

RESEARCH ARTICLE

Regularized Bayesian calibration and scoring of the WD-FAB IRT model improves predictive performance over marginal maximum likelihood

Joshua C. Chang¹, Julia Porcino¹, Elizabeth K. Rasch¹, Larry Tang^{1,2,3*}

1 Rehabilitation Medicine Department, NIH Clinical Center, Bethesda, Maryland, United States of America, **2** National Center for Forensic Science, University of Central Florida, Orlando, Florida, United States of America, **3** Department of Statistics and Data Science, University of Central Florida, Orlando, Florida, United States of America

* liansheng.tang@ucf.edu**OPEN ACCESS**

Citation: Chang JC, Porcino J, Rasch EK, Tang L (2022) Regularized Bayesian calibration and scoring of the WD-FAB IRT model improves predictive performance over marginal maximum likelihood. PLoS ONE 17(4): e0266350. <https://doi.org/10.1371/journal.pone.0266350>

Editor: Sheetal Kalyani, IIT Madras, INDIA

Received: March 15, 2021

Accepted: March 20, 2022

Published: April 8, 2022

Copyright: This is an open access article, free of all copyright, and may be freely reproduced, distributed, transmitted, modified, built upon, or otherwise used by anyone for any lawful purpose. The work is made available under the [Creative Commons CC0](https://creativecommons.org/licenses/by/4.0/) public domain dedication.

Data Availability Statement: The data used in this study were collected via an inter-agency agreement with the Social Security Administration. Data collected from claimants for disability benefits are owned by the SSA. We are legally prohibited from re-disclosing it, even if de-identified, without approval by the SSA. Researchers who are interested in accessing these data can seek permission for data use by reaching out to the NIH Principal Investigator (chanle@cc.nih.gov) who will seek approval from the appropriate SSA points of contact.

Abstract

Item response theory (IRT) is the statistical paradigm underlying a dominant family of generative probabilistic models for test responses, used to quantify traits in individuals relative to target populations. The graded response model (GRM) is a particular IRT model that is used for ordered polytomous test responses. Both the development and the application of the GRM and other IRT models require statistical decisions. For formulating these models (calibration), one needs to decide on methodologies for item selection, inference, and regularization. For applying these models (test scoring), one needs to make similar decisions, often prioritizing computational tractability and/or interpretability. In many applications, such as in the Work Disability Functional Assessment Battery (WD-FAB), tractability implies approximating an individual's score distribution using estimates of mean and variance, and obtaining that score conditional on only point estimates of the calibrated model. In this manuscript, we evaluate the calibration and scoring of models under this common use-case using Bayesian cross-validation. Applied to the WD-FAB responses collected for the National Institutes of Health, we assess the predictive power of implementations of the GRM based on their ability to yield, on validation sets of respondents, ability estimates that are most predictive of patterns of item responses. Our main finding indicates that regularized Bayesian calibration of the GRM outperforms the regularization-free empirical Bayesian procedure of marginal maximum likelihood. We also motivate the use of compactly supported priors in test scoring.

Introduction

Item response theory (IRT) encompasses a class of latent variable models for quantifying traits, such as abilities and attitudes, using questionnaires [1, 2]. Some of its highest-profile applications include the Graduate Record Exam (GRE) [3] and the Scholastic Aptitude Test (SAT)

Funding: This work is supported in part by US Social Security Administration and Intramural Research Program of the National Institutes of Health which participated in the study design, analysis, decision to publish and preparation of the manuscript. The authors received no specific funding for this work.

Competing interests: The authors have declared that no competing interests exist.

[4]. Besides its widespread use in high-stakes educational assessment, it is also heavily used in psychometrics [5, 6] and medical diagnostics [7].

Fundamentally, IRT models are generative non-linear factor analysis models [8]. These models yield predictions, in the form of probability mass functions, for how a particular person will respond to a particular test item. The key assumption in IRT is that the probability of one's response to any particular item on a test is a function composed of person and item-specific effects. As commonly implemented, the item-specific parameters inform each item's difficulty and discriminatory power. The person-specific parameter relates to an underlying ability relative to a target population. Altogether, these models assume that an individual can be characterized by a low-dimensional set of parameters known as traits or abilities, which are the embedded factors. Hence, these parameters constitute a representation of an individual's traits, or interchangeably, a person's responses.

Statistical choices in IRT model construction

IRT models are informed through a process known as item calibration. In calibration, one aims to train the item parameters in the IRT model using responses from a sample of the target population. In order to do so, however, one must either simultaneously infer or control for the person-specific parameters within the sample. Hence, calibration involves the parallel inference tasks of determining person and item-specific parameters.

Scoring is the process by which calibrated IRT models are applied to new responses, in order to obtain person-specific parameters applicable to the new respondents. In scoring, the new respondents are fit into the scale defined by the original calibration responses. Hence, the application of IRT models for use in prediction differs from usual statistical models in that another inference step is required after the initial training.

As high-dimensional models, the details of how one calibrates and scores IRT models are important. Unregularized maximum likelihood is known to be insufficient for calibration, due to parameter unidentifiability confounded by the nonlinear nature of the models. Similarly, scoring is sometimes done in settings where few test responses are known and maximum likelihood is unstable—this occurs regularly in computer adaptive testing for instance, where a test is scored in real-time, and this scoring informs the presentation of new items.

This particular drawback of the maximum likelihood method can be found in the Work Disability Functional Assessment Battery (WD-FAB) which is a relatively recent application of IRT to the assessment of work-related physical and behavioral function intended to inform processes related to disability determination and other potential applications [9].

Work Disability Functional Assessment Battery (WD-FAB)

We focus on the concrete application of improving the statistical properties of the IRT model underlying the Work Disability Functional Assessment Battery (WD-FAB) [9–13].

Background. The concept of work disability is evolving as represented by the World Health Organization's International Classification of Functioning, Disability and Health [14]. Modern models of work disability characterize the outcome of the interaction of a person's functional abilities within the work environment. Due to the complex nature of the interaction, a fundamental issue is how to identify and measure work disability within this contemporary framework [9].

The U.S. Social Security Administration (SSA) provides support to adults and children who qualify as disabled through the Supplemental Security Income (SSI) and Social Security Disability Insurance (SSDI) programs that provide health insurance and cash benefits to beneficiaries. As these are the two largest federal disability programs in the U.S. supporting millions

of Americans, accurate assessment of work disability is critical to applicants as an important safety net program and to the federal government to effectively allocate resources. The Social Security Administration (SSA) uses a statutory definition of work disability characterized as the inability to take part in “substantial gainful activity due to any medically determinable physical or mental impairment that can be expected to result in death or to last for a continuous period of not less than 12 months.” The Work-Disability Functional Assessment Battery (WD-FAB) was developed as an additional, comprehensive source of information about whole person function intended to support adjudicators when making disability determinations and/or re-determinations.

Prior development of the instrument. In this paper we evaluated the Work Disability Functional Assessment Battery (WD-FAB) which was developed by researchers at the Boston University Health and Disability Research Institute (BU) in conjunction with the National Institutes of Health (NIH). It is a computer-adaptive testing tool, backed by an IRT model, encompassing eight scales including four physical function scales and four mental function scales to identify self-reported function relative to work. The items within these scales consist of Likert-scaled multiple choice questions.

Work-related job function is a multifaceted concept, not easily summarized by any single quantitative factor. In the development of the WD-FAB, a combination of expert guidance and empirical evidence was used to inform the multidimensional nature of the instrument. An expert panel developed an item bank consisting of Likert scaled questions relating to physical and mental function. In a series of three surveys, a large-scale simple random sample of approximately 5000 SSA disability claimants in the United States was administered the entire item bank for use in calibrating the instrument. Separately, approximately 2000 individuals from the general population of both claimants and non-claimants were in placing the instrument in larger population context (in other words, a normative or reference sample).

The empirical portion of the creation of the WD-FAB, from individual responses, proceeded in two steps. The first step, exploratory factor analysis [15–18] (EFA) splits up an item bank into independent domains. The second step, confirmatory factor analysis [19] (CFA), verifies that a latent unidimensional trait is explanatory for the pattern of responses found for questions in each domain.

In EFA, step-wise item selection based on p-value cutoffs were used to factor the items into eight domains. Four of the domains pertain to physical function: Basic Mobility (BM), Fine motor function (FMF), Upper body function (UBF), and Community Mobility (CM). Four of the domains are used to evaluate psychological characteristics: Mood and emotions (ME), Resilience (RS), Self Regulation (SR), and Communication and Cognition (CC).

In CFA, heuristic (and ultimately arbitrary) values for fit statistics based on the PROMIS guidelines [2, 20] justified the instrument. In neither the EFA nor the CFA step were generalizability of the instrument considered. Additionally, both EFA and CFA, based on their own factorization models, do not guarantee consistency with nonlinear IRT models, which are themselves factor models. The specific IRT model adapted for the WD-FAB is the graded response model [21] (GRM). With the domains in place, marginal maximum likelihood (MML) was used to calibrate eight independent IRT models, and Warm’s weighted maximum likelihood [22] was adopted for test scoring.

Purpose of the present study

In this manuscript we improve on the predictive power of the graded response model underlying the WD-FAB by regularizing its calibration using fully-Bayesian methods. This approach is motivated by prior literature showing Bayesian calibration to be more accurate than MML

calibration for small datasets, even while using diffuse priors [23, 24], and the modern trend in applied Bayesian statistics towards utilizing stronger weakly informative priors for the purpose of statistical regularization [25–27].

As a metric for assessing predictive power, we adapt cross validation to estimate out-of-sample model accuracy in a manner that is consistent with how the WD-FAB instrument is used and interpreted. For predictive model assessment, cross validation and related methods are commonly used in machine learning and have found widespread adaption in the Bayesian statistical modeling world, but are not in common use for IRT model assessment.

Commonly-used alternatives to cross validation include information criteria such as the WAIC [28–30] and the AIC, which under certain conditions [31, 32] are asymptotically equivalent to leave-one-out cross validation. However, it is not straightforward to compare Bayesian versus non-Bayesian models using standard information criteria. Additionally, information criterion require two models to have identical prediction outputs for direct comparison and are therefore not flexible enough to handle common IRT model selection use cases like item selection. Finally, these criteria are not typically mindful of how models are interpreted.

To be mindful of common IRT-model use cases, where both model calibration and scoring require statistical decisions, and scores are interpreted alongside their inferred errors, we develop a custom variant of a cross validation metric based on the out-of-sample leave-K-out log marginal likelihood. We show that Bayesian IRT model calibration, coupled with regularized Bayesian scoring, outperforms the commonly-used procedures of marginal maximum likelihood (MML) calibration and weighted likelihood estimation (WLE) scoring.

Materials and methods

Although the methods in this paper generalize broadly to other item response theory models, we formulate our methods based on the unidimensional graded response model (GRM) for polytomous item responses. The GRM [21] applies to assessments where there is an intrinsic ordering in responses, as in Likert scales. According to the GRM, the probability of a response of j to item i for person p obeys the likelihood function

$$\begin{aligned} \Pr(X_{pi} = j | \theta_p, \boldsymbol{\tau}, \boldsymbol{\lambda}) &= \Pr(X_{pi} \geq j | \theta_p, \boldsymbol{\tau}, \boldsymbol{\lambda}) - \Pr(X_{pi} \geq j + 1 | \theta_p, \boldsymbol{\tau}, \boldsymbol{\lambda}) \\ &= \frac{1}{1 + \exp(\lambda_i(\tau_{ij} - \theta_p))} - \frac{1}{1 + \exp(\lambda_i(\tau_{i,j+1} - \theta_p))}, \end{aligned} \quad (1)$$

where λ_i are the item-specific discrimination parameters and $\tau_{ij} < \tau_{i,j+1}$ are the item-specific threshold parameters, and θ_p are person-specific ability parameters. The schematic of the GRM is presented in Fig 1, where for each of P people, predictions of their I item responses are determined by these parameters.

Eq 1, demonstrates that in the GRM, the person-specific parameter θ_p only has interpretability relative to the item-specific parameters. Mathematically speaking, the likelihood function in Eq 1 is unidentifiable, in both location and scale. To improve identifiability, one typically either imposes or encourages a scale and location on the person-specific θ_p , on the population level, so that a target population has θ_p follows roughly a Gaussian distribution of a given scale (typically unit).

The unidimensional formulation of Eq 1 can be extended to multidimensional traits (representing each individual using a vector of values rather than a scalar), by fitting models for each domain (dimension) separately, effectively partitioning tests into domains of items (as per the PROMIS guidelines [33]). Such item partitioning is typically completed with the aid of linear factor analysis methods [34] though factorization directly within the nonlinear IRT model is possible [8].

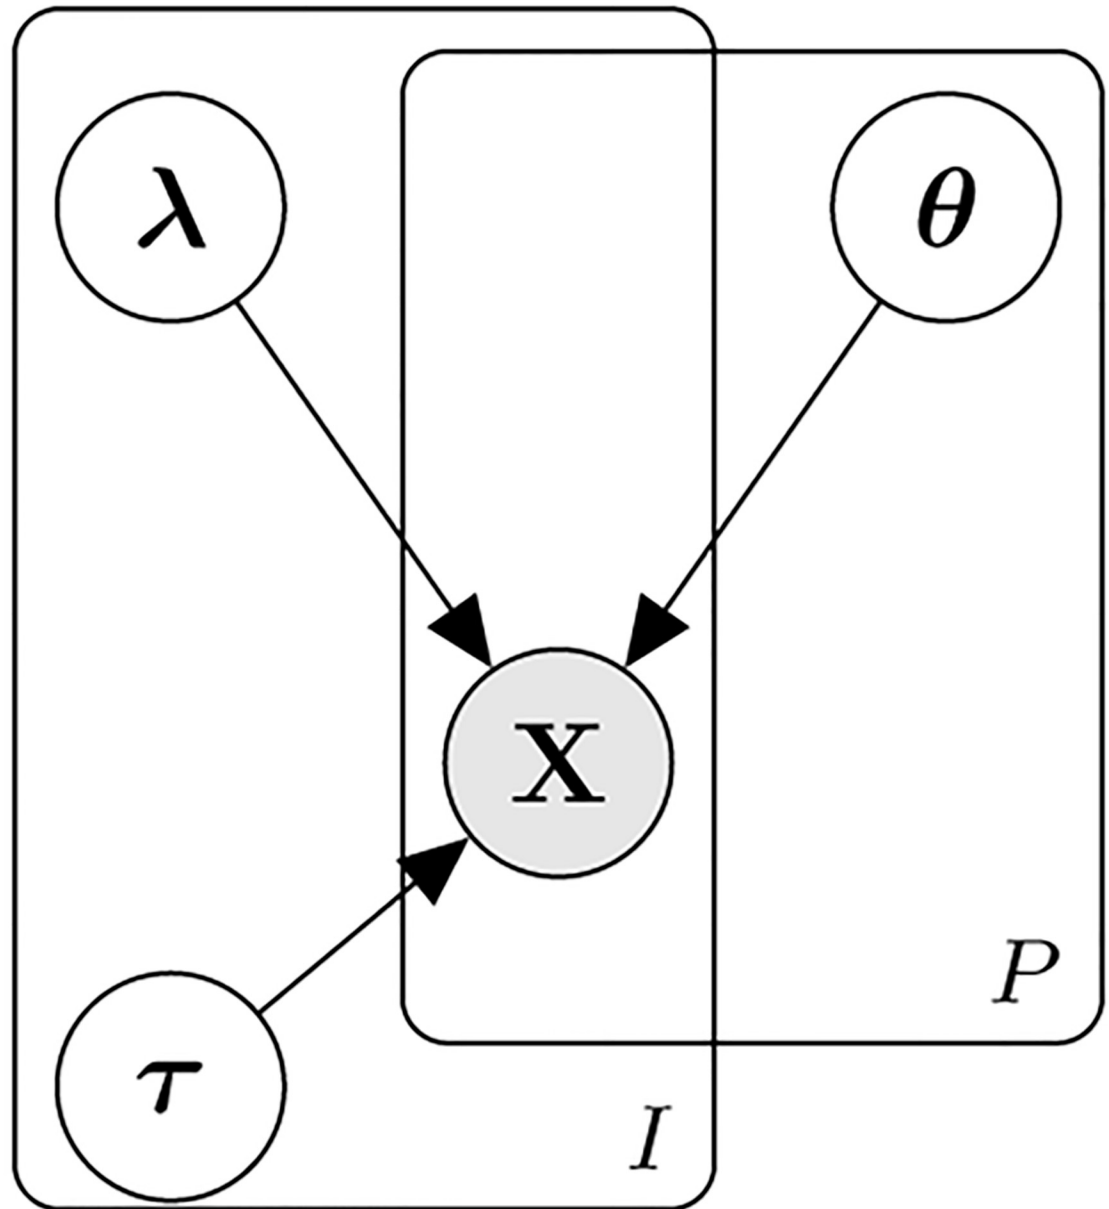


Fig 1. Graded response model. The Graded response model (GRM) with item parameters λ , τ , and person-specific parameters θ .

<https://doi.org/10.1371/journal.pone.0266350.g001>

In this manuscript, we assume that one has designed a test constituted of ordinal response items, intended to measure a latent construct in a population. We assume that the items in the test have already been pre-partitioned, through domain-specific knowledge, empirical methods, or a combination of both. Note however that the methods in this manuscript may be used to evaluate item partitioning schemes, in case multiple possibilities are being considered. Additionally, we assume that one has already administered the test to a representative sample of individuals within the target population, collecting responses that we will refer to as the calibration data. Fundamentally, we assume that one intends to use the data to create an instrument that will generalize to the wider target population. At this stage, one has the requisite materials needed for model calibration.

Model calibration

The model in Eq 1 can be learned in many ways. Non-Bayesian approaches typically involve maximum likelihood estimation (MLE) [35, 36] or the empirical Bayesian procedure of maximum marginal likelihood (MML) [37], where given a set of responses $\{x_{pi}\}_{p,i}$ the marginal likelihood

$$\prod_p \prod_i \Pr(X_{pi} = x_{pi} | \tau, \lambda) = \prod_p \prod_i \int \Pr(X_{pi} = x_{pi} | \theta_p, \tau, \lambda) d\mathcal{N}_{\theta_p}(\hat{\theta}_p, \hat{\sigma}_p) \tag{2}$$

is optimized recursively along with estimates of ability, typically using expectation maximization (EM). In Eq 2, the distribution of ability for a person p is approximated using a Gaussian distribution centered at $\hat{\theta}_p$ of variance $\hat{\sigma}_p^2$. Hence, model calibration yields estimates for both item-specific parameters and person-specific parameters.

Bayesian calibration. Fundamentally, the GRM is a high-dimensional nonlinear latent variable model informed using discrete observations. Hence, its fit must be constrained in order to ensure parametric identifiability. The standard methods for constraining its parameters typically lie in scaling of the ability parameters across the calibration sample so that they are unit scale. Beyond this imposition of scaling, most non-Bayesian methods do not use other regularization.

Modern high-dimensional statistical problems have necessitated the development of regularization techniques. In the Frequentist world, these regularization techniques usually involve penalty functions placed on the objects to be inferred, or hierarchical structure built into the problem as in the case of mixed effects linear regression. Using well-designed regularization, one obtains models that perform better at making predictions than those without regularization. This fact has to do with shrinkage and partial pooling properties of regularized models. Shrinkage is a statistical property where estimates of effect sizes (parameters) are shrunk towards zero. Partial pooling is a form of shrinkage where collectively the differences between parameters are shrunk so that overall group-level means are obeyed. These properties lead to more robust calibration of models in their response to noisy data. At that point, the lines between Frequentist and Bayesian methods are blurred as regularization can be interpreted as prior information.

In Bayesian modeling of IRT problems, prior distributions are placed on all model parameters [35]. Even when using diffuse priors, these distributions help regularize the overall inference problem [23, 24]. Bayesian modeling allows wide latitude in how one wishes to specify the IRT model. For the purposes of this manuscript, we consider modeling under the principle of using weakly informative prior distributions [27], for the purpose of parameter regularization. To put this principle in concrete terms, we consider the overall probabilistic model for generating X_{pid} , person p 's response to item i in domain d , on a test where each item has J possible responses,

$$\begin{aligned} X_{pid} &\sim \text{categorical}(p_{pi,1}^{(d)}, \dots, p_{pi,J}^{(d)}) \\ p_{pi,j}^{(d)} &= \Pr(X_{pid} = j | \{\theta_p^{(d)}\}_d, \{\{\tau_{ij}^{(d)}\}_j\}_d, \{\lambda_i^{(d)}\}_d) \quad (\text{see Eq. 1}) \\ \lambda_i^{(d)} &\sim \text{cauchy}^+(0, 5) \\ \tau_{i,j}^{(d)} &\sim \text{normal}^+(\mu_\tau^{(d)}, \sigma_\tau^{(d)}) \quad \text{subject to } \tau_{i,j}^{(d)} > \tau_{i,j-1}^{(d)} \\ \mu_\tau^{(d)} &\sim \text{normal}(0, 5) \\ \sigma_\tau^{(d)} &\sim \text{cauchy}^+(0, 1) \\ \theta_p^{(d)} &\sim \text{normal}(0, \sigma^{(d)}) \\ \sigma^{(d)} &\sim \text{cauchy}^+(0, 1), \end{aligned} \tag{3}$$

where normal^+ , cauchy^+ refer to the non-negative half-normal and half-cauchy distributions respectively. The model presented in Eq 3, features weakly-informative priors that are similar to those used in the prior literature [30, 38]. The purpose of weakly informative priors is to regularize the statistical and computational problem of inferring the parameters of the model. Weakly informative priors have been shown in many contexts to reduce type-I and type-M errors in statistical estimation problems [25, 27], including in IRT problems [26].

Bayesian inference involves computing the statistics of a hierarchical model. For IRT models, this inference is not analytically tractable and the computations are done typically through approximate methods like Markov Chain Monte Carlo (MCMC) or Automatic Differentiation Variational Inference (ADVI) [39, 40].

The end result of calibration are posterior distributions over the model parameters in the Bayesian setting, or point-estimates in the non-Bayesian setting. Point estimates can also be obtained from the Bayesian model through consideration of an appropriate loss function. For example, the posterior mean minimizes L^2 loss whereas the posterior median minimizes L^1 loss. In many applications, such as the one motivating this manuscript, point-estimate summaries of the model parameters are necessary.

Scoring

While calibration is performed off line, the model is usually intended for online use in the aim of determining the ability scores of new applicants. This process is known as scoring and pertains to estimation of the model's latent factors (person-specific ability parameters) conditional on learned posteriors for the item parameters.

In computer adaptive testing, the calibrated IRT model also guides the presentation of items. Classical methods for item selection have used the local Fisher information matrix, conditional on estimated score [41]. More-contemporary approaches also take score uncertainty into account [42–44]. Regardless of item selection approach, a method for scoring is necessary.

In scoring, full Bayesian treatment of the calibrated model parameters is often infeasible and optimization of the likelihood conditional on point-estimates of the item parameters is a common procedure. As opposed to calibration, in this step the item-specific parameters are assumed known and fixed. It is also computationally expensive to propagate posterior distributions of the item parameters so their uncertainty is ignored. Using their point estimates, the likelihood is maximized relative to the ability estimate of a new person, given his or her pattern of item responses [34].

To make meaningful comparisons between the ability of people, one must quantify the precision of the ability estimate. Given the fitted scores, it is common to perform an asymptotic approximation of the standard error using the Fisher information matrix $I(\hat{\theta}_p)$. This approximation is the Cramer-Rao lower bound,

$$\text{Var}(\hat{\theta}_p) \geq \frac{1}{I(\hat{\theta}_p)}, \quad (4)$$

an asymptotic bound for the variance for unbiased estimators based on applying Laplace's method on the posterior density. It is handy to interpret the ability estimate as if it were Gaussian, using the implied variance estimate of Eq 4, and the mapping

$$(\hat{\theta}_p, \hat{s}_p) \mapsto \mathcal{N}_{\theta_p}(\hat{\theta}_p, \hat{s}_p), \quad (5)$$

where $\hat{s}_p = 1/\sqrt{I(\hat{\theta}_p)}$.

Weighted Likelihood Estimation (WLE) scoring. The commonly-used weighted likelihood estimator [22] (WLE) removes the leading-order asymptotic bias of the maximum likelihood estimator. The asymptotic bias of this particular estimator is $\mathcal{O}(n^{-1})$. This estimator is often used in conjunction with the variance estimate of Eq 4.

Marginal Maximum Likelihood (MML) scoring. We also consider an estimator found by maximizing the marginal maximum likelihood of Eq 2 with respect to the score directly, hereby referred to as the MML scoring estimator. For this estimator, we optimize Eq 2 with respect to the score ($\hat{\theta}_p$) while using Eq 4 to impose the score variance. This estimator resembles a variational Bayesian estimator under interchange of expectations and logarithms within the objective.

Expected A Posteriori (EAP) scoring with a compactly-supported prior. Finally, we consider expected a-posteriori (EAP) estimation, using a fully Bayesian procedure where we regularize score inference using a prior distribution. In particular, we use an explicitly-truncated normal distribution to restrict score estimation within a compact interval surrounding zero. We believe this restriction to be well-motivated when one realizes that in scoring an individual, the model is interpolating that individual into the pre-calibrated model, finding a placement for that individual relative to the people in the original calibration sample.

In calibration, scores for a representative sample of a population are inferred. These scores follow some distribution, however, the usual assumption is that the scores are approximately normal. Since the scale of the distribution is arbitrary, let us assume without generality that the population follows a unit normal distribution on a given trait.

The tails for a normal distribution fall off rapidly. One should expect to observe someone with scores more extreme than four standard deviations once out of approximately sixteen thousand times. That ratio becomes one in 1.7 million outside of five standard deviations. Hence, for an IRT instrument calibrated using 1.7 million respondents, one would not expect to see scores more extreme than ± 5 , when placed on unit scale.

Predictive model evaluation

This manuscript evaluates model calibration and scoring methods for the GRM based on the predictive power of each model's corresponding Gaussian approximations of ability (Eq 6). In generality, one measures the predictive power of a model by approximating an appropriate risk function as computed by the model on new data. The typical ways of doing this task are cross-validation, and approximate cross-validation through the computation of information criterion.

Cross-validation involves the separation of datasets into training and testing sets, where the testing set is left out and the model is fit using the training set. The testing set is then used to test the model for predictive accuracy. Information criteria in the cases of the Akaike information criterion (AIC) and Watanabe-Akaike information criterion (WAIC) [28] are under certain conditions [32] asymptotic approximations of forms of cross validation. Regardless, the objective of each of these approaches is to approximate the log likelihood of the model for future data that is not available at training.

A limitation of each of the information criterion is that they can only be used when making comparisons between models based on the same data. For this reason, they cannot be used for looking at inclusion or exclusion of items since two models with different items use different data. Furthermore, the AIC and WAIC are from different statistical paradigms. The WAIC [29, 31, 45] is a Bayesian variant of the AIC, scaled to model deviance like the AIC. However, it operates under the assumption that one uses the full posterior of a Bayesian model in making predictions. In the scoring step at test administration, computational trade-offs must be made.

While calibration is performed off line, the model is intended for on-line use in determining the ability scores of new applicants. In this stage of computation, Bayes is often infeasible, and optimization of the likelihood is a common procedure. As opposed to calibration, in this step the item-specific parameters are assumed known and fixed. It is also computationally expensive to propagate posterior distributions of the item parameters, so their uncertainty is ignored.

Let $\Omega = \cup_{k=1}^K \Omega^{(k)}$, where $\Omega^{(j)} \cap \Omega^{(k)} = \{\}$ for $j \neq k$, represent a partition of the P people that responded to items for calibration. Leaving out one of the partitions $\Omega^{(k)}$ at a time, one calibrates (fits) K sets of model parameters. The outputs of these calibrations are discrimination parameters $\lambda^{|\Omega \setminus \Omega^{(k)}}$ and item threshold parameters $\tau^{|\Omega \setminus \Omega^{(k)}}$. Each calibration, applied to its corresponding left-out data, yields a set of ability estimates and estimates for the standard deviation of these ability estimates $\{(\hat{\theta}_p^{|\Omega \setminus \Omega^{(k)}}, \hat{s}_p^{|\Omega \setminus \Omega^{(k)}})\}_{p \in \Omega^{(k)}}$. These estimates self-consistently model the likelihood of the item responses of the left-out people by providing Gaussian approximations for their abilities θ_p through the mapping

$$(\hat{\theta}_p^{|\Omega \setminus \Omega^{(k)}}, \hat{s}_p^{|\Omega \setminus \Omega^{(k)}}) \mapsto \text{normal}_{\theta_p}(\hat{\theta}_p^{|\Omega \setminus \Omega^{(k)}}, \hat{s}_p^{|\Omega \setminus \Omega^{(k)}}), \tag{6}$$

where $\mathcal{N}_{\theta}(\mu, \sigma^2)$ is a Gaussian measure with mean μ and variance σ^2 . This mapping is crucial since it allows one to evaluate comparisons by evaluating quantities such as $\Pr(\theta_p - \theta_q > 0)$, while respecting uncertainty in ability estimates.

We wish to evaluate a model based on its predictive ability to forecast item response patterns in a way that is self consistent with such comparisons. To do so, we use the inferred approximations over the ability distributions and the GRM likelihood to formulate a prediction risk for the given model.

For risk, we consider an approximation of the information loss for a given model M , which is expressed by the deviance-scaled measure

$$\begin{aligned} D[M] &= -2 \sum_k \sum_{p \in \Omega^{(k)}} \log \prod_{i=1}^J \Pr(X_{pi} = x_{pi} | \tau^{|\Omega \setminus \Omega^{(k)}}, \lambda^{|\Omega \setminus \Omega^{(k)}}) \\ &\approx -2 \sum_k \sum_{p \in \Omega^{(k)}} \sum_{i=1}^J \log \left[J_{\hat{\theta}_p^{|\Omega \setminus \Omega^{(k)}}, \hat{s}_p^{|\Omega \setminus \Omega^{(k)}}}(\lambda_i^{|\Omega \setminus \Omega^{(k)}}, \tau_{i, x_{pi}}^{|\Omega \setminus \Omega^{(k)}}) \right. \\ &\quad \left. - J_{\hat{\theta}_p^{|\Omega \setminus \Omega^{(k)}}, \hat{s}_p^{|\Omega \setminus \Omega^{(k)}}}(\lambda_i^{|\Omega \setminus \Omega^{(k)}}, \tau_{i, x_{pi}+1}^{|\Omega \setminus \Omega^{(k)}}) \right]. \end{aligned} \tag{7}$$

The approximation

$$J_{\theta, \sigma}(\lambda, \tau) = \int \frac{d\mathcal{N}_{\phi}(\theta, \sigma)}{1 + \exp(-\lambda(\phi - \tau))} \approx \Phi \left(\frac{\pi\lambda(\theta - \tau)/8}{\sqrt{1 + (\pi\lambda\sigma/8)^2}} \right), \tag{8}$$

where \mathcal{N} is a Gaussian measure and Φ is the cumulative density function for the unit normal distribution, is well-known [46].

We use the criterion in Eq 7 to compare regularized Bayesian calibration of the GRM using the model of Eq 3 to calibration performed using marginal maximum likelihood (MML), coupled with the following scoring methods: expected a-posteriori (EAP), weighted maximum likelihood (WLE) [22], and marginal maximum likelihood (MML). The criterion being deviance-scaled has the same interpretation as the AIC, consisting of an estimate of the out-of-sample model likelihood.

Results

We performed all analyses in R 3.5, with empirical-Bayesian and non-Bayesian analyses performed using the R package `mirt`. For Bayesian analyses, we used the Stan probabilistic programming language [47], interfaced in R using the package `rstan`. We implemented the custom EAP scoring method in R, as well as all of the computations behind the cross-validation metric that we use to compare models. This section mainly shows the results of the comparisons performed on both cross-validation folds and on a control sample. The detailed discussion of the results is provided in the Discussion section.

Comparison of calibration methods on claimant sample

Our main objective is to compare the predictive performance of regularized Bayesian calibration to unregularized MML calibration across the different scoring methods (EAP, WLE, MML). To this end, we used response data collected from the target subpopulation of claimants and calibrated the WD-FAB using the Bayesian model of Eq 3 and using MML. To be specific, we used four-fold cross validation on each domain (BM, CC, CM, FMF, ME, RS, SR, UBF), leaving out one fold at a time and fitting the model on the remaining responses. From each model, we computed the cross-validation criteria for each left-out fold. The only exception was RS, where we instead used three-fold cross validation.

First, we computed the metric of Eq 7 for all pairings of calibration and scoring methods. In Fig 2, deviance values of Eq 7 are presented for each of the left-out groups. The criterion of Eq 7 takes uncertainty of the estimated scores into account. To look at the predictive accuracy of the point-estimate for the score, ignoring uncertainty, we considered the same deviance measure in Eq 7, with the variance taken to be zero, regardless of scoring method. The results of the point-wise criterion are shown in Fig 3.

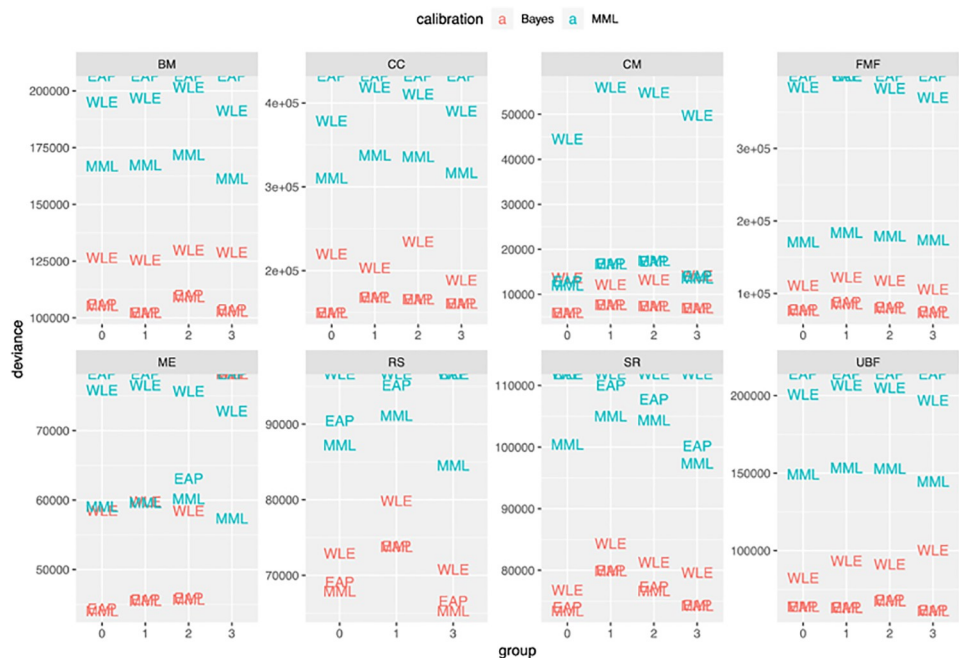


Fig 2. Four fold cross-validation comparison of calibration of and scoring choices for 8 domains of the WD-FAB (lower is better), performed on claimant sample. The deviance measure is computed for each left out fold using Eq 7. This measure is scaled in-line with the AIC so lower values are better.

<https://doi.org/10.1371/journal.pone.0266350.g002>

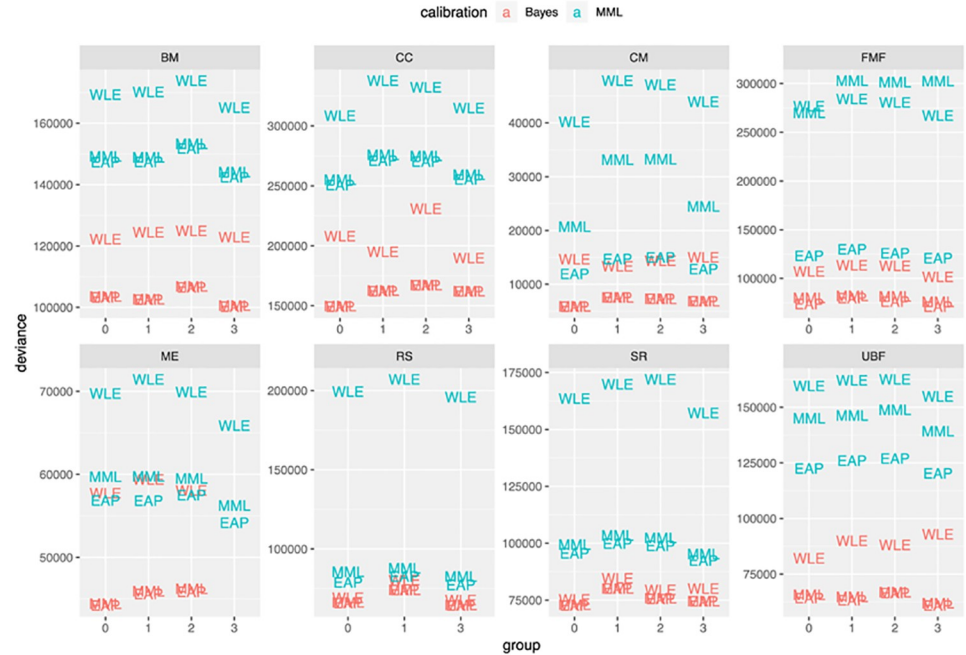


Fig 3. Point-estimated out-of-sample comparison of calibration and scoring choices on claimant responses for 8 domains of the WD-FAB (lower is better). The point-estimated deviance measure is computed by taking the variance parameters in Eq 7 to zero. This measure is scaled in-line with the AIC so lower values are better.

<https://doi.org/10.1371/journal.pone.0266350.g003>

Comparisons on an out-of-sample control population

For the WD-FAB, the calibration group was taken from a sample of disability claimants. To put work-related function for these people in context, a separate control sample of adults was also taken. This sample was meant to be representative of the population at-large. We evaluated the different model calibration and scoring methods on the data, using the same measure as in Fig 2, except without leaving out folds during calibration. Hence, we evaluated how well each model, trained solely on claimant data, predicts responses to the test items for the control population. The results of this evaluation are shown in Fig 4.

Comparing scoring methods

Finally, we evaluated each of the WLE, EAP, and MML scoring methods for consistency with each other and with scores inferred during model calibration. Recall that calibration also entails inference of abilities—the item parameters are found self-consistently with these. Fig 5 presents pairwise comparisons of scores obtained using the full Bayesian model, where the label “Calibration Bayes” corresponds to posterior means of ability estimates inferred at calibration. Shown are scores for the domain BM. We performed the same analysis on models calibrated using MML. These results are shown in Fig 6. Likewise, in this figure, “Calibration MML” corresponds to ability estimates for the MML model obtained at calibration.

Discussion

In this manuscript we compared full Bayesian inference of IRT models against marginal maximum likelihood (MML) based empirical Bayesian inference [48, 49]. Coupled with these these



Fig 4. Out-of-sample out-of-subpopulation comparison of calibration and scoring choices for the 8 domains of the WD-FAB (lower is better). The deviance measure is given in Eq 7. These are computed on a control group that was neither used nor intended for use in calibrating the WD-FAB.

<https://doi.org/10.1371/journal.pone.0266350.g004>

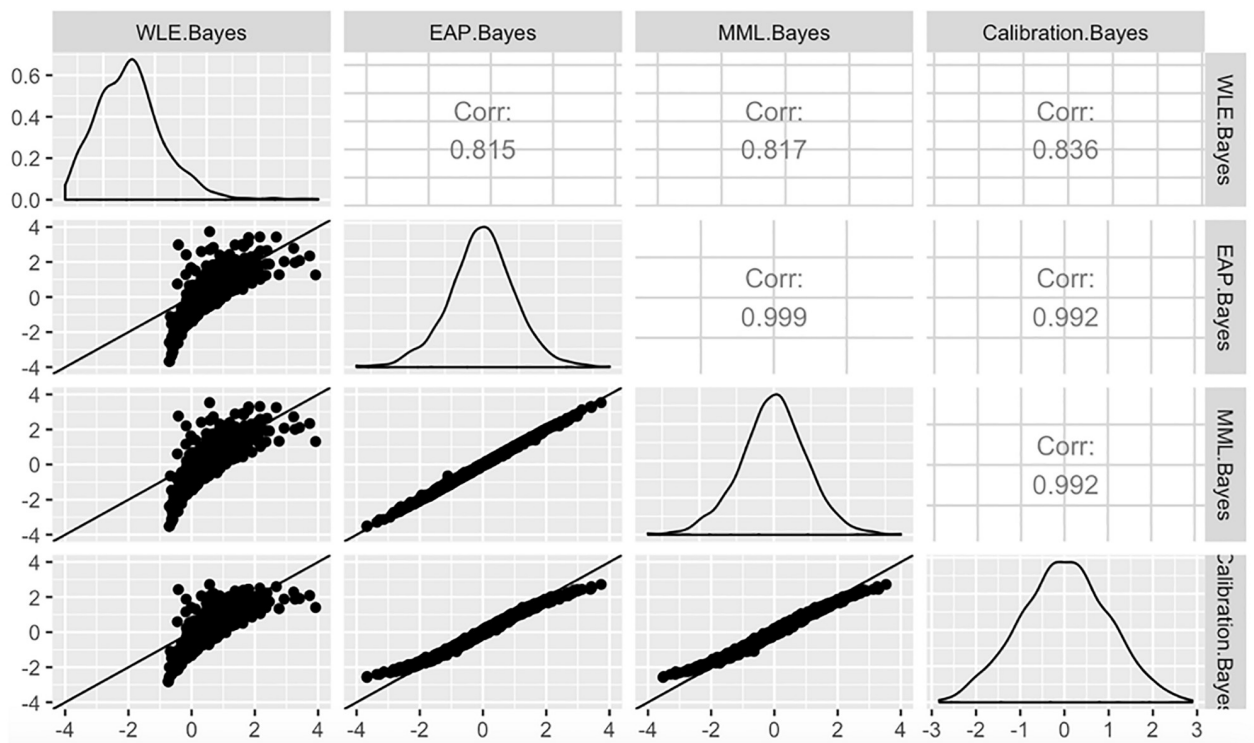


Fig 5. Pairwise comparison of scoring of the Basic Mobility (BM) scale calibrated using the full-Bayesian model. The methods compared are Warm’s weighted likelihood estimator (WLE), EAP, and MML. Calibration Bayes refers to the score inferred during item calibration.

<https://doi.org/10.1371/journal.pone.0266350.g005>

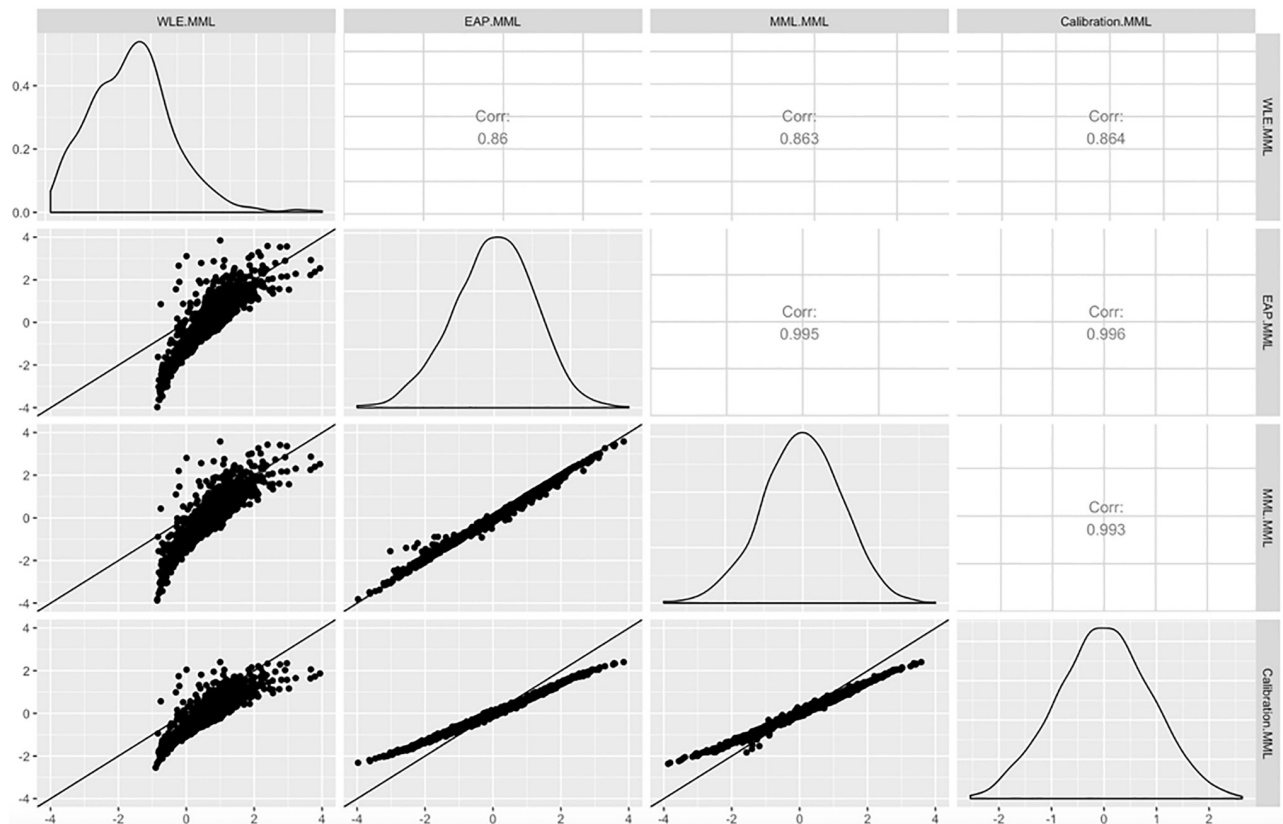


Fig 6. Pairwise comparison of scoring of the Basic Mobility (BM) scale calibrated using marginal maximum likelihood (MML). The methods compared are Warm's weighted likelihood estimator (WLE), EAP, and MML. Calibration MML refers to the score inferred during item calibration.

<https://doi.org/10.1371/journal.pone.0266350.g006>

choices for item calibration methodology, we also looked at compactly supported expectation a-posteriori (EAP) scoring compared to weighted likelihood and MML scoring.

Our evaluation metric, targeted at predictive accuracy, is rooted in the concrete and real life application of the assessment of work-related physical and mental function. In particular, the metric is consistent with how scores of such an instrument are typically interpreted—where the error in the ability estimate is interpreted as if it were the standard deviation of a Gaussian distribution. Hence, we defined the metric of Eq 7 to be consistent with such an interpretation. In defining the metric, we use the approximation of Eq 8. We note that using probit rather than logistic functions to model the item response functions would render such an approximation unnecessary. Using this metric we found that choices of calibration and scoring methodologies clearly matter.

Full Bayesian calibration consistently outperforms MML

In all of our analyses, models calibrated using full Bayesian inference outperformed MML, across all cross-validation splits and all item domains. This trend is visible in Fig 2, where the Bayesian models have lower estimated model deviance than the MML models, regardless of scoring methodology. In some applications, one does not care to interpret uncertainty in scoring. In these situations, one may refer to the analyses of Fig 3, where the results are consistent with those of Fig 2. The metric used here more-closely resembles Frequentist cross-validation measures like the AIC that do not account for parameter uncertainty.

Additionally, the Bayesian models transfer better than the MML models. When evaluating the resulting models on a true out-of-sample set of responses given by the control sample, the same trends held (Fig 4). While the WD-FAB instrument is meant to be calibrated relative to a claimant subpopulation, it is not known a-priori whether a given applicant should be a member of this population. Hence, it is important for the instrument to generalize and provide meaningful results for non-claimants as represented by the control sample.

The superiority of full Bayesian calibration was consistent across all scoring methods, though scoring methods also differed in terms of performance. The EAP and MML scoring methods exhibited similar performance in all of the cross-validation experiments of Figs 2 and 3. WLE, on the other hand, performed consistently poorly when paired with either MML calibration or full Bayesian calibration. For many of the domains, pairing the full Bayesian model with WLE scoring was sufficient to remove the performance advantage of Bayesian calibration over MML calibration.

The WLE approach gives scores inconsistent with item calibration

Focusing on the poor performance of WLE scoring, relative to the other methods, we compared the scores produced using WLE after calibration versus those of other methods. Figs 5 and 6 provide these comparisons using the full-Bayesian and MML calibrated IRT models respectively.

The item parameters produced in the calibration of these models are self-consistent with ability estimates produced during calibration. Hence, it is salient to compare the score obtained with each scoring method against ability estimates produced at calibration.

The WLE scores obtained, conditional on either model, are only weakly-correlated with the scores at calibration. The WLE is a correction to the first order term in an asymptotic approximation of the signed score bias [22]. Our results show that such a correction does more harm than good. Furthermore, we question the motivations behind attempting to correct this measure of bias, while ignoring other objectives such as absolute bias and estimator variance.

EAP regularization for scoring marginally improves predictive performance

In both Figs 5 and 6, both EAP and MML-based scoring had high correlations (>0.99) with scores computed at calibration. The MML-based scoring method does well—since the objective integrates over an estimate of the score uncertainty, which itself induces shrinkage in the ability estimates. Hence, it is more-regularized than unregularized maximum likelihood or the WLE.

Comparing EAP and MML directly, the EAP method has explicit regularization imposed by the truncated Gaussian prior. Hence, while the MML estimator performs some shrinkage, scores computed with EAP tend to be shrunk relative to MML, particularly at the tails. Imposing compact support on the scoring process guards against facetious extrapolation of the model beyond the range consistent with calibration. By definition, few calibration subjects fall into the extremes. Hence, IRT models are, by construction, less certain in estimating tail behavior within populations.

The shrinkage is most pronounced when performing the comparisons on the MML-calibrated IRT model. However, looking at the out-of-sample results of Fig 4, we see that EAP generally performs better than MML in terms of predictive performance as EAP tends to have smaller deviances. This behavior is expected because the general population has different characteristics than the subpopulation used at calibration. The regularization in the EAP method helps guard against instabilities induced by these differences. In Fig 5, using the MML-

calibrated IRT model, the effects of regularization are clear. EAP shrinks scores but better-preserved relative ordering of scores. In fact, the EAP scores have stronger correlation than the MML scores are with the scores obtained while calibrating the item parameters using MML.

In this manuscript we compared pairs of calibration and scoring methodologies as applied to assessing predictive ability of the WD-FAB. In this application, scores and their uncertainty are used to compare respondents. In-line with the interpretation of the IRT model, we developed the deviance metric of Eq 7. We found full-Bayesian item response calibration, coupled with regularized EAP scoring to provide for more-predictive self-consistent model interpretations.

Limitations and extensions

In this manuscript our goal was to evaluate calibration and scoring of the WD-FAB. For this reason, we did not formulate the cross validation criterion with item selection in mind—or other use cases where two candidate models would have a different set of items. A straightforward method to extend the enclosed methodology to these use cases would be to restrict the sums of Eq 7 to shared items. Future work will focus on looking at item selection through a predictive lens.

For process reasons, we used as a starting point of this work the previously-developed WD-FAB instrument. The development of WD-FAB followed the same procedure as the Patient-Reported Outcomes Measurement Information Systems (PROMIS) program. And all study design and sampling was contingent on the use of this program.

Each unidimensional factor in WD-FAB was developed from the exploratory and confirmatory factor analysis. It would be interesting to investigate how WD-FAB can be developed using truly multidimensional IRT models [8] which consider the correlation among multiple factors.

Author Contributions

Conceptualization: Joshua C. Chang.

Formal analysis: Joshua C. Chang, Larry Tang.

Funding acquisition: Elizabeth K. Rasch.

Methodology: Joshua C. Chang, Elizabeth K. Rasch, Larry Tang.

Resources: Elizabeth K. Rasch.

Supervision: Elizabeth K. Rasch.

Writing – original draft: Joshua C. Chang, Julia Porcino, Elizabeth K. Rasch, Larry Tang.

Writing – review & editing: Joshua C. Chang, Julia Porcino, Elizabeth K. Rasch, Larry Tang.

References

1. Chang CH, Reeve BB. Item Response Theory and Its Applications to Patient-Reported Outcomes Measurement; 28(3):264–282.
2. Fries JF, Witter J, Rose M, Cella D, Khanna D, Morgan-DeWitt E. Item Response Theory, Computerized Adaptive Testing, and PROMIS: Assessment of Physical Function; 41(1):153–158.
3. Kingston NM, Dorans NJ. The Feasibility of Using Item Response Theory as a Psychometric Model for the Gre Aptitude Test; 1982(1):i–148.
4. Carlson JE, von Davier M. Item Response Theory; 2013(2):i–69.
5. Altemeyer B. Enemies of Freedom: Understanding Right-Wing Authoritarianism. Enemies of Freedom: Understanding Right-Wing Authoritarianism. Jossey-Bass;

6. Goldberg LR. The Development of Markers for the Big-Five Factor Structure; 4(1):26–42.
7. Sullivan JT, Sykora K, Schneiderman J, Naranjo CA, Sellers EM. Assessment of Alcohol Withdrawal: The Revised Clinical Institute Withdrawal Assessment for Alcohol Scale (CIWA-Ar); 84(11):1353–1357.
8. Chang JC, Vattikuti S, Chow CC. Probabilistically-Autoencoded Horseshoe-Disentangled Multidomain Item-Response Theory Models;. Available from: <http://arxiv.org/abs/1912.02351>.
9. Marino ME, Meterko M, Marfeo EE, McDonough CM, Jette AM, Ni P, et al. Work-Related Measures of Physical and Behavioral Health Function: Test-Retest Reliability; 8(4):652–657.
10. Jette AM, Ni P, Rasch E, Marfeo E, McDonough C, Brandt D, et al. The Work Disability Functional Assessment Battery (WD-FAB); 30(3):561–572.
11. Marfeo EE, Ni P, Haley SM, Jette AM, Bogusz K, Meterko M, et al. Development of an Instrument to Measure Behavioral Health Function for Work Disability: Item Pool Construction and Factor Analysis; 94(9):1670–1678.
12. Marfeo E, Ni P, Meterko M, Marino M, Peterik K, McDonough C, et al. Development of a New Instrument to Assess Work-Related Function: Work Disability Functional Assessment Battery (WD-FAB); 70:7011500012p1–7011500012p1.
13. Marfeo EE, Ni P, McDonough C, Peterik K, Marino M, Meterko M, et al. Improving Assessment of Work Related Mental Health Function Using the Work Disability Functional Assessment Battery (WD-FAB); 28(1):190–199.
14. Organization WH. International Classification of Functioning, Disability, and Health. World Health Organization; 2001.
15. Ferguson E, Cox T. Exploratory Factor Analysis: A Users' Guide; 1(2):84–94.
16. Goretzko D, Pham TTH, Bühner M. Exploratory Factor Analysis: Current Use, Methodological Developments and Recommendations for Good Practice;.
17. Williams B, Onsman A, Brown T. Exploratory Factor Analysis: A Five-Step Guide for Novices; 8(3).
18. Howard MC. A Review of Exploratory Factor Analysis Decisions and Overview of Current Practices: What We Are Doing and How Can We Improve?; 32(1):51–62.
19. van Prooijen JW, van der Kloot WA. Confirmatory Analysis of Exploratively Obtained Factor Structures; 61(5):777–792.
20. DeWalt DA, Rothrock N, Yount S, Stone AA. Evaluation of Item Candidates: The PROMIS Qualitative Item Review; 45:S12–S21.
21. Samejima F. Estimation of Latent Ability Using a Response Pattern of Graded Scores; 34:100–100.
22. Warm TA. Weighted Likelihood Estimation of Ability in Item Response Theory; 54(3):427–450.
23. Kieftenbeld V, Natesan P. Recovery of Graded Response Model Parameters: A Comparison of Marginal Maximum Likelihood and Markov Chain Monte Carlo Estimation; 36(5):399–419. <https://doi.org/10.1177/0146621615577618> PMID: 29881020
24. Lord FM. Maximum Likelihood and Bayesian Parameter Estimation in Item Response Theory; 23(2):157–162.
25. Gelman A, Simpson D, Betancourt M. The Prior Can Often Only Be Understood in the Context of the Likelihood; 19(10):555.
26. Sheng Y. Investigating a Weakly Informative Prior for Item Scale Hyperparameters in Hierarchical 3PNO IRT Models; 8:123.
27. Lemoine NP. Moving beyond Noninformative Priors: Why and How to Choose Weakly Informative Priors in Bayesian Analyses; 128(7):912–928.
28. Watanabe S. Asymptotic Equivalence of Bayes Cross Validation and Widely Applicable Information Criterion in Singular Learning Theory; 11:3571–3594.
29. Watanabe S. A Widely Applicable Bayesian Information Criterion; 14:867–897.
30. Luo Y, Al-Harbi K. Performances of LOO and WAIC as IRT Model Selection Methods; 59(2):183.
31. Gelman A, Hwang J, Vehtari A. Understanding Predictive Information Criteria for Bayesian Models; 24(6):997–1016.
32. Stone M. An Asymptotic Equivalence of Choice of Model by Cross-Validation and Akaike's Criterion; 39(1):44–47.
33. The Promise of PROMIS; 59(2):77.
34. Hartig J, Höhler J. Multidimensional IRT Models for the Assessment of Competencies; 35(2):57–63.
35. Lord FM. Maximum Likelihood Estimation of Item Response Parameters When Some Responses Are Omitted; 48(3):477–482.

36. Wingersky MS, Lord FM. An Investigation of Methods for Reducing Sampling Error in Certain Irt Procedures*; 1983(2):i–52. PMID: [1772801](https://pubmed.ncbi.nlm.nih.gov/1772801/)
37. Bock RD, Thissen D, Zimowski MF. IRT Estimation of Domain Scores; 34(3):197–211.
38. Luo Y, Jiao H. Using the Stan Program for Bayesian Item Response Theory; 78(3):384–408.
39. Blei DM, Kucukelbir A, McAuliffe JD. Variational Inference: A Review for Statisticians; 112(518):859–877.
40. Natesan P, Nandakumar R, Minka T, Rubright JD. Bayesian Prior Choice in IRT Estimation Using MCMC and Variational Bayes; 7.
41. Wainer H, Dorans NJ, Flaugher R, Green BF, Mislevy RJ. Computerized Adaptive Testing: A Primer. Routledge;.
42. Veldkamp BP, Verschoor AJ. Robust Computerized Adaptive Testing. In: Veldkamp BP, Sluijter C, editors. Theoretical and Practical Advances in Computer-Based Educational Measurement. Methodology of Educational Measurement and Assessment. Springer International Publishing; p. 291–305.
43. Chang HH, Ying Z. A Global Information Approach to Computerized Adaptive Testing; 20(3):213–229.
44. Cheng PE, Liou M. Estimation of Trait Level in Computerized Adaptive Testing; 24(3):257–265.
45. Chang JC. Predictive Bayesian Selection of Multistep Markov Chains, Applied to the Detection of the Hot Hand and Other Statistical Dependencies in Free Throws; 6(3):182174.
46. Maragakis P, Ritort F, Bustamante C, Karplus M, Crooks GE. Bayesian Estimates of Free Energies from Nonequilibrium Work Data in the Presence of Instrument Noise; 129(2):07B609.
47. Carpenter B, Gelman A, Hoffman MD, Lee D, Goodrich B, Betancourt M, et al. Stan: A Probabilistic Programming Language; 76(1).
48. Harwell MR, Baker FB, Zwarts M. Item Parameter Estimation Via Marginal Maximum Likelihood and an EM Algorithm: A Didactic; 13(3):243–271.
49. Ogasawara H. Marginal Maximum Likelihood Estimation of Item Response Theory (IRT) Equating Coefficients for the Common-Examinee Design; 43(2):72–82.