**Investigating the Impact of Noneffortful Responses on Individual-Level Scores:
Can the Effort-Moderated IRT Model Serve as a Solution?**

Joseph A. Rios^a & James Soland^b

^aUniversity of Minnesota

^bUniversity of Virginia

Author's note

Joseph A. Rios: <https://orcid.org/0000-0002-1004-9946>

The first author holds departmental affiliation at the Department of Educational Psychology, University of Minnesota, Twin Cities. The second author is affiliated with the Department of Leadership, Foundations and Policy at the University of Virginia.

The authors would like to thank Hongwen Guo from the Educational Testing Service and Samuel Ihlenfeldt from the University of Minnesota for their helpful comments on an earlier draft.

Author contribution statement: The first author conceived of the presented idea and wrote the R syntax for the simulation analyses. The second author identified the datasets and conducted the analyses for the applied analyses. All authors interpreted findings, drafted the article, and conducted critical revisions of the article throughout the review process. Final approval of the version to be published was made by all authors.

Correspondence concerning this article should be sent to Joseph A. Rios, University of Minnesota, 56 E. River Road, 164 Education Sciences Building, Minneapolis, MN 55455.
Email: jrios@umn.edu

Abstract

Suboptimal effort is a major threat to valid score-based inferences. While the effects of such behavior have been frequently examined in the context of mean group comparisons, minimal research has considered its effects on individual score use (e.g., identifying students for remediation). Focusing on the latter context, this study addressed two related questions via simulation and applied analyses. First, we investigated how much including noneffortful responses in scoring using a three-parameter logistic (3PL) model affects person parameter recovery and classification accuracy for noneffortful responders. Second, we explored whether improvements in these individual-level inferences were observed when employing the Effort Moderated IRT (EM-IRT) model under conditions in which its assumptions were met and violated. Results demonstrated that including 10% noneffortful responses in scoring led to average bias in ability estimates and misclassification rates by as much as 0.15 *SDs* and 7% respectively. These results were mitigated when employing the EM-IRT model, particularly when model assumptions were met. However, once model assumptions were violated, the EM-IRT model's performance deteriorated, though still outperforming the 3PL model. Thus, findings from this study show that: (a) including noneffortful responses when using individual scores can lead to potential unfounded inferences and potential score misuse; and (b) the negative impact that noneffortful responding has on person ability estimates and classification accuracy can be mitigated by employing the EM-IRT model, particularly when its assumptions are met.

Keywords: noneffortful responding, rapid guessing, ability estimation, classification accuracy

Investigating the Impact of Noneffortful Responses on Individual-Level Scores: Can the Effort-Moderated IRT Model Serve as a Solution?

A fundamental assumption underlying most intended uses of test scores is that examinees employ their maximal effort. Yet, research shows this assumption is often not met, as examinees may disengage by either omitting a response or providing one without fully engaging with the item content (hereon referred to as “noneffortful” responding; Rios et al., 2017; Soland, 2018a). The latter is a common behavior that has been suggested to occur when an examinee perceives a low probability of success, is unaware or believes that their performance has minimal or no personal consequences, and/or is running out of time due to time constraints (Wise, 2017). Using rapid response time thresholds¹ to identify noneffortful responses, studies have found that as many as 25% of examinees can engage in noneffortful responding (e.g., DeMars, 2007). Such construct-irrelevant behavior can negatively impact estimates of measurement properties (e.g., item difficulty, item discrimination, test reliability; Rios & Soland, 2020) and aggregate-level inferences, such as teacher added-value estimates and subgroup comparisons (e.g., evaluations of achievement gaps; Soland 2018a; Soland 2018b). In short, low examinee effort has the potential to undermine valid uses of test scores in a host of ways.

To help reduce the bias introduced by low effort, researchers have proposed a handful of measurement models designed to address noneffortful responding. To date, the most frequently used in research and practice is the Effort-Moderated Item Response Theory (EM-IRT) model (Wise & DeMars, 2006). The EM-IRT model accounts for examinee effort by including a dichotomous indicator that specifies whether an item response is considered effortful or noneffortful. Any response classified as noneffortful is assumed an invalid indicator of examinee

¹ Low effort is typically quantified by using item response times to determine whether an examinee responded to an item so quickly, its content could not have been understood (so-called “rapid guessing”; [Wise & Kong, 2005]).

ability, which under the EM-IRT model is excluded from scoring as such responses are considered uninformative. Previous analyses of operational data support the view that these responses are uninformative as they have been found to possess proportion correct rates around chance (for a discussion, see Wise, 2017).

The popularity of this model can be attributed to three factors. For one, the EM-IRT model is general and flexible, because unlike other models (e.g., Bolt et al., 2002), it does not assume that all examinees engage in the same pattern of noneffortful responding across a given test. Second, dissimilar to competing models, the EM-IRT model is computationally simple. As an example, mixture IRT models have been shown to reduce the effects of low effort on ability estimates, but require specialized estimation procedures and longer convergence times compared to traditional models, making them less likely to be used in operational settings (Li et al., 2017). Third, the EM-IRT model has been shown to effectively recover true item and mean person parameters in the presence of noneffortful responses when compared to other commonly used models (e.g., three-parameter logistic [3PL] model; Wise & DeMars, 2006; Rios & Soland, 2020; Wise & Kingsbury, 2016). For these reasons, the EM-IRT model is quickly becoming the gold standard in measurement models that correct for low effort.

Yet, there are still gaps in our understanding of the EM-IRT model's ability to correct for low effort. One gap is that most related research assumes that the model's basic assumptions are met. Major model assumptions include that responses classified as effortful are representative of the range of item characteristics and content on the test, and that the true abilities of noneffortful responders are reflective of the sample distribution. However, prior research has demonstrated that these assumptions may be untenable in practice given that noneffortful responding has been found to be associated with item position, length, difficulty, and depth of knowledge required

(e.g., Wise, 2020). Evidence from operational tests have further illustrated that noneffortful responding can occur among examinees with low prior ability (Kuhfeld & Soland, 2020; Rios et al., 2017; Soland, 2018b). This may transpire as examinees perceive that they do not possess the requisite knowledge, skills, or abilities to successfully answer a particular item or set of items, and thus, expending effort is of minimal benefit (Wise, 2017). However, there are studies contending that any association between true ability and effort is minimal, given that effort is associated with the perceived consequences or values that examinees hold around the assessment results (for a summary, see Wise, 2015). For those that believe the assessment possesses little to no personal consequences or value, the cost of expending effort is seen to be too great when compared to the perceived benefits, which has been argued to occur independently of ability (Wise, 2015).

Another gap in the literature occurs because most studies examining the performance of the EM-IRT model (and the effects of low effort more generally) focus on assessments in which the primary inference is on comparing group means (e.g., Rios & Soland, 2020; Wise et al., 2020). In these assessments (e.g., the National Assessment of Educational Progress [NAEP] and the Program for International Student Assessment [PISA]), there are often instructions that explicitly tell examinees that the results have no personal consequences for them. This leads to a concern about low effort, and thus, is a major reason why many studies have concentrated on group-based assessments, such as NAEP and PISA (e.g., Lee & Jia, 2014). Given this focus in the literature, the effect of noneffortful responses on individual ability estimates and common uses of those individual scores is less understood. Such uses include determining instructional support (Townsend & Konold, 2010), grouping students in classrooms for differentiated instruction (Moon, 2005), and promoting students to the next grade (Jacob & Lefgren, 2009).

There is also little clarity on whether the EM-IRT model can sufficiently mitigate bias from noneffortful responses such that individual scores can still be used validly for these purposes.

To address these gaps in the literature, our simulation and empirical studies involve two objectives. First, we investigate how much including noneffortful responses in scoring affects ability parameter estimation² and classification accuracy (a common inference based on individual test scores) when estimating ability using a standard 3PL model. Second, we explore whether improvements in these individual-level inferences are observed when employing the EM-IRT model under conditions in which its assumptions are met and unmet. These objectives are addressed via the following research questions:

1. What is the impact of including noneffortful responses in scoring on:
 - a. Person parameter recovery for noneffortful responders?
 - b. Proficiency level misclassification of low performing noneffortful responders?
2. Under the same conditions, does the EM-IRT model improve person parameter and classification accuracy? If so, are improvements maintained when violating assumptions underlying the EM-IRT model?

The results from this study have the potential to provide practitioners with guidelines on how to address low effort when uses of individual scores are of interest.

Simulation

Data Generation

² While the effects on ability parameter estimation have been previously examined, those studies typically investigated the effects on group-based mean ability (e.g., Rios et al., 2017).

Noneffortful responding data were generated for a 50-item test administered to 5,000 simulees³ via a two-step process in *R*, version 3.5.0 (R Development Core Team, 2018). First, item responses were produced based on the 3PL model:

$$P_{ij}(\theta) = c_i + (1 - c_i) \left(\frac{e^{a_i(\theta_j - b_i)}}{1 + e^{a_i(\theta_j - b_i)}} \right), \quad (1)$$

where a_i is the discrimination parameter for item i , b_i is the difficulty parameter for item i , c_i is the lower asymptote of item i , and θ_j is the ability parameter for person j . Second, these data were manipulated by the conditions described in the next section to create the desired characteristics of noneffortful responding in the data matrix. This section will describe the former process.

To create item responses for each simulee, we first collected item and person generating parameters. The item parameters were taken from an operational administration of a National Assessment of Educational Progress (NAEP) math assessment. Across the 50 items, the mean discrimination, difficulty, and pseudo-guessing parameters were 1.17 ($SD = 0.43$, $\min = 0.45$, $\max = 1.91$), 0.06 ($SD = 1.19$, $\min = -2.14$, $\max = 2.17$), and 0.19 ($SD = 0.07$, $\min = 0.05$, $\max = 0.36$), respectively. Generating ability parameters were sampled from a normal distribution (more detail is provided in the next section). These generating parameters were then paired with the 3PL model to obtain item response probabilities. The probabilities obtained were compared to a random number sampled from a uniform distribution ranging from 0 to 1. For each simulee, if the random number was less than the probability, the item response was treated as correct.

Conditions

³ Prior research has demonstrated accurate 3PL model parameter estimation for a 50-item test with a minimum sample size of 1,000 when employing maximum likelihood estimation (Lord, 1968).

Below we describe how noneffortful responding was manipulated across four variables: (a) within-simulee noneffortful responding rate; (b) noneffortful responding pattern across the test; (c) percentage of noneffortful responders (i.e., simulees that engaged in at least one noneffortful response) in the sample; and (d) ability characteristics of noneffortful responders. These four variables were fully crossed producing a total of 48 conditions, with each condition replicated 100 times.

Within-simulee Noneffortful Responding Rate

The variable of greatest interest was the percentage of noneffortful responses engaged in by each noneffortful responder. The rates examined in this study were 10%, 30%, 50%, and 70% of the items seen by the simulee. All of these percentages except for 70% reflect those observed in both operational settings and prior simulation studies (DeMars & Wise, 2010; Rios et al., 2017; Wise & DeMars, 2006). The latter condition (70%) was included to evaluate estimation accuracy under an extreme rate.

Noneffortful Responding Pattern

This study examined two patterns of noneffortful responding: (a) difficulty-based (i.e., noneffortful responses occur only when simulees perceive the item to be too difficult); and (b) decreasing effort (i.e., examinees generally become less effortful as the test progresses due to cognitive fatigue). Across these two patterns, noneffortful responses were given a probability of a correct response equal to chance (.25; assuming each item possessed four response options). Specifically, for condition (a), true item probabilities for each simulee were rank ordered in descending order (ties were randomly ordered), and based on the rate of within-simulee noneffortful responding, the items with the lowest probabilities of success were replaced with the chance rate. To reflect effortfulness decreasing as the test progresses, noneffortful responding

was generated for condition (b) via a three-step process. First, the 50 items were split into five bins of 10 items each. Second, the number of noneffortful responses in each bin was specified. These numbers were determined based on the condition's specified within-simulee responding rate. As an example, when this rate was 50%, the number of noneffortful responses in each of the five bins was 3, 4, 5, 6, and 7. Third, once this distribution was determined, noneffortful responses were randomly selected in each bin and the true item probability was replaced with the chance rate. Given that items were not ordered in the simulated context, the decreasing effort pattern mimicked noneffortful responding that was unrelated to item difficulty.

Percentage of Noneffortful Responders

Although the primary interest of this study was evaluating inferences at the individual-level, it was important to examine the role of the percentage of noneffortful responders in the sample on individual ability estimates. This was of concern as prior research has shown that this percentage impacts the accuracy of item parameter estimates (e.g., Wise & DeMars, 2006), which in turn, influences ability parameter estimation accuracy (e.g., Feuerstahler, 2018). Thus, to largely reflect what is seen in operational settings (0% - 25%; Rios & Guo, 2020; Rios et al., 2017; Soland, 2018) and what has been studied in prior simulation studies (10%, 25%, 30%, 50%; Rios et al., 2017; Wise & DeMars, 2006), we examined three percentages of noneffortful responders in the sample: 10%, 30%, and 50%.

Ability Characteristics of Noneffortful Responders

There is some debate about the ability characteristics of noneffortful responding and whether this factor has an impact on ability parameter estimation accuracy (see Wise, 2015 for a discussion). To address this debate, conditions were included in which noneffortful responders possessed true abilities that were: (a) representative of the sample's ability distribution (i.e.,

noneffortful responders were sampled from across the ability continuum; hereon referred to as the representative condition); and (b) predominately of lower ability (hereon referred to as the low ability condition). Across both conditions, ability parameters were sampled separately for effortful and noneffortful simulees. For effortful simulees, ability parameters were sampled from a standard normal distribution, while for the latter group, the levels were differentiated by the sampling procedure. Specifically, for condition (a), ability parameters were sampled from a standard normal distribution. For (b), given that prior literature has demonstrated that noneffortful responding can occur more often for low ability examinees (Goldhammer et al., 2017; Rios et al., 2017; Soland & Kuhfeld, 2019), ability parameters were sampled from a normal distribution with a mean equal to -0.50 and *SD* of 0.50. This mean was chosen because Rios et al. (2017) found an average prior ability difference of 0.50 *SDs* (favoring effortful examinees) between effortful and noneffortful test takers.

Analyses

The analyses for this study consisted of estimating ability parameters from the data manipulated as described above, and then comparing these estimates for noneffortful responders to their parameters. Ability parameter estimation of the total sample (both effortful and noneffortful simulees) was conducted for two unidimensional models in the *R* package *mirt*, version 1.32.1 (Chalmers, 2012): (a) the standard 3PL (including noneffortful responses in scoring); and (b) the EM-IRT extension of the 3PL (hereon referred to as the EM-IRT model; i.e., treating all noneffortful responses as missing) model. The latter model is expressed as:

$$P_{ij}(\theta) = (SB_{ij}) \left(c_i + (1 - c_i) \left(\frac{e^{a_i(\theta_j - b_i)}}{1 + e^{a_i(\theta_j - b_i)}} \right) \right) + (1 - SB_{ij}) \left(\frac{1}{d_i} \right). \quad (2)$$

In this model, SB_{ij} is a dichotomous indicator specifying whether an item response is considered a valid ($SB_{ij} = 1$) or noneffortful ($SB_{ij} = 0$) response. If $SB_{ij} = 1$, $P_{ij}(\theta) = c_i +$

$(1 - c_i) \left(\frac{e^{a_i(\theta_j - b_i)}}{1 + e^{a_i(\theta_j - b_i)}} \right)$, which is equivalent to the standard 3PL model. However, if $SB_{ij} = 0$,

$P_{ij}(\theta) = \frac{1}{d_i}$, where d_i is roughly equal to the number of response options for item i .⁴ When

$SB_{ij} = 0$, an examinee's probability of correctly responding to item i does not depend on their underlying ability. Due to the uninformative nature of such responses to parameter estimation, scoring under the EM-IRT model excludes noneffortful responses (Wise & DeMars, 2006).⁵

Given the known issues in obtaining accurate estimates of the c parameter in the 3PL model (see Han, 2012), the c parameter was constrained to .25 for all items across the standard and EM-IRT models. For both models, the Bock-Aitkin expectation-maximization (EM) algorithm was applied to estimate item parameters. The EM convergence threshold was .0001 using the Broyden-Fletcher-Goldfarb-Shanno optimization algorithm with 61 quadrature points and the maximum number of cycles set to 1,000. Ability parameters were obtained via maximum likelihood estimation with the maximization accomplished using the Newton-Raphson algorithm. This estimation procedure was chosen given its popularity in practice, its desirable property of being asymptotically unbiased, and its documented performance of providing accurate ability estimates in the presence of aberrant responding (Kim & Moses, 2016). Standard errors of the ability parameter estimates were calculated based on the inverse of the square root of the

⁴ Correct rates of noneffortful responses may be beyond the chance level due to a function of the location of the correct answers on the assessments (see Pastor et al., 2019).

⁵ To run the EM-IRT model in R, for every item response, the user must first determine the value for SB_{ij} (for more information on this process, the reader is referred to Wise, 2017). If $SB_{ij} = 0$, the response must be treated as missing, given that it provides no information for the examinee on the underlying assessed construct. If $SB_{ij} = 1$, the observed item response is left as is. Upon completing this process for all items, a standard IRT model (e.g., the 3PL model) is estimated on the recoded data (i.e., recoding all item responses determined to be noneffortful responses as missing).

diagonal elements of the observed Fisher information matrix. Any replications that failed to converge were reanalyzed to ensure that each condition was based on the same number of converged replications.

Person Parameter Recovery

Upon estimating ability parameters, we examined person parameter recovery and classification accuracy. For ability recovery, our goal was to describe the estimation inaccuracies for noneffortful responders using both the 3PL and EM-IRT models. Our first approach involved examining both bias and root mean squared error (RMSE) for all noneffortful responders:

$$\text{Bias} = \sum_{i=1}^n (\hat{\theta}_i - \theta), \quad (3)$$

where $\hat{\theta}_i$ is the estimated parameter for replication i , θ is the known parameter, and n is the number of replications. RMSE was calculated as:

$$\text{RMSE: } \sqrt{\frac{\sum_{i=1}^n (\hat{\theta}_i - \theta)^2}{n-1}}. \quad (4)$$

At the individual score level, we examined how often bias led to statistical differences between true and estimated scores for both models. Statistical significance was determined by comparing the 95% confidence interval for simulees' estimated abilities (calculated based on the point estimate and associated standard error) to their known abilities. If the known ability was not included in the 95% confidence interval, it was categorized as statistically different.

Proficiency Level Misclassification

Proficiency level misclassification was examined by mirroring a scenario where a cut-score was used to make individual remediation decisions (Jacob & Lefgren, 2009). To generate this hypothetical, we chose a cut-score of 0.84 *SDs* below the average ability. This cut point was chosen based on the reading promotion policy implemented in Ohio during the 2017-2018 academic year, which recommended that any student scoring below -0.84 units on the theta scale

possessed limited proficiency, and thus, was eligible to receive additional reading instruction, including possible summer school attendance (American Institutes for Research, 2016; Ohio Department of Education, 2018).

The dependent variable for this analysis was the percentage of proficiency level type I errors. Specifically, we examined the extent to which simulees with a true ability parameter above -0.84 logits on the theta scale were estimated to possess an ability estimate at or below -0.84 logits. Although type II errors are also of importance, the focus on type I errors in this paper stems from research suggesting that approximately one-third of students in the United States are incorrectly placed into remedial courses (Jimenez et al., 2016). In the context of higher education, such incorrect placements costs students time and extra tuition dollars because most remedial courses do not count toward degree completion. Given the deleterious effects of misplacing students into remedial education and the fact that such classification decisions are often based on placement tests, our analyses focus solely on type I errors; however, for interested readers, type II error results can be obtained upon request from the corresponding author.

Results

As expected, the percentage of noneffortful responders in the sample had a significant impact on ability parameter recovery for the 3PL model (see Appendix A in the supplementary file; negligible effects on classification accuracy were observed for this model).⁶ In contrast, this factor negligibly influenced results for both dependent variables when employing the EM-IRT model (see Appendix B in the supplementary file). As these results echo prior findings in the literature, we provide results that are averaged across the percentage of noneffortful responders

⁴ In general, when noneffortful responders were of low ability, bias increased as the percentages of these simulees and noneffortful responses in the sample grew. This result was most likely associated with systematic inaccuracies in item parameter estimates.

and focus solely on the remaining three manipulation factors: (a) percentage of noneffortful responses; (b) ability characteristics of noneffortful responders; and (c) noneffortful responding patterns. We present these results separately by research question and outcome variable (person parameter recovery and proficiency level misclassifications).

What is the Impact of Including Noneffortful Responses in Scoring?

Person Parameter Recovery

In general, the proportion of ability estimates that differed significantly from their true scores paralleled patterns in person parameter recovery bias. Therefore, we do not discuss the former, though we do provide results in Appendix C of the supplementary file. Figure 1 presents a plot with average bias in the ability estimates for noneffortful responders on the vertical axis and percent of noneffortful responses on the horizontal (an identical figure is provided for the effortful subgroup in Appendix D of the supplementary file). Separate plots are provided by scoring model (3PL versus EM-IRT) and the ability characteristics of noneffortful responders. Furthermore, each plot disaggregates results by response pattern.

For the 3PL model, the findings demonstrated that when the ability characteristics of noneffortful responders were representative of the sample, ability estimates were consistently underestimated for both responding patterns, though the relationship between percent of noneffortful responses and bias was less pronounced for the difficulty-based pattern. For instance, the average ability for noneffortful responders was negatively biased by as little as 0.10 *SDs* when noneffortful responding was as high as 30% for the latter pattern. Meanwhile, for the decreasing effort pattern, the average bias indicated that ability was underestimated by a minimum of 0.40 *SDs* when the percent of noneffortful responses was $\geq 30\%$ and increased to approximately 1 *SD* under 70% noneffortful responding.

Next, we turn to conditions in which noneffortful responders' possessed low ability, which demonstrated distinctive directions of bias between noneffortful responding patterns. Specifically, for the difficulty-based pattern, ability was overestimated when up to 30% of noneffortful responses were present in the data matrix, with the degree of magnitude ranging from 0.15 to 0.12 *SDs*. However, when noneffortful responding increased to 50%, bias was near zero, while ability was negatively biased by 0.21 *SDs* under 70% noneffortful responses. This result indicated a beneficial effect for simulees engaging in noneffortful responding for a small percentage (up to 30%) of difficult items relative to simulees' ability, given the fact that the majority of simulees in this condition were predominately of below average ability, and thus, increased their probability of success by responding noneffortfully. However, as the number of noneffortful responses increased, their chance of correctly answering an item by guessing approached the theoretical probability, leading to underestimation of ability. In contrast, the decreasing effort pattern was observed to be associated with negligible bias (0.05 *SDs*) under 10% noneffortful responding and negative bias ranging from 0.16 to 0.55 *SDs* for noneffortful responding rates of 30% to 70%, respectively. These results are largely associated with the fact that the simulated test form possessed items with varying difficulty levels throughout, and thus, as the degree of noneffortful responding increased, the number of incorrect item responses increased in which simulees possessed high true probabilities of success, leading to greater underestimation of ability.

Proficiency Level Misclassification

In terms of proficiency level misclassifications, a large difference was noted when comparing conditions in which the ability characteristics of noneffortful responders differed, with the representative ability condition demonstrating drastically higher rates of misclassification

(Figure 2). This was due largely to the increased negative bias noted above when including noneffortful responses in scoring, which led to a growing number of simulees misclassified below the cut-point. For instance, across noneffortful responding patterns and percentages, on average type I errors were approximately three times greater when noneffortful responders possessed representative ability (23.13%) compared to when they were of predominately lower ability (7.88%; see Appendix E of the supplementary file for results based on the effortful subgroup). Furthermore, as expected, greater type I errors were generally observed for the decreasing effort pattern across noneffortful responding percentages for both ability characteristic conditions, with rates being more pronounced as the rate of noneffortful responding increased (see Figure 2). As an example, for the conditions in which noneffortful responders possessed representative ability, type I error percentages were higher by 1% to 17% for the decreasing effort responding pattern.

With that said, an interesting pattern arose for the difficulty-based pattern when noneffortful responders were of low ability. Specifically, as expected, type I errors decreased when noneffortful responding occurred on 10% to 50% of items given that ability estimates were positively biased under low rates of noneffortful responding (e.g., 10%), while the magnitude of bias approached zero for 50% of items, due to a cancellation effect of guessing on easy and difficult items relative to the simulees' ability. The interesting finding was that when noneffortful responding occurred on 70% of items, type I error rates increased to 7% (compared to 4% for the 50% condition). This was likely due to the fact that even when noneffortful responders possessed low ability, when disengaging on 70% of items, they likely guessed on a number of items in which they possessed a high probability of success, thereby leading to greater downward bias in

their ability estimates (Table 1), and ultimately, misclassification of their proficiency level as being below the cut-score.

Does the EM-IRT Model Improve Person Parameter Recovery and Classification Accuracy?

Person Parameter Recovery

Across all conditions, the EM-IRT model outperformed the 3PL model when noneffortful responders' underlying ability was representative of the sample. As shown in Figure 1, under decreasing effort, the average bias was equal to zero across all rates of noneffortful responding, and only increased to 0.07 *SDs* with progressive noneffortful rates as high as 70%. These results suggest that the EM-IRT model performed well even when violating the assumption that noneffortful responding must occur independently from item parameters/positioning.

By contrast, model performance deteriorated when violating the assumption that noneffortful responders are representative of the sample's ability distribution; however, the degree of bias was largely dependent on the responding pattern. Specifically, the EM-IRT model provided less biased ability parameter estimates than the 3PL under the decreasing effort responding pattern for larger percentages of noneffortful responding ($\geq 30\%$); though, for a small percentage of noneffortful responding (10%), bias was lower for the 3PL model (0.05 compared to 0.15 for the EM-IRT model). In addition, for the EM-IRT model, higher degrees of bias were observed under difficulty-based noneffortful responding that occurred on 30%, 50%, and 70% of items (Figure 1). These differences were greater for the EM-IRT model by 0.06 to 0.18 *SDs*.

Proficiency Level Misclassification

Across responding patterns, which were found to negligibly differ for the EM-IRT model, type I errors were roughly at or below 5% for noneffortful responding percentages as high as

70% for both ability characteristic conditions. Relative to including noneffortful responses in scoring (i.e., employing the 3PL model), the EM-IRT model demonstrated lower type I error rates across nearly all conditions, with observed reductions ranging from 1% to 45%. As can be seen in Figure 2, the largest improvements in classification accuracy for the EM-IRT model were observed for high percentages of noneffortful responding. In fact, type I errors were observed to decrease as the percentage of noneffortful responding increased, as bias for these conditions tended to be positive.

Empirical Study

To examine whether low effort affects rank orderings and classification accuracy in an empirical dataset, we examined test scores from ~500,000 7th and 8th grade students (509,369 in reading and 517,260 in math) who took Measures of Academic Progress (MAP) Growth tests in reading and mathematics in the spring of 2018. The MAP Growth assessments are computer adaptive tests that are aligned to state content standards. Test scores are reported on the Rasch Unit (RIT) scale, which is $200 + 10 \times \theta$ (θ refers to the logit scale units of the Rasch IRT model). Among the roughly half-million students in our sample, approximately 23% were black, 10% were Hispanic, 54% were white, and 50% were female.

Given that true noneffortful responses were unknown, we utilized item response times (i.e., the time that elapses between when an item is presented and answered) as a proxy of noneffortful responding. Specifically, response times were compared to response time thresholds indicating whether item responses were provided so rapidly that the examinee likely could not have understood the item's content (Wise & Kong, 2005). In this study, any response provided in less than 10% of the mean item response time was deemed a "rapid guess" and, therefore, treated

as noneffortful (Wise & Ma, 2012)⁷. This threshold-setting approach is supported by a range of validity evidence and is used in practice for MAP Growth, with thresholds based on a national norming sample (Wise & Ma, 2012). More detail on threshold setting and identification of low effort is provided in Appendix G of the supplemental material.

Using these thresholds, we investigated the effect of low effort on rank orderings and classification stability for students who responded noneffortfully on at least one item. In terms of the former, scatterplots of Rasch and EM-IRT Rasch (hereon referred to as EM-IRT) scores by subject were produced, and correlations between the two scores examined. To investigate classification stability, we used RIT scale cut scores that roughly correspond to the 5th and 10th percentiles for 7th and 8th grade students in reading and math nationwide (Thum & Hauser, 2015). These cut scores were used because percentiles in that range are utilized in practice with tests like MAP Growth to set cut scores for establishing policies on whether resources should be allotted to provide students with extra instruction (or even summer remediation; Schwerdt et al., 2017). We then examined the consistency with which students were identified as above or below the cut scores using Rasch and EM-IRT model scoring.

Results

Figure 3 shows scatterplots of Rasch versus EM-IRT scores by subject for any examinees with one or more noneffortful responses. While correlations between the two scores were quite high (.95 in reading and .97 in math), the figure makes clear that where students fell on the RIT scale was not always consistent between the two models. Specifically, Rasch scores were generally lower than EM-IRT scores, oftentimes by as many as 10 RIT. Such differences in

⁷ To mitigate the potential of overclassifying responses as noneffortful, a maximum threshold of 10 seconds was set.

scores can separate students who are well above the national average from those deemed well below (Thum & Hauser, 2015).

Meanwhile, low effort had a clear effect on which students were deemed low-performing based on RIT cut scores corresponding to the 5th and 10th percentiles. In reading, 37% of students that fell below both the 5th and 10th percentile cut score based on the Rasch model would have been above it using the EM-IRT model. In math, that figure increased to 44%. Complete cross-tabulations are available in Appendix H of the supplemental file. Overall, these results support our simulation findings by suggesting that a high rate of potential type I classifications errors can occur when including noneffortful responses in scoring.

Discussion

Our results provide several broad findings of interest to practitioners and researchers. First, including noneffortful responses when using individual scores can lead to unfounded inferences and potential score misuse. Second, the negative impact that noneffortful responding has on person ability estimates and classification accuracy can be mitigated by employing the EM-IRT model when its assumptions are met. However, once the ability characteristics of noneffortful responders were not representative of the sample's ability distribution (assumed to be normal), the EM-IRT model performed less well in most conditions.

Performance of the 3PL Model

When using the 3PL model, results demonstrated large degrees of mean bias in ability estimates even at low rates of noneffortful responding (10%). However, these results were moderated by the noneffortful responding pattern, and the ability characteristics of noneffortful responders. In general, bias in ability estimates was largest when noneffortful responders' ability was representative and noneffortful responding occurred due to decreasing effort. This likely

occurred because in the representative condition a larger percentage of noneffortful responders possessed above-average ability and more often engaged in noneffortful responding on items in which they possessed a high true probability of success, leading to a greater disconnect between estimated and true ability.

By contrast, when noneffortful responders possessed low ability, smaller bias in ability estimates and misclassification rates were observed, as many of the items that the noneffortful responders disengaged on had true probabilities of success that were low; thus, there was less of a difference between estimated and true ability. Regardless of these moderating effects, the simulation results demonstrated that including noneffortful responses in scoring can lead to deleterious results, such as misclassifications of individual proficiency levels. This was further corroborated by our applied analysis, which showed that more than one-third of examinees below the examined cut-score would be misclassified if including noneffortful responses in scoring by employing a 3PL model.

Performance of the EM-IRT Model

Across all conditions in which the ability of noneffortful responders was representative of the sample, the EM-IRT model outperformed the 3PL model in terms of both ability parameter recovery and classification accuracy. However, EM-IRT model performance deteriorated when noneffortful responders were primarily of low ability. For example, in these conditions, mean bias exceeded 0.15 *SDs*; though, this outcome was still better than including noneffortful responses in scoring and employing the 3PL model, particularly under decreasing effort. Furthermore, the EM-IRT model represented an improvement relative to the 3PL across virtually all conditions in terms of type I error classification rates. One potential reason for this is that when bias in ability estimates was present for the EM-IRT model it was consistently positive,

while bias for the 3PL was generally negative. Thus, higher type I errors were observed for the latter due to the model typically underestimating simulee ability.

Limitations and Future Research

First, misclassifications were evaluated for a testing context in which test information was not maximized around the cut-score. This approach is not ideal, as minimal standard errors would be desired when making classifications; however, in practice, this approach often does not occur (e.g., Koedel & Betts, 2010). Thus, the scenario presented in this study parallels many operational testing contexts. Second, in the simulation study, it was assumed that noneffortful responding was correctly identified with 100% accuracy. Although this approach has been taken in prior simulation research (e.g., Rios et al., 2017), findings in this paper should be interpreted as the best-case performance for the EM-IRT model. An area of future research is to examine how misclassifying noneffortful responses will impact parameter estimates for this model. Third, this study only investigated one approach to modeling noneffortful responding for ability parameter estimation. Although the EM-IRT model is one of the most popular approaches in the literature, other modeling methods have been proposed (see Wise & Kingsbury, 2016). These other approaches were not examined because they often make much stronger assumptions about how noneffortful responding occurs and unfolds over the course of a test, making their use more limited. In addition, there may be alternative scoring procedures that could be employed to downweight noneffortful responses without the requirement of log file information. One such approach is to modify maximum likelihood estimation equations by downweighting observations that are prone to response disturbances using Huber-Type weights (see Schuster & Yuan, 2011). Future research should compare these differing approaches to determine which one provides the most accurate individual-level scores. Lastly, our empirical example came from a single test.

Results might differ if other constructs and test designs were used. Such extensions are worthy of further investigation.

Recommendations for Practice

The findings from this study provide a number of implications for practitioners in educational and psychological measurement. First, including noneffortful responses in scoring can lead to deleterious results if individual scores are used for purposes like classifying examinees for instructional purposes. Therefore, if noneffortful responses are present for examinees, it is recommended that testing programs determine “...decision criteria regarding whether to include scores from individuals with questionable motivation...” (American Educational Research Association et al., 2014, p. 213). Establishing this criterion should be centered on gathering validity evidence based on test consequences for the score-based inferences of interest.

If test providers believe that providing a score is appropriate, they should make every effort to estimate ability accurately. As demonstrated in this paper, the EM-IRT model may serve as a feasible solution. However, the utility of this model is largely dependent on whether the ability characteristics of noneffortful responders are representative of the sample’s ability distribution. Previous research has attempted to evaluate this assumption by examining both prior ability measures (e.g., Rios et al., 2017) and the latent trait covariances between ability and effort (e.g., Liu et al., 2019). The former approach requires collateral information, which may not always be available to practitioners. Although the latter methods show some promise, they require certain model assumptions (e.g., noneffortful responding occurs independently of item characteristics) that may be untenable. As such, there is a need for more research on identifying

the ability characteristics of noneffortful responders prior to estimating a score when noneffortful responding is present.

References

- American Educational Research Association, American Psychological Association, & National Council on Measurement in Education (2014). *Standards for educational and psychological testing*. American Educational Research Association.
- American Institutes for Research (2016). *Recommending performance standards for Ohio's state tests*. https://oh.portal.airast.org/core/fileparse.php/3094/urlt/OST_Standard_Setting_Technical_Report_ELA_Math.pdf
- Bolt, D. M., Cohen, A. S., & Wollack, J. A. (2002). Item parameter estimation under conditions of test speededness: Application of a mixture Rasch model with ordinal constraints. *Journal of Educational Measurement*, 39, 331–348.
- Chalmers, R. P. (2012). Mirt: A multidimensional item response theory package for the R environment. *Journal of Statistical Software*, 48(6), 1-29.
- DeMars, C. E. (2007). Changes in rapid-guessing behavior over a series of assessments. *Educational Assessment*, 12, 23–45.
- DeMars, C. E., & Wise, S. L. (2010). Can differential rapid-guessing behavior lead to differential item functioning? *International Journal of Testing*, 10, 207–229.
- Feuerstahler, L. M. (2018). Sources of error in IRT trait estimation. *Applied Psychological Measurement*, 42, 359–375.
- Goldhammer, F., Martens, T., & Lüdtke, O. (2017). Conditioning factors of test-taking engagement in PIAAC: An exploratory IRT modelling approach considering person and item characteristics. *Large-scale Assessments in Education*, 5, 1–25.
- Han, K. T. (2012). Fixing the c parameter in the three-parameter logistic model. *Practical Assessment, Research, and Evaluation*, 17(1), 1-24.

- Jacob, B. A., & Lefgren, L. (2009). The effect of grade retention on high school completion. *American Economic Journal: Applied Economics*, 1, 33–58.
- Jimenez, L., Sargrad, S., Morales, J., & Thompson, M. (2016). *Remedial education: The cost of catching up*. Center for American Progress. https://cdn.americanprogress.org/content/uploads/2016/09/29120402/CostOfCatchingUp2-report.pdf?_ga=2.161063806.2143483662.1614893802-719964981.1614893802
- Kim, S., & Moses, T. (2016). *Investigating robustness of item response theory proficiency estimators to atypical response behaviors under two-stage multistage testing* (ETS RR-16-22). Educational Testing Service.
- Koedel, C., & Betts, J. (2010). Value added to what? How a ceiling in the testing instrument influences value added estimation. *Education Finance and Policy*, 5, 54–81.
- Kuhfeld, M., & Soland, J. (2020). Using assessment metadata to quantify the impact of test disengagement on estimates of educational effectiveness. *Journal of Research on Educational Effectiveness*, 13, 147–175.
- Lee, Y., & Jia, Y. (2014). Using response time to investigate students' test-taking behaviors in NAEP computer-based study. *Large-scale Assessments in Education*, 2(1), 8.
- Li, Z., Banerjee, J., & Zumbo, B. D. (2017). Response time data as validity evidence: Has it lived up to its promise and, if not, what would it take to do so. In B. D. Zumbo & A. M. Hubley (Eds.), *Understanding and Investigating Response Processes in Validation Research* (pp. 159–177). Springer.
- Liu, Y., Li, Z., Liu, H., & Luo, F. (2019). Modeling test-taking non-effort in MIRT models. *Frontiers in psychology*, 10, 145.
- Lord, F. M. (1968). An analysis of the Verbal Scholastic Aptitude Test using Birnbaum's three-parameter logistic model. *Educational and Psychological Measurement*, 28, 989-1020.

Moon, T. R. (2005). The role of assessment in differentiation. *Theory into Practice*, 44, 226–233.

Ohio Department of Education (2018). *Understanding Ohio's state tests score reports: 2017-2018*. https://oh.portal.airast.org/core/fileparse.php/3094/urlt/Understanding_State_Tests_Reports_2017-2018.pdf

R Development Core Team (2018). *R: A language and environment for statistical computing* [Computer software]. R Foundation for Statistical Computing. <https://www.R-project.org/>.

Rios, J. A., & Guo, H. (2020). Can culture be a salient predictor of test-taking engagement? An analysis of differential noneffortful responding on an international college-level assessment of critical thinking. *Applied Measurement in Education*, 33(4), 263-279.

Rios, J. A., Guo, H., Mao, L., & Liu, O. L. (2017). Evaluating the impact of noneffortful responses on aggregated scores: To filter unmotivated examinees or not? *International Journal of Testing*, 1, 74–104.

Schwerdt, G., West, M. R., & Winters, M. A. (2017). The effects of test-based retention on student outcomes over time: Regression discontinuity evidence from Florida. *Journal of Public Economics*, 152, 154–169.

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.

Soland, J. (2018a). Are achievement gap estimates biased by differential student test effort? Putting an important policy metric to the test. *Teachers College Record*, 120(12).

Soland, J. (2018b). The achievement gap or the engagement gap? Investigating the sensitivity of gaps estimates to test motivation. *Applied Measurement in Education*, 31, 312–323.

- Soland, J., & Kuhfeld, M. (2019). Do students rapidly guess repeatedly over time? A longitudinal analysis of student test disengagement, background, and attitudes. *Educational Assessment, 24*, 327–342.
- Thum, Y. M., & Hauser, C. (2015). *NWEA 2015 MAP Norms for Student and School Achievement Status and Growth*. NWEA.
- Townsend, M., & Konold, T. R. (2010). Measuring early literacy skills: A latent variable investigation of the phonological awareness literacy screening for preschool. *Journal of Psychoeducational Assessment, 28*, 115–128.
- Wise, S. L. (2015). Effort analysis: Individual score validation of achievement test data. *Applied Measurement in Education, 28*, 237–252.
- Wise, S. L. (2017). Rapid-guessing behavior: Its identification, interpretation, and implications. *Educational Measurement: Issues and Practice, 36*(4), 52–61.
- Wise, S. L. (2020). The impact of test-taking disengagement on item content representation. *Applied Measurement in Education, 33*, 1–12.
- Wise, S. L., & DeMars, C. E. (2006). An application of item response time: The effort-moderated IRT model. *Journal of Educational Measurement, 43*, 19–38.
- Wise, S. L., & Kingsbury, G. G. (2016). Modeling student test-taking motivation in the context of an adaptive achievement test. *Journal of Educational Measurement, 53*, 86–105.
- Wise, S. L., & Kong, X. (2005). Response time effort: A new measure of examinee motivation in computer-based tests. *Applied Measurement in Education, 18*, 163–183.
- Wise, S. L., & Ma, L. (2012, April 12–16). Setting response time thresholds for a CAT item pool: The normative threshold method. *Annual Meeting of the National Council on Measurement in Education, Vancouver, Canada*.

Wise, S. L., Soland, J., & Bo, Y. (2020). The (non) impact of differential test taker engagement on aggregated scores. *International Journal of Testing*, 20(1), 57-77.

Appendix A**Person Parameter Recovery Bias for the 3PL model by Percentage of Noneffortful****Responders**

Appendix B**Person Parameter Recovery Bias for the EM-IRT model by Percentage of Noneffortful****Responders**

Appendix C

Percentage of Statistically Different Ability Estimates by Model for Noneffortful Responders

Note. Within each condition, average standardized bias values are aggregated across percentage of unmotivated simulees in the sample.

Appendix D

Average Bias for Effortful Responders by Model

Note. Within each condition, average standardized bias values are aggregated across percentage of unmotivated simulees in the sample.

Appendix E

Percentage of Misclassified Simulees by Model for Effortful Responders

Note. Within each condition, average standardized bias values are aggregated across percentage of unmotivated simulees in the sample.

Appendix F

Density Plot of Theta Parameter Distributions for Noneffortful Responders by Ability Characteristic Condition

Note. The panel on the left represents the ability distribution of the representative condition, while the right-panel is reflective of the low ability condition. The horizontal line represents the cut-score used for classification accuracy.

Appendix G

Approach to Identifying Noneffortful Responses in the Empirical Data

With the increasing prevalence of computer-based tests, methods for detecting noneffortful responding have expanded (Eklöf, Pavešič, & Grønmo, 2014; Guo et al., 2016; Rios, Liu, & Bridgeman, 2014; Wise, 2015). These methods often rely on student metadata captured when a test is administered. Specifically, most detection methods use response time, or the seconds that elapse between when an item is presented and answered. Research suggests that in some cases students respond to an item so quickly that its content could not have been understood (Eklöf et al., 2014; Guo et al., 2016; Rios et al., 2014; Wise, 2015). This behavior is often termed “rapid guessing” because responses are, when effectively detected, correct at a rate no better than chance (Wise & Kong, 2005). Further, once a response is deemed a rapid guess, the student’s test effort can be summarized using the proportion of item responses on the test that were not rapid (Wise & Kong, 2005). Wise (2015) catalogued nearly a decade’s worth of research making a validity argument for the use of such metrics to identify noneffortful item responses. In our empirical analyses, we too use rapid guessing methods to identify noneffortful responses.

To identify rapid guesses, we followed the empirical approach developed by Wise and Kong (2005) and Wise and Ma (2012) to set item thresholds. These thresholds are the response time below which a response is deemed rapid. Under the Wise and Ma (2012) approach, the threshold for rapid responding is 10% of the average time students take to answer an item with a maximum time of 10 seconds. These thresholds were previously set using a nationally representative sample of students, which helps ensure they are not sensitive to our sample. We compared a subset of items deemed rapid under the Wise and Ma (2012) method to the one developed by Guo et al. (2016), but found very high overlap, likely because the frequency with

which students answer an item correctly was used as a criterion to validate the thresholds set using the Wise and Ma (2012) approach. After the item thresholds were established, each item response was flagged as a rapid response if the response time was less than the item threshold.

References

- Eklöf, H., Pavešič, B. J., & Grønmo, L. S. (2014). A cross-national comparison of reported effort and mathematics performance in TIMSS Advanced. *Applied Measurement in Education, 27*(1), 31-45.
- Guo, H., Rios, J. A., Haberman, S., Liu, O. L., Wang, J., & Paek, I. (2016). A new procedure for detection of students' rapid guessing responses using response time. *Applied Measurement in Education, 29*(3), 173-183.
- Rios, J. A., Liu, O. L., & Bridgeman, B. (2014). Identifying low-effort examinees on student learning outcomes assessment: A comparison of two approaches. *New Directions for Institutional Research, 2014*(161), 69-82.
- Wise, S. L. (2015). Effort analysis: Individual score validation of achievement test data. *Applied Measurement in Education, 28*(3), 237-252.
- Wise, S. L., & Kong, X. (2005). Response time effort: A new measure of examinee motivation in computer-based tests. *Applied Measurement in Education, 18*(2), 163–183.
- Wise, S. L., & Ma, L. (2012). Setting response time thresholds for a CAT item pool: The normative threshold method. *Annual Meeting of the National Council on Measurement in Education, Vancouver, Canada*. <https://nwea.org/content/uploads/2012/04/Setting-Response-Time-Thresholds-for-a-CAT-Item-Pool.pdf>

Appendix H

Complete Cross Tabulations of Students above and below Cut Scores on MAP Growth by Model

Reading

AboveRIT_r	AboveEM_r		Total
	0	1	
0	40,249	23,655	63,904
	62.98	37.02	100.00
	96.65	5.06	12.55
1	1,394	444,071	445,465
	0.31	99.69	100.00
	3.35	94.94	87.45
Total	41,643	467,726	509,369
	8.18	91.82	100.00
	100.00	100.00	100.00

AboveRIT_r 2	AboveEM_r2		Total
	0	1	
0	55,723	31,741	87,464
	63.71	36.29	100.00
	96.95	7.02	17.17
1	1,753	420,152	421,905
	0.42	99.58	100.00
	3.05	92.98	82.83
Total	57,476	451,893	509,369
	11.28	88.72	100.00
	100.00	100.00	100.00

Math

AboveRIT_m	AboveEM_m		Total
	0	1	
0	36,017	28,564	64,581
	55.77	44.23	100.00
	95.91	5.95	12.49
1	1,537	451,142	452,679
	0.34	99.66	100.00
	4.09	94.05	87.51
Total	37,554	479,706	517,260
	7.26	92.74	100.00
	100.00	100.00	100.00

AboveRIT_m 2	AboveEM_m2		Total
	0	1	
0	52,112	41,555	93,667
	55.64	44.36	100.00
	96.41	8.97	18.11
1	1,939	421,654	423,593
	0.46	99.54	100.00
	3.59	91.03	81.89
Total	54,051	463,209	517,260
	10.45	89.55	100.00
	100.00	100.00	100.00

Note. The left-side panel represents examinees in grade 7, while the right reflects data for grade 8.

Acknowledgements

The authors would like to thank Hongwen Guo from the Educational Testing Service and Samuel Ihlenfeldt from the University of Minnesota for their helpful comments on an earlier draft.