

UC Berkeley

UC Berkeley Previously Published Works

Title

Computational Strategies and Estimation Performance With Bayesian Semiparametric Item Response Theory Models

Permalink

<https://escholarship.org/uc/item/4d21g17c>

Journal

Journal of Educational and Behavioral Statistics, 48(2)

ISSN

1076-9986

Authors

Paganin, Sally
Paciorek, Christopher J
Wehrhahn, Claudia
[et al.](#)

Publication Date

2023-04-01

DOI

10.3102/10769986221136105

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Computational strategies and estimation performance with Bayesian semiparametric Item Response Theory models

Sally Paganin ^{*1}, Christopher J. Paciorek¹, Claudia Wehrhahn², Abel Rodríguez³, Sophia Rabe-Hesketh¹, and Perry de Valpine¹

¹University of California, Berkeley

²University of California, Santa Cruz

³University of Washington, Seattle

Sally Paganin Postdoctoral Researcher at UC Berkeley University of California, Berkeley, 201 Wellman Hall, Berkeley, CA 94720-3100 US; e-mail: sallypaganin@gmail.com. Her research interests are in Bayesian methods, latent variables models and Bayesian nonparametrics.

Christopher Paciorek is an Adjunct Professor in the Statistics Department at UC Berkeley, 367 Evans Hall, CA 94720-3860 US; email: paciorek@stat.berkeley.edu. His research interests are in Bayesian methods, environmental statistics, and statistical computing.

Claudia Wehrhahn is a Visiting Assistant Professor at University of California, Santa Cruz, 1156 High St, Santa Cruz, CA 95064 US; e-mail: cwehrhah@ucsc.edu. Her research focuses on the development of Bayesian nonparametric models for single density estimation and regression modeling on compact spaces, time dynamic point processes, and diagnostic test validation. She is also the main developer of the nonparametric functionality of the NIMBLE statistical software.

Abel Rodríguez is a Professor of Statistics at the University of Washington, Padelford Hall, RM B-313, Box 354322, Seattle, WA 98195-4322 US; email: abelrod@uw.edu. His research interests are in Bayesian statistical methods and their applications in the natural and social sciences.

Sophia Rabe-Hesketh is a Professor, Graduate School of Education, University of California, Berkeley, CA 94720; sophiarh@berkeley.edu. Her research interests are in multilevel and latent variable modeling, estimation, and missing data.

*sallypaganin@gmail.com

Perry de Valpine Perry de Valpine is a Professor at the University of California, Berkeley, 130 Mulford Hall #3114, Berkeley CA 94720-3114 US; e-mail: pdevalpine@berkeley.edu. His research interests include computational statistical methods applied to biology.

Acknowledgements

This work was supported by funding from grants ACI-1550488, DMS-1622444, DMS-2114727 and DMS-2023495 from the National Science Foundation.

Abstract

Item response theory (IRT) models typically rely on a normality assumption for subject-specific latent traits, which is often unrealistic in practice. Semiparametric extensions based on Dirichlet process mixtures offer a more flexible representation of the unknown distribution of the latent trait. However, the use of such models in the IRT literature has been extremely limited, in good part because of the lack of comprehensive studies and accessible software tools. This paper provides guidance for practitioners on semiparametric IRT models and their implementation. In particular, we rely on NIMBLE, a flexible software system for hierarchical models that enables the use of Dirichlet process mixtures. We highlight efficient sampling strategies for model estimation and compare inferential results under parametric and semiparametric models.

Keywords: binary IRT models, Dirichlet process mixture, MCMC strategies, NIMBLE.

1 Introduction

Traditional approaches in item response theory (IRT) modeling rely on the assumption that subject-specific latent traits follow a normal distribution. This assumption is often considered for computational convenience, but there are many situations in which it may be unrealistic (Samejima, 1997). For example, Micceri (1989) gives a comprehensive review of many psychometric datasets where the distribution of latent individual trait does not respect the normality assumption and presents instead asymmetries, heavy-tails or multimodality. In addition, estimation of IRT parameters in the presence of non-normal latent traits has been shown to produce biased estimates of the parameters (see, for example Finch & Edwards, 2016; Kirisci, chi Hsu, & Yu, 2001; Schmitt, Mehta, Aggen, Kubarych, & Neale, 2006; Seong, 1990).

Different proposals have been made in the general IRT literature for relaxing this normality assumption, using either Markov chain Monte Carlo (MCMC) or Marginal Maximum Likelihood (MML) estimation methods. One option is to rely on more general parametric assumptions. For example, Azevedo, Bolfarine, and Andrade (2011) considered a skew-normal distribution (Azzalini, 1985), while others have suggested finite mixtures of normal distributions (Bambirra Gonçalves, da Costa Campos Dias, & Machado Soares, 2018; Bolt, Cohen, & Wollack, 2001). Alternatively, one can refrain from making a distributional assumption on the latent abilities by using nonparametric maximum likelihood estimation (Laird, 1978; Mislevy, 1984), B-splines (Johnson, 2007; Woods & Thissen, 2006) or empirical histograms (Woods, 2007).

This paper considers a Bayesian nonparametric approach that uses a Dirichlet process mixture (Escobar & West, 1995; Ferguson, 1973; Lo, 1984) as a nonparametric prior on the distribution of the subject-specific latent trait. Dirichlet process mixtures are often used as flexible models to describe the unknown distribution of a heterogeneous population of interest. These models are sometimes interpreted as mixture models with an infinite number of components. In practice these models treat the number of groups as an unknown parameter and estimate it from the data, so that the model can easily account for multimodality, asymmetries or outliers in the latent trait distribution. We focus in particular on semiparametric extensions of logistic IRT models for binary responses. Such models are semiparametric because they retain other, parametric, assumptions of binomial mixed models, such as the functional form of the link function.

Even though some Bayesian nonparametric extensions of binary IRT models have been presented in the literature, they have been given only limited consideration. Within this approach, the semiparametric 1PL model has been the focus of more effort as well as software (Jara, Hanson, Quintana, Müller, & Rosner, 2011, `DPpackage`, no longer actively maintained). San Martín, Jara, Rolin, and Mouchart (2011) investigated a semiparametric generalization of the 1PL model from a theoretical perspective, while Finch and Edwards (2016) provided results from simulation studies. An example using the semiparametric 2PL model is given in Duncan and MacEachern (2008). However, such semiparametric models have not received much attention in applied IRT modeling, in good part because of the lack of comprehensive studies and accessible software tools.

The goal of this paper is to provide a practical guide to semiparametric IRT models for both (i) applied researchers interested in using Dirichlet process mixtures, and (ii) those fa-

miliar with Bayesian nonparametrics concepts who are interested in IRT models. To achieve these goals, we fill three major gaps that hinder the widespread application of semiparametric Bayesian IRT models.

First, we implement the semiparametric 1PL, 2PL and 3PL models in NIMBLE (de Valpine et al., 2017) (R package `nimble`, de Valpine et al., 2020), a flexible R-based system for hierarchical modeling. In particular, NIMBLE provides functionality for fitting hierarchical models that involve Dirichlet process priors either via a Chinese Restaurant Process (CRP) (Aldous, 1985; Blackwell & MacQueen, 1973; Pitman, 1996) or a truncated stick-breaking (SB) (Sethuraman, 1994) representation of the prior. Hence, NIMBLE supports a much wider class of models than those that are implemented in standard software packages. Code is provided for all examples in a publicly accessible GitHub repository (<https://github.com/salleuska/IRTnimblecode>).

Second, focusing on the 2PL model, we study the efficiency of several MCMC sampling strategies in both simulated and real-data scenarios. We define sampling strategies as the combination of model parameterization, identifiability constraints and sampling algorithms, focusing on general MCMC algorithms available in easy-to-access software tools for Bayesian hierarchical models. We find that some choices of parameterization and identifiability constraints can yield order-of-magnitude differences in sampling efficiency compared to others. This approach also allows us to compare various random walk Metropolis-Hastings MCMC strategies to the Hamiltonian Monte Carlo (HMC) strategy implemented in the widely used Stan package (Stan Development Team, 2018). Although there is Stan support for many parametric IRT models (Bürkner, 2021; Furr, 2017), HMC algorithms are not readily available for Dirichlet process prior models, since HMC cannot sample discrete parameters (the component indicators), which cannot be easily integrated out in infinite mixture models.

Finally, we present a comparison of inferential results for item and subject parameters under parametric and semiparametric specifications. To make these comparisons fair, we carefully elicit prior distributions for the models by matching the prior predictive distribution of the data to a common distribution (Berger & Pericchi, 1996). We also illustrate how to estimate the entire distribution of latent traits and its functionals under the two specifications. As expected, we find that the semiparametric model improves recovery of item and individual latent trait parameters in the case of non-normal latent traits. More surprisingly, there seems to be little inferential penalty in using a semiparametric model when a parametric model would be correct, supporting the benefit of greater robustness to mis-specification. These conclusions are based on analyses carried out on simulated data as well as two real datasets related to education and medical assessments: the 2007 Trends in International Mathematics and Science Study (TIMSS) and the 1996 Health Survey for England. For both the real data examples, the semiparametric model performs better than the parametric counterpart, with the semiparametric model identifying distinct modes in the distribution of the latent trait missed by the parametric model.

We note that other authors have considered nonparametric IRT models that rely on a general monotonic function in place of the logistic/probit link function. These models are sometimes referred as NIRT models. Some work using the Dirichlet process falls in this class of models (Karabatsos, 2017; Miyazaki & Hoshino, 2009; Qin, 1998). While we do not pursue this direction in this paper, focusing instead on nonparametric modeling of the latent trait distribution, such an extension is relatively straightforward.

The remainder of the paper is organized as follows. In Section 2 we present the standard IRT model and the Bayesian semiparametric extension along with considerations for identifiability. We then present different potential sampling strategies (Section 3) and discuss the goals of our experiments. To fairly compare the different strategies, we give guidance on selecting prior distributions in Section 4. We introduce simulated and real-world data in Section 5. Comparison of the results in terms of MCMC efficiency and statistical inference is presented in Sections 6 and 7. In Section 8, we conclude that having access to semiparametric models can be broadly useful, as it allows inference on the entire underlying latent trait distribution and its functionals, with NIMBLE being a flexible framework for estimation of such models.

2 IRT models and background

IRT models are widely used in various social science disciplines to scale binary responses into continuous constructs. For conciseness, in this section we introduce model notation in the context of educational assessment, where typically data are answers to exam questions from a set of individuals and the latent trait is interpreted as an individual’s ability. In particular, let y_{ij} denote the answer of individual j to item i for $j = 1, \dots, N$ and $i = 1, \dots, I$, with $y_{ij} = 1$ when the answer is correct and 0 otherwise. Responses from different individuals are assumed to be independent, while responses from the same individual are assumed independent conditional on the latent trait (this is sometimes called the *local independence assumption* in the psychometric literature).

2.1 Binary logistic IRT models

Let π_{ij} denote the probability that individual j answers item i correctly, given the model parameters $\eta_j, \lambda_i, \beta_i$; i.e., $\pi_{ij} = \Pr(y_{ij} = 1 \mid \eta_j, \lambda_i, \beta_i)$ for $i = 1, \dots, I$ and $j = 1, \dots, N$. The parameter η_j represents the latent ability of the j -th individual, while β_i and λ_i encode the item characteristics for the i -th item. In the two-parameter logistic (2PL) model, the probability π_{ij} is determined using the logistic function as

$$\text{logit}(\pi_{ij}) = \lambda_i(\eta_j - \beta_i), \quad i = 1, \dots, I, \quad j = 1, \dots, N. \quad (1)$$

A further assumption for the latent abilities is that they are independently and identically distributed according to some distribution G ,

$$\eta_j \stackrel{iid}{\sim} G, \quad j = 1, \dots, N, \quad (2)$$

with G traditionally a standard normal distribution. The parameter $\lambda_i > 0$ is often referred to as *discrimination*, since items with a large λ_i are better at discriminating between subjects with similar abilities, while β_i is called *difficulty* because for any fixed η_j the probability of a correct response to item i is decreasing in β_i .

Often, the log-odds in (1) are reparameterized as

$$\lambda_i \eta_j + \gamma_i,$$

with $\gamma_i = -\lambda_i\beta_i$. The two parameterizations are sometimes referred to as *IRT parameterization* and *slope-intercept (SI) parameterization*, respectively. While the slope-intercept parameterization is often considered for computational convenience, the IRT parameterization is the most traditional in terms of interpretation. In exploring different strategies for Bayesian estimation, we will consider both alternatives and investigate potential differences in terms of computational performance.

Alternative models can be obtained by considering a different number of item parameters. When $\lambda_i = 1$ for all $i = 1, \dots, I$, the model in (1) reduces to the one-parameter logistic (1PL) model, also known as Rasch model (Rasch, 1990). In some settings one may wish to account for the probability of answering correctly by chance, by introducing a third set of item parameters, $v_i, i = 1, \dots, I$, referred to as *guessing* parameters, so that

$$\Pr(y_{ij} = 1 \mid \eta_j, \lambda_i, \beta_i, v_i) = v_i + (1 - v_i)\text{expit}\{\lambda_i(\eta_j - \beta_i)\},$$

where $\text{expit}\{\cdot\}$ denotes the inverse of the logistic function. This model is typically referred to as the three-parameter logistic model (3PL), and is often relevant in educational assessments.

2.2 Semiparametric IRT models

The classical formulation of IRT models assumes that the latent abilities in (2) follow a normal distribution. This assumption can be relaxed, modeling the distribution of ability as a mixture of normal distributions, where the number of mixture components does not need to be specified in advance but rather is learned from the data. This can be achieved using a Dirichlet process mixture (DPM) model for the distribution of ability. In particular, the distribution of ability G in (2) can be specified as a convolution involving a Dirichlet process (DP) prior, i.e.

$$G = \int \mathcal{K}(\eta_j \mid \theta)F(d\theta), \quad F \sim \text{DP}(\alpha, G_0), \quad (3)$$

where $\mathcal{K}(\cdot \mid \theta)$ is a suitable probability kernel indexed by the parameter θ , while α and G_0 are, respectively, the concentration parameter and the base distribution of the Dirichlet process.

In the context of binary IRT models, it seems natural to choose a normal kernel for $\mathcal{K}(\cdot \mid \theta)$, indexed by parameters $\theta = \{\mu, \sigma^2\}$. This means the distribution of ability is a mixture of normal distributions, where the number of mixture components and their means and variances are unknown. Furthermore, under this choice, taking $\alpha \rightarrow 0$ leads to the original parametric model discussed in Section 2.1, in this case a single normal. Parameters characterizing each mixture component are drawn from the base distribution, G_0 . For computational convenience the base distribution is typically the product of conjugate distributions, e.g., a normal distribution for μ and an inverse-gamma distribution for σ^2 .

We proceed now to discuss the Dirichlet process prior in more detail. There are two main representations of the Dirichlet process, each leading to a different MCMC posterior sampling strategy, namely the stick-breaking representation (SB) (Sethuraman, 1994) and the Chinese Restaurant Process (CRP) (Aldous, 1985; Blackwell & MacQueen, 1973; Pitman, 1996). In

this work we use the CRP representation. The CRP representation is derived from (3) integrating out the random measure F . More specifically, let $\theta_1, \dots, \theta_N$ be an independent sample from F , with some values possibly repeated. Integrating over F one can obtain the joint prior distribution on $(\theta_1, \dots, \theta_N)$, which can be written as the product of a sequence of conditional distributions, where

$$(\theta_j \mid \theta_{j-1}, \dots, \theta_1) \sim \frac{\alpha}{\alpha + j - 1} G_0 + \sum_{l=1}^{j-1} \frac{1}{\alpha + j - 1} \delta_{\theta_l}, \quad (4)$$

for $j = 1, \dots, N$, where δ_a is the Dirac probability measure concentrated at a . The second term in (4) represents the probability that a new observation is equal to one of the previous ones, while the first term captures the possibility that we observe a new value, which would be drawn from the base measure G_0 .

The CRP name comes from an analogy often used to describe the process in (4). Consider a Chinese restaurant with an infinite number of tables, each table serving one dish shared by all customers sitting at that table. In this metaphor, each table represents a possible mixture component, while each dish represents the parameter indexing the distribution associated with the mixture component. Customers entering the restaurant can seat themselves at a previously occupied table and share the same dish (with probability proportional to the number of customers already sitting at the table), or go to a new table and order another dish (with probability proportional to α). The dishes are selected according to the centering distribution G_0 .

One way to make the Chinese restaurant analogy clearer is by reparameterizing the model. Denote by θ_k^* the dish served in table k (which is a draw from G_0) and let z_j be the variable denoting the table chosen by the j th customer. Then

$$p(z_j = k \mid z_{j-1}, \dots, z_2, z_1, \alpha) = \begin{cases} \frac{n_k^{j-1}}{\alpha + j - 1}, & k = 1, \dots, K^{j-1}, \\ \frac{\alpha}{\alpha + j - 1}, & k = K^{j-1} + 1, \end{cases} \quad (5)$$

where K^{j-1} is the total number of occupied tables by the first $j - 1$ customers, and n_k^{j-1} is the number of customers at table k among the first $j - 1$. This new parameterization can be related to the old one by noting that $\theta_i = \theta_{z_i}^*$. The concentration parameter α controls the distribution of the number of tables (components), with larger values favoring more tables. Using the indicators $\mathbf{z} = \{z_j, j = 1, \dots, N\}$, we can denote by $\mathbf{z} \mid \alpha \sim \text{CRP}(\alpha)$ the joint distribution induced by (5), and rewrite the DPM model for the distribution of ability in (3) using

$$\begin{aligned} \eta_j \mid z_j, \theta_1^*, \theta_2^*, \dots &\stackrel{ind}{\sim} \mathcal{K}(\cdot \mid \theta_{z_j}^*), \quad j = 1, \dots, N, \\ \mathbf{z} \mid \alpha &\sim \text{CRP}(\alpha), \\ \theta_k^* &\stackrel{iid}{\sim} G_0, \quad k = 1, 2, \dots \end{aligned}$$

Together, the probability kernel, base measure, and CRP form a Dirichlet process mixture. Alternatively, the distribution F can be written using the stick-breaking representation:

$$F(\cdot) = \sum_{k=1}^{\infty} w_k \delta_{\tilde{\theta}_k},$$

where $\tilde{\theta}_1, \tilde{\theta}_2, \dots$ is a sequence of independent draws from G_0 and the weights are constructed by letting $w_k = v_k \prod_{l=1}^{k-1} (1 - v_l)$, with v_1, v_2, \dots being a sequence of independent draws from a $\text{Beta}(1, \alpha)$ distribution. This construction makes it clear that, as long as the kernel $\mathcal{K}(\cdot | \theta)$ is continuous, the distribution of ability G is also continuous, but F is almost surely discrete, naturally inducing clustering via repeats in the parameter indexing the distribution of ability.

Note that an alternative to the formulation described above is to model the distribution of the ability G directly using a DP, e.g., centered around a normal distribution. Such a model also comprises the standard parametric model as a limiting case (now, when $\alpha \rightarrow \infty$) and leads to slightly simpler computational algorithms. However, we believe that such an approach has some serious drawbacks in the context of most IRT applications. By definition, realizations from a Dirichlet process are almost surely discrete. This property has made the Dirichlet process a useful tool in clustering applications. However, in our context, it implies that we believe that two (or more) individuals potentially have exactly the same ability. Not only is this assumption not realistic, but it potentially prevents us from distinguishing individuals based on their abilities, which is one common goal in IRT modeling. The use of a Dirichlet process mixture with a continuous kernel (Gaussian, in this case) sidesteps this issue.

2.3 Identifiability and constraints

Without additional constraints, the parameters of the models presented in Section 2 are not identifiable (e.g., see Bafumi, Gelman, Park, & Kaplan, 2005; Geweke & Singleton, 1981, as well as Section A in the Supplementary Materials). For example in the 2PL and 3PL models, increasing all η_j and β_i values by the same amount yields the same probabilities in (1) for all i and j . More generally, the ability parameters are known up to a linear transformation, and constraints are needed to identify them. To address this problem, traditional work on parametric IRT models assumes that latent abilities in (2) come from a standard normal distribution, i.e., $G \equiv \mathcal{N}(0, 1)$, and constrains the discrimination parameters λ_i for $i = 1, \dots, I$ to be positive.

Alternative constraints can also establish identifiability and could yield different computational performance for MCMC sampling. A common alternative considers sum-to-zero constraints for the item parameters (Fox, 2010)

$$\sum_{i=1}^I \beta_i = 0, \quad \left(\text{or } \sum_{i=1}^I \gamma_i = 0 \right), \quad \sum_{i=1}^I \log(\lambda_i) = 0. \quad (6)$$

Centering the difficulty parameters addresses the invariance to translations, while centering the log of the discrimination parameters (setting their product to one) accounts for the invariance to rescalings of the latent space. Another potential set of constraints, popular in political science applications, involves fixing the value of the latent traits for two individuals (e.g., see Clinton, Jackman, & Rivers, 2004). Whatever the set of constraints, it is worthwhile to note that they can be either directly incorporated in the model as part of the prior (and, therefore in the structure of the sampling algorithms), or they can be applied as a postprocessing step (after running an unconstrained MCMC). This last approach is typical of parameter-expanded algorithms, which embed the target model in a larger specification.

Parameter expansion has been proposed in the literature to accelerate EM (C. Liu, Rubin, & Wu, 1998) and Gibbs sampler (J. S. Liu & Wu, 1999) convergence, as well as to induce new classes of priors (Gelman, 2004). Although targeting the same posterior, constrained priors and parameter expansion can lead to very different results in terms of convergence and mixing of the MCMC algorithms.

Similar arguments apply for the semiparametric extensions using the Dirichlet process mixture. In that setting, one identifiability strategy may be to constrain the base distribution G_0 , e.g., by letting $G_0 \sim \mathcal{N}(0, 1)$ (for example, see Duncan & MacEachern, 2008). However, even if the prior expectation and variance of G_0 are zero and one, the corresponding posterior quantities can deviate substantially from these values, leading to biased inference (Yang & Dunson, 2010). More general centering approaches have been proposed in the literature when a DP distribution is used to model random effects or latent variables in a hierarchical model (Li, Müller, & Lin, 2011; Yang & Dunson, 2010; Yang, Dunson, & Baird, 2010). These approaches rely on parameter expansion by sampling from the unconstrained DP model and then applying a post-processing procedure to the posterior samples. This post-processing procedure requires the analytical evaluation of the posterior mean and variance of the DP random measure, with Li et al. (2011) providing results under the CRP representation and Yang et al. (2010) under the stick-breaking one. Although such strategies are useful for general hierarchical models to avoid identifiability issues, for the semiparametric 1PL, 2PL, and 3PL models it is simpler to use the sum-to-zero constraints on the item parameters in (6), and that is the approach we adopt in this work. As in the parametric case, we can either include these constraints in the prior or use the parameter expansion approach for sampling and then center and rescale the posterior samples as appropriate.

3 Sampling strategies for logistic IRT models

In this work we explore different *sampling strategies* for Bayesian estimation of the logistic parametric and semiparametric IRT models. We define a *sampling strategy* to include the combination of model parameterization, identifiability constraints and sampling algorithms. We focus on the case of the 2PL model, as it contains the 1PL as a special case and presents the same identifiability challenges as the 3PL model. The strategies considered are summarized in Table 1.

We explore both parameterizations of the 2PL model mentioned in Section 2.1: the IRT and the slope-intercept parameterization. To compare estimates obtained from different parameterizations on a common scale, we post-process posterior samples (using transformations described in the Supplementary Materials, Section A) to respect the following base parameterization

$$\begin{aligned} \text{logit}(\pi_{ij}) &= \lambda_i(\eta_j - \beta_i), \quad i = 1, \dots, I, \\ \sum_{i=1}^I \log(\lambda_i) &= 0 \quad \sum_{i=1}^I \beta_i = 0, \\ \eta_j &\sim G, \quad j = 1, \dots, N, \end{aligned} \tag{7}$$

Identifiability constraints	Parametric		Semi-parametric	
	Slope-intercept	IRT	Slope-intercept	IRT
Constrained abilities	MH/conjugate Centered	MH/conjugate HMC (Stan)		
Constrained item parameters	MH/conjugate	MH/conjugate*	MH/conjugate	MH/conjugate*
Unconstrained	MH/conjugate Centered	MH/conjugate	MH/conjugate Centered	MH/conjugate

Table 1: Summary of the 14 sampling strategies considered for the parametric and semi-parametric 2PL model. Each of the 14 entries is a different strategy with “MH/conjugate”, “Centered”, and “HMC (Stan)” referring to three different sampling algorithms discussed below. The asterisk symbol denotes the sampling strategies that lead directly to samples parameterized as model (7). Others need post-processing to correspond to model (7).

where G denotes a general distribution for the latent abilities, either parametric or nonparametric. The model in (7) follows the IRT parameterization with sum-to-zero identifiability constraints, which is typically the target one for inference for interpretability reasons.

As discussed in Section 2.3, our target inferential model in (7) can be estimated directly, accounting for identifiability constraints. This can be achieved by introducing in the model formulation a set of auxiliary item parameters, $\{\lambda'_i, \beta'_i\}$ for each $i = 1, \dots, I$ and defining $\{\lambda_i, \beta_i\}$ as

$$\log(\lambda_i) = \log(\lambda'_i) - \frac{1}{I} \sum_{i=1}^I \log(\lambda'_i) \quad \beta_i = \beta'_i - \frac{1}{I} \sum_{i=1}^I \beta'_i, \quad i = 1, \dots, I. \quad (8)$$

In Table 1 we label this model as the *constrained item parameters model*. Unconstrained priors are then placed on the auxiliary parameters $\{\lambda'_i, \beta'_i\}$. The same formulation applies under the slope-intercept parameterization, where the sum-to-zero constraints are placed on the pairs $\{\log(\lambda_i), \gamma_i\}$ for $i = 1, \dots, I$.

Alternatively, we can consider *unconstrained* versions of the 2PL model, treating the unconstrained model as a parameter-expanded version of the target inferential model in (7), where the redundant parameters are the means of the difficulty and log discrimination parameters. Hence we conduct MCMC sampling with known model unidentifiability, and before using the results for inference, we transform samples to follow the model in (7), as described in the Supplementary Materials, Section A.

Finally, for the parametric case only, we consider the traditional version of the 2PL model that assumes the ability parameters η_j for $j = 1, \dots, N$ follow a standard normal distribution (*constrained abilities model*).

3.1 Sampling algorithms

We focus on general MCMC sampling algorithms available in easy-to-access software tools for Bayesian hierarchical models, such as NIMBLE and Stan, that can flexibly accommodate different choices of prior distributions and link functions. In this section we give an overall

description of the algorithms considered for the different sampling strategies. We refer to the Supplementary Materials, Section B, for a detailed summary of the samplers used for each parameter.

For both the parametric and semiparametric models we consider NIMBLE’s default sampling configuration (*MH/conjugate algorithm*). NIMBLE’s MCMC uses an overall one-at-a-time sampling strategy, cycling over individual parameters, or parameter blocks for parameters with a multivariate prior. By default, specific sampler types are assigned to the parameters or parameter blocks, but the user can choose to change sampler types, control blocking strategies, and modify details of sampling algorithm behavior. NIMBLE’s default MCMC configuration assigns a conjugate (sometimes called “Gibbs”) sampler where possible, sampling from the corresponding full conditional posterior distribution. For non-conjugate continuous-valued parameters, NIMBLE’s default sampler assignment is an adaptive random walk Metropolis-Hastings. For the parametric versions of the 2PL model, the strategies using the default NIMBLE assignments (*MH/conjugate algorithm*) correspond to these conjugate and adaptive random walk Metropolis-Hastings samplers, with the latter also used for most parametric components of the semiparametric 2PL. Specialized samplers are assigned when Bayesian nonparametric priors are considered in the semiparametric 2PL.

In the case of the slope-intercept parameterization, we take advantage of NIMBLE’s flexibility to include user-programmed custom samplers (*centered sampler*). The proposed centered sampler uses an adaptive random walk Metropolis-Hastings sampler with a joint proposal for each pair of item parameters $\{\lambda_i, \gamma_i\}$ for $i = 1, \dots, I$, thereby accounting for their posterior correlation. The proposal is made under a reparameterization of the model that centers the abilities to have mean zero. Implementation details are provided in the Supplementary Materials, Section B.

Finally, in the parametric setting only, we consider a Hamiltonian Monte-Carlo (HMC) algorithm, as implemented in the Stan software (Carpenter et al., 2017). Stan implements an adaptive HMC sampler (Betancourt, Byrne, Livingstone, & Girolami, 2017) based on the No-U-Turn sampler (NUTS) of Hoffman and Gelman (2014). HMC algorithms are known to produce samples that are much less autocorrelated than those of other samplers but at more computational cost given the need to calculate the gradient of the log-posterior. In this work, we limit the comparison to the IRT parameterization with constraints on the abilities distribution, as that is the model provided in the `edstan` R package (Furr, 2017).

3.2 Aims

In the remainder of the paper, we study the efficiency of the MCMC sampling strategies in Table 1 to fit binary logistic IRT models, and we compare inferential results under parametric and semiparametric specifications. Using both simulated and real-world data, we aim to answer the following questions:

- Q.1 For the parametric binary logistic IRT model, which of the Metropolis-Hastings-based MCMC sampling strategies in Table 1 are most efficient? Do different strategies work better in different scenarios for the distribution of ability?
- Q.2 For the parametric binary logistic IRT model, how does efficiency of random walk Metropolis-Hastings sampling compare to Hamiltonian Monte Carlo (HMC), as imple-

mented in the popular Stan package? This question is of interest because HMC is not readily available for semiparametric models using Dirichlet process priors.

- Q.3 How does MCMC efficiency of a semiparametric model compare to that of a parametric one? Does this comparison differ when the parametric model is correctly vs. incorrectly specified?
- Q.4 To what degree does the use of a parametric model when its assumptions are violated yield bad inference? Does the use of a semiparametric model change inference even when a parametric model would be valid?
- Q.5 How much do results differ between the semiparametric and parametric models for the real data examples?

In Section 6 we discuss the choice of efficiency metrics to address Q.1-Q.3 and present the results obtained using simulated and real-world data. In Section 7, we discuss differences in the inferential results to investigate Q.4-Q.5.

4 Choice of prior distributions

Past research on Bayesian IRT models has warned about the use of either vague priors or highly informative priors when there is little information about the parameters (Natesan, Nandakumar, Minka, & Rubright, 2016; Sheng, 2010). In particular Natesan et al. (2016) investigated the use of different prior choices in 1PL and 2PL models using MCMC and variational Bayes algorithms and found that the use of vague priors tends to produce biased inference or convergence issues. Similarly, it is well known that highly informative prior distributions on parameters can strongly affect model comparison procedures.

To ensure a fair comparison between results from different strategies, we chose the parameters of the priors in such a way that the induced prior predictive distribution of the data is similar across all the different model parameterizations. This “predictive matching approach” has been widely used to guide prior elicitation in model comparison settings (Bedrick, Christensen, & Johnson, 1996; Berger & Pericchi, 1996; Ibrahim, 1997).

In the context of binary logistic IRT models, we aim to match the prior marginal predictive distribution of a response y_{ij} , which in turn can be achieved by matching the induced prior distribution on the marginal prior probability of a correct response, $\pi_{ij} = \text{expit}\{\lambda_i(\eta_j - \beta_i)\}$. Note that all the priors discussed in this paper are separately exchangeable, which means that this prior marginal will be the same for any values of i and j . In particular, we attempt to match a Beta(0.5, 0.5) distribution, which is both the reference and the Jeffreys prior for the Bernoulli likelihood in the fully exchangeable case (Berger, Bernardo, & Sun, 2009; Bernardo, 1979). A similar approach to prior elicitation in the context of latent space models for networks can be found in Guhaniyogi and Rodriguez (2020) and Sosa and Rodríguez (2021). Because there are no analytical expressions available for the prior distribution of π_{ij} , we use simulations to estimate the shape of the prior distribution and obtain an approximate match. This is facilitated by our implementation in NIMBLE. Indeed, one of the advantages of the NIMBLE system is that it provides a seamless way to

simulate from the model of interest. Histograms of samples from the resulting induced priors can be seen in Figure 1 for a set of parametric and semiparametric models. Further details are presented in the following subsections.

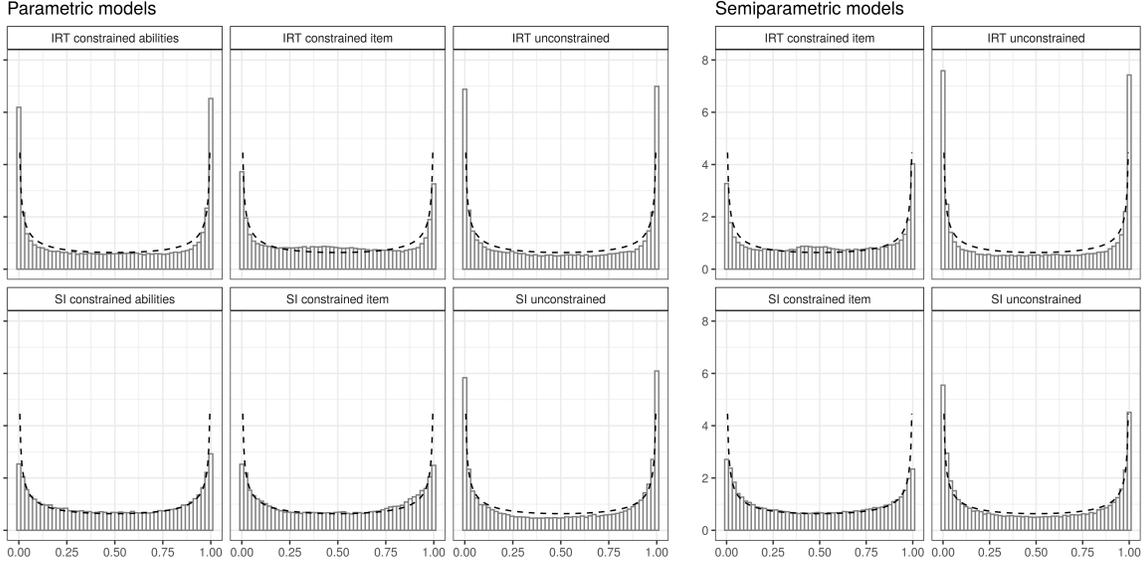


Figure 1: Histogram of samples from the induced prior on π_{ij} under each of the considered models. Dashed line indicates the density function of a $\text{Beta}(0.5, 0.5)$ distribution. Samples for the semiparametric models use a prior distribution $\text{Gamma}(2, 4)$ for the DP concentration parameter α , but similar results are obtained under the other settings presented Section 4.

4.1 Priors for the item parameters

In Bayesian IRT modeling, normal distributions are typically chosen as priors for the item parameters. This is true under both parameterizations. In addition, the discrimination parameters, $\{\lambda_i\}_{i=1}^I$, are typically assumed positive, so we consider a normal distribution on the log-scale. To summarize, priors on the item parameters are:

$$\log \lambda_i \sim \mathcal{N}(\mu_\lambda, \sigma_\lambda^2), \quad \beta_i \sim \mathcal{N}(0, \sigma_\beta^2), \quad \gamma_i \sim \mathcal{N}(0, \sigma_\gamma^2) \quad i = 1, \dots, I.$$

By default, we center on the difficulty parameters β_i (or the reparameterized version γ_i) on 0 for $i = 1, \dots, I$, while we set $\sigma_\beta^2 = \sigma_\gamma^2 = 3$. For the discrimination parameters, we set $\mu_\lambda = \sigma_\lambda^2 = 0.5$ such that the prior probability mass on the original scale is mostly in the range $(0.5, 2.5)$.

4.2 Priors for the distribution of ability

In choosing priors for the abilities, we distinguish between the parametric and semiparametric cases. In the parametric case, excluding the strategies in which the distribution is a standard normal, we assume $G \equiv \mathcal{N}(\mu_\eta, \sigma_\eta^2)$. We specify hyperpriors for the unknown mean and variance, using a normal distribution for the mean $\mu_\eta \sim \mathcal{N}(0, 3)$, and an inverse-gamma

distribution for the variance, $\sigma_\eta^2 \sim \text{InvGamma}(2.01, 1.01)$ as in [Paulon, De Iorio, Guglielmi, and Ieva \(2018\)](#), with hyperparameter values implying an a priori marginal expected value of 1 and an a priori variance equal to 100.

In the semiparametric case, we need to specify the base distribution G_0 of the DP mixture prior along with the hyperparameters. We choose $G_0 \equiv \mathcal{N}(0, \sigma_0^2) \times \text{InvGamma}(\nu_1, \nu_2)$ where $\text{InvGamma}(\nu_1, \nu_2)$ denotes an inverse-gamma distribution with shape parameter ν_1 and mean $\nu_2/(\nu_1 - 1)$. In choosing values for the hyperparameters $\{\sigma_0^2, \nu_1, \nu_2\}$, we first considered the concentration parameter α as fixed and evaluated the induced prior distribution on $\boldsymbol{\pi}$ for values of $\alpha \in \{0.01, 0.05, 0.5, 1, 1.5, 2\}$. Recall that α controls the prior expectation and variance of the number of clusters induced by the DP, which are both of the order $\alpha \log(N)$. We discuss prior choice for the α in [Section 4.3](#). As in the parametric case, we center the normal distribution for the mixture component means on 0 with $\sigma_0^2 = 3$ and set $\nu_1 = 2.01$ and $\nu_2 = 1.01$ for the inverse-gamma distribution. Given these settings, we found that choosing $\alpha \in \{0.01, 0.05, 0.5, 1, 1.5, 2\}$ does not have much effect on the marginal prior distribution of the π_{ij} s.

4.3 Prior on the DP concentration parameter

One may be interested in placing a prior distribution on the concentration parameter α of the Dirichlet process. A typical choice for the DP concentration parameter is a $\text{Gamma}(a, b)$, with shape $a > 0$ and scale $b > 0$, due to its computational convenience ([Escobar & West, 1995](#)). As previously stated, the concentration parameter controls the prior distribution of the number of clusters ([Escobar & West, 1995](#); [S. J. Liu, 1996](#)). In choosing values a and b , we considered the implied prior mean and variance of the number of clusters.

Let K_N denote the number of clusters for a sample of size N . Results from [Antoniak \(1974\)](#) and [S. J. Liu \(1996\)](#) show that the expected value and variance of K_N given α is

$$\mathbb{E}(K_N | \alpha) = \sum_{i=1}^N \frac{\alpha}{\alpha + N - i}, \quad \text{Var}(K_N | \alpha) = \sum_{i=1}^N \frac{\alpha(i-1)}{(\alpha + N - i)^2}. \quad (9)$$

We exploit these results to choose values a and b that lead to reasonable a priori values for the moments of the number of clusters for each of our applications. For a given N and for different values of a and b , we evaluated the marginal expectation and variance of the quantities in (9) via Monte Carlo approximation. We sample α_r for $r = 1, \dots, R$ from its prior and compute

$$\widehat{\mathbb{E}}(K_N) = \frac{1}{R} \sum_{r=1}^R \mathbb{E}[K_N | \alpha_r], \quad \widehat{\text{Var}}(K_N) = \frac{1}{R} \sum_{r=1}^R \text{Var}(K_N | \alpha_r) + \widehat{\text{Var}}(\mathbb{E}[K_N | \alpha]),$$

where $\widehat{\text{Var}}(\mathbb{E}[K_N | \alpha]) = R^{-1} \sum_{r=1}^R [\mathbb{E}[K_N | \alpha_r] - \widehat{\mathbb{E}}[K_N]]^2$.

We explored a few prior choices and tabulate approximated moments in [Table 4.3](#), for the values of N in our datasets. We consider the popular choice of $a = 2, b = 4$ for the hyperparameters as in [Escobar and West \(1995\)](#) along with values favoring a small number of clusters ($a = 1, b = 3$) and values leading to a more vague prior ($a = 1, b = 1$). For our applications we decided to favor a relatively small number of clusters, choosing $a = 2, b = 4$ as hyperparameters for the simulated data, and $a = 1, b = 3$ for the real-world data.

$\alpha \sim \text{Gamma}(a, b)$	$\mathbb{E}[\alpha]$	$\text{Var}(\alpha)$	$\widehat{\mathbb{E}}(K_{2,000})$	$\widehat{\text{Var}}(K_{2,000})$	$\widehat{\mathbb{E}}(K_{14,525})$	$\widehat{\text{Var}}(K_{14,525})$	$\widehat{\mathbb{E}}(K_{7,377})$	$\widehat{\text{Var}}(K_{7,377})$
$a = 2, b = 4$	0.5	0.12	4.7	9.3	5.6	14.12	5.3	12.2
$a = 1, b = 3$	0.3	0.11	3.5	7.6	4.2	11.76	3.9	10.2
$a = 1, b = 1$	1	1	7.8	43.7	9.8	73.35	9.2	64.5

Table 2: Approximate expectation and variance of the a priori number of clusters, K_N , under different choices of the concentration parameter distribution, for $N = 2,000, 14,525, 7,377$ as in our data example presented in Section 5.

5 Data examples

5.1 Synthetic data

We specify three different simulation scenarios for the distribution of ability: *unimodal*, *bimodal* and *multimodal* distributions. For all the scenarios, we simulate responses from $N = 2,000$ individuals on $I = 15$ binary items. Values for the discrimination parameters $\{\lambda_i\}_{i=1}^{15}$ are sampled from a Uniform(0.5, 1.5) distribution, while values for difficulty parameters $\{\beta_i\}_{i=1}^{15}$ are taken to be equally spaced in $(-3, 3)$. We center the log of the discrimination parameters on zero, while the difficulty parameters are already centered based on how they are generated. We consider three different underlying distributions for the latent abilities, η_j for $j = 1, \dots, 2,000$. In the *unimodal* scenario, latent abilities are generated from a normal distribution with mean 0 and variance $(1.25)^2$. In the *bimodal* scenario, we use an equal-weights mixture of two normal distributions with means $\{-2, 2\}$ and common variance $(1.25)^2$. Finally, for the *multimodal* scenario, latent abilities are generated from the following mixture

$$\frac{1}{5}\mathcal{N}(-2, 1) + \frac{2}{5}\mathcal{N}(0, (0.5)^2) + \frac{2}{5}\mathcal{SN}(3, 1, -3),$$

where $\mathcal{SN}(\xi, \omega, \zeta)$ indicates a skew-normal distribution (Azzalini, 1985) with location parameter ξ , scale parameter $\omega > 0$ and parameter ζ that controls the asymmetry of the distribution.

As a sensitivity analysis, we considered other values of N and I , simulating data for each scenario following a factorial design with $I \in \{10, 30\}$ and $N \in \{1,000, 5,000\}$. We discuss efficiency results of the different sampling strategies in Section 6 and report results in the Supplementary Materials, Section D.

5.2 Real world data

The first example is a subset of data from the 1996 Health Survey for England (Joint Health Surveys Unit of Social and Community Planning Research and University College London, 2017), a survey conducted yearly to collect information concerning the health and behavior of households in England. In particular, we have data for 10 items measuring Physical Functioning (PF-10), which is a sub-scale of the SF-36 Health Survey (Ware, 2003) administered to people aged 16 and above. In this case the latent trait quantifies the physical status of a given individual (Hays, Morales, & Reise, 2000; McHorney, Haley, & Ware Jr, 1997).

Participants in the survey were asked whether they perceived limitations in a variety of physical activities (e.g., running, walking, lifting heavy objects) and if so the degree of

limitation. We list the original questions in the Supplementary Materials, Section C. Answers to items comprised three possible responses (“yes, a lot”, “yes, limited a little”, “no, not limited at all”); however, in our analysis we consider the dichotomous indicator for not being limited at all. The left panel of Figure 2 shows the distribution of raw scores, i.e. the total of correct answers. We consider the 2PL model for this data, as it is reasonable to assume that some of the questions are more informative in defining individuals with high physical impairment (see Supplementary Materials, Section C). For simplicity, we analyzed complete case data from 14, 525 individuals out of 15, 592 respondents, although the model can easily accommodate missing data.

The second example uses data from the TIMSS (Trends in International Mathematics and Science Study) survey, which is an international comparative educational survey dedicated to improving teaching and learning in mathematics and science for students around the world (<http://timssandpirls.bc.edu/TIMSS2007/about.html>). We used data from the 2007 eighth-grade mathematics assessment for the United States ($N = 7, 377$), publicly available at <https://timssandpirls.bc.edu/TIMSS2007/idb.ug.html>. The dataset comprises 214 items, with 192 of them dichotomous, while the remaining 22 have three category responses (“incorrect”, “partially correct”, “correct”). We dichotomized these latter questions, considering partially correct answers as incorrect ones. Like other large-scale assessments, participants in TIMSS only received a subset of the items according to a booklet design, resulting in 28-32 item responses per student. Distribution of the raw scores for the data are shown in the right panel of Figure 2. As for the previous example, it is reasonable to assume that some items discriminate differently between students with high and low ability. However, in the context of educational testing, the 3PL model is often considered because it accounts for the probability of answering correctly by chance. Hence we consider both the 2PL and the 3PL models when evaluating the different strategies.

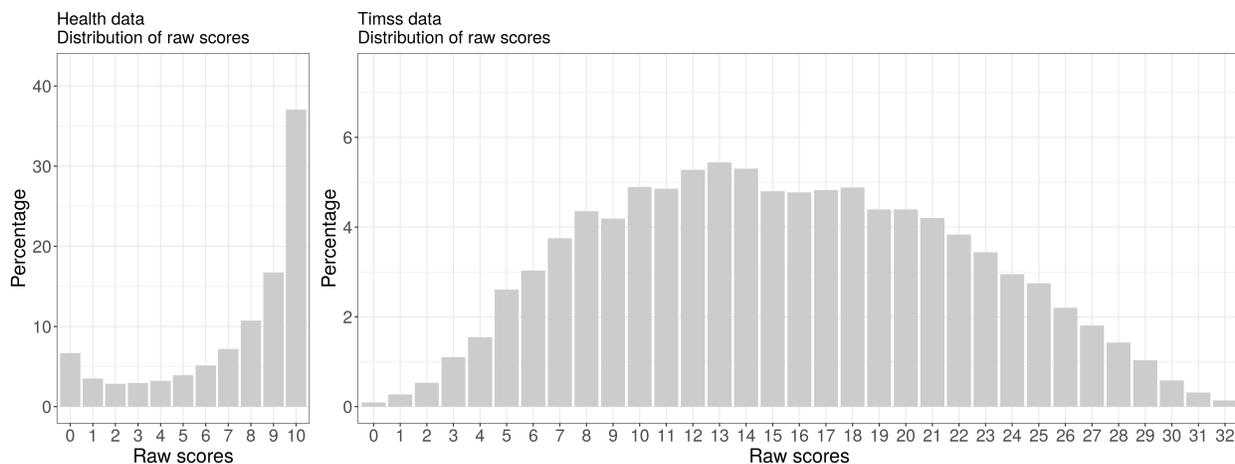


Figure 2: Distribution of the raw scores (total of correct answers) for the real data examples: health data (left panel) and TIMSS data (right panel).

Finally, we note that participants in the survey were sampled according to a complex two-stage clustered sampling design that we did not consider in our application. In other contexts the design is typically taken into account using sampling weights for model estimation, as

discussed for example by [Rutkowski, Gonzalez, Joncas, and von Davier \(2010\)](#).

6 Comparing results in terms of efficiency

MCMC performance is often evaluated in terms of mixing, often by calculating the effective sample size (ESS), which is the equivalent number of independent samples that would contain the same statistical information as the actual non-independent samples. However, comparison between different MCMC algorithms based solely on mixing can be misleading, as different samplers can vary greatly in terms of computational cost ([Nguyen et al., 2020](#)). Hence, it is appropriate to consider ESS per computation time (in seconds), the rate at which effectively independent samples are generated. A second issue is how to combine ESS results for multiple parameters. For this purpose, we used the *multivariate ESS* (mESS) recently introduced by [Vats, Flegal, and Jones \(2019\)](#), which accounts for cross-correlations among parameters.

Computation time is typically measured for the actual MCMC run, not accounting for steps to prepare for a run, thereby focusing on the algorithms of interest rather than unrelated aspects of the software. Comparison between HMC and MCMC algorithms raises the question of how to fairly account for computation times, given that these two classes of algorithms use different tuning phases. Since there is not an established way to compare these two algorithms in the literature, we decided to consider different timings when using the two algorithms: (i) *sampling time*, which accounts only for the time to draw the posterior samples, hence discarding the time needed for the burn-in and warm-up phases of the two algorithms and (ii) *total time* comprising also the burn-in and warm-up phases. Although one can use alternative metrics for the comparison, this choice can provide interesting and useful insights. When computing efficiency based on the sampling time, we can assess pure efficiency of sampling from the posterior. Using total time accounts for potentially different times needed for warm-up/burn-in by the different algorithms but introduces the difficulty of determining the optimal burn-in/warm-up time, which we avoided here in favor of using basic defaults.

Estimate of the mESS is based on the multivariate batch means estimator as described in [Vats et al. \(2019\)](#) and implemented in the `mcmcse` ([Flegal et al., 2021](#)) package. Since we used different specifications for the distribution of ability across the different sampling strategies, we calculate the mESS considering only the common parameters (i.e., the item parameters and sampled abilities) after transforming samples to the parameterization of our target inferential model (7). The mESS provides a single scalar measure of joint mixing for all the parameters of interest in a model, but it does not necessarily reflect univariate ESS values of each parameter. For example, mESS can be larger than all the univariate ESS values (see Supplementary Materials, Section D, Figures 1-3 for some insights). Given this, it can be useful as a simple overall performance metric but does not replace ESS for specific parameters of interest.

Given the large number of examples and sampling strategies, we performed a preliminary experiment to choose the number of iterations and number of burn-in or warm-up samples. In particular, for a portion of the simulations we used multiple runs to determine the number of iterations and samples needed to obtain reliable estimates of the mESS. For all MCMCs

using NIMBLE, we decided to use a total of 50,000 iterations, with a 10% burn-in of 5,000 for all examples. When running the HMC algorithm via Stan, we used a total of 15,000 iterations, with the first 5,000 iterations as warm-up steps. However, we observed highly variable values of the mESS estimates when using the HMC algorithm. We also found that values of mESS are correlated with those of the HMC tuning parameters (see Supplementary Materials, Section D). Given this variability in mixing performance, we decided to limit comparisons with HMC to the simulation scenarios, reporting results from the run with the median ESS across multiple replications. All the models were estimated using a Linux cluster with 4 nodes having 24 cores and 128 GB RAM per node (Intel(R) Xeon(R) CPU E5-2643 v2 @ 3.50GHz). Across simulations, the running times ranged between 15-94 minutes for parametric models and between 37-124 minutes for semiparametric ones. For the data applications, running times ranged between 85-384 minutes for the parametric models, and 410-610 minutes for semiparametric ones.

6.1 Efficiency results for simulated data

For the three simulation scenarios we estimated the 2PL parametric model using the different sampling strategies summarized in Table 1. Figure 3 compares efficiency for all these strategies using the multivariate ESS per second, computed with respect to both the total and sampling time.

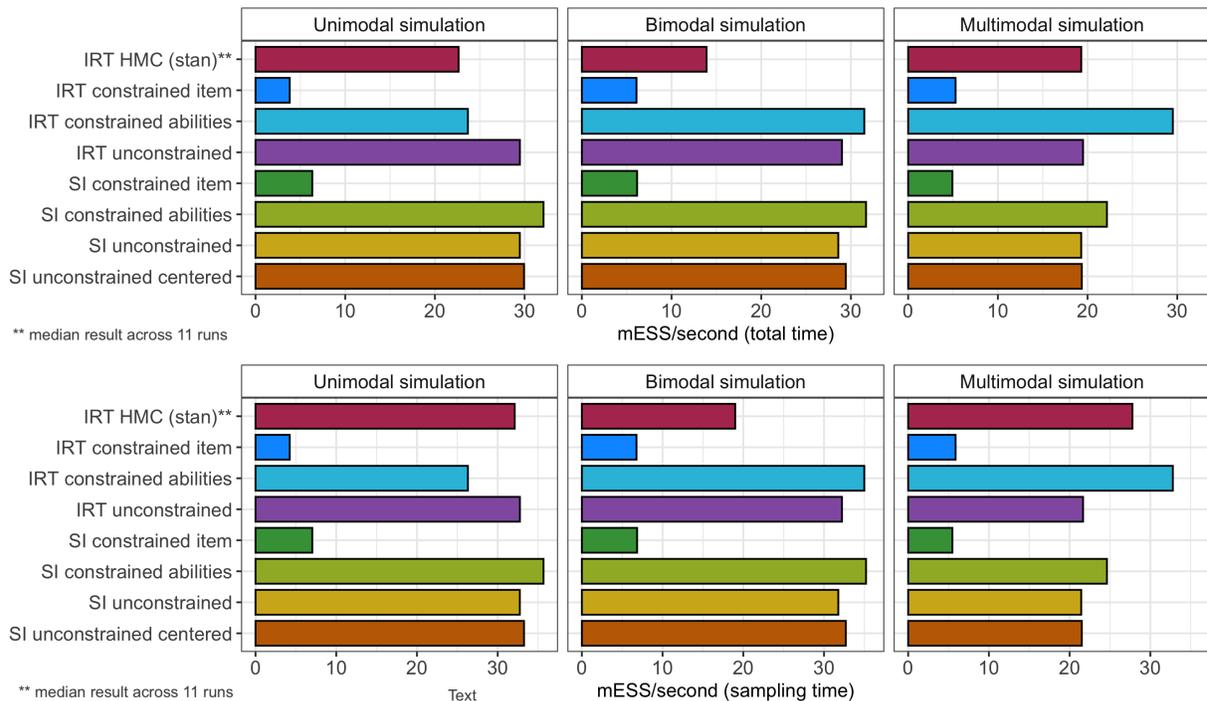


Figure 3: Multivariate ESS per second for various sampling strategies used to estimate the 2PL parametric model for the unimodal (left column), bimodal (middle column), and multimodal (right column) scenarios. Results are computed using total time (top row) and sampling time (bottom row). Note that SI stands for “slope-intercept”.

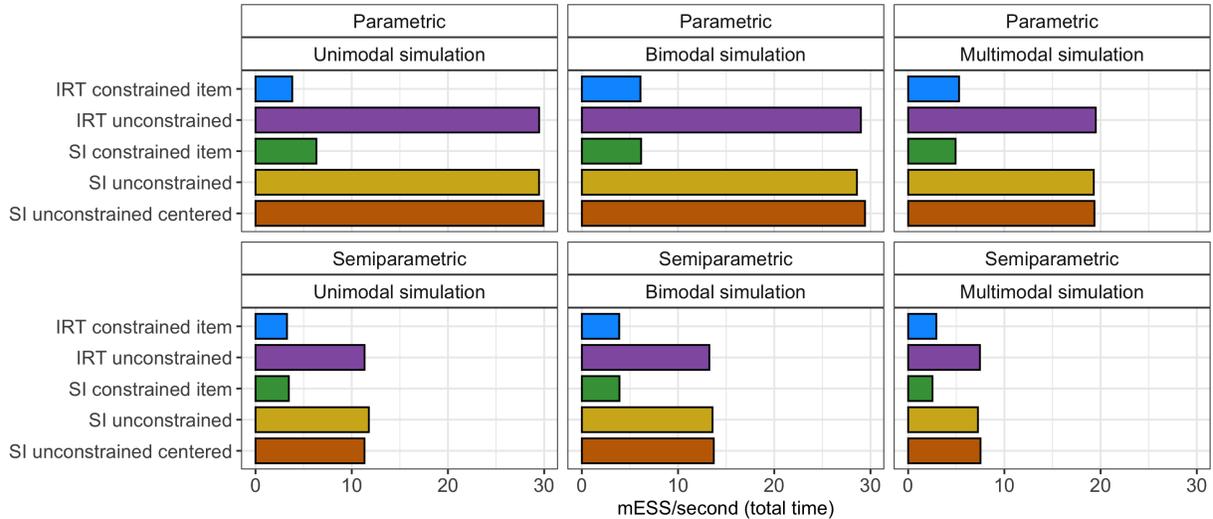


Figure 4: Multivariate ESS per second (computed using the total time) for semiparametric 2PL models (bottom row) in comparison with their parametric version (top row) under the two simulation scenarios. Note that SI stands for “slope-intercept”.

Using different time baselines when computing efficiency changes the ranking of the MCMC algorithms for the unimodal and multimodal simulation, highlighting the trade-off for the HMC algorithm between sampling efficiency and computational cost. While the HMC is highly efficient in producing samples with low correlation, warm-up steps are computationally expensive. Recall that efficiency values presented for the HMC strategy are relative to a median performance across multiple runs. Amongst the non-HMC strategies, unconstrained scenarios generally mix well, as do scenarios with constraints on the abilities. However, imposing constraints on the item parameters directly in the sampling performs poorly because obtaining each sample is time-consuming. This is because the constraints in (8) require calculation of all the likelihood terms for each parameter update, whereas for other strategies only the likelihood terms for individuals’ responses on the item under consideration need to be calculated. While incorporating constraints on the item parameters is time-consuming, such a strategy could be useful in more complicated hierarchical models, in particular when it is unclear how to rescale posterior samples. The centered strategy for the slope-intercept parameterization seems to have little impact across the scenarios.

Moving to the semiparametric models, recall that we did not consider identifiability constraints on the abilities. We either included identifiability constraints on the item parameters in the sampling or sampled from the unconstrained model and rescaled the posterior samples.

As expected when using a more complicated model, we observed some reduction in efficiency in the semiparametric model compared to the parametric model, but not a drastic one (Figure 4). Results for the semiparametric case are similar in relative terms, but not in absolute magnitudes, when comparing the different parameterizations and constraints.

For the non-HMC strategies, we also looked at how different combinations of numbers of items and individuals affect efficiency of the different sampling strategies (as discussed in Section 5.1), with results shown in the Supplementary Materials, Section D. We found

that the ranking across strategies is generally stable across the different scenarios; strategies using constraints on items are the worst overall, and the benefit of using other strategies is most evident when the number of individuals N is low.

6.2 Efficiency results for real-world data

We did similar comparisons using the real-data examples (Figure 5), noting that we excluded the constrained items strategy, given its poor performance on the simulated datasets, and the HMC strategy, because of the high variability in mixing performance.

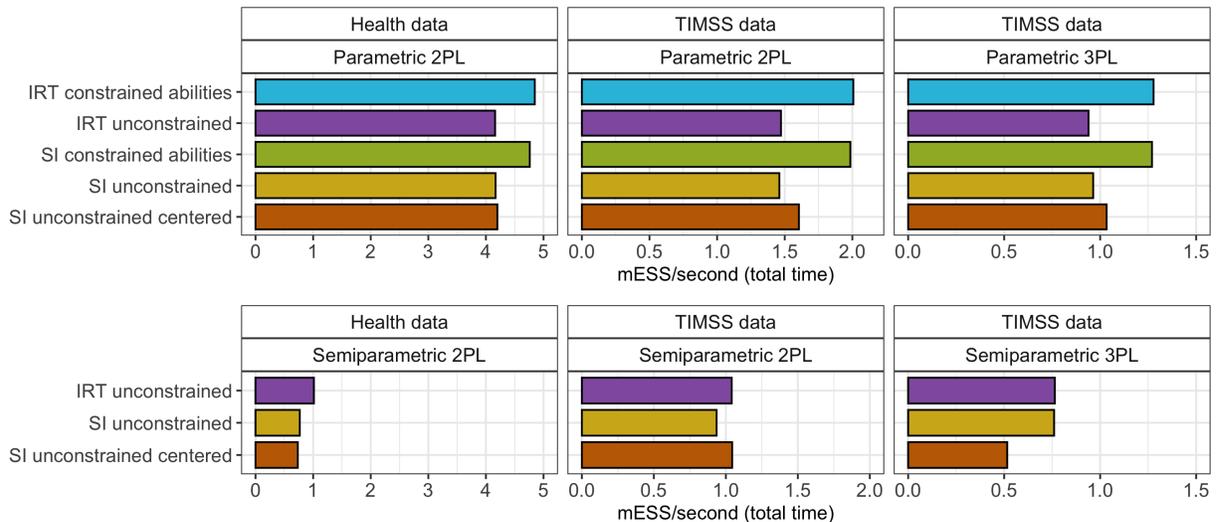


Figure 5: Multivariate ESS per second (computed using the total time) for the health data (left column) parametric and semiparametric 2PL models, and TIMSS data for the 2PL (middle column) and 3PL (right column) models. Note the scales are different for the different rows and that SI stands for “slope-intercept”.

The efficiency is lower than for the simulated datasets because the real data have many more individuals or items, and therefore more parameters. The same consideration applies when comparing efficiency of the 2PL and 3PL models for the TIMSS data. We also noted that for the health data, when considering the posteriors for the abilities in the semiparametric model, we saw evidence for multimodality and some difficulty moving between modes for individuals with high raw scores. The multimodality is likely related to it being difficult for the semiparametric model to identify the exact magnitude of the ability for such individuals. The use of more informative priors, with careful elicitation of the prior distribution, may be important in such cases.

7 Comparing results in terms of statistical inference

In this section we compare results for the parametric and semiparametric models in terms of statistical inference, regardless of the sampling strategy used to obtain posterior samples. All posterior samples follow the parameterization in (7), and we use samples from the most

efficient sampling strategy for each dataset. We use posterior means as point estimates for the item and ability parameters. For the simulated datasets, we measure how well the models recover the (known) true value of the parameters using absolute error, e.g., $|\hat{\beta}_i - \beta_i|$, and squared error, e.g., $(\hat{\beta}_i - \beta_i)^2$.

A crucial point of this paper is to make inference on the distribution of latent abilities. An estimate of this distribution is sometimes based on the posterior means of the abilities (Bambirra Gonçalves et al., 2018; Duncan & MacEachern, 2008), and histograms or kernel density plots are reported. Such an estimate ignores uncertainty in the estimates of individual abilities. Instead, one should directly obtain the point estimate of the posterior distribution of the latent abilities $p(\eta \mid \mathbf{Y})$ (for any value of η) using the posterior samples. In the parametric case, this reduces to:

$$p(\widehat{\eta \mid \mathbf{Y}}) = \frac{1}{T} \sum_{t=1}^T \mathcal{N}(\eta; \mu^{(t)}, \sigma^{2(t)}), \quad (10)$$

with $\mathcal{N}(\cdot; \mu, \sigma^2)$ indicating the probability density function of a normal distribution with mean μ and variance σ^2 , and $t = 1, \dots, T$ denoting an MCMC iteration. In the semiparametric case, a point estimate of $p(\eta \mid \mathbf{Y})$ is the posterior mean of the mixing measure G of the Dirichlet process. This can be obtained using posterior samples, averaging over the DP conditional distribution in (4) computed for each iteration $t = 1, \dots, T$,

$$p(\widehat{\eta \mid \mathbf{Y}}) = \frac{1}{T} \sum_{t=1}^T \left\{ \left[\sum_{k=1}^{K^{(t)}} \frac{n_k^{(t)}}{\alpha^{(t)} + N} \mathcal{N}(\eta; \mu_k^{(t)}, \sigma_k^{2(t)}) \right] + \frac{\alpha^{(t)}}{\alpha^{(t)} + N} \mathcal{N}(\eta; \mu_{K^{(t)+1}, \sigma_{K^{(t)+1}}^2} \right\}, \quad (11)$$

with $n_k^{(t)}$ the number of observations in cluster k at iteration t , $K^{(t)}$ the total number of clusters at iteration t , and $\mu_{K^{(t)+1}$ and $\sigma_{K^{(t)+1}}^2$ sampled from G_0 (conditional on the data). We graphically compare estimates for the distribution of ability resulting from (10)–(11) with the estimates obtained using the posterior means.

It is possible to make full inference on $p(\eta \mid \mathbf{Y})$ in the semiparametric setting; this requires sampling from the posterior of the mixing distribution F . A computational approach to obtain the entire posterior distribution has been presented in Gelfand and Kottas (2002), a version of whose algorithm is implemented in NIMBLE in the function `getSamplesDPMeasure`. This function provides samples of a truncated version of the infinite mixture to a level L . The value of L varies at each iteration of the MCMC's output when α is random, while it is the same at each iteration when α is fixed. In our case, for every MCMC iteration, we can obtain samples of the vector of mixture weights $\{w_1^{(t)}, \dots, w_{L^{(t)}}^{(t)}\}$ and parameters of the mixture components. We can use these samples to make inference on functionals of the distribution, such as the percentile for an individual, $100 \times p_j$, where $p_j = \int_{-\infty}^{\eta_j} p(\eta \mid \mathbf{Y}) d\eta$, typically paired with test scores when giving results for educational assessments. For an individual j for $j = 1, \dots, N$ we estimate p_j at each MCMC iteration as

$$p_j^{(t)} = \sum_{l=1}^{L^{(t)}} w_l^{(t)} F_{\mathcal{N}}(\eta_j^{(t)}; \mu_l^{(t)}, \sigma_l^{2(t)}), \quad (12)$$

where $F_{\mathcal{N}}$ denotes the distribution function of the normal distribution. For comparison, we define the parametric counterpart as $p_j^{(t)} = F_{\mathcal{N}}(\eta_j^{(t)}; \mu^{(t)}, \sigma^{2(t)})$.

	Unimodal Simulation				Bimodal Simulation				Multimodal Simulation			
	Parametric		Semi-parametric		Parametric		Semi-parametric		Parametric		Semi-parametric	
	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE
Difficulty parameters	0.0996	0.0185	0.0988	0.0183	0.0895	0.0137	0.0673	0.0083	0.0647	0.0080	0.0737	0.0103
Discrimination parameters	0.0734	0.0069	0.0731	0.0070	0.0832	0.0105	0.0397	0.0020	0.0721	0.0082	0.0677	0.0062
Ability parameters	0.4836	0.3719	0.4836	0.3720	0.5944	0.5571	0.5501	0.4775	0.5477	0.4753	0.5212	0.4444

Table 3: MAE and MSE for the item and ability parameters estimates, under the three simulation scenarios, using samples from most efficient MCMC-based sampling strategies.

7.1 Inferential results for the simulated datasets

We report results using posterior samples from the unconstrained sampling strategy under the IRT parameterization. Using the absolute error and the squared error for each parameter, we report in Table 3 the mean absolute error (MAE) and the mean squared error (MSE) across item and ability parameters.

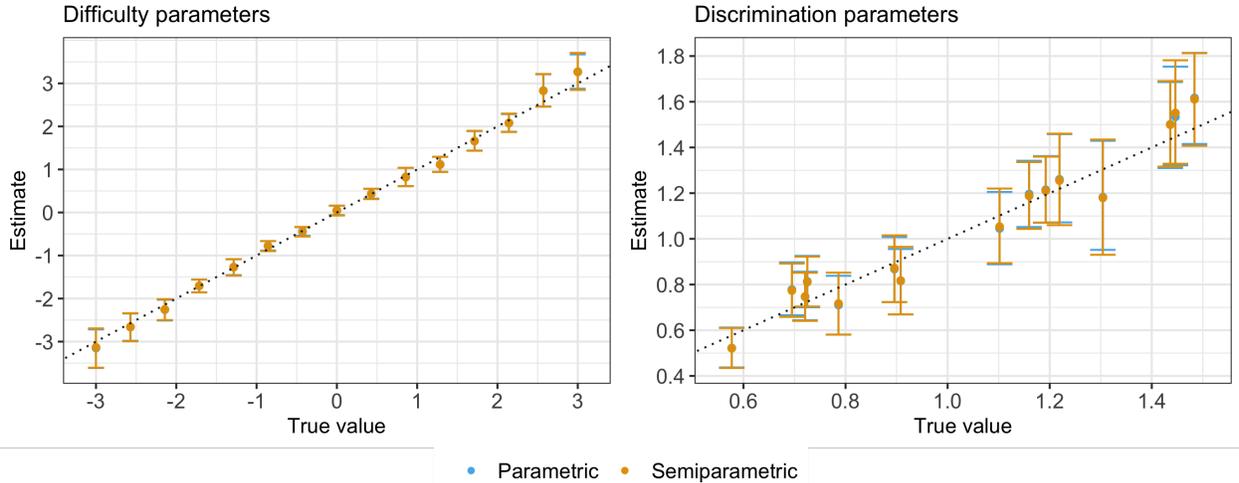


Figure 6: Unimodal simulated data. Comparison of the posterior mean estimates (with 95% credible interval) of item parameters (difficulties β and discriminations λ) for parametric and semiparametric 2PL models and true simulated values. Note that some estimates from the parametric model overlap almost exactly with semiparametric ones.

In the unimodal scenario, we observe similar performance for the parametric and semi-parametric 2PL in estimating item parameters (Figure 6). As expected, we observe some differences when considering the bimodal and multimodal scenario (Figure 7). The use of a model with a more flexible distribution improves recovery of both item and ability parameters in the bimodal scenarios (Table 3). This is especially evident when comparing estimates of the discrimination parameters, in particular for larger values. Instead, in the multimodal case the inference for the item parameters seems relatively insensitive to the specification.

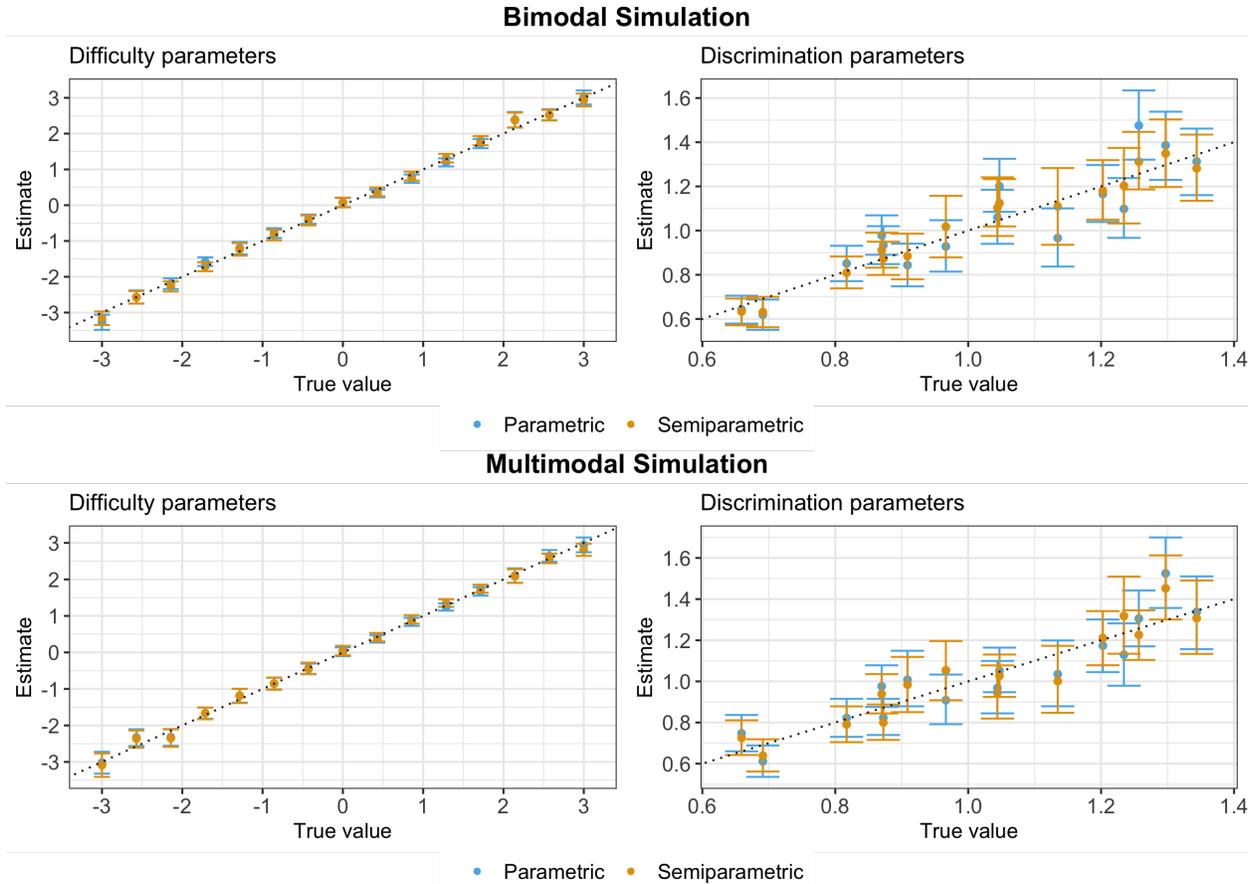


Figure 7: Bimodal Simulation (top row) and multimodal simulation (bottom row). Comparison of the posterior mean estimates (with 95% credible interval) of item parameters (difficulties β and discriminations λ) for parametric and semiparametric 2PL models and true simulated values.

Results for the ability parameters are similar when estimating abilities using the posterior means of the individual abilities (Figure 8). However, results are very different when one looks at estimates of the distribution of ability (Figure 9).

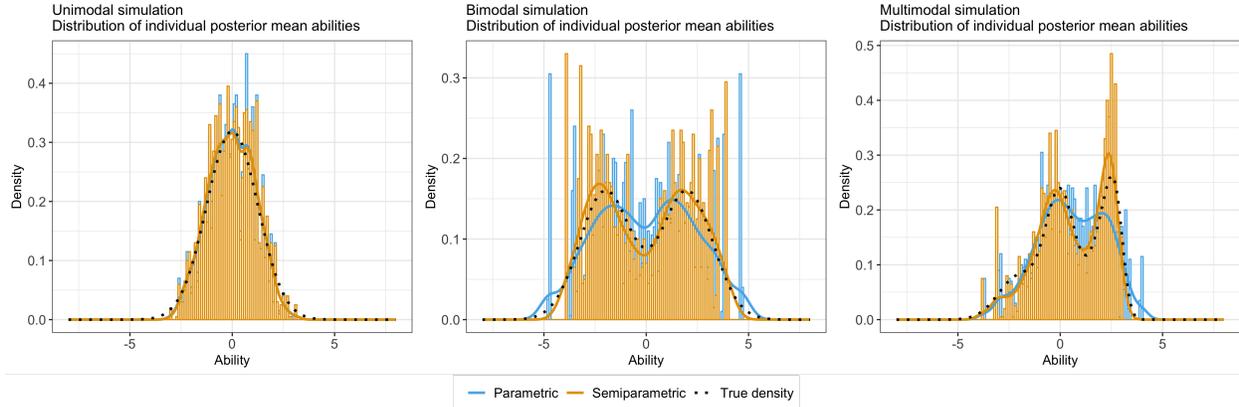


Figure 8: Histogram and density estimate of individual posterior mean abilities, under unimodal (left column), bimodal (middle column), and multimodal (right column) scenarios compared with the true density (dotted line).

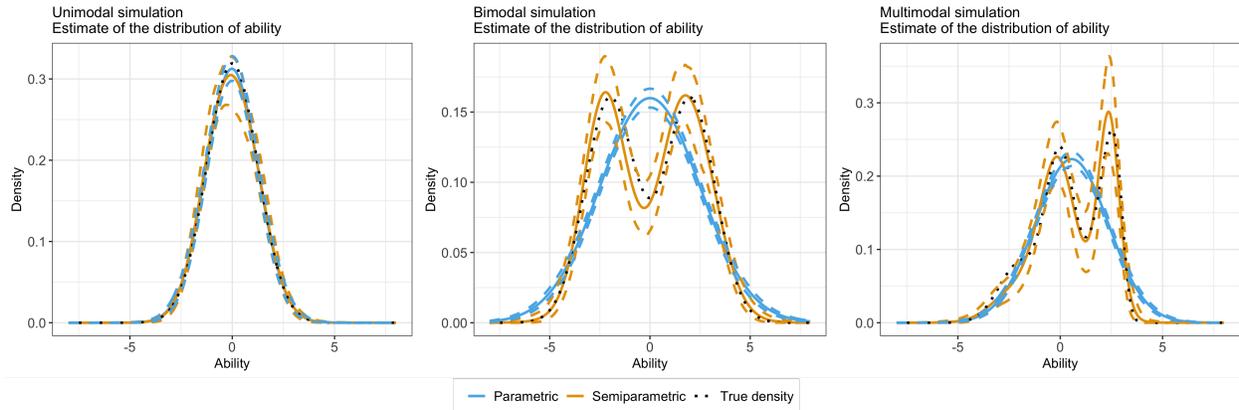


Figure 9: Distribution of ability estimated under the unimodal (left column), bimodal (middle column), and multimodal (right column) scenarios, compared with the true density (dotted line). Dashed lines indicate 95% credible intervals for the estimated distributions.

The normality assumption of the parametric model leads to unimodal density estimates, inconsistent with the true distribution, whereas the semiparametric model can recover it. The posterior means of individual abilities in Figure 8 are a compromise between the inferred distribution of ability and the information in the data, so with sufficient observations, one can obtain estimates of the distribution that are reasonable even with severe model misspecification. In other words, for misspecified parametric models, the in-sample predictions for observed individuals can be reasonable, while the out-of-sample predictions based on (10) for new individuals are poor. Note that when using the parametric model, inspection of the posterior means of individual abilities can be used to assess model misspecification relative to the assumed parametric distribution.

Mis-specification of the distribution of ability has limited effect when estimating individual percentiles. In Figure 10 we compare the posterior mean estimates of individual

percentiles with the percentiles calculated using the true distribution assumed for the simulation, for a subset of 50 individuals. Overall, the parametric and semiparametric estimates produce similar results even when the estimated density is largely different. This is because estimation of percentile is basically ranking the individuals; it makes sense that ranking is relatively insensitive to the estimation of the distribution of ability.

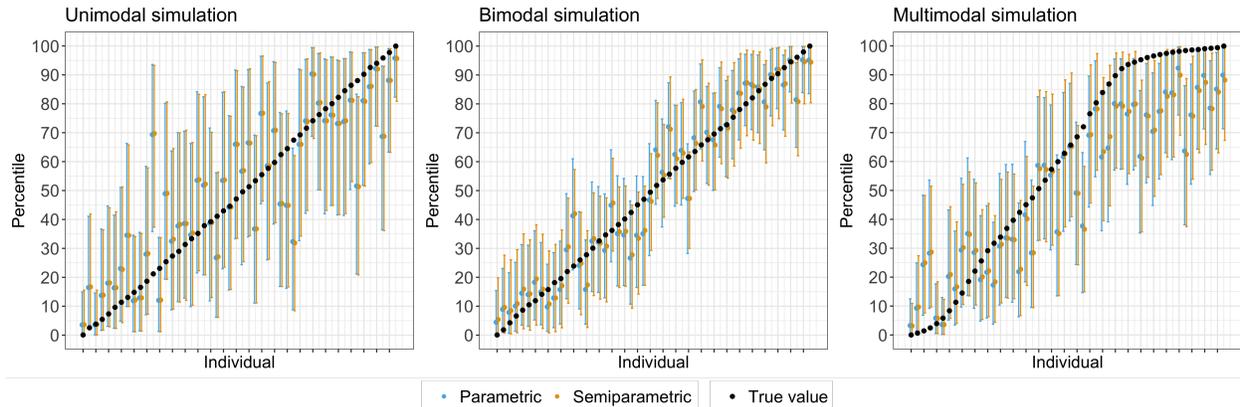


Figure 10: Estimates of individual percentiles (with 95% credible interval) for a subset of 50 individuals with varying (true) ability levels under the unimodal (left column), bimodal (middle column), and multimodal (right column). Black dots correspond to true percentiles.

7.2 Inferential results for real-world data

For the real data examples we graphically inspect results from the parametric and semiparametric models. To compare the overall model fit in the parametric and semiparametric cases, we computed the Widely Applicable Information Criterion (WAIC) (Watanabe & Opper, 2010). We refer to the NIMBLE user manual for details about WAIC calculation (de Valpine et al., 2022, see Section 7.7).

We found that for the health data the semiparametric model performs better than the parametric one (WAIC of 68,527 versus 70,414), while for the TIMSS data the best model is the semiparametric 3PL (WAIC of 228,871; parametric 3PL: 229,123; semiparametric 2PL: 229,529; parametric 2PL: 229,575). Hence in this Section, we show inferential results using samples from the IRT unconstrained model for both the health and TIMSS data, using the 3PL model for TIMMS. We also report inferential results for the 2PL model in the Supplementary Materials (Section E).

In Figure 11 we compare item parameter estimates from the two models for the health data application, while Figure 12 shows estimates for the distribution of abilities. Recall that, in this case, we interpret the latent ability as physical ability, with high values characterizing healthy individuals. As with the bimodal simulation, estimates from the parametric model of the distribution of physical ability are quite different than the distribution of individual posterior mean abilities. It is clear that the parametric model is badly mis-specified and would produce bad out-of-sample predictions. In contrast, the semiparametric model seems to nicely characterize multi-modality in the latent distribution. We observe in Figure 12 large credible intervals for high values of this distribution that can be explained by the presence of

many individuals with high raw scores, (i.e., 9 or 10 out of 10, Figure 2) for whom the model can clearly determine that their physical abilities are high, but with the exact magnitudes being difficult to identify.

The two modeling assumptions yield different estimates of the item parameters (Figure 11), with this difference being higher for extreme values. However, the relative ranking of the items is roughly the same in both cases, with for example item 1 (*Vigorous activities*) being the most difficult item and the one with the lowest value of the discrimination parameter. According to the parametric model, discrimination parameters for item 3 (*Lift/carry*) and item 10 (*Bathing/dressing*) should have similar values, while the semiparametric model separates them.

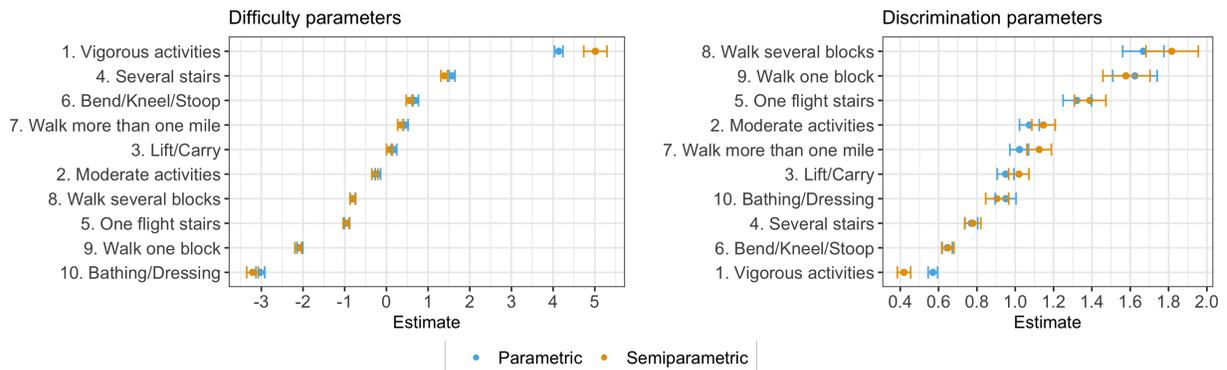


Figure 11: Health data. Comparison of item parameter estimates from the parametric and semiparametric models. In each panel items are ordered by increasing values of the parameter estimate under the semiparametric model.

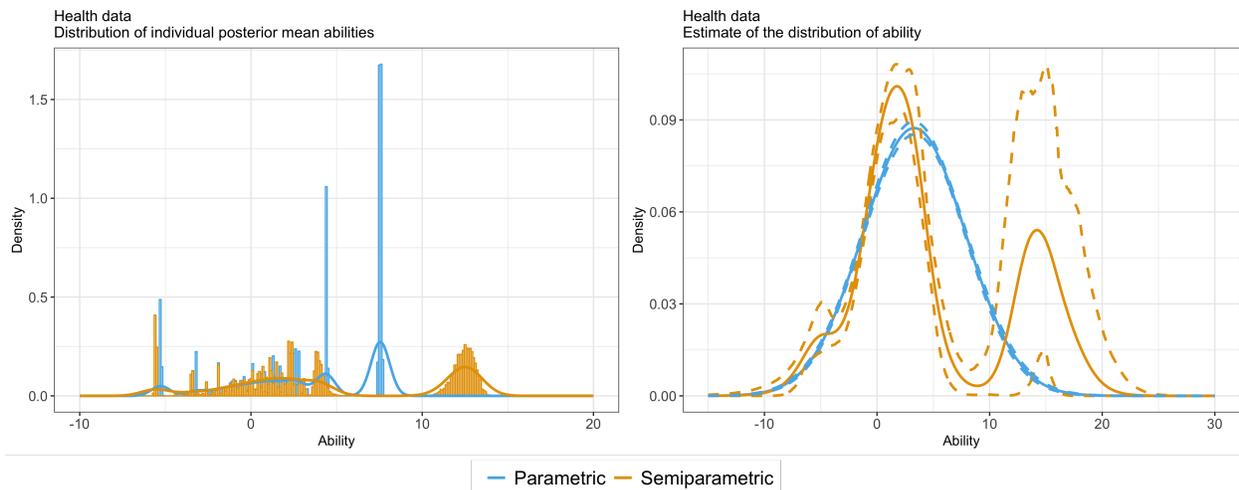


Figure 12: Health Data. Histogram and density estimate of the posterior means of the latent abilities (left panel) and estimate of the posterior distribution for the latent abilities (right panel). Dashed lines indicate 95% credible intervals for the estimated distributions.

For the TIMSS data we only compare point estimates of the item parameters (Figure 13),

due to the large number of parameters. Estimates for the difficulty parameters and discrimination parameters have the largest difference for low/high values of the parameters, while estimates of the guessing parameters are quite different. Figure 14 shows estimates of the distribution of ability. In this case the semiparametric model estimate shows departure from the normal parametric assumption, with multimodality in the estimated distribution. Differences in the distribution of ability between the semiparametric and parametric model may explain the large discrepancies between the guessing parameter estimates; however, further investigation of this matter is outside the scope of the paper. We found also that the estimate of the ability distribution is quite different under the semiparametric 3PL and semiparametric 2PL model with the distribution under the 2PL model being unimodal and right-skewed instead of multimodal (see Supplementary Materials, Section F).

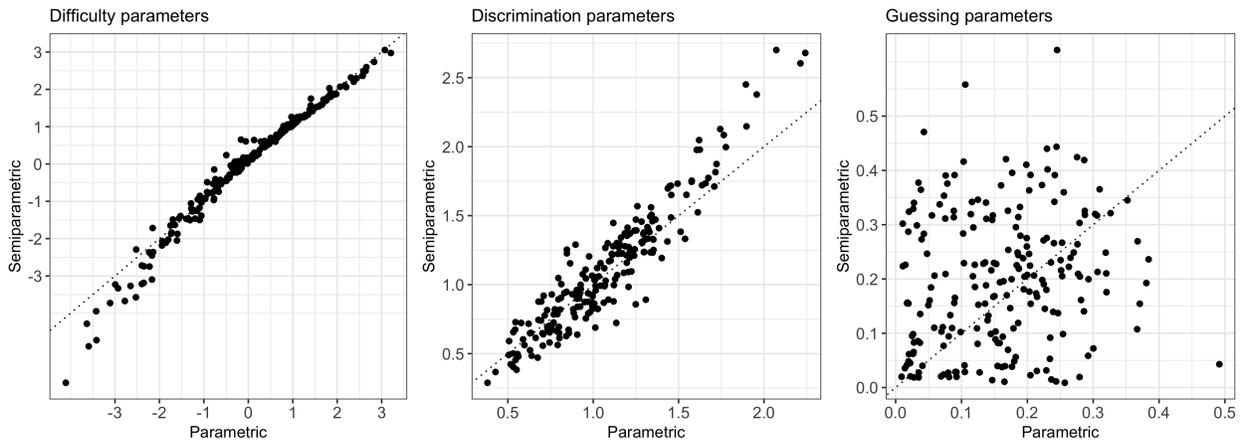


Figure 13: TIMSS data. Comparison of posterior estimates of the item parameters between the parametric and semiparametric 3PL model both using the SI unconstrained centered sampling strategy.

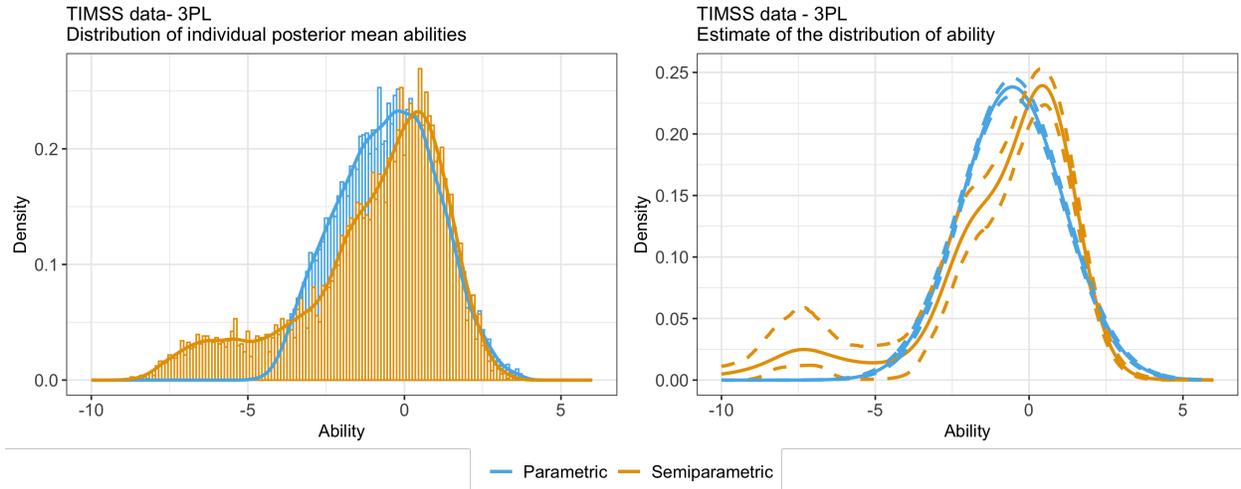


Figure 14: TIMSS Data. Histogram and density estimate of the posterior means of the latent abilities (left panel), and estimate of the posterior distribution for the latent abilities (right panel). Dashed lines indicate 95% credible intervals for the estimated distributions.

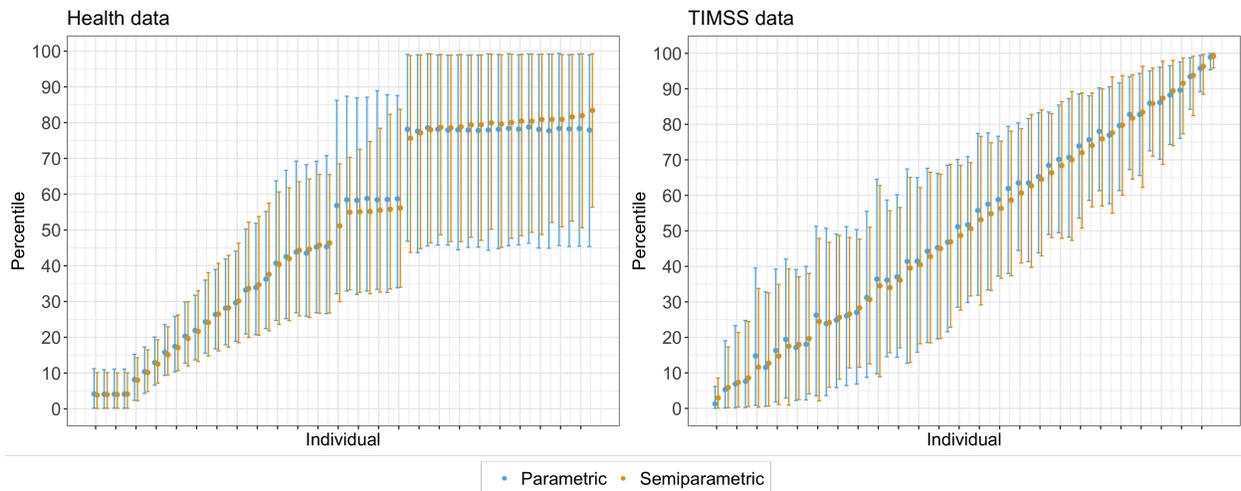


Figure 15: Estimates of individual percentiles (with 95% credible interval) for a subset of 50 individuals, for the health data (left panel) and TIMSS data (right panel).

Figure 15 compares estimates of the percentiles for both the health and TIMSS data for a sample of 50 individuals sorted according to the point estimates of the abilities from the semiparametric model. There are moderate differences in percentile values and individual ordering between the parametric and semiparametric models; in particular some estimates are associated with larger intervals than in the semiparametric case.

8 Discussion

In this paper, we consider a semiparametric extension for binary logistic IRT models, using Dirichlet process mixtures as a nonparametric prior to flexibly characterize the distribution of ability. We provide an overview of these models and study how different sets of constraints can address identifiability issue and lead to different MCMC estimation strategies.

Focusing on the 2PL and 3PL models, we compare efficiency and inferential results under different sampling strategies based on model parametrization, constraints and sampling algorithms. We find that MCMC performance across strategies can vary in relation to the underlying shape of the latent distribution and the total number of parameters.

When moving to semiparametric modeling, the computational cost can be high for large datasets, given that sampling from the Dirichlet process requires iteration through all individuals. However we find computational costs to be reasonable in our applications in light of the better inferential results.

In particular under model mis-specification, inference for item parameters worsens noticeably in the parametric model compared to the semiparametric model. With sufficient data, inference for the abilities of observed individuals can be decent even under mis-specification of the distribution of ability, but inference for the unknown latent distribution (i.e., the predictive distribution for new individuals) as a whole can be quite bad. Although parametric IRT models can work well in applications in educational assessment, having access to semiparametric models can be broadly useful as it allows inference on the entire underlying distribution of ability and its functionals. This is particularly relevant in contexts where the distribution of the individual latent trait is more complicated, for example, when measuring health (Smits, Ögreden, Garnier-Villarreal, Terwee, & Chalmers, 2020) or psychological outcomes (Reise & Rodriguez, 2016).

Results of this work potentially can be generalized to versions of binary IRT models using different prior distributions or link functions (e.g. probit), since we considered general MCMC sampling algorithms as opposed to algorithms tailored to specific choices of such model components. As a general recommendation for IRT models, we found that sampling strategies using parameter expansion are more efficient than those embedding the constraints.

In this work we extensively use the NIMBLE software for hierarchical modeling, with code reproducing results in the paper available at https://github.com/salleuska/IRT_nimble_code. Although there are other software solutions enabling Bayesian nonparametric modeling, these are often limited in the type of algorithms or in the class of models available. NIMBLE offers a high degree of flexibility in that the models considered in this paper could be easily embedded in more complicated ones. Sampler assignment can be highly customized by the user, including user-defined sampling algorithms. This customizability makes NIMBLE a powerful platform for comparing different sampling strategies. At the same time, NIMBLE allows easy sharing of the most successful strategies as black-box implementations for end users.

References

Aldous, D. J. (1985). Exchangeability and related topics. In *École d'été de probabilités de*

- saint-flour xiii—1983* (pp. 1–198). Springer.
- Antoniak, C. E. (1974, 11). Mixtures of Dirichlet Processes with applications to Bayesian nonparametric problems. *Annals of Statistics*, 2(6), 1152–1174. Retrieved from <https://doi.org/10.1214/aos/1176342871> doi: 10.1214/aos/1176342871
- Azevedo, C. L., Bolfarine, H., & Andrade, D. F. (2011). Bayesian inference for a skew-normal IRT model under the centred parameterization. *Computational Statistics & Data Analysis*, 55(1), 353–365.
- Azzalini, A. (1985). A class of distributions which includes the normal ones. *Scandinavian Journal of Statistics*, 171–178.
- Bafumi, J., Gelman, A., Park, D. K., & Kaplan, N. (2005). Practical issues in implementing and understanding Bayesian ideal point estimation. *Political Analysis*, 13(2), 171–187.
- Bambirra Gonçalves, F., da Costa Campos Dias, B., & Machado Soares, T. (2018). Bayesian item response model: a generalized approach for the abilities' distribution using mixtures. *Journal of Statistical Computation and Simulation*, 88(5), 967–981.
- Bedrick, E. J., Christensen, R., & Johnson, W. (1996). A new perspective on priors for generalized linear models. *Journal of the American Statistical Association*, 91(436), 1450–1460.
- Berger, J. O., Bernardo, J. M., & Sun, D. (2009). The formal definition of reference priors. *The Annals of Statistics*, 37(2), 905–938. Retrieved from <http://www.jstor.org/stable/30243652>
- Berger, J. O., & Pericchi, L. R. (1996). The intrinsic Bayes factor for model selection and prediction. *Journal of the American Statistical Association*, 91(433), 109–122.
- Bernardo, J. M. (1979). Reference posterior distributions for Bayesian inference. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(2), 113–128.
- Betancourt, M., Byrne, S., Livingstone, S., & Girolami, M. (2017). The geometric foundations of Hamiltonian Monte Carlo. *Bernoulli*, 23(4A), 2257–2298.
- Blackwell, D., & MacQueen, J. B. (1973). Ferguson distributions via Pólya urn schemes. *The Annals of Statistics*, 1(2), 353–355.
- Bolt, D. M., Cohen, A. S., & Wollack, J. A. (2001). A mixture item response model for multiple-choice data. *Journal of Educational and Behavioral Statistics*, 26(4), 381–409. Retrieved from <http://www.jstor.org/stable/3648167>
- Bürkner, P.-C. (2021). Bayesian item response modeling in R with brms and Stan. *Journal of Statistical Software*, 100(5), 1–54. Retrieved from <https://www.jstatsoft.org/index.php/jss/article/view/v100i05> doi: 10.18637/jss.v100.i05
- Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., . . . Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software, Articles*, 76(1), 1–32. Retrieved from <https://www.jstatsoft.org/v076/i01> doi: 10.18637/jss.v076.i01
- Clinton, J., Jackman, S., & Rivers, D. (2004). The statistical analysis of roll call data. *American Political Science Review*, 355–370.
- de Valpine, P., Paciorek, C., Turek, D., Michaud, N., Anderson-Bergman, C., Obermeyer, F., . . . Hug, J. (2022). Nimble user manual [Computer software manual]. Retrieved from <https://r-nimble.org> doi: 10.5281/zenodo.1211190
- de Valpine, P., Paciorek, C. J., Turek, D., Michaud, N., Anderson-Bergman, C., Obermeyer, F., . . . Paganin, S. (2020). *NIMBLE: MCMC, particle filtering, and programmable hi-*

- erarchical modeling*. Retrieved from <https://cran.r-project.org/package=nimble> (R package version 0.9.1) doi: 10.5281/zenodo.1211190
- de Valpine, P., Turek, D., Paciorek, C. J., Anderson-Bergman, C., Lang, D. T., & Bodik, R. (2017). Programming with models: Writing statistical algorithms for general model structures with nimble. *Journal of Computational and Graphical Statistics*, *26*(2), 403–413. Retrieved from <https://doi.org/10.1080/10618600.2016.1172487> doi: 10.1080/10618600.2016.1172487
- Duncan, K. A., & MacEachern, S. N. (2008). Nonparametric Bayesian modelling for item response. *Statistical Modelling*, *8*(1), 41–66.
- Escobar, M. D., & West, M. (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, *90*(430), 577–588.
- Ferguson, T. S. (1973, 03). A Bayesian analysis of some nonparametric problems. *The Annals of Statistics*, *1*(2), 209–230. doi: 10.1214/aos/1176342360
- Finch, H., & Edwards, J. M. (2016). Rasch model parameter estimation in the presence of a nonnormal latent trait using a nonparametric Bayesian approach. *Educational and Psychological Measurement*, *76*(4), 662–684. Retrieved from <https://doi.org/10.1177/0013164415608418> doi: 10.1177/0013164415608418
- Flegal, J. M., Hughes, J., Vats, D., Dai, N., Gupta, K., & Maji, U. (2021). mcmcse: Monte Carlo standard errors for MCMC [Computer software manual]. Riverside, CA, and Kanpur, India. (R package version 1.5-0)
- Fox, J.-P. (2010). *Bayesian item response modeling: Theory and applications*. Springer Science & Business Media.
- Furr, D. C. (2017). edstan: Stan models for Item Response Theory [Computer software manual]. Retrieved from <https://CRAN.R-project.org/package=edstan> (R package version 1.0.6)
- Gelfand, A. E., & Kottas, A. (2002). A computational approach for full nonparametric Bayesian inference under Dirichlet Process mixture models. *Journal of Computational and Graphical Statistics*, *11*(2), 289–305. Retrieved from <https://doi.org/10.1198/106186002760180518> doi: 10.1198/106186002760180518
- Gelman, A. (2004). Parameterization and Bayesian modeling. *Journal of the American Statistical Association*, *99*(466), 537–545.
- Geweke, J. F., & Singleton, K. J. (1981). Maximum likelihood “confirmatory” factor analysis of economic time series. *International Economic Review*, 37–54.
- Guhaniyogi, R., & Rodriguez, A. (2020). Joint modeling of longitudinal relational data and exogenous variables. *Bayesian Analysis*, *15*(2), 477–503.
- Hays, R. D., Morales, L. S., & Reise, S. P. (2000). Item response theory and health outcomes measurement in the 21st century. *Medical Care*, *38*(9 Suppl), II28.
- Hoffman, M. D., & Gelman, A. (2014). The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, *15*(1), 1593–1623.
- Ibrahim, J. G. (1997). On properties of predictive priors in linear models. *The American Statistician*, *51*(4), 333–337. Retrieved from <https://www.tandfonline.com/doi/abs/10.1080/00031305.1997.10474408> doi: 10.1080/00031305.1997.10474408
- Jara, A., Hanson, T. E., Quintana, F. A., Müller, P., & Rosner, G. L. (2011). DPpackage: Bayesian semi-and nonparametric modeling in R. *Journal of Statistical Software*, *40*(5),

1. Johnson, M. S. (2007). Modeling dichotomous item responses with free-knot splines. *Computational Statistics & Data Analysis*, 51(9), 4178–4192.
- Joint Health Surveys Unit of Social and Community Planning Research and University College London. (2017). *Health survey for england, 1996. [data collection]*. 5th Edition. UK Data Service. SN: 3886. Retrieved from <http://doi.org/10.5255/UKDA-SN-3886-2>
- Karabatsos, G. (2017). Bayesian nonparametric response models. In W. J. Van der Linden (Ed.), *Handbook of item response theory, volume one: Models* (p. 323-336). CRC Press. Retrieved from <https://www.taylorfrancis.com/books/e/9781466514423/chapters/10.1201/9781315374512-32>
- Kirisci, L., chi Hsu, T., & Yu, L. (2001). Robustness of item parameter estimation programs to assumptions of unidimensionality and normality. *Applied Psychological Measurement*, 25(2), 146-162. Retrieved from <https://doi.org/10.1177/01466210122031975> doi: 10.1177/01466210122031975
- Laird, N. M. (1978). Nonparametric maximum likelihood estimation of a mixing distribution. *Journal of the American Statistical Association*, 73, 805-807.
- Li, Y., Müller, P., & Lin, X. (2011). Center-adjusted inference for a nonparametric Bayesian random effect distribution. *Statistica Sinica*, 21(3), 1201–1223.
- Liu, C., Rubin, D. B., & Wu, Y. N. (1998). Parameter expansion to accelerate EM: the PX-EM algorithm. *Biometrika*, 85(4), 755–770.
- Liu, J. S., & Wu, Y. N. (1999). Parameter expansion for data augmentation. *Journal of the American Statistical Association*, 94(448), 1264–1274.
- Liu, S. J. (1996). Nonparametric hierarchical Bayes via sequential imputations. *The Annals of Statistics*, 24(3), 911 – 930. Retrieved from <https://doi.org/10.1214/aos/1032526949> doi: 10.1214/aos/1032526949
- Lo, A. Y. (1984). On a class of Bayesian nonparametric estimates: I. density estimates. *The Annals of Statistics*, 351–357.
- McHorney, C. A., Haley, S. M., & Ware Jr, J. E. (1997). Evaluation of the MOS SF-36 physical functioning scale (PF-40): II. comparison of relative precision using Likert and Rasch scoring methods. *Journal of Clinical Epidemiology*, 50(4), 451–461.
- Micceri, T. (1989). The unicorn, the normal curve, and other improbable creatures. *Psychological bulletin*, 105(1), 156.
- Mislevy, R. J. (1984). Estimating latent distributions. *Psychometrika*, 49(3), 359–381. Retrieved from <https://doi.org/10.1007/BF02306026> doi: 10.1007/BF02306026
- Miyazaki, K., & Hoshino, T. (2009). A Bayesian semiparametric item response model with Dirichlet Process priors. *Psychometrika*, 74(3), 375–393.
- Natesan, P., Nandakumar, R., Minka, T., & Rubright, J. D. (2016). Bayesian prior choice in IRT estimation using MCMC and variational Bayes. *Frontiers in Psychology*, 7, 1422.
- Nguyen, D., de Valpine, P., Atchade, Y., Turek, D., Michaud, N., & Paciorek, C. J. (2020, 12). Nested adaptation of mcmc algorithms. *Bayesian Analysis*, 15(4), 1323–1343. Retrieved from <https://doi.org/10.1214/19-BA1190> doi: 10.1214/19-BA1190
- Paulon, G., De Iorio, M., Guglielmi, A., & Ieva, F. (2018, 07). Joint modeling of recurrent events and survival: a Bayesian non-parametric approach. *Biostatistics*, 21(1), 1-14.

- Retrieved from <https://doi.org/10.1093/biostatistics/kxy026> doi: 10.1093/biostatistics/kxy026
- Pitman, J. (1996). Some developments of the blackwell-macqueen urn scheme. In T. S. Ferguson, L. S. Shapley, & J. B. MacQueen (Eds.), *Statistics, probability and game theory* (Vol. Volume 30, pp. 245–267). Hayward, CA: Institute of Mathematical Statistics. Retrieved from <https://doi.org/10.1214/lnms/1215453576> doi: 10.1214/lnms/1215453576
- Qin, L. (1998). *Nonparametric Bayesian models for item response data* (Unpublished doctoral dissertation). The Ohio State University.
- Rasch, G. (1990). *Probabilistic models for some intelligence and attainment tests*. Copenhagen: Danish Institute for Educational Research.
- Reise, S., & Rodriguez, A. (2016). Item response theory and the measurement of psychiatric constructs: some empirical and conceptual issues and challenges. *Psychological Medicine*, *46*(10), 2025–2039.
- Rutkowski, L., Gonzalez, E., Joncas, M., & von Davier, M. (2010). International large-scale assessment data: issues in secondary analysis and reporting. *Educational Researcher*, *39*(2), 142–151.
- Samejima, F. (1997). Departure from normal assumptions: a promise for future psychometrics with substantive mathematical modeling. *Psychometrika*, *62*(4), 471–493.
- San Martín, E., Jara, A., Rolin, J.-M., & Mouchart, M. (2011, Jul 01). On the Bayesian nonparametric generalization of IRT-type models. *Psychometrika*, *76*(3), 385–409. Retrieved from <https://doi.org/10.1007/s11336-011-9213-9> doi: 10.1007/s11336-011-9213-9
- Schmitt, J. E., Mehta, P. D., Aggen, S. H., Kubarych, T. S., & Neale, M. C. (2006). Semiparametric methods for detecting latent non-normality: A fusion of latent trait and ordered latent class modeling. *Multivariate Behavioral Research*, *41*(4), 427–443.
- Seong, T. (1990). Sensitivity of marginal maximum likelihood estimation of item and ability parameters to the characteristics of the prior ability distributions. *Applied psychological measurement*, *14*(3), 299–311.
- Sethuraman, J. (1994). A constructive definition of Dirichlet priors. *Statistica Sinica*, *4*(2), 639–650.
- Sheng, Y. (2010). A sensitivity analysis of gibbs sampling for 3PNO IRT models: Effects of prior specifications on parameter estimates. *Behaviormetrika*, *37*(2), 87–110.
- Smits, N., Öğreden, O., Garnier-Villarreal, M., Terwee, C. B., & Chalmers, R. P. (2020). A study of alternative approaches to non-normal latent trait distributions in item response theory models used for health outcome measurement. *Statistical Methods in Medical Research*, *29*(4), 1030–1048.
- Sosa, J., & Rodríguez, A. (2021). A latent space model for cognitive social structures data. *Social Networks*, *65*, 85 - 97.
- Stan Development Team. (2018). *Stan modeling language users guide and reference manual, version 2.18.0*. Retrieved from <http://mc-stan.org>
- Vats, D., Flegal, J. M., & Jones, G. L. (2019). Multivariate output analysis for Markov chain Monte Carlo. *Biometrika*, *106*(2), 321–337.
- Ware, J. E. (2003). *Sf-36 health survey: Manual and interpretation guide..*
- Watanabe, S., & Opper, M. (2010). Asymptotic equivalence of Bayes cross validation and

- widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, 11(12).
- Woods, C. M. (2007). Empirical histograms in item response theory with ordinal data. *Educational and Psychological Measurement*, 67(1), 73–87.
- Woods, C. M., & Thissen, D. (2006). Item response theory with estimation of the latent population distribution using spline-based densities. *Psychometrika*, 71(2), 281.
- Yang, M., & Dunson, D. B. (2010). Bayesian semiparametric structural equation models with latent variables. *Psychometrika*, 75(4), 675–693.
- Yang, M., Dunson, D. B., & Baird, D. (2010). Semiparametric Bayes hierarchical models with mean and variance constraints. *Computational Statistics & Data Analysis*, 54(9), 2172–2186.

Supplementary Material to “Computational strategies and estimation performance with Bayesian semiparametric Item Response Theory models”

A. Identifiability

The 2PL model is not identifiable based on the likelihood. Here we demonstrate the non-identifiability for the two parameterizations we consider, showing how different linear transformations lead to the same probabilities. Note that these transformations are defined for every parameter associated with each item $i = 1, \dots, I$ and individual $j = 1, \dots, N$.

Under the IRT parameterization:

1. $\eta'_j = \eta_j/s$ and $\lambda'_i = s\lambda_i$

$$\lambda'_i(\eta'_j - \beta_i) = s\lambda_i(\eta_j/s - \beta_i) = \lambda_i\eta_j - \lambda_i\beta_i = \lambda_i(\eta_j - \beta_i),$$

2. $\eta'_j = \eta_j + c$ and $\beta'_i = \beta_i + c$,

$$\lambda_i(\eta'_j - \beta'_i) = \lambda_i(\eta_j + c - (\beta_i + c)) = \lambda_i(\eta_j - \beta_i).$$

Under the slope-intercept parameterization:

1. $\eta'_j = \eta_j/s$ and $\lambda'_i = s\lambda_i$,

$$\lambda'_i\eta'_j + \gamma_i = s\lambda_i\eta_j/s + \gamma_i = \lambda_i\eta_j + \gamma_i,$$

2. $(\lambda_i\eta_j)' = \lambda_i\eta_j + c$ and $\gamma'_i = \gamma_i - c$, or $\eta'_j = \eta_j + c$ and $\gamma'_i = \gamma_i - \lambda_i c$

$$\lambda_i\eta'_j + \gamma'_i = \lambda_i(\eta_j + c) + \gamma_i - \lambda_i c = \lambda_i\eta_j + \gamma_i.$$

Post-processing to satisfy identifiability constraints

This section reports the transformations we apply to item and ability parameters in order to satisfy the identifiability constraints in our base parameterization (??). These transformations are applied to each posterior sample.

Under the IRT parameterization, the set of transformations for each posterior sample of $\{\lambda_i, \beta_i, \eta_j\}$ for $i = 1, \dots, I, j = 1, \dots, N$ takes these forms:

$$\lambda_i^* = s\lambda_i, \quad \beta_i^* = \frac{\beta_i - b}{s}, \quad \eta_j^* = \frac{\eta_j - b}{s},$$

subject to $\prod_{i=1}^I \lambda_i^* = 1$, $\sum_{i=1}^I \beta_i^* = 0$. By solving the system of equations given by the transformations and the set of identifiability constraints, we obtain

$$s = \exp \left\{ \sum_{i=1}^I \log(\lambda_i) / I \right\}, \quad b = \frac{\sum_{i=1}^I \beta_i}{I}. \quad (1)$$

Under the slope-intercept parameterization, the set of transformations for each posterior sample of $\{\lambda_i, \gamma_i, \eta_j\}$ for $i = 1, \dots, I, j = 1, \dots, N$ takes these forms:

$$\tilde{\lambda}_i = s\lambda_i, \quad \tilde{\gamma}_i = \gamma_i - \lambda_i c, \quad \tilde{\eta}_j = \frac{\eta_j + c}{s},$$

subject to $\prod_{i=1}^I \tilde{\lambda}_i = 1$, $\sum_{i=1}^I \tilde{\gamma}_i = 0$. Similarly, by solving the system of equations given by the transformations and the set of identifiability constraints, we obtain

$$s = \exp \left\{ \sum_{i=1}^I \log(\lambda_i) / I \right\}, \quad c = \frac{\sum_{i=1}^I \gamma_i}{\sum_{i=1}^I \lambda_i}. \quad (2)$$

Finally, to get from the slope-intercept parameterization to the IRT parameterization, we define $\tilde{\beta}_i := -\tilde{\gamma}_i / \tilde{\lambda}_i$ and then calculate $\beta_i^* = \tilde{\beta}_i - \sum_i \tilde{\beta}_i / I$.

Rescaling the DP density

We can obtain posterior samples from the mixing distribution F via NIMBLE's `getSamplesDPmeasure` function, allowing us to estimate the density for the latent ability distribution. However, when comparing these estimated densities between models, for some of the sampling strategies, we need to transform the estimated density to account for the transformations of the abilities from the scale on which sampling is done to the scale in (??).

As an example, consider the IRT parameterization without constraints. From the MCMC output we can obtain $p(\tilde{\eta})$ evaluated for different values of $\tilde{\eta}$, but we want $p(\tilde{\eta}^*)$ with $\tilde{\eta}^* = (\tilde{\eta} - b) / s$. To do so we need the Jacobian of the transformation, which is simply s . Then, we obtain $p(\tilde{\eta}^*)$

$$p(\tilde{\eta}^*) = p_{\tilde{\eta}}(s\tilde{\eta}^* + b) \left| \frac{\partial(s\tilde{\eta}^* + b)}{\partial\tilde{\eta}^*} \right| = p_{\tilde{\eta}}(s\tilde{\eta}^* + b)s. \quad (3)$$

B. Details about the sampling algorithms

Tables 1-2 summarize the sampling algorithms used for each sampling strategy in Table ??.

Table 1: Summary of the sampling algorithms used for each parameter under the sampling strategies considered for the parametric 2PL model.

Model constraints	IRT parameterization		SI parameterization	
	MH/conjugate	HMC (Stan)	MH/conjugate	Centered
Constrained abilities	Adaptive MH $\{\log(\lambda_i), \beta_i, \eta_j\}$	HMC $\{\{\log(\lambda_i)\}, \{\beta_i\}, \{\eta_j\}\}$	Adaptive MH $\{\log(\lambda_i), \gamma_i, \eta_j\}$	Centered sampler for pairs $\{\log(\lambda_i), \gamma_i\}$ Adaptive MH $\{\eta_j\}$
Constrained item	Adaptive MH $\{\log(\lambda_i^*), \beta_i^*, \eta_j\}$	-	Adaptive MH $\{\log(\lambda_i^*), \gamma_i^*, \eta_j\}$	-
Unconstrained	Adaptive MH $\{\log(\lambda_i), \beta_i, \eta_j\}$ Conjugate $\{\mu, \sigma^2\}$	-	Adaptive MH $\{\log(\lambda_i), \gamma_i, \eta_j\}$ Conjugate $\{\mu, \sigma^2\}$	Centered sampler for pairs $\{\log(\lambda_i), \gamma_i\}$ Adaptive MH $\{\eta_j\}$ Conjugate $\{\mu, \sigma^2\}$

Table 2: Summary of the sampling algorithms used for each parameter under the sampling strategies considered for the semiparametric 2PL model.

Model constraints	IRT parameterization	SI parameterization	
	MH/conjugate	MH/conjugate	Centered
Constrained item	Adaptive MH $\{\log(\lambda_i^*), \beta_i^*, \eta_j\}$ CRP sampler $\{\alpha, \{z_j\}\}$ Conjugate $\{\{\mu_k^*\}, \{\sigma_k^{2*}\}\}$	Adaptive MH $\{\log(\lambda_i^*), \gamma_i^*, \eta_j\}$ CRP sampler $\{\alpha, \{z_j\}\}$ Conjugate $\{\{\mu_k^*\}, \{\sigma_k^{2*}\}\}$	-
Unconstrained	Adaptive MH $\{\log(\lambda_i), \beta_i, \eta_j\}$ CRP sampler $\{\alpha, \{z_j\}\}$ Conjugate $\{\{\mu_k^*\}, \{\sigma_k^{2*}\}\}$	Adaptive MH $\{\log(\lambda_i), \gamma_i, \eta_j\}$ CRP sampler $\{\alpha, \{z_j\}\}$ Conjugate $\{\{\mu_k^*\}, \{\sigma_k^{2*}\}\}$	Centered sampler for pairs $\{\log(\lambda_i), \gamma_i\}$ Adaptive MH $\{\eta_j\}$ CRP sampler $\{\alpha, \{z_j\}\}$ Conjugate $\{\{\mu_k^*\}, \{\sigma_k^{2*}\}\}$

Description of the sampling algorithms

- **Conjugate sampler:** for the models used in the paper, we exploit conjugacy results for the Normal distribution with Normal Inverse-gamma priors for the mean and variance.
- **Adaptive MH** (Metropolis-Hastings): uses a normal proposal distribution, with initial proposal variance equal to 1 and an adaptation interval of 200 iterations. The adaptation routine is implemented as given in [Shaby and Wells \(2010\)](#).
- **Centered sampler:** this is a custom defined sampler implemented by the authors. Details are given in the below.
- **CRP sampler:** under the CRP specification, the random measure G is integrated out from the model and NIMBLE assigns a collapsed sampler.
 - clustering indicators \mathbf{z} are updated as in described in [Neal \(2000\)](#);
 - the DP concentration parameter α is sampled as described in [Escobar and West \(1995\)](#) (Section 6) when a Gamma prior is used, as in the models considered in the paper. If another prior is considered, NIMBLE uses a random walk Metropolis-Hastings.

Centered sampler

We consider a custom sampler for the 2PL model under the slope-intercept parameterization. Intuition for this sampling strategy comes from the resemblance to a linear model. In order to sample $\{\lambda_i, \gamma_i\}$ efficiently, we propose centering the implied covariate, η_j , to have mean zero. This is analogous to centering covariates in a linear model, but in this case the "covariate" values are not fixed, so the centering needs to be done in each iteration. For a given item i for $i = 1, \dots, I$ we can rewrite

$$\begin{aligned}\lambda_i \eta_j + \gamma_i &= \lambda_i (\eta_j - \bar{\eta}) + \lambda_i \bar{\eta} + \gamma_i, \\ &= \lambda_i \eta_j^c + \gamma_i^c,\end{aligned}$$

such that the quantity $\eta_j^c = \eta_j - \bar{\eta}$ is centered. The idea is to propose a new value λ_i^* in this new parameterization at each MCMC iteration, using a random walk on the log scale. Translating to the original parameterization, we have:

$$\begin{aligned}\lambda_i^* \eta_j^c + \gamma_i^c &= \lambda_i^* (\eta_j - \bar{\eta}) + \lambda_i \bar{\eta} + \gamma_i, \\ &= \lambda_i^* \eta_j - \lambda_i^* \bar{\eta} + \lambda_i \bar{\eta} + \gamma_i.\end{aligned}$$

This means that we are proposing $\gamma_i^* = \gamma_i + \bar{\eta}(\lambda_i - \lambda_i^*)$. Thus we have a joint proposal $(\lambda_i^*, \gamma_i^*)$ that accounts for the usual correlation in a regression between intercept and slope. Apart from accounting for sampling λ_i on the log scale, the proposal is symmetric, so no Hastings correction is needed. The original sampler for γ_i can stay the same. This is because in the reparameterization with γ_i^c above, shifting γ_i by a certain amount is equivalent to shifting γ_i^c .

C. Health data questions

The following items are about activities you might do during a typical day. Does your health now limit you in these activities? If so, how much?

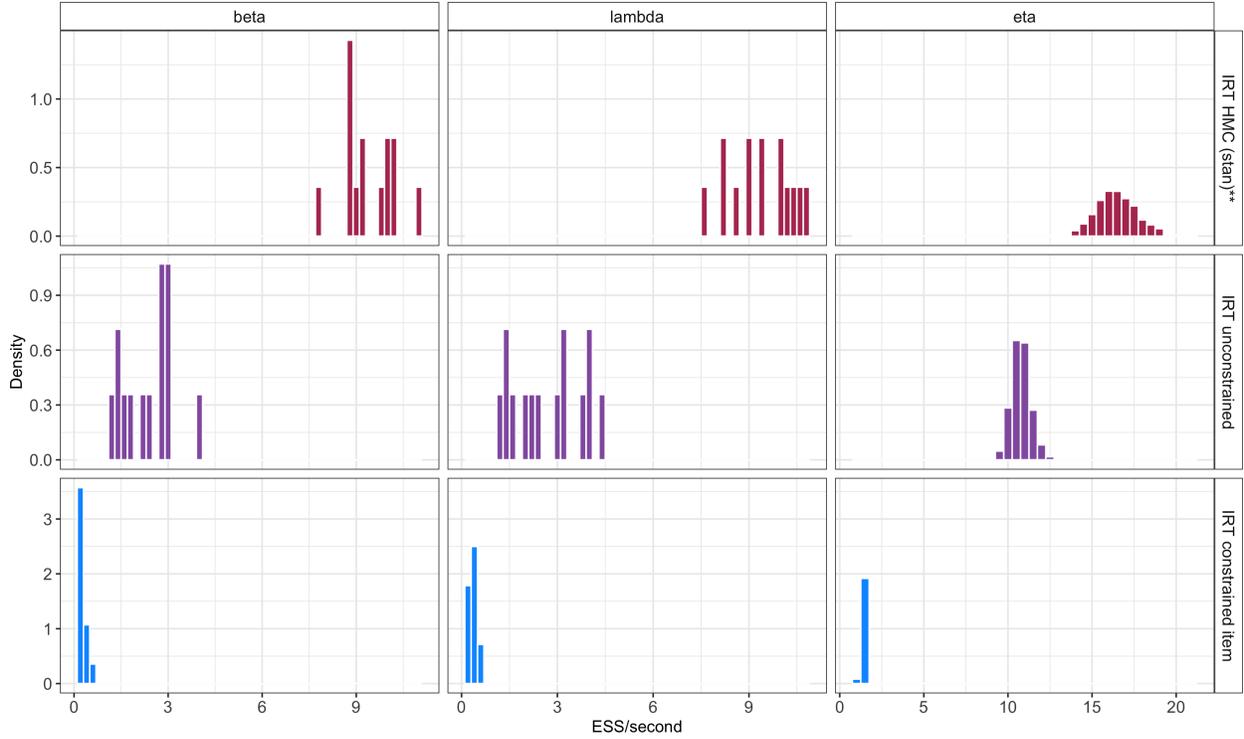
1. Vigorous activities: Vigorous activities, such as running, lifting heavy objects, participating in strenuous sports.
2. Moderate activities: Moderate activities, such as moving a table, pushing a vacuum cleaner, bowling or playing golf.
3. Lift/Carry: Lifting or carrying groceries.
4. Several stairs: Climbing several flights of stairs.
5. One flight stairs: Climbing one flight of stairs.
6. Bend/Kneel/Stoop: Bending, kneeling, or stooping.
7. Walk more mile: Walking more than a mile.
8. Walk several blocks: Walking several blocks.

9. Walk one block: Walking one block.
10. Bathing/Dressing: Bathing or dressing yourself.

D. A note on efficiency comparisons

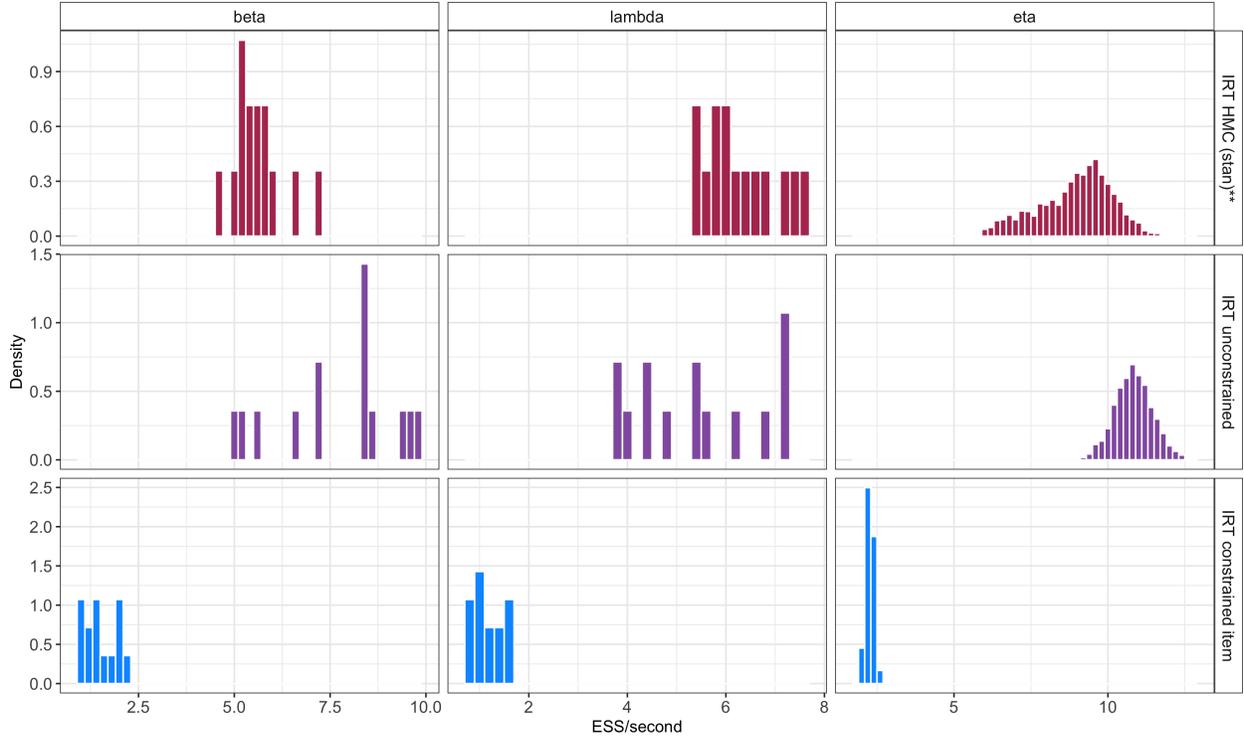
Comments on the multivariate ESS

In this section we compare univariate and multivariate efficiency metrics. In particular, we compare efficiency values based on the mESS (multivariate efficiency), with the distribution of efficiencies calculated for each parameter (univariate efficiency) using the total time. We report these metrics for each simulation scenario, selecting three representative strategies under the IRT parameterization for the parametric 2PL model. Figures 1-3 show the distribution of univariate efficiency for difficulty, discrimination and ability parameters, along with a table reporting information for the multivariate efficiency. There are some differences between the multivariate and univariate efficiency results. This is expected because the mESS provides a single scalar measure of mixing performance that accounts for cross-correlation among the parameters and does not necessarily reflect the distribution of univariate ESSs. In fact, values of the mESS can be larger than all the univariate ESSs. For example, the IRT HMC strategy has larger univariate ESS values compared to other sampling strategies, but lower mESS.



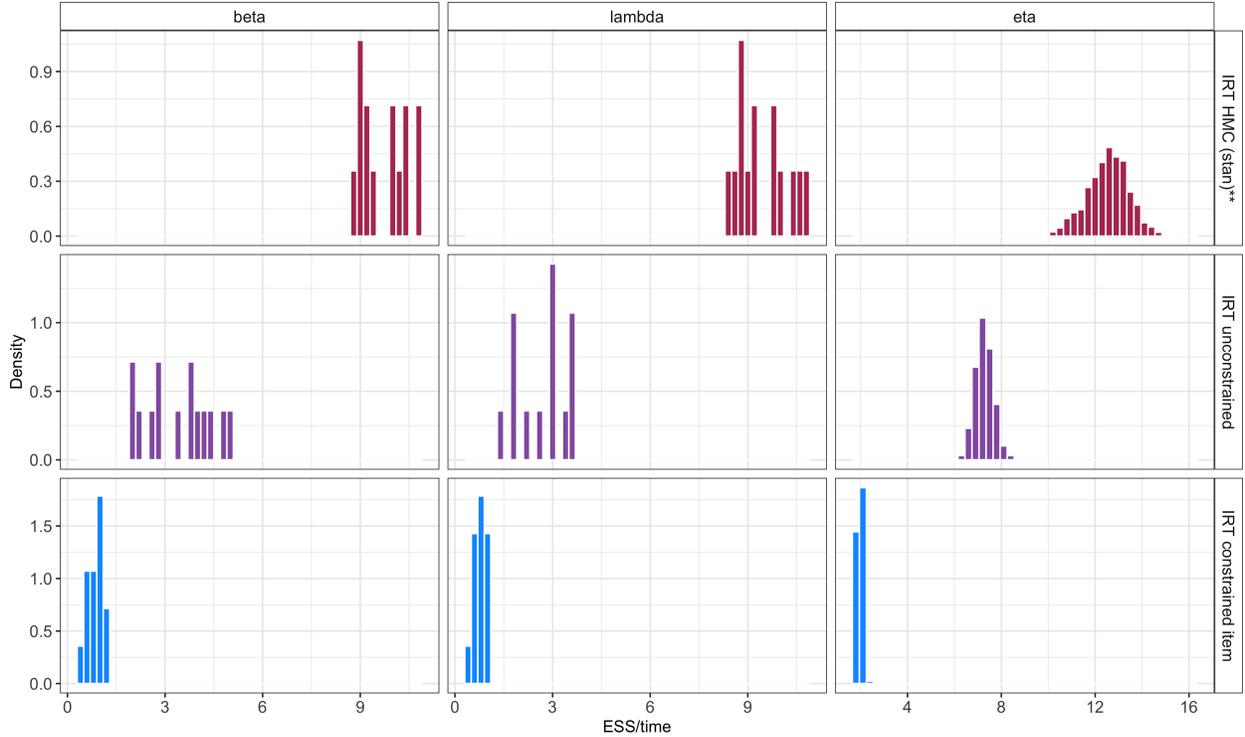
Strategy	mESS	total time (second)	efficiency (mESS/second)
IRT HMC (Stan)**	31802	1405	23
IRT unconstrained	27540	934	29
IRT constrained item	27904	7318	4

Figure 1: Unimodal simulation ($N = 2000, I = 15$). Distribution of univariate efficiencies (univariate ESS/seconds) for each group of parameters used to compute the multivariate efficiency (mESS/second) using the total time. The symbol ** denotes median results across 11 runs.



Strategy	mESS	total time (second)	efficiency (mESS/second)
IRT HMC (Stan)**	30473	2191	14
IRT unconstrained	27167	937	29
IRT constrained item	27463	4508	6

Figure 2: Bimodal simulation ($N = 2000, I = 15$). Distribution of univariate efficiencies (univariate ESS/seconds) for each group of parameters used to compute the multivariate efficiency (mESS/second) using the total time. The symbol ** denotes median results across 11 runs.



Strategy	mESS	total time (second)	efficiency (mESS/second)
IRT HMC (Stan)**	22698	1176	19
IRT unconstrained	27442	1407	19
IRT constrained item	27756	5248	5

Figure 3: Multimodal simulation ($N = 2000, I = 15$). Distribution of univariate efficiencies (univariate ESS/seconds) for each group of parameters used to compute the multivariate efficiency (mESS/second) using the total time. The symbol ** denotes median results across 11 runs.

Variability of mESS when using HMC from the Stan software

When deciding on the number of posterior, burn-in and warm-up samples, we tried to obtain a reliable estimate of the multivariate ESS. To ensure the chains were long enough, we used multiple runs for some of the experiments. We found that mESS estimates based on the chosen settings of the MCMC algorithm (i.e., number of iterations, number of burn-in or warm-up samples) have negligible variability across multiple runs for all strategies with the exception of using HMC as implemented in Stan. As an example, Figure 4 shows the distribution of multivariate ESS for the IRT constrained abilities approach.

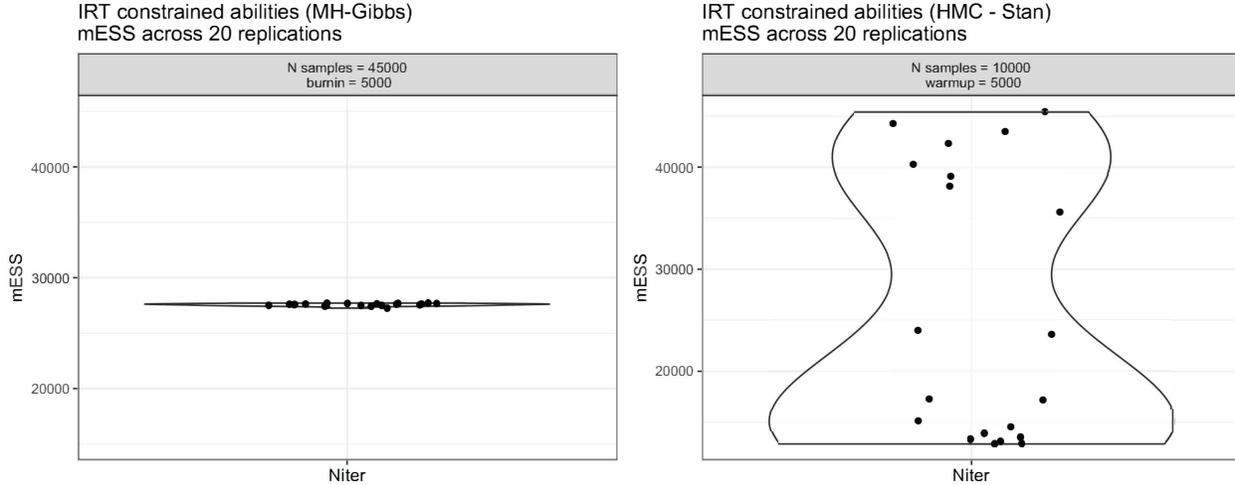


Figure 4: Comparison of mESS estimates across multiple runs for the IRT constrained abilities approach using the unimodal simulation scenario.

We also found that lower and higher values of ESS for the strategies using the HMC are highly correlated with values of the tuning parameters of the HMC (i.e., leapfrog and step-size parameters), which are typically estimated during the warm-up phase (see Figure 5).

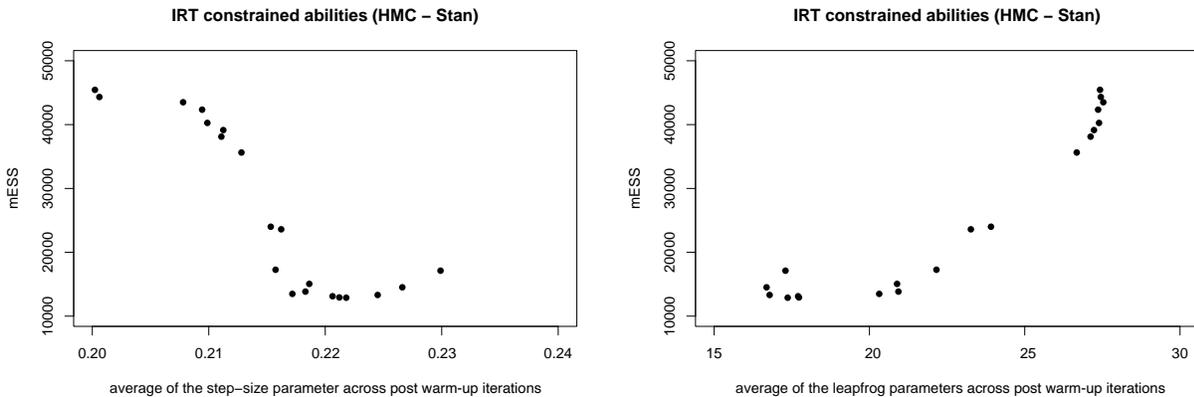


Figure 5: Estimates of the mESS versus average values of the HMC tuning parameters (step-size and leapfrog) across post warm-up iterations.

E. Results from additional simulations

We investigated how different combinations of numbers of items and individuals affect efficiency of the different sampling strategies. In particular, we simulated data under the three scenarios presented in Section ?? following a factorial design with $I \in \{10, 30\}$ and individuals $N \in \{1, 000, 5, 000\}$. We omit results for the strategy using HMC due to the high variability in estimating the mESS.

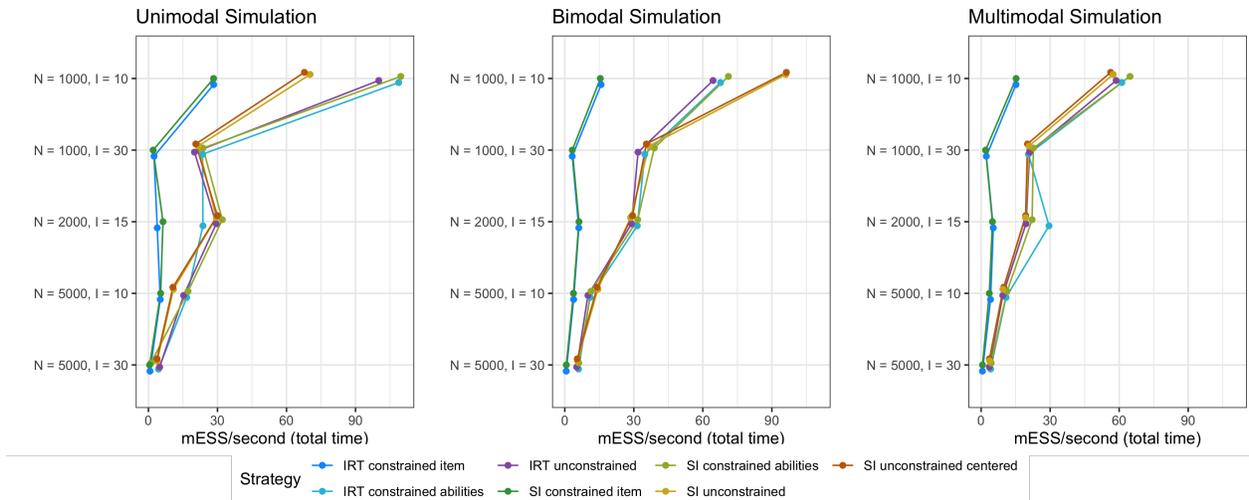


Figure 6: Multivariate ESS per second of the parametric sampling strategies (excluding HMC) across different combinations of numbers of items and individuals.

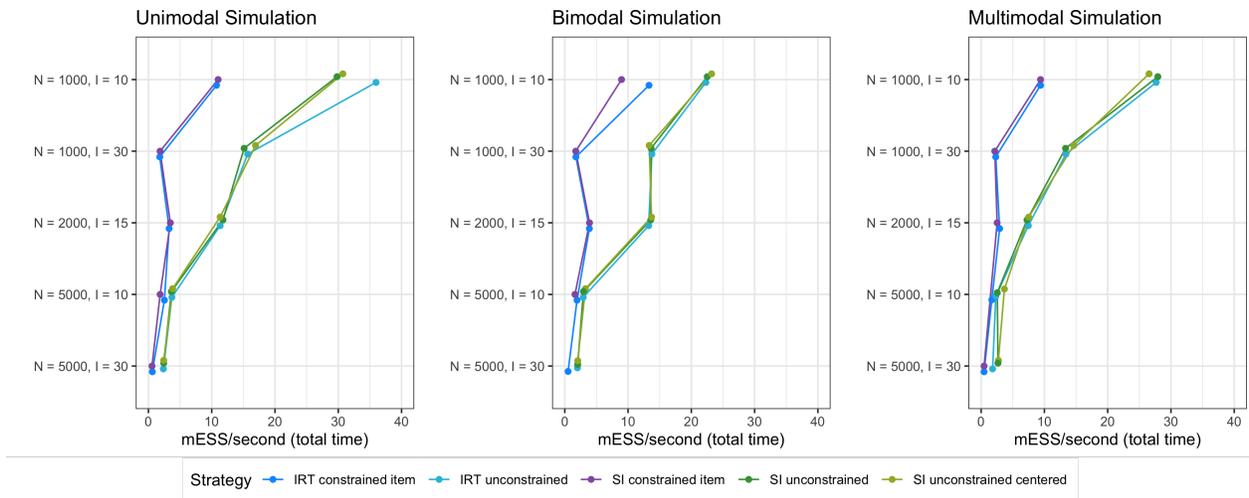


Figure 7: Multivariate ESS per second of the semiparametric sampling strategies (excluding HMC) across different combinations of numbers of items and individuals.

F. Inferential results for the TIMSS data using the 2PL model

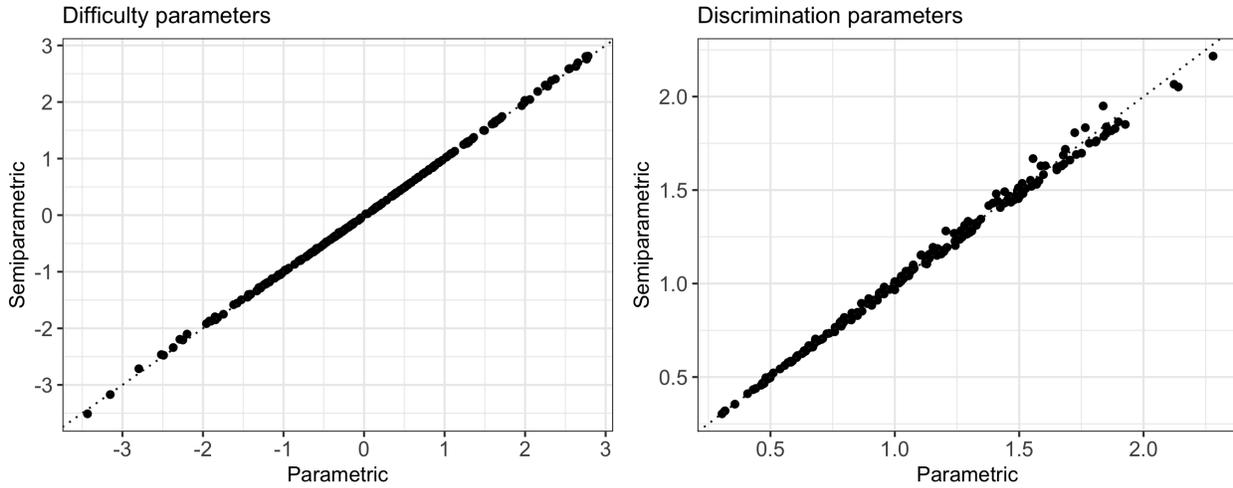


Figure 8: TIMSS data, 2PL model. Comparison of posterior estimates of the item parameters between the parametric and semiparametric model both using the SI unconstrained centered sampling strategy.

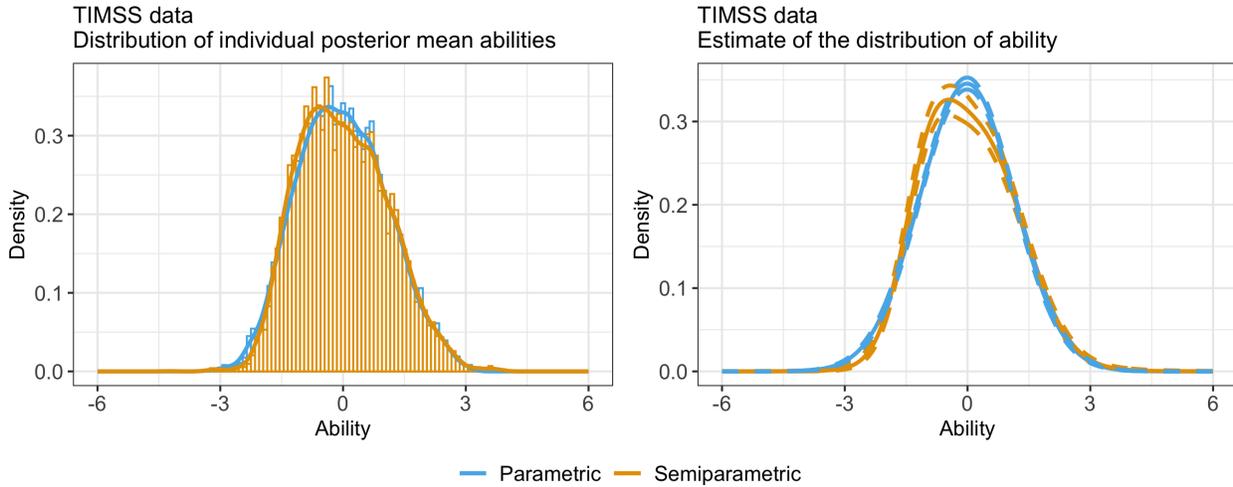


Figure 9: TIMSS Data, 2PL model. Histogram and density estimate of the posterior means of the latent abilities (left panel), and estimate of the posterior distribution for the latent abilities (right panel). Dashed lines indicate 95% credible intervals for the estimated distributions.

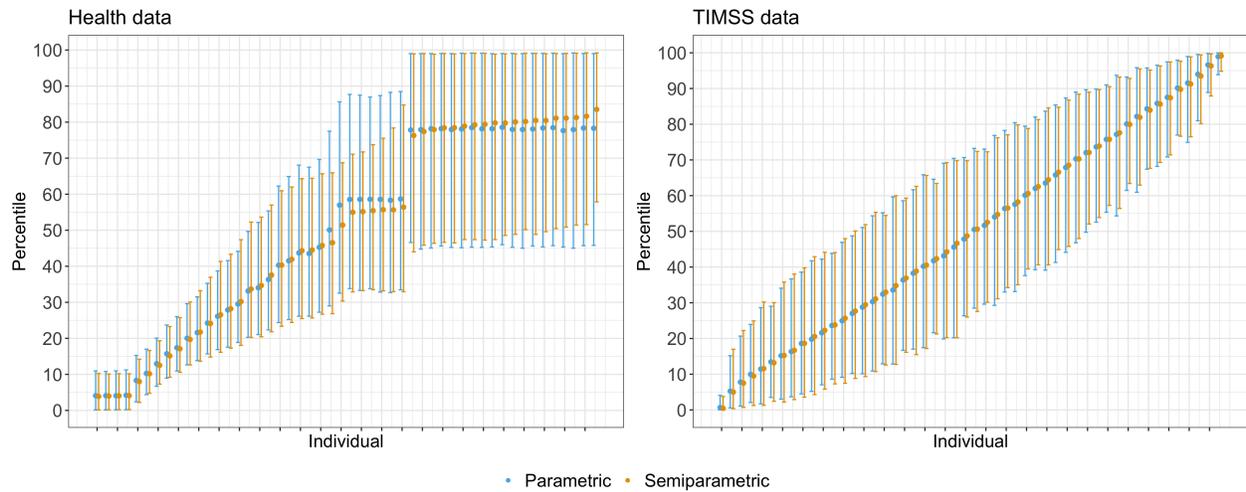


Figure 10: Estimates of individual percentiles (with 95% credible interval) for a subset of 50 individuals, for the health data (left panel) and TIMSS data (right panel) under the 2PL model.

References

- Escobar, M. D., & West, M. (1995). Bayesian density estimation and inference using mixtures. *Journal of the American Statistical Association*, *90*(430), 577–588.
- Neal, R. M. (2000). Markov chain sampling methods for Dirichlet Process mixture models. *Journal of Computational and Graphical Statistics*, *9*(2), 249–265.
- Shaby, B., & Wells, M. (2010). Exploring an adaptive Metropolis algorithm (Technical Report No. 1011-14). Duke University Department of Statistical Science.