Robust Maximum Marginal Likelihood (RMML) Estimation for Item Response Theory Models

Maxwell R. Hong

University of Notre Dame

Ying Cheng

University of Notre Dame

This is a Preprint

Author Note

Please address your correspondence to Ying Cheng, Department of Psychology,
University of Notre Dame, Notre Dame, IN 46556. Email: ycheng4@nd.edu

Abstract

Self-report data are common in psychological and survey research. Unfortunately, many of these samples are plagued with careless responses due to unmotivated participants. The purpose of this study is to propose and evaluate a robust estimation method in order to detect careless, or unmotivated, responders while leveraging Item Response Theory (IRT) person fit statistics. First, we outline a general framework for robust estimation specific for IRT models. Subsequently, we conduct a simulation study covering multiple conditions to evaluate the performance of the proposed method. Ultimately, we show how robust maximum marginal likelihood (RMML) estimation significantly improves detection rates for careless responders and reduce bias in item parameters across conditions. Furthermore, we apply our method to a real dataset to illustrate the utility of the proposed method. Our findings suggest that robust estimation coupled with person fit statistics offers a powerful procedure to identify careless respondents for further review, and to provide more accurate item parameter estimates in presence of careless responses.

Keywords: Item Response Theory, Careless Responses, Person fit, Robust Estimation, Robust Maximum Marginal Likelihood

Robust Maximum Marginal Likelihood (RMML) Estimation of Item Response Theory Models

With recent technological advances, researchers have access to large samples that are rapidly collected via online platforms, such as Amazon Mechanical Turk (Sakaluk, 2016). These new rapid data collection methods improve one's ability to efficiently collect large samples, which in turn improves the ability to detect an effect in a more precise way. However, crowdsourcing data collection poses a trade-off when we consider "noisy data" as researchers have little control over the data collection environment (Pauszek, Sztybel, & Gibson, 2017). Noisy data, or data with large amounts of measurement error, hinders the researchers' ability to detect any true effect due to bad quality data (Loken & Gelman, 2017; Maniaci & Rogge, 2014). Oftentimes, measurement error can arise from aberrant responses, where participants respond independently of item level content (Huang, Curran, Keeney, Poposki, & DeShon, 2012; Meade & Craig, 2012). The range of inattentive responses in one sample can vary, where researchers reported the proportion of compromised samples from 3.8% up to 50.0% (Berry et al., 1992; Johnson, 2005).

Different types of aberrant responses in psychological literature has been studied extensively. For instance, aberrant responses can manifest themselves as a form of careless or insufficient effort response (Curran, 2016; Meade & Craig, 2012). One manifestation of carelessness is through random responses, where a participant endorses every response option available within an item with equal probability, i.e., the probability of endorsing each response option follows the discrete uniform distribution (Beach, 1989). However, it is very difficult for a participant to repeatedly produce truly discrete uniform responses across multiple items (Meade & Craig, 2012). Instead, participants may still answer independent of item-level content by consecutively endorsing for example "strongly agree" within a set of responses (Huang et al.,

2012). Therefore, in this study we differentiate between two types of careless responses where participants can *randomly* respond versus *systematic* carelessness. In contrast to random responses, systematic carelessness is where participants respond in a non-random fashion such as choosing the same response category over multiple items.

Careless responses can lead to many serious consequences in research, such as resulting in biased correlations and attenuated or inflated scale reliability (Attali, 2005; Clark, Girona, & Young, 2003; Wise & DeMars, 2009). Reducing the effect of careless responses is no trivial problem to solve. Some researchers have proposed ways to reduce the harmful effect of careless responses by removing response patterns associated with self-reported insufficient efforts or implausible responses on bogus items (Maniaci & Rogge, 2014; Meade & Craig, 2012).

However, these approaches require planning on the researcher's part to include bogus items or self-report on response effort. For instance, some researches work with pre-collected data, which may not have the luxury of such pre-planned outlier detection. In addition, the efficacy of such approaches is questionable at best (Meade & Craig, 2012).

Some researchers control for aberrant responses by directly modeling these types of complex processes (Shao, Li & Cheng 2016; Wang, Xu & Shang, 2016; Yamamoto & Everson, 2003). In the psychological literature, others consider careless responses as a "special case" of a more general modeling framework of response styles (Böckenholt, 2017; Falk & Cai, 2016; Wetzel & Carstensen, 2017). Although these methods are attractive, directly modeling carelessness or low motivation can be challenging. First, the researcher has to assume they know the aberrant mechanism underlying the data. Many different types of carelessness have already been identified; correctly choosing what type is in a dataset would be challenging because the researcher would never know (Meade & Craig, 2012). Moreover, the general goal when

modeling response styles is to gain a better understanding of the cognitive process of item responding. Modeling carelessness or response styles is certainly an interesting line of inquiry, but qualitatively different when carelessness is considered a nuisance factor.

Our proposed method belongs to another category, i.e., post hoc procedures, to identify careless responses and minimize the impact of careless responders. These methods require no prior planning on the researchers' part, thus providing a much more flexible solution. To this end, researchers have proposed many procedures of this nature, such as the popular Mahalanobis Distance (MD) method for outlier detection (Curran, 2016). Unfortunately, the computation of MD is prohibitively intensive for long surveys (Meade & Craig, 2012), in which case careless responses are more likely to occur, because each item is essentially one dimension. Latent variable modeling that allows dimension reduction and model-based outlier detection therefore became desirable (Yuan, Fung & Reise, 2004).

Person fit statistics based on Item Response Theory (IRT) is such an example. Given a chosen IRT model, outliers, or outlying response patterns, can be identified by comparing person fit statistics to established cutoffs. This approach has been shown in the literature as a viable way to classify careless respondents (Niessen, Meijer, & Tendeiro, 2016). Person fit statistics have a rich literature in the context of education literature in order to detect aberrant responses, (e.g. Karabatsos, 2003; Meijer & Sijtsma, 2001), and have also begun to emerge in psychological research when evaluating scores from psychological scales. For instance, person fit statistics have been used to refine scales for health outcomes, anxiety and depression, mindfulness, personality and attitudes (Chien, Shao & Kuo, 2017; Edelen & Reeve, 2007; Ferrando, 2004; Pallant & Tennant, 2007; Van Dam, Earleywine & Broders, 2010). Beyond identifying person misfit, person fit statistics have been used in moderated multiple regression analysis to evaluate

validity estimates (McGrath et al., 2010). Others have shown that person fit scores add information when calculating criterion scores in a validation sample (Schmitt, Cortina & Whitney, 1993). Person fit statistics also provide more nuance into identifying participant subtypes, such as screening people of high risk for suicide who would remain undetected otherwise (Conrad et al., 2010).

When carelessness is viewed as a nuisance variable, detecting careless responses is crucial for the research process. Carelessness can lead to unreliable measures, underestimated effect size, and attenuated correlation between predictor and outcome variables, which can contribute to replication failure (Stanley & Spence, 2014). Recently, a special issue in the *Journal of Experimental Social Psychology* aimed to identify rigorous and replicable methods in social psychology in replication studies. The very first article in the issue highlights the importance of detecting carelessly invalid responses in survey data, where person fit statistics are a clear contender (Curran, 2016). As a field, we must first identify powerful methods that can correctly classify careless responders in data.

Most, if not all, person fit statistics require estimated parameters based on the data. One would normally estimate an IRT model from the full sample, which include both normal and careless responders. Researchers can then use the estimated item and person parameters to calculate person fit statistics to classify careless responders. However, the existence of a non-ignorable portion of outliers could bias the model parameter estimates (Kim & Moses, 2016; Meijer & Sijtsma, 2001; Oshima, 1994) to the extent of “masking” outliers, i.e., reducing the probability of detecting the outliers. This phenomenon is well known in the literature of outlier detection as the “masking effect” (Yuan, Fung, & Reise, 2004; Yuan & Zhong, 2008).

Removing careless respondents can potentially overcome the “masking effect” in outlier detection and some researchers have even proposed to do so iteratively to achieve the best results (Cheng & Patton, 2014). However, removing data can be a very controversial practice. In addition, removing outliers could potentially remove participants who exhibit only partial carelessness, i.e., only some of his or her responses are affected by carelessness. By removing all of their responses, we may lose useful information. Given these considerations, we propose a flexible robust estimation framework based on IRT to improve item parameter estimation, which in turn improves the detection of careless respondents. This method is flexible in that different weighting mechanisms can be used in the same general framework. In this study the weights are derived from the normalized p -value of Sinharay’s (2016a) person fit statistic l_z^* , but it does not necessarily have to be so. There is a large and diverse set of person fit statistics to choose from. Moreover, our parameter estimation procedure is demonstrated through the simple unidimensional IRT model case. However, the framework is by no means limited to the unidimensional case. It can also work with multidimensional IRT models or be applied to each individual subscale in a test battery (Felt, Castaneda, Tiemensma, & Depaoli, 2017).

The rest of our paper will be organized as follows. We will provide a review of the literature pertaining to careless responses and person fit statistics. Then, we will describe a robust procedure to estimate IRT models given response data contaminated by careless responses. By obtaining more accurate IRT model parameters we will demonstrate that we will be able to achieve better classification rates (of careless vs. non-careless response patterns) via a simulation study. Finally, we apply our methodology to the Feelings Scale in the National Longitudinal Study of Adolescent Health (AddHealth). Sample codes are provided as an appendix. We will conclude with discussion and future directions.

Method

Carelessness is one type of content-independent behavior, i.e., participants responding independently of item content. As discussed earlier, carelessness under that definition can manifest itself in several ways. For instance, for a K -point Likert-scale item, participants can endorse any of the K response categories with equal probability. This would constitute a random response. Meanwhile, carelessness may also manifest itself as having invariant responses to multiple items consecutively on the survey. This type of response behavior ties into our definition of content-independent responses, yet the careless response itself is not necessarily random.

In this study we focus on carelessness that does inflict additional randomness in item responses, which is less straightforward to detect (Meade & Craig, 2012). The most common, or “off the shelf,” statistic for the detection of outliers is the well-known Mahalanobis distance (MD):

$$D(\mathbf{u}_i) = \sqrt{(\mathbf{u}_i - \bar{\mathbf{u}})' \mathbf{S}_u^{-1} (\mathbf{u}_i - \bar{\mathbf{u}})},$$

where the MD for response vector $\mathbf{u}_i$, $D(\mathbf{u}_i)$, is evaluated based on a general distance metric between a response vector and average item scores, for person i (Curran, 2016). Assuming that $\mathbf{u}_i$ is multivariate normally distributed with average item scores $\bar{\mathbf{u}}$ and the covariance-covariance matrix $\mathbf{S}_u$, the null distribution for $D(\mathbf{u}_i)$ should follow a central chi-square distribution with the degrees of freedom equal to the number of items. Participants whose $D(\mathbf{u}_i)$ is larger than a critical value from the central chi-square distribution are considered aberrant and can be flagged at a nominal type-I error level. Unfortunately, the computation of MD is prohibitively intensive for long surveys (Meade & Craig, 2012). Methods that incorporate scale dimension reduction therefore become more attractive.

For many questionnaires that have items based on a Likert rating scale, participants rate an item on a fixed number of response options such as from *Strongly Disagree* to *Strongly Agree*. Item responses are typically coded as values one up to the total number of response options. Arguably, the most widely used IRT model for such item response data is the Graded Response Model (GRM) (Samejima, 1969). Under a GRM, each item subsumes $K-1$ boundary functions specified between item steps for K response options. The GRM represents the probability of responding either above or below the boundary as a function of latent ability, θ . The probability of endorsing the score category k or above can be characterized by a simple 2-parameter logistic model (Birnbaum, 1968):

$$P_{jk}^*(\theta) = \frac{1}{1 + \exp [-Da_j(\theta - b_{jk})]} \quad (1)$$

Note that D is a scaling constant often fixed to 1.7, so that item parameters are approximately put on the same scale as of a normal ogive model. Here b_{jk} is the location parameter for the boundary that separates the $(k-1)$ th and the k th response category of item j ; a_j is the discrimination parameter for item j for all boundary functions within item j . The probability of endorsing response option k can be expressed by taking the difference between adjacent boundary functions:

$$P_{jk}(\theta) = P_{jk}^*(\theta) - P_{j(k+1)}^*(\theta) \quad (2)$$

The probability responding below the first and highest option is set to 0.0 and 1.0, respectively. This model subsumes the widely used two-parameter logistic (2PL) model for dichotomously scored items, i.e., $K = 2$. It further subsumes the one-parameter logistic (1PL) model which sets all the a_j to be equal across items.

For content-independent responses, the GRM would not be able to capture them. Person fit statistics provide researchers a statistical method to decide whether or not a participant's response vector is aberrant in the sense that it cannot be sufficiently explained by the IRT model. Careless responses being one type of aberrant responses can therefore be possibly detected by person fit statistics. A commonly used person fit statistic is the standardized log-likelihood person fit index l_z (Drasgow, Levine, & Williams, 1985; Van Krimpen-Stoop & Meijer, 2002). Denote the item parameter matrix as $\boldsymbol{\gamma}$, where each column $\boldsymbol{\gamma}_j = (a_j, b_{j1}, b_{j2}, \dots, b_{j(K-1)})'$. The likelihood function for the i th respondent's response vector $\mathbf{u}_i$ can be written as:

$$L_i(\mathbf{u}_i | \theta; \boldsymbol{\gamma}) = \prod_{j=1}^J \prod_{k=1}^K P_{ijk}(\theta)^{\delta_{ijk}}, \quad (3)$$

where δ_{ijk} is an indicator function that the i th respondent endorses the k th response category of the j th item:

$$\delta_{ijk} = \begin{cases} 1, & \text{if } u_{ij} = k, \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

where u_{ij} is the response of the i th respondent to the j th item. The standardized person fit statistic l_z for the i th respondent takes the following form:

$$l_{zi} = \frac{l_i - E(l_i)}{\sqrt{\text{Var}(l_i)}} \quad (5)$$

where l_i is the logarithm of the likelihood function of the GRM for person i , $E(l_i)$ is the expectation of l_i and $\text{Var}(l_i)$ is the corresponding variance. Respondents who respond carelessly should exhibit large negative residuals of the likelihood, i.e. large negative l_{zi} . The standardized likelihood function should asymptotically follow a standard normal distribution under the null hypothesis of person-model fit. It is therefore common practice to flag response patterns with $l_{zi} < -1.64$, which corresponds to a nominal type-I error of 5%.

A well-known issue to this strategy is that the computation of the l_z statistic requires the true latent trait value. In practice, this value is never known and is estimated based on item responses. Estimated latent trait estimates would shrink the variance of l_z so the statistic would not follow a standard normal distribution. In consequence, the type-I error rate has been shown to be over conservative. To overcome this issue, researchers have proposed corrections to the original l_z statistic, such as Snijders (2001)'s l_{z-d}^* for dichotomous items, and Sinharay (2016a)'s correction for polytomous items, l_{z-p}^* . In this study we will consider l_{z-p}^* as we are dealing with polytomous items, or l_z^* for short. Appendix A includes the formulas required to compute l_z^* .

Using the proposed correction addresses the uncertainty introduced by the latent trait estimate. However, the person fit statistic still relies on estimated item parameter, which are biased due to careless responders in the sample. The magnitude of the bias would depend on the manifestation of carelessness and proportion of careless responses. It is noteworthy that the biased item parameter estimates would also lead to bias in latent trait estimates (Oshima, 1994). Together the biases in the structural and person parameters will distort the l_z^* , and hence the “masking effect” may result.

In the literature, some suggest that responses should be removed once identified as aberrant (e.g. Curran, 2016; Niessen, Meijer, & Tendeiro, 2016). However, data removal can be controversial. Such controversy has been well documented in the context of detection and treatment of outliers (e.g. Orr, Sackett, & Dubois, 1991). Furthermore, complete removal of a participant's response vector may not be ideal when a careless responder exhibits partial carelessness, i.e., he or she responds carelessly to some but not all items. Partial carelessness can manifest in many forms, but the most commonly observed form is called “back random

responding” (BRR), meaning careless responses occur towards the later part of a questionnaire due to cognitive fatigue or boredom etc. (Clark et al., 2003). Full removal would cause loss of information when participants exhibit partial carelessness.

We propose a robust estimation procedure that addresses the masking effect while overcoming the two limitations. This is done by weighting response patterns according to their fit in the process of item parameter estimation, i.e., misfitting patterns will be down-weighted. Note that misfitting patterns are not removed anywhere in the analysis. Meanwhile, partial carelessness is taken into account as a response pattern with higher proportion of careless responses is expected to result in smaller weight, while a response pattern with smaller proportion of careless responses should receive higher weight. A normal response pattern is expected to receive close to a full weight. Next, we will introduce details of the robust estimation procedure.

Robust Estimation of GRM Through Robust Maximum Marginal Likelihood

Robust estimation has been previously examined in other measurement models, such as confirmatory factor analysis (Yuan & Zhong, 2008; Yuan & Zhong, 2013). It has also been previously applied in different applications of IRT, such as scale equating and latent trait estimation (e.g. Bejar & Wingersly, 1981; Linn, Levine, Hastings & Wardrop, 1980; Stocking & Lord, 1983; Schuster & Yuan, 2011). However, these applications in IRT concerned themselves with robust estimation of the incidental parameters, i.e., latent trait of individuals, instead of robust estimation of the structural parameters, i.e., the item parameters.

There are at least two reasons why we choose to focus on robust estimation of the item parameters. First of all, without accurate item parameter estimates it will be difficult to accurately detect response patterns affected by carelessness, due to the masking effect.

Previous studies have investigated the utility of robust estimation for the latent trait estimator when conducting outlier detection with person fit (Sinharay, 2016b). Researchers found that using robust estimation led to better detection rates for aberrant responders compared to non-robust estimates, at the price of inflated type -I error rates. The inflation became negligible for very long tests. We therefore expect to be able to improve the detection of carelessness with robust estimation of item parameters under certain conditions. Second, obtaining accurate item parameters by itself is important. Item parameters are often treated as if they are known and without error when used in many IRT applications such as latent trait estimation (see Baker & Kim, 2004; Cheng & Yuan, 2010), scale linking and equating (van der Linden & Barrett, 2016), and computerized adaptive testing (Patton, Cheng, Yuan & Diao, 2013). If not estimated accurately, using item parameter estimates in those applications can cause many undesirable consequences (Cheng, Liu, & Behrens, 2015; Patz & Junker, 1999; Tsutakawa & Johnson, 1990).

Robust procedures typically involve assigning a weight to each individual case “according to its distance from the center of the majority of the data” (Yuan & Zhang, 2012). For our purposes when estimating item parameters, robust estimation entails assigning a weight to each respondent when applying the common Maximum Marginal Likelihood estimation method. More specifically, it entails assigning a weight to each respondent in computing the overall marginal likelihood for the GRM model (Equations 1 and 2) using the expectation maximization (EM) algorithm proposed in Bock & Aitkin (1981). Assuming that $\theta \sim f(\theta)$, the marginal probability of observing the item response vector $\mathbf{u}_i$ can be written as:

$$P(\mathbf{u}_i|\boldsymbol{\gamma}) = \int L_i(\theta|\mathbf{u}_i, \boldsymbol{\gamma})f(\theta)d\theta, \quad (6)$$

where $f(\theta)$ is often assumed to be the PDF of the standard normal distribution. The overall marginal likelihood of observing the entire response data matrix is therefore

$$P(\mathbf{u}|\boldsymbol{\gamma}) = \prod_{i=1}^N P(\mathbf{u}_i|\boldsymbol{\gamma}). \quad (7)$$

The MML procedure seeks to computationally find the $\hat{\boldsymbol{\gamma}}$ that maximizes $P(\mathbf{u}|\boldsymbol{\gamma})$ after the data are observed, but the observed data are regarded as random sample from a population. The persons are considered a random effect, whereas “the items are still considered fixed” (de Ayala, 2009, p. 69). This leads to maximizing over the log-likelihood of the response data in MML:

$$\log P(\mathbf{u}) = \sum_{i=1}^N \log P(\mathbf{u}_i|\boldsymbol{\gamma}). \quad (8)$$

Estimation based on Equation 8 implicitly assumes equal weight for each respondent. We can modify Equation 8 by adding a weight, w_i , in order to weight individuals differentially given the plausibility of their response patterns:

$$\log P(\mathbf{u}) = \sum_{i=1}^N w_i \log P(\mathbf{u}_i|\boldsymbol{\gamma}). \quad (9)$$

The EM algorithm proceeds in the same manner as usual with just a modified log-likelihood function (Baker & Kim, 2004, p. 285-291). Weighting the likelihood function based on groups of individuals is similar to the so-called pseudolikelihood approach for complex survey data (Thomas & Cyr, 2002). Note that if we remove participants from the data set, we are essentially fixing their weights to be 0.

By incorporating different weighting mechanisms into the likelihood function we can devise different robust estimation procedures. Weighting the likelihood function by some weight function follows the general form of M-estimation in the robust literature for general linear models (Carroll & Pederson, 1993; Künsch, Stefanski & Carroll, 1989). For an accessible

introduction to properties of general robust estimation, we defer readers to Chapter 2 in Wilcox (2016). In general, robust estimation requires two parts: a residual, r , and weight function, $w(\cdot)$. The residual requires one to specify which observations will be considered inconsistent. Fortunately, l_z^* functions as a standardized residual which we use in this study. Using a standardized residual is similar to the approach used in Yang and Yuan (2016). Next, we need to choose a weight function. One possibility is to choose common weight functions used in the robust statistic literature, such as Huber weights or Tukey biweights (Beaton & Tukey, 1974; Huber, 1964). However, prior work for a 2PL model found that these common weights in the robust literature to have undesirable properties for IRT model estimation (Cheng & Patton, 2014). The Huber-type weights were recommended in some previous research (e.g., Yuan, Bentler & Chan, 2004; Yuan & Zhang, 2012) on robust estimation of SEM, when the majority of the data are normally distributed but some moderate outlying cases exist (not sitting too far away from the center). In our case, the Huber-type weights may not work well as with item response data (essentially ordinal data) we can hardly expect the majority of data to be normally distributed. Meanwhile, the type of “outliers” are not ones that lead to fatter tails than a normal distribution but ones that result from model misspecification: for the inattentive respondents their responses no longer follow the usual IRT model. Hence weights based on multivariate t -distributions are not suitable either.

Given the wide use of the person-fit statistic in detecting response aberrance, the pervious success of robust estimation for the latent trait estimator when conducting outlier detection with person fit (Sinharay, 2016b), and the superiority of the l_z^* over l_z , we propose a weighting scheme based on l_z^* . Note that l_z^* asymptotically follows a standard normal distribution, we therefore propose to set

$$w_i = \frac{\Phi(l_{zi}^*)}{\sum_i \Phi(l_{zi}^*)}, \quad (10)$$

where $\Phi(\cdot)$ is the CDF of the standard normal distribution, and l_{zi}^* is the l_z^* statistic computed for the i th respondent. This essentially sets w_i equal to the normalized p -value of the l_z^* statistic under a one-sided hypothesis testing framework. The following discussion will be based on such a weight, although this does not prevent other researchers from choosing a different person fit statistic or weight function.

By setting the weight to be equal to the normalized p -value of the l_z^* statistic, we expect to give smaller weights to aberrant responses, such as careless responses, while maintaining larger weights for normal response patterns. We provide a brief illustration to show the utility of the proposed weight function. Item responses were simulated from a GRM with six response categories. The discrimination parameter was sampled from a uniform distribution with a lower bound of .5 and upper bound of 2. The location parameters were also sampled from a uniform distribution with a lower bound of -2 and upper bound of 2. In order to avoid the confounding effect of sparse categories, we resampled location parameters if the difference between adjacent categories was smaller than .3, a similar design to previous studies (Jiang, Wang, & Weiss, 2016). We generated 1,000 latent traits drawn from a standard normal distribution for a 50-item test. However, 100 participants gave random responses to 25 of the items randomly chosen from the test. Carelessness was simulated by generating responses from the discrete uniform distribution $U[0, 5]$. These participants were considered aberrant responders. Data generation and analysis were conducted in R (R Core Team, 2017).

In Figure 1, the left panel plots a respondent's l_z^* statistic against their true latent ability. Note that participants who respond aberrantly, or in this case for 50% of a survey, their l_z^* statistic is more negative compared to the normal responder counterpart. The line drawn in the

graph shows when participants are typically flagged in studies, at a cut-off value of -1.64 given that the asymptotic distribution of l_z^* is a standard normal. The right panel shows the corresponding weights if one integrates the normal distribution from $-\infty$ to q , where q is the l_z^* statistic, compared to the corresponding l_z^* statistic. This integrand is the associated p-value for each l_z^* statistic. Note that the weights in this figure are rescaled from Eq. 9 by multiplying each weight by N so that the weights sum to the total sample size. We can observe that aberrant responders are now given a smaller weight compared to the normal responders. Again, the rationale behind applying a gradient weight instead of data removal (which is essentially assigning a weight of 0 for aberrant response patterns) is to take into account the severity, or proportion of careless responses within a single participant. If we plug in these weights from our given sample into Equation 9, we conjecture that our robust procedure will produce less biased parameter estimates compared to the unweighted (or equal-weight) estimation procedure. Compared to complete data removal method we expect our method to yield smaller standard errors because we still glean valuable information from partial valid responses, and we do not decrease our sample size compared to data-removal.

Our robust method can then be described as a two-stage approach to estimate a given IRT model. First, we estimate the item parameters from the full data, which would include equal weights for all participants. Next, we calculate the person fit statistic, l_z^* , for each participant given the current estimated item parameters from the first step. Finally, we use normalized l_z^* 's p -value as weights to plug into Equation 9 and re-estimate the model. Smaller p -values would indicate larger severity of response aberrance and the corresponding response pattern would be assigned a smaller weight. Such a weighting mechanism would allow us a gradient to differentiate between partial and complete carelessness. The item parameter estimates should be

less biased, which would in turn improve detection of careless responses. A summary for our RMML estimation procedure is provided in Algorithm 1.

Algorithm 1: Robust Maximum Marginal Likelihood (RMML) Estimation for Item Response Theory Models

1. Estimate an IRT model using an unweighted maximum marginal likelihood estimation procedure;
 2. Calculate a person fit statistic based on parameters estimated in the previous step and convert the statistic to an associated weight by a prespecified function, e.g., the normalized p -value;
 3. Reestimate the IRT model using the weights from Step 2;
 4. Calculate a new person fit statistic based on updated parameters from Step 3;
 5. Identify participants of suspected aberrant behavior based on a predetermined cut-off. ...
-

Note that this general algorithm can be adapted in multiple ways. For example, one can devise different weighting mechanisms and incorporate them in step 2. In addition, Steps 2 - 4 can be iterative. That is, person-fit statistics, weights, and item parameters can be iteratively updated until some convergence criteria are met. Here we only consider one iteration, i.e., the item parameters are updated only once. If the end goal is only to obtain item parameters then one may remove step 5.

To evaluate the performance of the proposed robust procedure, we conducted a simulation study by generating a dataset containing both regular and careless responses. Accuracy and precision of parameter estimation was examined. The power and type-I error of the proposed procedure in detecting carelessness were also investigated. It is important to note that our method should accommodate various forms of aberrant response pattern. Carelessness, given its prevalence in low-stakes testing, is one example we are most concerned with.

Simulation Study

Item and person responses were simulated the same way as the data generation scheme previously described with minor changes. For accurate parameter recovery of the GRM, previous

literature suggested that the minimum required sample size ranges from 375 to 750 depending on the length of the scale (De Ayala, 1994) as well as the number of response categories per item. This is true when the responses do follow the GRM. With misfitting response patterns in the data, it is unclear what would be an adequate sample size. Hence we included two sample sizes in our study, 500 or 3000, with the former being within the range of recommended sample size, while the latter being way more than adequate in traditional settings.

Meade and Craig (2012) suggested that the efficacy of detection methods of carelessness depends strongly on the nature of the response data. Following their study, we simulated both *random* and *midpoint* carelessness. Random carelessness was simulated by generating responses from the discrete uniform distribution $U[0, 5]$. Each response option has equal chance of endorsement, $1/6$. Midpoint carelessness was simulated where participants are more likely to select the middle categories, in our case categories 2 and 3. This was done by drawing from a Binomial distribution where the number of categories equal to 5 and probability of success fixed to .5. This way the expected probability of endorsing 0 or 5 is roughly 3%, endorsing 1 or 4 is around 16%, and endorsing 2 or 3 is around 31%.

It is also important to consider the prevalence and severity of careless responses in the response data, because large amount of careless responses can cause the “masking effect”. Prevalence refers to the proportion of participants who responded carelessly to some or all the items. Severity refers to the proportion of items affected by carelessness within a response vector. Together they indicate the proportion of responses in an N by J data matrix affected by carelessness. In this study prevalence in the total sample was set at either 10% or 30%, with the former being consistent with recent research of prevalence rate, which reported the prevalence rate to be 5 – 20% (Curran, Kotrba, & Denison, 2010) and 11% (Meade & Craig, 2012), whereas

the latter being a more extreme rate to challenge our proposed procedure. With respect to severity, in past research people self-reported to respond carelessly to 50% or more items on long surveys that involve multiple subscales (Baer, Ballenger, Berru, & Wetter, 1997; Berry et al., 1992). Given that our study focuses on shorter, unidimensional measures, we simulated severity rate to be either 30 or 50%. In other words, for a careless participant, we simulated 30 or 50% of his or her responses to be careless responses intermittently throughout the survey. We also varied the test length to be either 30 or 50.

In sum, we had a total of 32 conditions for data generation with 100 replications for each condition, with five factors fully crossed (2 test lengths $\times$ 2 sample sizes $\times$ 2 types of careless responses $\times$ 2 prevalence rates $\times$ 2 severity rates). For each data generation condition, we compared the performance of three procedures that involves the use of l_z^* in terms of the detection of careless respondents and quality of item parameter estimates. The first procedure (denoted as full-sample l_z^*) used the full sample to both estimate the GRM and identify participants who aberrantly responded using l_z^* based on estimated parameters. Participants whose l_z^* are smaller than -1.64 are flagged as aberrant. The second procedure (denoted by data removal) involves data removal, i.e., participants whose l_z^* are smaller than -1.64 are removed. Then the model is re-estimated and participants' l_z^* are updated with new parameter estimates. Data removal implies that during model re-estimation each un-flagged participant receives a weight of one and each flagged participant a weight of 0 based on Equation 9. Those with updated l_z^* smaller than -1.64 are flagged as aberrant (data removal). The third procedure (denoted by robust estimation) is our RMML estimation procedure detailed in previous sections. All programming was done using the R language. For the RMML procedure, the number of quadrature points during the E step was fixed to 60. Maximization over the log-likelihood was

done using a Newton-Raphson procedure during the M step. Convergence was established if the log-likelihood between adjacent iterations did not change beyond .0001.

We evaluated the power and type-I error of the three procedures in identifying careless respondents. Power is defined as the proportion of flagged participants over the total number of participants affected by carelessness. Type-I error is defined as the proportion of participants unaffected by carelessness that were flagged. For ability estimation we used maximum likelihood estimation procedures. When participants endorsed 0 for all items or obtained a full score, we fixed the ability estimates to be -3 and 3, respectively. We compared these three procedures in power and Type-I error against two scenarios: a) the baseline scenario when item parameters are known and l_z^* is used to flag participants based on the true parameter values; and b) when outlier detection uses MD on all of the data. This is the same procedure as the full-sample l_z^* procedure except that MD is used for flagging participants instead of l_z^* . The motivation for including the full-sample MD conditions is due to the ubiquity of MD in the outlier detection literature. The known item parameters condition serves as an ideal case, which never holds in practice.

We evaluated average bias of the location parameter estimates, $\sum_{j=1}^J (\widehat{b}_{jk} - b_{jk}) / J$ and root mean square error (RMSE), $(\sum_{j=1}^J (\widehat{b}_{jk} - b_{jk})^2 / J)^{-.5}$, for $k = 1, 2, \dots, (K-1)$. For the discrimination parameter, we evaluated the parameter using conditional bias and RMSE where we first bin the true discrimination parameters into six equally spaced intervals between .5 and 2, then evaluate the parameter bias and RMSE. The average parameter bias and RMSE were also evaluated across only the three estimation procedures using l_z^* . Bias, RMSE and classification accuracy were averaged across 100 replications.

Results

Preliminary results showed little differences between sample sizes of 500 vs. 3000 and test length. We therefore only present results for the sample size of 500 and a test length of 50. Table 1, and 2 show the average power, type-I error, and classification accuracy rates for those conditions, respectively.

The full sample condition using either l_z^* or MD is the common practice in reality. Both the data removal method and the robust estimation method try to take the inaccurate item parameter estimates into account in detecting careless responses. Between the two methods, data removal and robust estimation, the former is a special case of the latter when flagged participants are given a weight of 0. The known parameter condition is an ideal condition. Evidently the power does not appear to be high even in the ideal situation, especially when the prevalence rate is low. This is consistent with existing literature, reporting traditional methods with power well below .5 in classifying random responders as such (Huang et al., 2012; Marjanovic et al., 2015), even when multiple indices are employed as suggested by Meade and Craig (2012). Our simulation conditions particularly challenge the detection methods, because we simulated partial carelessness (in contrast to careless responses on every item) and midpoint carelessness, both have been shown in previous studies to hurt the power (Meade & Craig, 2012; Cheng & Patton, 2014). In Table 2 we also observe that the power is uniformly higher for detecting random careless responses than midpoint carelessness.

Compared to the known item parameter condition, the full sample condition leads to an evident drop of power when using either both l_z^* or MD. The data removal method does help improve the power slightly. Data removal and the full sample conditions result in conservative type-I error rates. On the other hand, the robust estimation procedure increased power across conditions compared to the baseline and other estimation methods. However, we do experience

an inflated type-I error in the robust condition compared to all other methods. Further discussions did not evaluate procedures based on MD due to its poor performance in detecting simulated careless respondents compared to methods that utilized l_z^* .

Figures 2 and 3 present the empirical bias and RMSE for the item location and discrimination parameters for three procedures that involve l_z^* (full-sample l_z^* , data removal, and robust estimation), respectively. For item location parameters, a clear trend emerges when we look across groups. The largest and smallest location parameters are the most biased estimates across replications. Data removal reduces the bias in many scenarios. However, the robust estimation either out-performs or operates at the same level compared to the data removal procedure. A different pattern emerges when we look at RMSE. Data removal reduces the RMSE when compared to the full estimation. The robust estimation procedure has slightly larger RMSE when compared to data removal across conditions, and even exceeds the RMSE for the full data sample occasionally. This perhaps could explain the inflated type-I error rates for the robust estimation procedure.

For the conditional bias and RMSE for the discrimination parameters, the full sample approach creates (negative) bias for the discrimination parameters. The magnitude of the bias increases when the true discrimination parameter increases. The bias and RMSE also increase with more severity and prevalence of careless responses. Data removal improves estimation for larger discrimination values when we look at both RMSE and bias. However, robust estimation outperforms both full estimation and data removal across the majority of the conditions, whether we look at RMSE or bias.

Taken together, results from the simulation study suggest that the robust estimation method improves item parameter estimation in presence of careless responses, even when the

prevalence and severity rates are high. It also improves the power of detecting careless responses. However, the improvement in power seems to come at a cost of inflated type-I error. This suggests that we need to be careful with the next steps once we flag the participants. How to handle suspicious responses is an evolving practice and our goal is not to provide a strict guideline for practitioners. For instance, Allalouf, Gutentag, and Baumer (2017) recommended that following statistical quality control procedures a human review of suspicious cases should be conducted. Other precautionary steps may be taken as outlined in Wainer (2014) for suspicious responses. We argue that robust estimation's high power makes it a powerful tool for data driven methods when identifying suspicious responses. We strongly caution against data removal on the sole basis of results from statistical analysis in testing.

To further exemplify our key findings, we built upon our simulation study with an empirical analysis to illustrate the robust estimation procedure.

Empirical Analysis

We applied both the data removal and robust procedures to data from the National Longitudinal Study of Adolescent Health (AddHealth), which is an education and health-related study of adolescents in grades 7-12 (Harris & Udry, 2010). We restricted our analysis to the base year, which was an in-home questionnaire administered during 1994-1995. We analyzed the Feelings scale, which collects information about the respondents' current emotional state. The scale contains 19 items in which participants respond to Likert-scale items with 4 response categories on how often each of the statements apply to the past week (e.g. "You were bothered by things that usually don't bother you"). However, participants rarely endorse the category 4, "most of the time or all of the time." Therefore, we collapsed this category with response

category 3. This approach mirrors similar studies on this scale where a unidimensional IRT model was fit to the data and found to have adequate fit (Edelen & Reeve, 2007).

The original sample contains 6,504 respondents, but respondents with any missing data in the Feelings scale were removed, leaving 6,457 respondents. Using this final sample, item parameter estimates were obtained without cleansing (i.e., using all 6,457 participants), with data removal using l_z^* , and with the robust estimation method. We did not utilize cleansing based on MD due to its suboptimal performance compared to procedures that used l_z^* in our simulation conditions. We evaluated all three methods based on the classification of respondents and differences among the resulting item parameter estimates and their standard errors. Without any cleansing, 7.6% of the sample was flagged. With data removal, 11.18% of the sample was flagged, and with robust estimation, 20.59% of the sample was flagged. This is consistent with our simulation studies, where using the entire data set resulted in reduced power and overly conservative type-I error rate. According to the simulation study, using data removal increases power and here we see that data removal flags 3.5% more participants. Robust estimation flagged an additional 13% of the total sample than using the non-weighted full sample.

Item parameter estimates based on the full, data removal, and robust cleansed samples with their standard errors are presented in Table 3. Data removal resulted in little change or slightly larger discrimination parameter estimates. For item 1, the full sample yielded an estimate of 1.42, whereas data removal yielded 1.48 and the robust estimation procedure produced 1.67. Larger discrimination estimates were observed in the majority of the other items as well. Moreover, the effect of careless responses on b parameter estimates varied across items. Some parameter estimates increased or decreased depending on the estimation procedure. Potentially different kinds of aberrant responses may attribute to different amounts of bias for the location

parameter, which can explain why there are different manifestations of item parameter bias. It would be difficult to generalize the overall effect of careless responses specifically for the location parameter. All standard error estimates for the item parameters were also comparable, which could be attributed to the large sample size.

Discussion

This paper proposed a robust estimation framework to improve item parameter estimation in presence of aberrant responses. We provide a general method that could leverage person fit statistics, such as l_z^* , to detect careless responders and down-weight them in the estimation process. This procedure was compared to what is common in practice, the unweighted full data estimation procedure. We also applied a data cleansing procedure, which is essentially a special case of the robust estimation procedure, when one fixes the weights to be zero for flagged participants. We compared all three procedures that involve the use of l_z^* using a simulation study, where different careless response behavior, percentage of severity and prevalence, sample size, and item bank size were considered. In general, the impact of careless responses was regulated mostly by the severity and prevalence, or the number of careless responses. Test length and sample size made minimal differences when considering the impact of careless responses. Moreover, the type of careless response had similar impacts across conditions, with varying magnitudes.

Results from our simulation study suggest that using the entire data set when fitting an IRT model severely biases both location and discrimination parameters, at least in the context with the GRM polytomous model. Furthermore, we found that the biased item parameter estimates significantly reduces the ability of detecting outliers with the person fit statistic, l_z^* . This finding is consistent with previous research, where the masking effect reduces the ability of

finding outliers due to biased structural parameter estimates (Yuan & Zhong, 2008). In order to assuage the impact of bias item parameters, our goal was to compare two procedures that could potentially improve both item parameter estimates and detection rates: a data-removal process and a more general robust estimation procedure. Both procedures improved item calibration, i.e., they resulted in less biased item parameter estimates. In turn, both procedures were able to better detect outliers due to a reduction in the masking effect.

In spite of the encouraging results, there are several limitations and room for improvement concerning our robust procedure. One confounding element for our weight function is that we require an estimate of latent ability. For a finite test length, l_z^* may not follow a normal distribution. Sinharay (2016b) and Snijder (2001) found that l_z^* may have smaller variances for extreme values of latent ability. This would in turn influence our weight function. Second, aberrant responses contribute their own form of bias when estimating latent trait estimates (Oshima, 1994). In order to correct this form of bias in the latent trait estimate, robust estimators for latent ability can certainly be used (e.g. Schuster & Yuan, 2011). Coupling robust estimators for latent ability and item parameters could be studied in the future. The performance of our method may be evaluated under more simulation conditions. For instance, the tendency to respond carelessly can also be correlated with latent trait ability (Falk & Cai, 2016), or carelessness may lead to responses that are not completely independent of item content (Shao & Cheng, 2017; Yu & Cheng, 2017). These scenarios are not covered in this paper but certainly warrant attention in future studies.

Moreover, robust estimators of an IRT model can have a breakdown point. A breakdown point for a robust estimator corresponds to the proportion of aberrant observations an estimator can accommodate before giving a severely biased or inconsistent estimate. For an IRT model, the

breakdown point for robust estimation would become most evident concerning the discrimination parameter estimate. The robust estimator for slope parameter estimates could explode after several iterations of down weighting outliers (Croux, Flandre & Haesbroeck, 2002). This occurs when it is easy to distinguish between high and low ability examinees or where a vertical line can separate when a person would or would not endorse an item. This would cause discrimination parameter estimate to approach infinity. Outliers can be thought of a “safe-guard” to minimize this effect.

An unforeseen side effect of the robust-estimation procedure was inflated RMSE compared to the two other estimation procedures, especially in the largest location parameters in absolute terms. A trade-off between bias and efficiency is not unusual, where a less bias estimate using robust estimation may be worth the loss of efficiency compared to non-robust estimation (Carrol & Pederson, 1993). Moreover, larger RMSE could potentially explain why our type-I error rates for outlier detection were significantly larger in the case of robust estimation procedures. A similar finding when using robust estimators for latent ability further validates our argument (Sinharay, 2016b). Perhaps coupling our robust estimators with a resampling procedure, such as the bootstrap, could build an empirical null distribution for the l_z^* statistic (Efron & Tibshirani, 1994). The bootstrap procedure would retain the nice properties of unbiased item-parameter estimates and could potentially reduce the type-I error rates to a nominal level.

These limitations aside, both other researchers and applied psychometricians should find our study meaningful. First, our method provides a general framework for robust estimation. Our study provides the first initial steps that outlines a general methodology, but can be improved upon or modified. For instance, our choice of utilizing the l_z^* statistic was based on its popularity, familiarity and simplicity of use. Researchers could potentially choose other weighting

procedures, such as using the MD (Meade & Craig, 2012). Other Person fit statistics can also be used, such as non-parametric person fit statistics such as *Gnorm* (Emons, 2008).

Our overall approach in this paper was done by maximizing a so-called psuedolikelihood function which directly maps onto literature estimating IRT models with complex survey designs (Thomas & Cyr, 2002) and robust generalized linear model (Carrol & Pederson, 1993; Künsch, Stefanski, & Carroll, 1989). However, other robust methods may be employed such as as those in the general structural equation modeling framework by iteratively reweighted least squares (Yuan & Zhong, 2008).

In addition, our study could be applied to other measurement situations, such as using IRT models suited for educational testing scenario as opposed to a psychological battery. In that scenario, other common IRT models such as the 2PL or 3PL for dichotomous items may be more appropriate. As GRM subsumes 2PL and 1PL, our RMML procedure applies directly to those models. With proper adaptation, the procedure can be used with 3PL as well. Other considerations could be to understand our method in the context of more complex IRT models that are common in practice for psychology, such as multidimensional IRT models. Our method can address this situation two ways. First, one can apply the robust estimation procedure to subscales in a psychological battery. This approach would certainly be the easiest computationally, and provide a rich picture on which scales a participant does not perform at their true ability level. Our method could also be framed using a multidimensional IRT model, where one would just add a weighting procedure to the likelihood when estimating a multidimensional IRT model.

Our method is also easy to adopt in practice. Example code to calculate conduct robust estimation is provided in Appendix B using R. Moreover, one could also use other R packages to calculate person fit statistics, such as *PerFit*, or create their own weights to employ in *mirt*,

which has a weight function argument can be used to implement our method (Chalmers, 2012; Tendiero, 2015).

Finally, our study contributes to the literature more broadly in psychology and other sciences have large sources of measurement error (Maxwell, Lau, & Howard, 2015). More specifically, we address issues when identifying powerful statistics to detect outliers in a data set that could potentially confound replication studies (Loken & Gelman, 2017). Detecting outliers is an important step psychologists should take during data preprocessing; our method provides one improvement on pre-existing methods.

References

- Allalouf, A., Gutentag, T., & Baumer, M. (2017). Quality Control for Scoring Tests Administered in Continuous Mode: An NCME Instructional Module. *Educational Measurement: Issues and Practice*, 36(1), 58–68. doi: 10.1111/emip.12140
- Attali, Y. (2005). Reliability of Speeded Number-Right Multiple-Choice Tests. *Applied Psychological Measurement*, 29(5), 357–368. doi: 10.1177/0146621605276676
- Baker, F.B., Kim, S.H. *Item response theory: Parameter estimation techniques*. 2nd Eds. New York: Marcel Dekker; 2004.
- Baer, R. A., Ballenger, J., Berru, D., & Wetter, M. W. (1997). Detection of random responding on the MMPI-A. *Journal of Personality Assessment*, 68, 139–151.
- Beach, D. A. (1989). Identifying the random responder. *Journal of Psychology: Interdisciplinary and Applied*, 123(1), 101–103. doi: 10.1080/00223980.1989.10542966
- Beaton, A. E., & Tukey, J. W. (1974). The Fitting of Power Series, Meaning Polynomials, Illustrated on Band-Spectroscopic Data. *Technometrics*, 16(2), 147–185. doi: 10.1080/00401706.1974.10489171
- Bejar, I. & Wingersky, M.S. *An application of item response theory to equating the Test of Standard Written English* (College Board Report No. 81-8). Princeton NJ: Educational Testing Service, 1981. (ETS No. 81-35)
- Berry, D., Wetter, M., Baer, R. A., Larsen, L., Clark, C., & Monroe, K. (1992). MMPI-2 Random Responding Indices: Validation Using a Self-Report Methodology. *Psychological Assessment*, 4(3), 340–345. doi: 10.1037/1040-3590.4.3.340
- Birnbaum, A. (1968). Some latent trait models and their use in inferring an examinee's ability. *Statistical Theories of Mental Test Scores*.

- Bock, R. D., & Aitkin, M. (1981). EM Solution of the Marginal Likelihood Equations. *Psychometrika*, 46(4), 443–459.
- Böckenholt, U. (2017). Measuring response styles in likert items. *Psychological Methods*, 22(1), 69–83. doi: 10.1037/met0000106
- Carroll, R. J., & Pederson, S. (1993). On robustness in the logistic regression model. *Journal of the Royal Statistical Society Series B*, 84(3), 693–706.
- Chalmers, R. P. (2012). mirt: a Multidimensional Item Response Theory Package for R Environment. *Journal of Statistical Software*, 28(6), 1-29. doi: 10.18637/jss.v048.i06
- Cheng, Y., & Patton, J. M. (2014). *Detection and treatment of careless responses in survey data*. Poster presented at the annual convention of the Association for Psychological Science, San Francisco, CA.
- Cheng, Y., Liu, C., & Behrens, J. (2015). Standard error of ability estimates and the classification accuracy and consistency of binary decisions. *Psychometrika*, 80(3), 645–664. doi: 10.1007/s11336-014-9407-z
- Cheng, Y., & Yuan, K. (2010). The impact of fallible item parameter estimates on latent trait recovery. *Psychometrika*, 75(2), 280–291. doi: 10.1007/s11336-009-9144-x
- Chien, T. W., Shao, Y., & Kuo, S. C. (2017). Development of a Microsoft Excel tool for one-parameter Rasch model of continuous items: An application to a safety attitude survey. *BMC Medical Research Methodology*, 17(1). doi: 10.1186/s12874-016-0276-2
- Clark, M. E., Gironda, R. J., & Young, R. W. (2003). Detection of back random responding: effectiveness of MMPI-2 and Personality Assessment Inventory validity indices. *Psychological Assessment*, 15(2), 223–234. doi: 10.1037/1040-3590.15.2.223

- Conrad, K. J., Bezruczko, N., Chan, Y. F., Riley, B., Diamond, G., & Dennis, M. L. (2010). Screening for atypical suicide risk with person fit statistics among people presenting to alcohol and other drug treatment. *Drug and Alcohol Dependence*, *106*(2–3), 92–100. doi: 10.1016/j.drugalcdep.2009.07.023
- Croux, C., Flandre, C., & Haesbroeck, G. (2002). The breakdown behavior of the maximum likelihood estimator in the logistic regression model. *Statistics & Probability Letters*, *60*(4), 377–386.
- Curran, P. G. (2016). Methods for the detection of carelessly invalid responses in survey data. *Journal of Experimental Social Psychology*, *66*, 4–19. doi: 10.1016/j.jesp.2015.07.006
- Curran, P. G., Kotrba, L., & Denison, D. (2010, April). *Careless responding in surveys: Applying traditional techniques to organizational settings*. Poster presented at the 25th Annual Conference of the Society for Industrial and Organizational Psychology, Atlanta.
- de Ayala, R. J. (1994). The influence of multidimensionality on the graded response model. *Applied Psychological Measurement*, *18*(2), 155–170. doi: 10.1177/014662169401800205
- de Ayala, R. J. (2009). *The theory and practice of item response theory*. New York: The Guilford Press.
- Drasgow, F., Levine, M. V., & Williams, E. A. (1985). Appropriateness measurement with polychotomous item response models and standardized indices. *British Journal of Mathematical and Statistical Psychology*, *38*(1), 67–86. doi: 10.1111/j.2044-8317.1985.tb00817.x
- Edelen, M. O., & Reeve, B. B. (2007). Applying item response theory (IRT) modeling to questionnaire development, evaluation, and refinement. *Quality of Life Research*, *16*, 5–18. doi: 10.1007/s11136-007-9198-0

- Efron, B., & Tibshirani, R. J. (1994). *An introduction to the bootstrap*. New York: Chapman & Hall.
- Emons, W. H. M. (2008). Nonparametric Person-Fit Analysis of Polytomous Item Scores. *Applied Psychological Measurement*, 32(3), 224–247. doi: 10.1177/0146621607302479
- Falk, C. F., & Cai, L. (2016). A flexible full-information approach to the modeling of response styles. *Psychological Methods*, 21(3), 328–347. doi: 10.1037/met0000059
- Felt, J. M., Castaneda, R., Tiemensma, J., & Depaoli, S. (2017). Using person fit statistics to detect outliers in survey research. *Frontiers in Psychology*, 8(863). doi: 10.3389/fpsyg.2017.00863
- Ferrando, P. J. (2004). Person Reliability in Personality Measurement: An Item Response Theory Analysis. *Applied Psychological Measurement*, 28(2), 126–140. doi: 10.1177/0146621603260917
- Harris, K. M., & Udry, J. R. (2010). National Longitudinal Study of Adolescent Health (Add Health), 1994-2008: Core Files Restricted Use. *Inter-University Consortium for Political and Social Research*. doi: 10.3886/ICPSR27021
- Huang, J. L., Curran, P. G., Keeney, J., Poposki, E. M., & DeShon, R. P. (2012). Detecting and Deterring Insufficient Effort Responding to Surveys. *Journal of Business and Psychology*, 27(1), 99–114. doi: 10.1007/s10869-011-9231-8
- Huber, P. J. (1964). Robust Estimation of a Location Parameter. *The Annals of Mathematical Statistics*, 35(1), 73–101. doi: 10.1214/aoms/1177703732
- Jiang, S., Wang, C., & Weiss, D. J. (2016). Sample Size Requirements for Estimation of Item Parameters in the Multidimensional Graded Response Model. *Frontiers in Psychology*, 7. doi: 10.3389/fpsyg.2016.00109

- Johnson, J. A. (2005). Ascertaining the validity of individual protocols from Web-based personality inventories. *Journal of Research in Personality*, 39(1), 103–129. doi: 10.1016/j.jrp.2004.09.009
- Karabatsos, G. (2003). Comparing the Aberrant Response Detection Performance of Thirty-Six Person-Fit Statistics, 16(4), 277–298. doi: 10.1207/S15324818AME1604
- Kim, S., & Moses, T. (2016). *ETS GRE ® Board Research Report Investigating Robustness of Item Response Theory Proficiency Estimators to Two-Stage Multistage Testing*.
- Künsch, H. R., Stefanski, L. A., & Carroll, R. J. (1989). Conditionally unbiased bounded-influence estimation in general regression models, with applications to generalized linear models. *Journal of the American Statistical Association*, 84(406), 460–466. doi: 10.1080/01621459.1989.10478791
- Loken, E., & Gelman, A. (2017). Measurement error and the replication crisis. *Science*, 355(6325), 584–585. doi: 10.1126/science.aal3618
- Linn, R.L., Levine, M. V., Hasings, C. N., & Wardrop, J. L. *An investigation of item bias in a test of reading comprehension* (Technical Report No. 163). Urbana IL: Center for the Study of Reading, University of Illinois, 1980.
- Maniaci, M. R., & Rogge, R. D. (2014). Caring about carelessness: Participant inattention and its effects on research. *Journal of Research in Personality*, 48(1), 61–83. doi: 10.1016/j.jrp.2013.09.008
- Marjanovic, Z., Holden, R., Struthers, W., Cribbie, R., & Greenglass, E. (2015). The inter-item standard deviation (ISD): An index that discriminates between conscientious and random responders. *Personality and Individual Differences*, 84, 79–83. doi: 10.1016/j.paid.2014.08.021

- Maxwell, S. E., Lau, M. Y., & Howard, G. S. (2015). Is psychology suffering from a replication crisis? What does “failure to replicate” really mean? *American Psychologist*, 70(6), 487–498. doi: 10.1037/a0039400
- McGrath, R. E., Mitchell, M., Kim, B. H., & Hough, L. (2010). Evidence for Response Bias as a Source of Error Variance in Applied Assessment. *Psychological Bulletin*, 136(3), 450–470. doi: 10.1037/a0019216
- Meade, A. W., & Craig, S. B. (2012). Identifying careless responses in survey data. *Psychological Methods*, 17(3), 437–455. doi: 10.1037/a0028085
- Meijer, R. R., & Sijtsma, K. (2001). Methodology Review: Evaluating Person Fit. *Applied Psychological Measurement*, 25(2), 107–135. doi: 10.1177/01466210122031957
- Niessen, A. S. M., Meijer, R. R., & Tendeiro, J. N. (2016). Detecting careless respondents in web-based questionnaires: Which method to use? *Journal of Research in Personality*, 63, 1–11. doi: 10.1016/j.jrp.2016.04.010
- Orr, J. M., Sackett, P. R., & Dubois, C. L. Z. (1991). Outlier Detection and Treatment in I/O Psychology: A Survey of Researcher Beliefs and an Empirical Illustration. *Personnel Psychology*, 44(3), 473–486. doi: 10.1111/j.1744-6570.1991.tb02401.x
- Oshima, T. C. (1994). The Effect of Speededness on Parameter Estimation in Item Response Theory. *Journal of Educational Measurement*, 31(3), 200–219. doi: 10.1111/j.1745-3984.1994.tb00443.x
- Pallant, J. F., & Tennant, A. (2007). An introduction to the Rasch measurement model: An example using the Hospital Anxiety and Depression Scale (HADS). *British Journal of Clinical Psychology*, 46(1), 1–18. doi: 10.1348/014466506X96931
- Patton, J. M., Cheng, Y., Yuan, K.-H., & Diao, Q. (2013). The influence of item calibration error

- on variable-length computerized adaptive testing. *Applied Psychological Measurement*, 37, 24-40.
- Patz, R.J., & Junker, B.W. (1999). A straightforward approach to Markov chain Monte Carlo methods for item response models. *Journal of Educational and Behavioral Statistics*, 24, 146–178.
- Pauszek, J. R., Sztybel, P., & Gibson, B. S. (2017). Evaluating Amazon’s Mechanical Turk for psychological research on the symbolic control of attention. *Behavior Research Methods*, 1969–1983. doi: 10.3758/s13428-016-0847-5
- R Core Team. (2017). R Development Core Team. *R: A Language and Environment for Statistical Computing*. Retrieved from <https://www.r-project.org/>
- Sakaluk, J. K. (2016). Exploring Small, Confirming Big: An alternative system to The New Statistics for advancing cumulative and replicable psychological research. *Journal of Experimental Social Psychology*, 66, 47–54. doi: 10.1016/j.jesp.2015.09.013
- Samejima, F. (1969). Estimation of latent ability using a response pattern of graded scores. *Psychometrika*. doi: 10.1007/BF02290599
- Schmitt, N., Cortina, J. M., & Whitney, D. J. (1993). Appropriateness Fit and Criterion-Related Validity. *Applied Psychological Measurement*, 17(2), 143–150. doi: 10.1177/014662169301700204
- Schuster, C., & Yuan, K.-H. (2011). Robust Estimation of Latent Ability in Item Response Models. *Journal of Educational and Behavioral Statistics*, 36(6), 720–735. doi: 10.3102/1076998610396890
- Shao, C., & Cheng, Y. (2017, May). *Detection of Test Speededness Using Change-Point Analysis with Response Time Data*. Paper presented at the Annual Meeting of National

Council for Measurement in Education, San Antonio, TX.

Shao, C., Li, J., & Cheng, Y. (2016). Detection of test speededness using change-point analysis.

Psychometrika, 81(4), 1118-1141. doi: 10.1007/s11336-015-9476-7

Sinharay, S. (2016a). Asymptotically Correct Standardization of Person-Fit Statistics Beyond

Dichotomous Items. *Psychometrika*, 81(4), 992-1013. doi: 10.1007/s11336-015-9465-x

Sinharay, S. (2016b). The choice of the ability estimate with asymptotically correct standardized

person-fit statistics. *British Journal of Mathematical and Statistical Psychology*, 69(2),

175–193. doi: 10.1111/bmsp.12067

Snijders, T. A. B. (2001). Asymptotic null distribution of person fit statistics with estimated

person parameter. *Psychometrika*, 66(3), 331–342. doi: 10.1007/BF02294437

Stanley, D. J., & Spence, J. R. (2014). Expectations for Replications. *Perspectives on*

Psychological Science, 9(3), 305–318. doi: 10.1177/1745691614528518

Stocking, M. L., & Lord, F. M. (1983). Developing a Common Metric in Item Response Theory.

Applied Psychological Measurement, 7(2), 201–210. doi: 10.1177/014662168300700208

Tendeiro, J. N. (2015). Perfit (R package version 1.4) [computer software] Available from

<<http://cran.r-project.org/web/packages/PerFit/index.html>>.

Thomas, D. R., & Cyr, A. (2002). Applying item response theory methods to complex survey

data. In *Proceedings of the Survey Methods* (pp. 17-26).

Tsutakawa, R.K., & Johnson, J.C. (1990). The effect of uncertainty of item parameter estimation

on ability estimates. *Psychometrika*, 55, 371–390.

Van Dam, N. T., Earleywine, M., & Borders, A. (2010). Measuring mindfulness? An Item

Response Theory analysis of the Mindful Attention Awareness Scale. *Personality and*

Individual Differences, 49(7), 805–810. doi: 10.1016/j.paid.2010.07.020

- van der Linden, W. J., & Barrett, M. D. (2016). Linking Item Response Model Parameters. *Psychometrika*, 81(3), 650–673. doi: 10.1007/s11336-015-9469-6
- van Krimpen-Stoop, E. M., & Meijer, R. R. (2002). Detection of Person Misfit in Computerized Adaptive Tests with Polytomous Items. *Applied Psychological Measurement*, 26(2), 164–180. doi: 10.1177/01421602026002004
- Wainer, H. (2014). Cheating: Some ways to detect it badly. In N. M. Kingston & A. K. Clark (Eds.), *Test fraud: Statistical detection and methodology* (pp. 8-20). New York, NY: Taylor and Francis.
- Wang, C., Xu, G., & Shang, Z. (2016). A Two-Stage Approach to Differentiating Normal and Aberrant Behavior in Computer Based Testing. *Psychometrika*, 1–32. doi: 10.1007/s11336-016-9525-x
- Wetzel, E., & Carstensen, C. H. (2017). Multidimensional modeling of traits and response styles. *European Journal of Psychological Assessment*, 33(5), 352–364. doi: 10.1027/1015-5759/a000291
- Wilcox, R. R. (2016) *Introduction to robust estimation and hypothesis testing* (4th ed.). San Diego, CA: Academic Press.
- Wise, S. L., & DeMars, C. E. (2009). A Clarification of the Effects of Rapid Guessing on Coefficient : A Note on Attali’s Reliability of Speeded Number-Right Multiple-Choice Tests’. *Applied Psychological Measurement*, 33(6), 488–490. doi: 10.1177/0146621607304655
- Yamamoto, K., & Everson, H. (2003). Estimating the Effects of Test Length and Test Time on Parameter Estimation Using the Hybrid Model. *ETS Research Report Series*, 1995(1), 277–298. doi: 10.1002/j.2333-8504.1995.tb01637.x

- Yu, X., & Cheng, Y. (2017, Sept). *Using change point analysis to detect inattentiveness in polytomous survey response data*. Paper presented at the 2017 Conference on Test Security, Madison, WI.
- Yang, M., & Yuan, K. H. (2016). Robust Methods for Moderation Analysis with a Two-Level Regression Model. *Multivariate Behavioral Research*. doi: 10.1080/00273171.2016.1235965
- Yuan, K.-H., Bentler, P. M., & Chan, W. (2004). Structural equation modeling with heavy tailed distributions. *Psychometrika*, 69(3). doi: 10.1007/BF02295644
- Yuan, K.-H., Fung, W. K., & Reise, S. P. (2004). Three Mahalanobis distances and their role in assessing unidimensionality. *The British Journal of Mathematical and Statistical Psychology*, 57, 151–65. doi: 10.1348/000711004849231
- Yuan, K. H., & Zhang, Z. (2012). Robust Structural Equation Modeling with Missing Data and Auxiliary Variables. *Psychometrika*, 77(4), 803-826. doi: 10.007/s11336=012-9282-4
- Yuan, K. H., & Zhong, X. (2008). Outliers, leverage observations, and influential cases in factor analysis: Using robust procedures to minimize their effect. *Sociological Methodology*, 38(1), 329–368. doi: 10.1111/j.1467-9531.2008.00198.x
- Yuan, K. H., & Zhong, X. (2013). Robustness of fit indices to outliers and leverage observations in structural equation modeling. *Psychological Methods*, 18(2), 121–136. doi: 10.1037/a0031604

Table 1: Type-1 error for aberrant and normal response patterns

Aberrant Type	Sev.	Prev.	Full- Sample l_z^*	Data Removal	Robust Estimation	Known Item parameters	Full-Sample MD
Random	10%	30%	0.025	0.036	0.115	0.036	0.026
Random	10%	50%	0.018	0.034	0.111	0.035	0.021
Random	30%	30%	0.011	0.019	0.078	0.035	0.011
Random	30%	50%	0.005	0.014	0.067	0.034	0.007
Midpoint	10%	30%	0.030	0.039	0.123	0.037	0.038
Midpoint	10%	50%	0.028	0.039	0.120	0.038	0.039
Midpoint	30%	30%	0.021	0.029	0.098	0.036	0.036
Midpoint	30%	50%	0.019	0.028	0.096	0.035	0.043

Note: Sev. = Severity, Prev. = Prevalence. Both columns l_z^* , MD and Known Item Parameters are based on the full data set.

Table 2: Power for aberrant and normal response patterns

Aberrant Type	Sev.	Prev.	Full- Sample l_z^*	Data Removal	Robust Estimation	Known Item parameters	Full-Sample MD
Random	10%	30%	0.332	0.389	0.626	0.408	0.189
Random	10%	50%	0.607	0.700	0.857	0.729	0.354
Random	30%	30%	0.223	0.293	0.539	0.415	0.117
Random	30%	50%	0.390	0.562	0.787	0.726	0.183
Midpoint	10%	30%	0.169	0.201	0.405	0.213	0.022
Midpoint	10%	50%	0.278	0.342	0.570	0.381	0.015
Midpoint	30%	30%	0.120	0.150	0.340	0.216	0.016
Midpoint	30%	50%	0.158	0.208	0.422	0.388	0.014

Note: Sev. = Severity, Prev. = Prevalence. Both columns l_z^* , MD and Known Item Parameters are based on the full data set.

Table 3: Item parameter estimates for the *Feelings Scale*.

Item	a _{Full}	a _{Remove}	a _{Robust}	b _{1-Full}	b _{1-Remove}	b _{1-Robust}	b _{2-Full}	b _{2-Remove}	b _{2-Robust}
1	1.42 (0.04)	1.48 (0.05)	1.67 (0.05)	0.40 (0.03)	0.41 (0.03)	0.47 (0.02)	2.27 (0.06)	2.39 (0.06)	2.53 (0.06)
2	1.04 (0.04)	1.08 (0.04)	1.25 (0.04)	0.69 (0.03)	0.71 (0.04)	0.79 (0.03)	2.67 (0.09)	2.79 (0.09)	2.82 (0.08)
3	2.20 (0.06)	2.31 (0.07)	2.58 (0.08)	0.72 (0.02)	0.74 (0.02)	0.79 (0.02)	1.8 (0.04)	1.89 (0.04)	1.99 (0.04)
4	0.86 (0.03)	0.94 (0.03)	1.04 (0.03)	-0.76 (0.04)	-0.79 (0.04)	-0.85 (0.04)	1.01 (0.05)	0.92 (0.04)	0.82 (0.04)
5	1.24 (0.04)	1.27 (0.04)	1.41 (0.04)	-0.40 (0.03)	-0.40 (0.03)	-0.42 (0.03)	1.64 (0.05)	1.70 (0.05)	1.82 (0.05)
6	2.74 (0.08)	2.80 (0.08)	3.08 (0.09)	0.34 (0.02)	0.34 (0.02)	0.35 (0.02)	1.54 (0.03)	1.61 (0.03)	1.73 (0.03)
7	1.01 (0.03)	1.04 (0.04)	1.16 (0.04)	-0.38 (0.03)	-0.40 (0.03)	-0.40 (0.03)	2.28 (0.07)	2.38 (0.08)	2.50 (0.07)
8	0.78 (0.03)	0.85 (0.03)	0.95 (0.03)	-1.16 (0.05)	-1.19 (0.05)	-1.31 (0.05)	0.89 (0.05)	0.79 (0.04)	0.63 (0.04)
9	2.13 (0.07)	2.33 (0.08)	2.62 (0.09)	1.29 (0.03)	1.33 (0.03)	1.40 (0.03)	2.35 (0.06)	2.50 (0.06)	2.67 (0.06)
10	1.25 (0.04)	1.28 (0.05)	1.52 (0.05)	1.00 (0.04)	1.03 (0.04)	1.07 (0.03)	3.20 (0.10)	3.57 (0.12)	3.62 (0.11)
11	1.28 (0.04)	1.36 (0.04)	1.51 (0.04)	-0.56 (0.03)	-0.60 (0.03)	-0.72 (0.03)	1.29 (0.04)	1.26 (0.04)	1.25 (0.03)
12	0.86 (0.03)	0.89 (0.03)	1.04 (0.04)	0.33 (0.03)	0.36 (0.04)	0.41 (0.03)	2.90 (0.10)	3.10 (0.11)	3.13 (0.10)
13	1.96 (0.06)	2.02 (0.06)	2.25 (0.06)	0.47 (0.02)	0.47 (0.02)	0.50 (0.02)	1.86 (0.04)	1.98 (0.05)	2.11 (0.04)
14	0.97 (0.04)	0.98 (0.04)	1.19 (0.04)	0.83 (0.04)	0.87 (0.04)	0.94 (0.04)	3.38 (0.12)	3.79 (0.14)	3.82 (0.13)
15	1.37 (0.04)	1.46 (0.04)	1.60 (0.04)	-0.07 (0.02)	-0.09 (0.02)	-0.15 (0.02)	1.33 (0.04)	1.31 (0.04)	1.34 (0.03)
16	2.27 (0.06)	2.40 (0.07)	2.61 (0.07)	0.07 (0.02)	0.05 (0.02)	0.03 (0.02)	1.88 (0.04)	1.98 (0.04)	2.13 (0.04)
17	1.38 (0.04)	1.39 (0.05)	1.63 (0.05)	0.62 (0.03)	0.65 (0.03)	0.73 (0.03)	2.55 (0.07)	2.82 (0.08)	2.95 (0.08)
18	0.96 (0.03)	0.97 (0.04)	1.07 (0.04)	-0.10 (0.03)	-0.12 (0.03)	-0.12 (0.03)	2.85 (0.09)	3.03 (0.10)	3.18 (0.10)
19	2.14 (0.08)	2.43 (0.09)	2.71 (0.11)	1.55 (0.04)	1.58 (0.04)	1.68 (0.03)	2.46 (0.06)	2.56 (0.06)	2.75 (0.06)

Note: Standard Error estimates are in parenthesis.

Figure 1: l_z^* statistic and weight values for aberrant and clean responses. Note that the weights are rescaled so they sum to the total sample size.

Figure 2: Bias and RMSE for the location parameters for three estimation procedures based on I_z^* (Full = Full Data, Rem = Data Removal; Rob = Robust Estimation). The top two rows are estimated bias and the bottom two rows are RMSE. Solid and dashed lines represent different prevalence rates and gray and black colors represent different severity rates.

Figure 3: Conditional Bias and RMSE for the discrimination parameters for three estimation procedures based on I_z^* (Full = Full Data, Rem = Data Removal; Rob = Robust Estimation). The top two rows are estimated bias and the bottom two rows are RMSE. Solid and dashed lines represent different prevalence rates and gray and black colors represent different severity rates.

Appendix A

Suppose $P_{jk}(\theta)$ follow the graded response model for $j = 1, \dots, J$ items and $k = 1, \dots, K$ response options. Its first order derivative is denoted as $P'_{jk}(\theta)$. The derivation and notation in Sinharay (2016a), $l_z^*(\hat{\theta})$ can be expressed as

$$l_z^*(\hat{\theta}) = \frac{[T(\hat{\theta}) + c'_J(\hat{\theta})s_0(\hat{\theta})]}{[\sqrt{J}\zeta_J(\hat{\theta})]},$$

where

$$T(\hat{\theta}) = \sum_{j=1}^J \sum_{k=1}^K [\delta_{jk} - P_{jk}(\hat{\theta})][\tilde{w}_{jk}(\hat{\theta})],$$

$$\delta_{jk} = \begin{cases} 1, & \text{if } u_j = k \\ 0, & \text{otherwise} \end{cases}$$

$$c'_J(\hat{\theta}) = \frac{\sum_{j=1}^J \sum_{k=1}^K P'_{jk}(\hat{\theta}) \log P_{jk}(\hat{\theta})}{\sum_{j=1}^J \sum_{k=1}^K \frac{P'_{jk}(\hat{\theta})^2}{P_{jk}(\hat{\theta})}}$$

For $s_0(\hat{\theta})$, it depends on the estimation method of θ . For example, $s_0(\hat{\theta}) = 0$ for ML estimate of θ . In addition,

$$\zeta_J^2(\hat{\theta}) = \frac{1}{J} \sum_{j=1}^J \mathbf{v}'_j(\hat{\theta}) \mathbf{D}_j(\hat{\theta}) \mathbf{v}_j(\hat{\theta}),$$

where $\mathbf{D}_j(\hat{\theta})$ is the variance-covariance matrix of $(\delta_{j1}, \delta_{j2}, \dots, \delta_{jK})'$:

$$\mathbf{D}_j(\hat{\theta}) = \begin{pmatrix} P_{j0}(\hat{\theta})(1 - P_{j0}(\hat{\theta})) & -P_{j0}(\hat{\theta})P_{j1}(\hat{\theta}) & \dots & -P_{j0}(\hat{\theta})P_{jK}(\hat{\theta}) \\ -P_{j1}(\hat{\theta})P_{j0}(\hat{\theta}) & P_{j1}(\hat{\theta})(1 - P_{j1}(\hat{\theta})) & \dots & \dots \\ \dots & \dots & \dots & \dots \\ -P_{jK}(\hat{\theta})P_{j0}(\hat{\theta}) & -P_{jK}(\hat{\theta})P_{j1}(\hat{\theta}) & \dots & P_{jK}(\hat{\theta})(1 - P_{jK}(\hat{\theta})) \end{pmatrix},$$

and

$$\mathbf{v}'_j(\hat{\theta}) = (\tilde{w}_{j1}(\hat{\theta}), \tilde{w}_{j2}(\hat{\theta}), \dots, \tilde{w}_{jK}(\hat{\theta})) \text{ where } \tilde{w}_{jk}(\hat{\theta}) = \log(P_{jk}(\hat{\theta})) - c'_J(\hat{\theta}) \frac{P'_{jk}(\hat{\theta})}{P_{jk}(\hat{\theta})}.$$

Appendix B

```
# Load package and example data set  
require(mirt)  
data(Science)  
  
# Initial Model estimation  
mod <- mirt(Science, 1, optimizer = 'NR')  
  
# Calculating person fit  
per.fit <- personfit(mod, method = 'ML')$Zh  
  
# Calculating weight  
weight <- pnorm(per.fit)*nrow(Science)/sum(pnorm(per.fit))  
  
# Robust estimation  
robust.mod <- mirt(Science, 1, survey.weights = weigh, optimizer = 'NR')
```