed-score “equatings”. *Applied Psychological Measurement*. 1984; 8:453–461.
- Meng X-L, Rubin DB. Using EM to obtain asymptotic variance-covariance matrices: The SEM algorithm. *Journal of the American Statistical Association*. 1991; 86:899–909.
- Mislevy, RJ.; Wingersky, MS.; Sheehan, KM. Dealing with uncertainty about item parameters: Expected response functions. Educational Testing Service; Princeton, NJ: 1993. Research Report No. 94-28
- Mislevy, RJ.; Yan, D. Dealing with uncertainty about item parameters: Multiple imputations and SIR. Annual Meeting of the Psychometric Society; Princeton, NJ. 1991.
- Patz RJ, Junker BW. A straightforward approach to Markov chain Monte Carlo methods for item response models. *Journal of Educational and Behavioral Statistics*. 1999; 24:146–178.
- Raudenbush, SW.; Bryk, AS. *Hierarchical linear models: Applications and data analysis methods*. 2nd ed.. Sage; Thousand Oaks, CA: 2002.
- Rubin, DB. *Multiple imputation for nonresponse in surveys*. Wiley; New York: 1987.
- Schafer, JL. *Analysis of incomplete multivariate data*. Chapman & Hall/CRC; New York: 1997.
- Thissen D, Wainer H. Confidence envelopes for item response theory. *Journal of Educational Statistics*. 1990; 15:113–128.
- Thissen, D.; Wainer, H. *Test scoring*. Lawrence Erlbaum Associates; Mahwah, NJ: 2001.
- Tsutakawa RK, Johnson JC. The effect of uncertainty of item parameter estimation on ability estimates. *Psychometrika*. 1990; 55:371–390.
- Tsutakawa RK, Soltys MJ. Approximation for Bayesian ability estimation. *Journal of Educational Statistics*. 1988; 13:117–130.
- van der Vaart, AW. *Asymptotic statistics*. Cambridge University Press; Cambridge, UK: 1998.
- Zhang J, Xie M, Song X, Lu T. Investigating the impact of uncertainty about item parameters on ability estimation. *Psychometrika*. 2011; 76(1):97–118.




Figure 1. Item trace lines and posterior distributions for three response patterns. Solid lines represent the tracelines and posteriors based on the MML parameter estimates; those based on M=20 randomly imputed parameter sets appear as dotted lines.

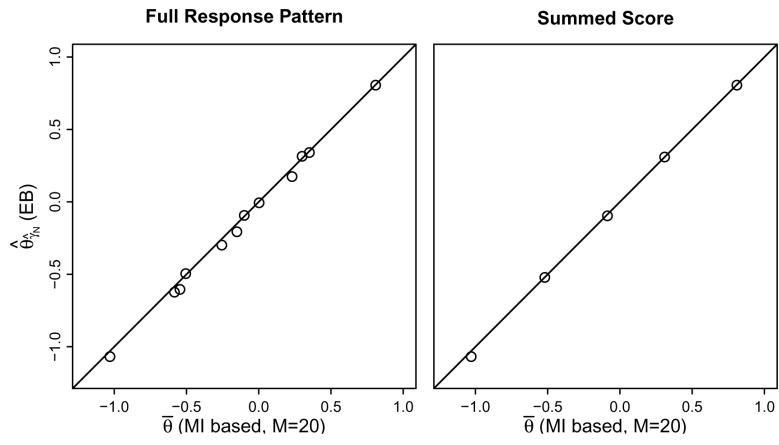


Figure 2. Traditional and MI based EAP estimates.

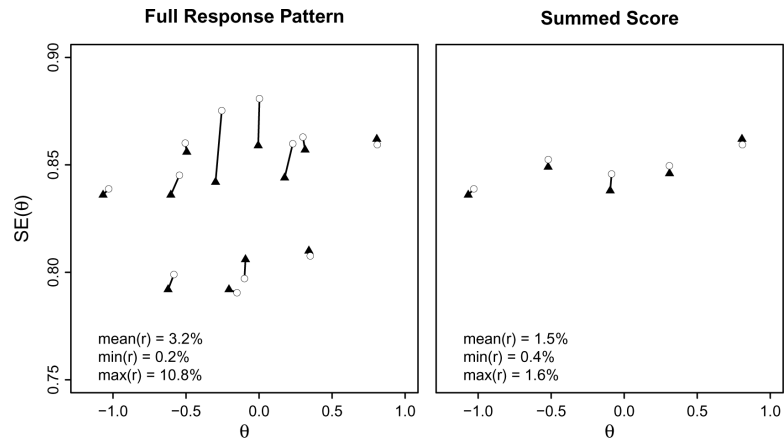


Figure 3. MI based scale score and SEM corrections. Traditional estimates are represented by the closed triangles. MI based estimates are shown as open circles.

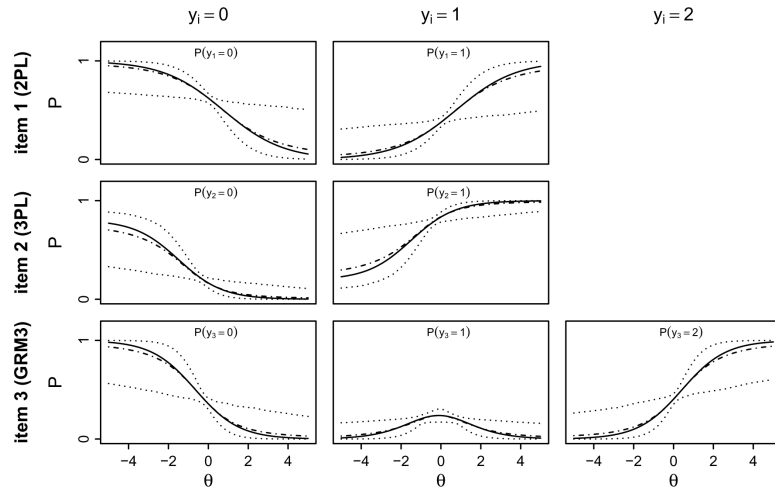


Figure 4. Confidence envelopes for item tracelines. The solid curves represent the item characteristic curves based on the MML item parameter estimates. The dotted curves represent the 95% confidence limits. The dashed curves show the expected response functions. The confidence limits and expected functions are generated based on $M = 1000$ imputations

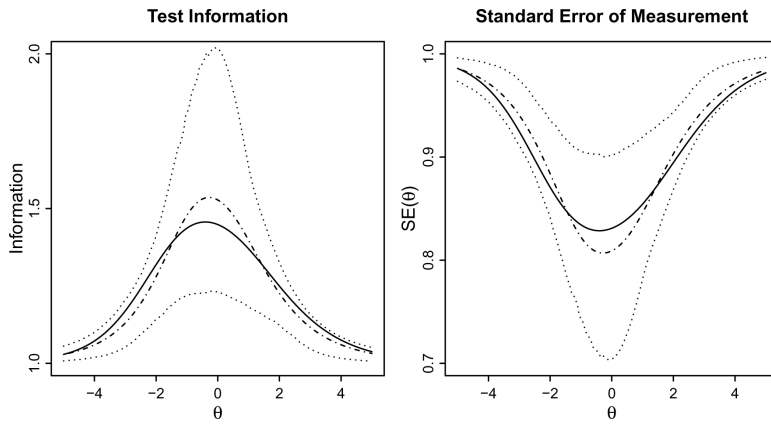


Figure 5. Confidence envelopes for test information and the conditional SEM. Solid curves show test information and the conditional SEM based on the MML item parameter estimates. The dotted curves represent the 95% confidence limits. The expected information and conditional SEM curves are shown as dashed lines. The confidence limits and expected functions are based on $M = 1000$ imputations.

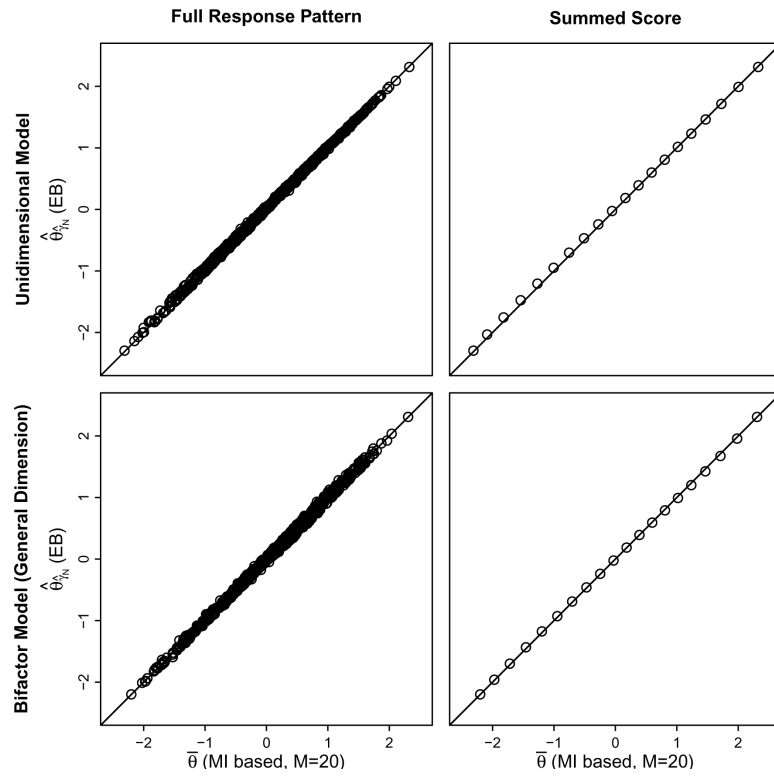


Figure 6.
 Estimated IRT Scale Scores for PISA Math Booklet 8.

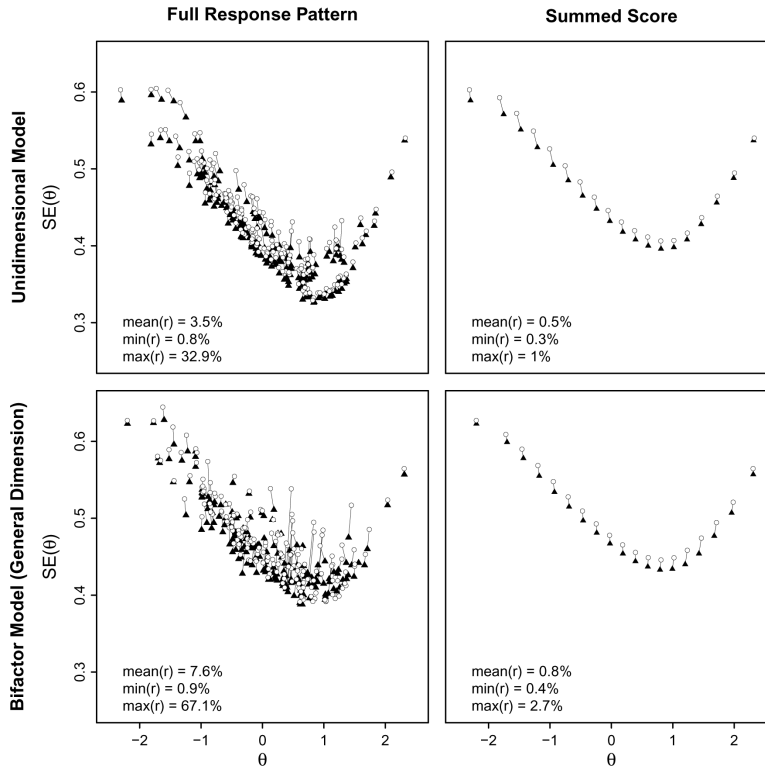


Figure 7. Traditional and MI based response pattern and summed score SEM estimates plotted against EAPs for PISA Math Booklet 8. For clarity, values for 200 randomly selected individuals in the scoring sample are plotted; the mean, minimum, and maximum values of r are based on the full scoring sample of 1000 individuals. Traditional estimates and SEMs based on the MML item parameters are shown as triangles. MI based results are shown as circles.




Figure 8. Confidence envelopes for test information and the conditional SEM for PISA Math Booklet 8. The solid curves show test information and the conditional SEM based on the MML item parameter estimates. The dotted curves represent the 95% confidence limits. The expected information and conditional SEM curves are shown as dashed lines. The confidence limits and expected functions are based on $M = 1000$ imputations.

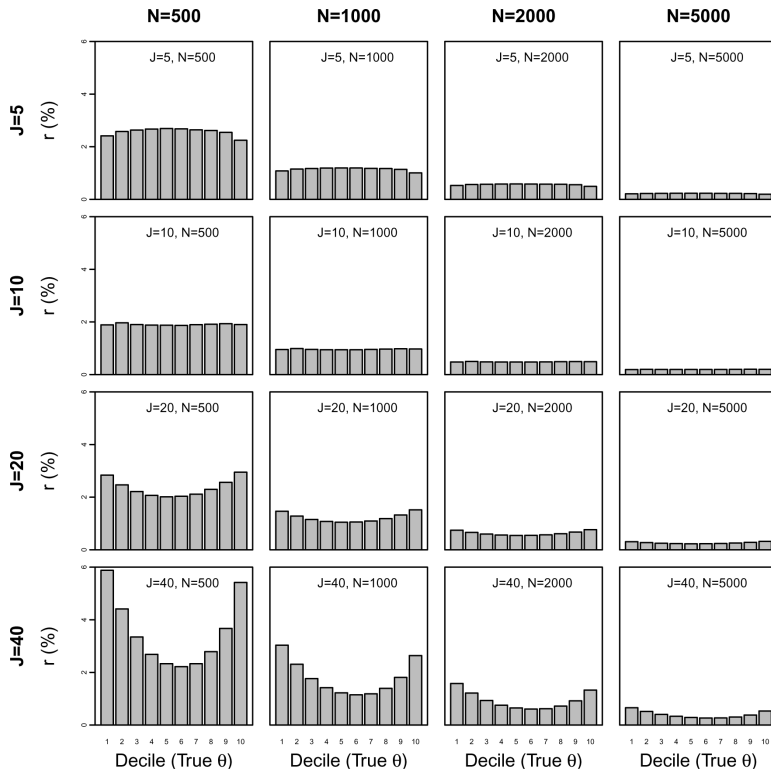


Figure 9. Mean relative increase in variance (by deciles) under various simulated test lengths and calibration sample sizes for the logistic graded response model with 5 categories.

Table 1

Item parameter estimates and the asymptotic covariance matrix for the 3-item artificial data set

Item	Parameter	True Value	$\hat{\gamma}_N$ (SE)	Parameter Error Covariance Matrix (Γ_N^{-1})							
				Item 1 (2PL)		Item 2 (3PL)		Item 3 (GRM3)			
				c	a	c	a	c	a		
1 (2PL)	intercept, c	-.54	-.50 (.11)	.01							
	slope, a	.62	.67 (.29)	-.01	.08						
2 (3PL)	guessing, logit(g)	-1.39	-1.41 (.50)	.00	.00	.25					
	intercept, c	1.34	1.35 (.24)	.00	.00	-.07	.06				
	slope, a	.88	.90 (.40)	.00	.00	.02	.05	.16			
3 (GRM3)	intercept, c ₀	.66	.59 (.13)	.00	-.01	.00	.00	-.01	.02		
	intercept, c ₁	-.54	-.39 (.12)	.00	.01	.00	.00	.01	.00	.01	
	slope, a	1.01	.91 (.40)	.01	-.05	.00	-.03	-.07	.03	-.02	.16

Table 2

MML parameter estimates and random imputations for the 3-item artificial data set

Item	Parameter	γ_N (SE)	Imputed Parameters ($M = 20$)						
			γ_1	γ_2	γ_3	...	γ_{20}		
1 (2PL)	intercept, c	-.50 (.11)	-.50	-.68	-.60	-.38	
	slope, a	.67 (.29)	.61	.75	.7132	
2 (3PL)	guessing, $\text{logit}(g)$	-1.41 (.50)	-2.09	-.86	-.95	-1.84	
	intercept, c	1.35 (.24)	1.331	1.33	1.31	1.00	
	slope, a	.90 (.40)	.76	1.10	.9135	
3 (GRM3)	intercept, c_0	.59 (.13)	.64	.67	.3279	
	intercept, c_1	-.39 (.12)	-.50	-.26	-.56	-.51	
	slope, a	.91 (.40)	1.16	.81	.23	1.36	

Table 3

EAP scores based on imputed item parameters

Response Pattern	Estimates Based on Imputed Item Parameters ($M = 20$)										MI Summary Statistics		
	$\hat{\theta}_1$ (SE)	$\hat{\theta}_2$ (SE)	$\hat{\theta}_3$ (SE)	...	$\hat{\theta}_{20}$ (SE)	$\bar{\theta}$	\bar{V}	B	$V(\bar{\theta})$				
0 0	-1.07 (.83)	-1.10 (.82)	-.86 (.89)	...	-.90 (.85)	-1.03	.69	.01	.70				
0 0	-.51 (.77)	-.72 (.78)	-.74 (.88)	...	-.26 (.76)	-.58	.61	.02	.64				
0 1	-.58 (.83)	-.47 (.86)	-.21 (.92)	...	-.67 (.84)	-.51	.72	.02	.74				
1 0	-.65 (.83)	-.59 (.82)	-.30 (.89)	...	-.68 (.84)	-.55	.68	.03	.71				
0 0	-.10 (.83)	-.44 (.82)	-.64 (.89)30 (.84)	-.26	.69	.07	.76				
0 1	-.08 (.77)	-.12 (.82)	-.09 (.91)	...	-.07 (.76)	-.10	.63	.00	.63				
1 0	-.15 (.76)	-.26 (.79)	-.18 (.89)	...	-.08 (.76)	-.15	.61	.01	.62				
1 1	-.15 (.84)	.09 (.86)	.38 (.92)	...	-.44 (.84)	.00	.72	.05	.77				
0 1	.42 (.84)	.26 (.86)	.02 (.92)54 (.85)	.30	.72	.03	.74				
1 0	.33 (.83)	.08 (.83)	-.08 (.89)53 (.84)	.23	.69	.04	.74				
1 1	.29 (.77)	.39 (.82)	.50 (.91)11 (.76)	.35	.64	.01	.65				
1 1	.85 (.84)	.81 (.86)	.61 (.92)77 (.85)	.81	.73	.01	.74				

Table 4

Response pattern and summed score EAPs

		Response pattern EAP			Summed score EAP				
y_1	y_2	y_3	$\hat{\theta}_{\gamma_N}$ (SE)	$\bar{\theta}$ (SE)	$r(\%)$	sum	$\hat{\theta}_{\gamma_N}$ (SE)	$\bar{\theta}$ (SE)	$r(\%)$
0	0	0	-1.07 (.84)	-1.03 (.84)	1.6	0	-1.07 (.84)	-1.03 (.84)	1.6
0	0	1	-.62 (.79)	-.58 (.80)	4.2	1	-.52 (.85)	-.52 (.85)	1.5
0	1	0	-.50 (.86)	-.51 (.86)	2.5				
1	0	0	-.60 (.84)	-.55 (.85)	4.3				
0	0	2	-.30 (.84)	-.26 (.88)	10.8	2	-.10 (.84)	-.09 (.85)	1.5
0	1	1	-.09 (.81)	-.10 (.80)	.2				
1	0	1	-.21 (.79)	-.15 (.79)	2.5				
1	1	0	-.01 (.86)	.00 (.88)	7.6				
0	1	2	.32 (.86)	.30 (.86)	3.7	3	.31 (.85)	.31 (.85)	1.5
1	0	2	.18 (.84)	.23 (.86)	6.6				
1	1	1	.34 (.81)	.35 (.81)	2.1				
1	1	2	.81 (.86)	.81 (.86)	1.4	4	.81 (.86)	.81 (.86)	1.4

Note. $\hat{\theta}_{\gamma_N}$ is the "plug-in" estimator that uses maximum likelihood estimates of item parameters, and $\bar{\theta}$ is the MI estimator.

Table 5Relative increase in variance, $r(\%)$

Items, J	Calibration sample size (N)			
	500	1000	2000	5000
2-parameter logistic model				
5	2.6	1.2	.6	.2
10	2.2	1.1	.5	.2
20	2.7	1.3	.7	.3
40	14.7	3.5	1.2	.4
3-parameter logistic model				
5	4.7	2.1	.9	.4
10	2.9	1.4	.8	.3
20	3.5	1.8	.9	.4
40	8.8	2.4	1.1	.5
graded response model				
5	2.3	1.1	.5	.2
10	1.8	.9	.5	.2
20	2.3	1.2	.6	.3
40	3.3	1.7	.9	.4