

Using Odds Ratios to Detect Differential Item Functioning

Applied Psychological Measurement
2018, Vol. 42(8) 613–629
© The Author(s) 2018
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/0146621618762738
journals.sagepub.com/home/apm



Kuan-Yu Jin¹ , Hui-Fang Chen²  and
Wen-Chung Wang³ 

Abstract

Differential item functioning (DIF) makes test scores incomparable and substantially threatens test validity. Although conventional approaches, such as the logistic regression (LR) and the Mantel–Haenszel (MH) methods, have worked well, they are vulnerable to high percentages of DIF items in a test and missing data. This study developed a simple but effective method to detect DIF using the odds ratio (OR) of two groups' responses to a studied item. The OR method uses all available information from examinees' responses, and it can eliminate the potential influence of bias in the total scores. Through a series of simulation studies in which the DIF pattern, impact, sample size (equal/unequal), purification procedure (with/without), percentages of DIF items, and proportions of missing data were manipulated, the performance of the OR method was evaluated and compared with the LR and MH methods. The results showed that the OR method without a purification procedure outperformed the LR and MH methods in controlling false positive rates and yielding high true positive rates when tests had a high percentage of DIF items favoring the same group. In addition, only the OR method was feasible when tests adopted the item matrix sampling design. The effectiveness of the OR method with an empirical example was illustrated.

Keywords

odds ratio, differential item functioning, Mantel–Haenszel, logistic regression, scale purification, missing data

The presence of differential item functioning (DIF) violates the assumption of measurement invariance and makes test scores incomparable across different groups of participants. In DIF assessments, participants from different groups (e.g., gender or ethnicity groups) are placed on the same metric via a matching variable, and the performances of a focal group (e.g., minority) and a reference group (e.g., majority) on the studied item are compared. The performance on the studied item is conditional on participants' ability levels, which serves as a foundation to distinguish differences in the item functioning and ability level. When matched participants

¹The University of Hong Kong, Pokfulam, Hong Kong

²City University of Hong Kong, Kowloon, Hong Kong

³The Education University of Hong Kong, New Territories, Hong Kong

Corresponding Author:

Hui-Fang Chen, Department of Applied Social Sciences, City University of Hong Kong, Tat Chee Ave., Kowloon, Hong Kong.

Email: hfchen@cityu.edu.hk

from different groups demonstrate disparate probabilities of endorsing or accurately answering the studied item, the studied item is deemed as having DIF. The present study addresses the drawbacks of two standard methods that rely on a matching variable and proposes a new, robust method to detect DIF items for practical use.

Several approaches have been developed to detect DIF items, and these methods can be classified as item response theory (IRT)-based and non-IRT-based approaches (Magis, Béland, Tuerlinckx, & De Boeck, 2010). Approaches based on IRT, such as the likelihood ratio test (Cohen, Kim, & Wollack, 1996), Lord's chi-square test (Lord, 1980), and Raju's (1988) signed area method, are implemented to compare item parameters (or derived item characteristic curves) among different groups. Once the parameter estimates of an item are significantly different between groups, this item is flagged as a DIF item. Although IRT-based approaches function well in DIF detection, their applications are somewhat challenging for practitioners who are not familiar with IRT models.

Non-IRT-based approaches to DIF assessment include the Mantel-Haenszel (MH; Holland & Thayer, 1988), logistic regression (LR; Rogers & Swaminathan, 1993), delta (Angoff & Ford, 1973), standardization (Dorans & Kulick, 1986), and Simultaneous Item Bias Test (SIBTEST) methods (Shealy & Stout, 1993). These procedures require neither specific forms of item response functions nor large sample sizes (Narayanon & Swaminathan, 1996) and demonstrate computational simplicity. Among these approaches, the MH and LR methods function well when the percentage of DIF items in a test is not high and there is no impact (no mean ability difference) between groups (French & Maller, 2007; Narayanon & Swaminathan, 1996). Moreover, these two methods are readily available in most commercial (e.g., SPSS, STATA, and SAS) or free statistical software (e.g., R).

In the MH method, examinees or participants are stratified by test scores, and item performance is compared for two groups across all strata. Suppose there are two groups of participants (e.g., boys and girls) and items are scored dichotomously (i.e., 0 and 1). A 2×2 contingency table can be created for each test score stratum on the studied item i . In stratum k ($k = 1, \dots, K$), T_{ik} denotes the total number of examinees who answered item i . The number of examinees in the reference group who answered item i correctly and incorrectly are denoted as t_{R1ik} and t_{R0ik} , respectively, and t_{F1ik} and n_{F0ik} denote the numbers of examinees in the focal group who answered item i correctly and incorrectly, respectively. Accordingly, the null hypothesis that item i does not have DIF is tested by computing the following odds ratio (OR) across the strata for item i :

$$\alpha_i = \frac{\sum_k (t_{R1ik}t_{F0ik}/T_{ik})}{\sum_k (t_{F1ik}t_{R0ik}/T_{ik})}, \quad (1)$$

which indicates the general association between the grouping variable and the item response in a series of contingency tables. To test the null hypothesis (e.g., $\alpha_i = 1$), the MH chi-square statistic is computed as follows:

$$\chi_1^2 = \frac{\left\{ \sum_k [t_{R1ik} - E(t_{R1ik})] - 0.5 \right\}^2}{\sum_k \text{Var}(t_{R1ik})}, \quad (2)$$

where

$$E(t_{R1ik}) = \frac{(t_{R1ik} + t_{R0ik}) \times (t_{R1ik} + t_{F1ik})}{T_{ik}}, \quad (3)$$

$$\text{Var}(t_{R1ik}) = \frac{(t_{R1ik} + t_{R0ik}) \times (t_{R1ik} + t_{F1ik}) \times (t_{F1ik} + t_{F0ik}) \times (t_{R0ik} + t_{F0ik})}{T_{ik}^2 (T_{ik}^2 - 1)}. \quad (4)$$

When the null hypothesis of no DIF on the studied item holds, the MH chi-square statistic will follow the chi-square distribution with one degree of freedom asymptotically. Consequently, item i is flagged as having DIF if the chi-square statistic is statistically significant.

In practice, all items in a studied test should be assessed for DIF, as we do not know which one is DIF-free. Holland and Thayer (1988) observed a special relationship between the MH method and the Rasch model-based DIF detection method (Rasch, 1960), in that the hypothesis of the MH method is equivalent to that of the Rasch model-based method when (a) the matching score includes the studied item, (b) all items but the studied item are DIF-free, and (c) both the reference and focal groups are random samples. In other words, when the Rasch model fits the data and the three conditions are met, the MH method will be the best way to detect DIF compared with other non-IRT-based approaches.

In the LR method, the test score X and the binary grouping variable G (e.g., $G = 0$ for the reference group and $G = 1$ for the focal group) are used to predict the log-odds of success over failure on the studied item i as follows:

$$\log\left(\frac{P_{i1}}{P_{i0}}\right) = \beta_0 + \beta_1 X + \beta_2 G + \beta_3 XG, \quad (5)$$

where P_{i1} and P_{i0} are the probabilities of success and failure on item i , respectively; β_1 corresponds to the influence of observed test scores on the studied item; and β_2 and β_3 refer to the group difference and the interaction between the observed test score and group membership, respectively. The studied item is deemed as having uniform DIF if $\beta_2 \neq 0$ and $\beta_3 = 0$, while it has nonuniform DIF if $\beta_3 \neq 0$ (regardless of whether $\beta_2 = 0$).

Both the MH and the LR methods are promising in detecting uniform DIF in dichotomous items when tests do not contain too many DIF items (e.g., less than 15%) and DIF magnitudes are large (Narayanan & Swaminathan, 1996; Rogers & Swaminathan, 1993). The LR method outperforms the MH method in detecting nonuniform DIF (Hidalgo & López-Pina, 2004) and/or small magnitude of DIF (Hidalgo & López-Pina, 2004). The power of the MH and LR methods increases as the sample size or DIF magnitude increases (Narayanan & Swaminathan, 1996).

A common feature of the MH and the LR methods is that different groups of examinees are placed on the same metric based on the test score. Ideally, a matching variable should consist of exclusively DIF-free items so that the matching is correct. Prior studies have shown that a test score can serve as a matching variable to yield satisfactory DIF detection only if the test score is a sufficient statistic for latent ability or is highly related to latent ability and the group ability distribution across groups are identical or extremely similar (Magis & De Boeck, 2014). In other words, the use of a matching variable is critical for DIF detection (Kopf, Zeileis, & Strobl, 2015). If a contaminated matching variable (i.e., consisting of DIF items) is used, examinees with the same ability levels will not be matched, and the subsequent DIF detection will be biased (Clauser, Mazor, & Hambleton, 1993). It has been found that the MH or LR method loses control of false positive rates (FPRs) in DIF detection when the matching variable (i.e., test score) consists of many DIF items (French & Maller, 2007).

It is challenging to identify a set of DIF-free items to serve as a clean matching variable. Various scale purification procedures have been proposed (Holland & Thayer, 1988; Kopf et al., 2015; Wang & Su, 2004); however, if the percentage of DIF items in a test is not low (e.g., more than 20%) and the DIF magnitude is large, scale purification procedures often fail, and the resulting FPRs are severely inflated (French & Maller, 2007). Even if the percentage is low, conducting scale purification procedures makes DIF assessment tedious, especially for practitioners (Magis & De Boeck, 2012, 2014). It is desirable to develop DIF detection methods that do not require specification of DIF-free items to serve as a matching variable and that can yield well-controlled FPRs and high true positive rates (TPRs), even when a test consists of a high percentage of DIF items (e.g., 30% or higher), which is the major goal of this study.

Recent developments in the outlier approach appear promising for addressing some limitations in non-IRT-based DIF detection. In the outlier approach, each item has its own value on a DIF statistic, and those items with extreme values on this statistic (termed outliers) are deemed to have DIF (Magis & De Boeck, 2012, 2014). Given the assumption that a well-constructed test should have fewer DIF items than DIF-free items, it seems reasonable to declare the outliers with DIF. For instance, Magis and De Boeck (2012) transformed the MH statistics across items into standardized z scores and classified items with extreme z scores as outliers with DIF. Because real tests usually have multiple DIF items, the sample median will be a better indicator for representing the central tendency than the sample mean, with which the outliers are identified; this method is called the robust MH method (Magis and De Boeck, 2012). It has been found that the robust MH method outperforms the conventional MH method when tests consist of DIF items.

The robust MH method is not widely used by many practitioners because it is relatively new, and the concept and computation of the MH statistic are complex. To facilitate robust DIF detection methods, a statistic that is much easier to understand and compute than the MH statistic is needed. In this study, the ORs of the reference and focal groups were proposed. Because, in practice, uniform DIF is more of concern than nonuniform DIF, and the MH method is more appropriate in detecting uniform DIF than nonuniform DIF, the authors focus on uniform DIF in this study. The study proceeds as follows: an introduction to the OR method is provided; three simulation studies are described, which were conducted to evaluate the performance of the OR method under various conditions; an empirical example is given to demonstrate the feasibility and advantages of the robust method; and conclusions and suggestions for future study are provided.

The OR Method

Conventionally, item difficulty (easiness) is defined as the passing rate for a group of examinees on an item. Let n_{R1i} and n_{R0i} be the number of examinees in the reference group who answer item i correctly and incorrectly, respectively, and n_{F1i} and n_{F0i} be the number of examinees in the focal group who answer item i correctly and incorrectly, respectively. The passing rate of item i is $n_{R1i}/(n_{R0i} + n_{R1i})$ for the reference group and $n_{F1i}/(n_{F0i} + n_{F1i})$ for the focal group. The passing rate can be also represented as the odds of success, n_{R1i}/n_{R0i} and n_{F1i}/n_{F0i} , for the reference group and the focal group, respectively. Let $\hat{\lambda}_i$ denote the logarithm of the OR of success over failure on item i for the reference group and the focal group:

$$\hat{\lambda}_i = \log \left(\frac{n_{R1i}/n_{R0i}}{n_{F1i}/n_{F0i}} \right), \quad (6)$$

which follows a normal distribution asymptotically (Agresti, 2002), with a mean of λ and standard deviation of

$$\sigma(\hat{\lambda}_i) = (n_{R1i}^{-1} + n_{R0i}^{-1} + n_{F1i}^{-1} + n_{F0i}^{-1})^{1/2}. \tag{7}$$

When data follow the Rasch model and there is no DIF in the test, λ represents the mean ability difference between the reference and focal groups (often referred to as impact). A positive λ indicates that the reference group has a higher mean than the focal group, and a negative λ indicates the reverse.

When tests are perfect (not containing any DIF item), every $\hat{\lambda}_i$ ($i = 1, \dots, I$) is an unbiased estimator of λ . The sample mean $\tilde{\lambda} = \sum_{i=1}^I \hat{\lambda}_i / I$ is more efficient than individual $\hat{\lambda}_i$ values in estimating λ because its sampling variance is smaller. When item i has DIF, the expected value of $\hat{\lambda}_i$ will not be equal to λ . In other words, if $\hat{\lambda}_i$ is far away from λ , item i is deemed as having DIF. Because λ is unknown in reality, the sample mean $\tilde{\lambda}$ was used instead. If $\hat{\lambda}_i \pm z_{\alpha/2} \times \sigma(\hat{\lambda}_i)$ does not contain $\tilde{\lambda}$, item i will be deemed as having DIF. If the studied item is the only item that might have DIF, this approach will be appropriate; however, tests are not always perfect, and they usually contain multiple DIF items, such that the sample mean $\tilde{\lambda}$ is no longer an unbiased estimator of λ . The sample median $\tilde{\lambda}$ was used then to replace the sample mean $\tilde{\lambda}$ to estimate λ , which is more resistant to multiple DIF items. Specifically, if $\hat{\lambda}_i \pm z_{\alpha/2} \times \sigma(\hat{\lambda}_i)$ does not include $\tilde{\lambda}$, item i will be deemed as having DIF.

As an example, a dataset was generated with 100 examinees in each of the reference and focal group, and they answered 10 dichotomous items. Items 1 to 7 were generated as non-DIF items, and Items 8, 9, and 10 were DIF items, which favors the reference group by 1 logit. There was no impact (i.e., mean ability difference) between the two groups. For Item 1, Examinees 21 and 28 in the reference and focal groups had a correct answer, respectively. According to Equations 6 and 7, the authors found that

$$\hat{\lambda}_1 = \log\left(\frac{n_{R11}/n_{R01}}{n_{F11}/n_{F01}}\right) = \log\left(\frac{21/79}{28/72}\right) = -0.380,$$

$$\sigma(\hat{\lambda}_1) = (n_{R11}^{-1} + n_{R01}^{-1} + n_{F11}^{-1} + n_{F01}^{-1})^{1/2} = (21^{-1} + 79^{-1} + 28^{-1} + 72^{-1})^{1/2} = 0.331.$$

$$\hat{\lambda}_1 \pm 1.96 \times \sigma(\hat{\lambda}_1) = -1.029, 0.269.$$

$\hat{\lambda}_i$, $\sigma(\hat{\lambda}_i)$, and $\hat{\lambda}_i \pm z_{\alpha/2} \times \sigma(\hat{\lambda}_i)$ were computed for the other items, and the sample median $\tilde{\lambda}$ (the mean of $\hat{\lambda}_2$ and $\hat{\lambda}_7$) was 0.022. When $\alpha = .05$, if an item's 95% confidence interval does not cover 0.022, the corresponding item was deemed as having DIF. Finally, Items 8, 9, and 10 are successfully flagged as DIF (favoring the reference group; see Table 1 for details).

If the DIF detection is based on the sample mean, the accurate classification rate was reduced. In this case, the sample mean was 0.363. Items 1, 3, 4, 8, and 10 had a 95% confidence interval not covering 0.363; thus, they were deemed as having DIF, and the other items were deemed as non-DIF. Only six out of 10 items were correctly classified, and the correct classification rate was 60%. It is noticeable that using the median over the mean is not salient under the balanced DIF conditions, in which some DIF items favor the reference group and the other items favor the focal group; thus, the DIF effects are canceled out between groups. No matter how test developers try to balance the effects of potential DIF items, it is unlikely that the DIF effects will be canceled out completely between groups.

Table 1. An Example of the OR Method Where Items 8, 9, and 10 of the 10 Items Favor the Reference Group.

Items	Reference group		Focal group		OR statistics		
	Success	Fail	Success	Fail	$\hat{\lambda}_i$	$\sigma(\hat{\lambda}_i)$	$\hat{\lambda}_i \pm z_{\alpha/2} \times \sigma(\hat{\lambda}_i)$
1	21	79	28	72	-0.380	0.331	-1.029, 0.269
2	28	72	28	72	0.000	0.315	-0.617, 0.617
3	74	26	80	20	-0.340	0.338	-1.002, 0.322
4	52	48	63	37	-0.452	0.288	-1.016, 0.112
5	18	82	13	87	0.385	0.395	-0.389, 1.159
6	33	67	35	65	-0.089	0.299	-0.675, 0.497
7	35	65	34	66	0.044	0.298	-0.540, 0.628
8	42	58	12	88	1.670	0.368	0.948, 2.392
9	71	29	52	48	0.851	0.298	0.232, 1.399
10	58	42	16	84	1.981	0.340	1.315, 2.647

Note. $\alpha = .05$. OR = odds ratio.

The OR method is easy to understand and implement; there is no need to identify DIF-free items. The performance of the OR method especially depends on two conditions. First, a long test is required to justify the OR method. The longer the test, the smaller the sampling variation in $\hat{\lambda}_i$, making the OR method more reliable. Second, as a basic concept of outliers, the number of DIF items should not exceed the number of DIF-free items; otherwise, the classification of outliers becomes tricky.

The OR method has two distinct features that differentiate it from the conventional MH and robust MH methods. The first is with regard to stratification. The OR method does not require stratification, and the log OR is directly computed for all items (see Equation 6). In contrast, both the conventional and robust MH methods assign examinees to strata based on examinees' test scores, calculate an OR in each stratum, and generate a composite OR across strata (see Equation 1). The second feature is the principle of identifying DIF items. In the conventional MH method, an item is flagged as having DIF if its OR significantly deviates from 1 (or equivalently, the MH chi-square statistic is statistically significant). In the OR method, an item whose log OR is an outlier is deemed as having DIF. Therefore, although both methods use ORs to detect DIF, they differ in nature.

The OR method has several advantages. First, it is easy and simple for practitioners. The computation task in the OR method is minimized, and it can easily be implemented with a portable calculator. What is more, interpreting its result is intuitive and does not require any advanced statistical knowledge or techniques. Second, unlike conventional DIF assessment approaches, the OR method does not require matching variables; thus, there is no need for the identification of DIF-free items. The OR method uses the median of the log ORs of all items as the reference point for DIF detection. In contrast, the MH and LR methods rely on test scores to match examinees, which is biased when tests have multiple DIF items. Even if tests do not contain any DIF items, test scores cannot match examinees well when the impact is large (Narayanan & Swaminathan, 1996). Third, the OR method can accommodate missing data much more easily than the conventional or robust MH method. When there are missing data, the test score is usually no longer a valid matching variable, so the conventional or robust MH method will be adversely affected (Robitzsch & Rupp, 2009). In contrast, the OR method does not rely on test score and will be less affected by missingness. Fourth, scale purification procedures can easily be incorporated into the OR method; all that is necessary is the precomputation

of the sample median based on presumably DIF-free items. Take a test with 20 items as an example. Suppose Items 1 to 5 are classified as having DIF by the OR method; a new sample median $\tilde{\lambda}$ can be computed then based on Items 6 to 20. Again, for each item, the authors checked whether $\hat{\lambda}_i \pm z_{\alpha/2} \times \sigma(\hat{\lambda}_i)$ contains the new $\tilde{\lambda}$. If not, the corresponding item is deemed as having DIF. Suppose that, this time, Items 1 to 4 are classified as having DIF; a new $\tilde{\lambda}$ based was computed on Items 5 to 20, and for each item the authors checked whether $\hat{\lambda}_i \pm z_{\alpha/2} \times \sigma(\hat{\lambda}_i)$ contains the new $\tilde{\lambda}$. Again, suppose that the same set of items (Items 1 to 4) is classified as having DIF. In this case, the scale purification stops; otherwise, a new $\tilde{\lambda}$ will be computed, and the DIF detection continues until either the same set of items is classified as having DIF in two successive iterations or the maximum number of iterations is reached. These arguments are verified by the simulations below.

Simulation Study I

Design

Study 1 concerned the performance of the OR method for complete data. A total of 1,000 examinees answered 20 dichotomous items. Five factors were manipulated as follows: (a) equal and unequal sample sizes of the reference and focal groups: 500/500 and 800/200; (b) percentages of DIF items: 0%, 10%, 20%, 30%, and 40%; (c) DIF patterns: balanced and unbalanced; (d) impact: 0 and 1; and (e) purification procedure: with or without. Item responses were generated from the Rasch model. Item difficulties were generated from a uniform distribution between -1.5 and 1.5 . The abilities of the reference group were generated from $N(0, 1)$. When impact = 0, the abilities of the focal group were also generated from $N(0, 1)$; when impact = 1, they were generated from $N(-1, 1)$. The differences in the item difficulty for DIF items were set at a constant of 0.5 logits. Under the unbalanced DIF conditions, all DIF items were set to favor the reference group; under the balanced DIF conditions, half of the DIF items favored the reference group and the other half favored the focal group. It has been shown that the MH and LR methods perform well under balanced DIF conditions because the DIF effects are canceled out between groups, as if there are no DIF items in the test. In contrast, they perform poorly under unbalanced DIF conditions because the test score is seriously contaminated (Wang & Su, 2004). Because the OR method does not rely on test scores to match examinees, it was anticipated that the OR method would not suffer much from unbalanced DIF patterns.

As in other DIF studies, the nominal level was set at .05. For completeness, scale purification procedures were also implemented on the three methods. Unlike the scale purification procedures for the MH and LR methods, which involve intensive computation and are often implemented in specific computer programs, the scale purification procedure for the OR method involves only recalculation of the sample median and can be easily implemented using portable calculators.

A total of 100 replications were carried out under each condition. The performance of the OR method was evaluated and compared with the performances of the MH and LR methods in terms of the FPR, in which a DIF-free item was misclassified as having DIF, and the TPR, in which a DIF item was correctly classified as having DIF. The MH and LR methods were implemented via the *difMH* and *difLogistic* functions in the *difR* package (Magis et al., 2010), and the OR method was implemented in R version 3.2.5 (R Core Team, 2016) and is available upon request.

Table 2. Averaged FPR and TPR (%) Under the Conditions With Impact = 0 in Simulation I.

	FPR						TPR					
	500/500			800/200			500/500			800/200		
	OR	LR	MH	OR	LR	MH	OR	LR	MH	OR	LR	MH
0%	3.20 (3.45)	4.75 (3.75)	4.10 (3.80)	3.55 (3.70)	4.70 (4.90)	3.50 (3.65)	—	—	—	—	—	—
Balanced DIF												
10%	3.33 (3.50)	5.05 (4.56)	4.17 (4.06)	3.50 (3.44)	5.28 (5.06)	4.06 (4.17)	87.50 (87.50)	94.50 (91.50)	94.00 (89.50)	79.00 (78.50)	85.50 (83.00)	84.50 (79.50)
20%	3.25 (4.00)	4.81 (4.63)	3.94 (4.06)	3.56 (3.69)	5.19 (5.06)	4.00 (3.63)	87.75 (87.25)	92.25 (89.25)	91.75 (88.50)	68.75 (70.00)	79.75 (76.00)	74.75 (70.75)
30%	3.14 (3.57)	5.50 (5.21)	4.14 (4.00)	2.86 (3.29)	4.71 (4.93)	3.93 (4.00)	83.50 (82.50)	89.33 (85.67)	87.50 (82.33)	70.67 (70.50)	78.33 (75.17)	77.00 (71.83)
40%	3.00 (3.58)	4.25 (4.83)	3.67 (3.67)	2.58 (3.08)	4.58 (5.08)	3.17 (3.42)	87.50 (86.75)	93.75 (90.38)	91.63 (88.50)	74.75 (74.38)	82.88 (76.75)	80.50 (73.38)
Unbalanced DIF												
10%	3.22 (3.11)	6.39 (4.50)	5.22 (3.78)	4.28 (4.50)	6.33 (5.78)	5.00 (4.89)	86.00 (87.00)	88.50 (90.50)	84.50 (88.50)	75.50 (77.00)	76.00 (81.50)	72.50 (78.00)
20%	4.44 (4.25)	11.00 (5.31)	9.44 (4.56)	4.75 (4.31)	8.50 (6.13)	7.63 (4.25)	80.50 (83.75)	80.25 (89.25)	74.50 (85.00)	61.00 (65.50)	60.75 (69.50)	56.25 (67.75)
30%	6.29 (5.14)	17.50 (6.14)	15.57 (5.43)	5.93 (5.93)	12.21 (7.58)	10.07 (6.71)	73.00 (82.00)	66.50 (85.17)	63.33 (82.33)	47.83 (54.83)	48.17 (58.17)	42.50 (52.50)
40%	10.75 (7.42)	26.08 (10.17)	22.42 (8.00)	11.17 (11.00)	20.75 (16.75)	18.42 (13.58)	59.88 (72.75)	54.63 (74.50)	50.88 (73.13)	34.75 (42.50)	36.75 (43.88)	31.50 (38.63)

Note. The values in parentheses are obtained with the purification procedure. Inflated FPRs ($\geq 7.5\%$) are marked in bold. FPR = false positive rate; TPR = true positive rate; OR = odds ratio; LR = logistic regression; MH = Mantel-Haenszel; DIF = differential item functioning.

Results

Due to space constraints, the FPR and TPR for individual items were not shown; rather, the averaged FPR across DIF-free items and averaged the TPR across DIF items were computed. Table 2 shows the averaged FPRs and TPRs under all conditions when $\text{impact} = 0$. As expected, all methods yielded well-controlled FPRs under the no-DIF and balanced DIF conditions, although the OR and MH methods were slightly conservative. Under the unbalanced DIF conditions, the MH and LR methods yielded inflated FPRs when tests had 20% or more DIF items; fortunately, the scale purification procedures could bring the inflation back to normal if the DIF items represented no more than 30%. The scale purification did not work in the MH and the LR approaches when a test included 40% or more DIF items. In contrast, the OR method yielded slightly inflated FPRs only when tests had 40% or more DIF items, and the inflation was back to normal when the scale purification procedure was incorporated.

When the FPRs for all methods were well controlled, as under the no-DIF and balanced DIF conditions, in general, the TPRs were slightly higher in the LR method than in the MH and OR methods. Under the unbalanced DIF conditions, the FPRs were inflated and the TPRs were deflated for the MH and LR methods, while the TPRs were lower than those for the OR method. In other words, once the FPRs were inflated, the TPRs were reduced, and the MH and LR methods suffered a greater loss in TPRs than the OR method did. In addition, the TPRs were higher in equal sample sizes than in unequal sample sizes given the fixed total sample size.

Table 3 summarizes the averaged FPRs and TPRs when $\text{impact} = 1$. The findings were comparable with those in Table 2 when $\text{impact} = 0$, except that the inflation of FPRs and deflation of TPRs in the MH and LR methods under the unbalanced DIF conditions became worse when $\text{impact} = 1$ than when $\text{impact} = 0$.

Simulation Study 2

Design

Study 2 mimicked a horizontal equating design and aimed to demonstrate the advantages of the OR method over the MH and LR methods in accommodating missing data by design. There were 30 items in two 20-item booklets. The first booklet contained Items 1 to 20, and the second booklet contained Items 1 to 10 and 21 to 30. Thus, Items 1 to 10 were common between the two booklets. There were 600 examinees each in the reference and focal groups. In each group, half of the examinees received the first booklet, and the other half received the second booklet. Each examinee answered only 20 out of 30 items, so the missing rate was 33% by design. The data generation procedures and the three manipulated factors, including percentage of DIF items, impact, and purification procedure, were the same as those in Study 1. For example, in the condition of 10% DIF items, Items 10, 20, and 30 had DIF, so that each booklet included two DIF items and 18 DIF-free items. The DIF magnitude was set at 0.5 logits, and all DIF items favored the reference group. An examinee's test score on his or her booklet (ranging from 0-20) was used as the matching variable in the MH and LR methods. A total of 100 replications were carried out under each condition. FPRs and TPRs were the outcome variables.

Results

Table 4 summarizes the averaged FPRs and TPRs. As in Simulation Study 1, the OR method yielded acceptable FPRs and satisfactory TPRs when tests had 30% or fewer DIF items, but it had lower TPRs in assessing the unique items than the common items because the number of

Table 3. Averaged FPR and TPR (%) Under the Conditions With Impact = 1 in Simulation I.

	FPR						TPR					
	500/500			800/200			500/500			800/200		
	OR	LR	MH	OR	LR	MH	OR	LR	MH	OR	LR	MH
0%	2.70 (2.80)	4.50 (4.50)	3.40 (3.30)	2.75 (3.15)	4.60 (4.90)	3.65 (3.60)	—	—	—	—	—	—
Balanced DIF												
10%	3.72 (4.11)	5.67 (5.33)	4.56 (4.50)	3.56 (3.78)	4.56 (4.56)	3.17 (3.00)	85.00 (84.50)	89.00 (84.00)	86.50 (83.00)	73.00 (73.00)	79.50 (75.50)	74.50 (71.00)
20%	3.25 (3.56)	4.69 (4.81)	3.81 (3.63)	3.13 (3.44)	5.19 (5.50)	4.31 (4.13)	83.50 (83.25)	86.75 (83.25)	85.50 (81.25)	66.00 (65.75)	70.00 (66.25)	65.00 (61.00)
30%	3.93 (4.36)	5.29 (5.50)	4.29 (4.43)	3.14 (4.14)	5.36 (6.29)	3.79 (3.71)	90.83 (90.50)	92.83 (87.83)	91.83 (86.50)	69.17 (68.33)	77.50 (71.33)	72.67 (66.83)
40%	3.42 (4.17)	4.92 (4.75)	4.08 (3.83)	3.33 (4.17)	5.67 (6.42)	4.00 (5.00)	87.63 (86.13)	90.25 (86.50)	88.00 (84.50)	70.50 (69.75)	77.13 (69.38)	71.50 (64.88)
Unbalanced DIF												
10%	3.56 (3.50)	5.78 (4.83)	4.94 (3.83)	3.50 (4.44)	7.33 (6.22)	5.11 (4.56)	84.50 (85.00)	82.50 (83.00)	80.00 (83.50)	69.00 (71.50)	71.50 (73.00)	67.50 (69.00)
20%	5.19 (4.56)	10.63 (5.69)	8.31 (4.00)	3.88 (4.44)	8.19 (6.00)	5.56 (4.13)	76.75 (82.00)	69.00 (78.75)	67.50 (79.25)	58.50 (61.25)	53.00 (56.25)	48.50 (53.75)
30%	7.86 (5.57)	17.86 (8.86)	13.57 (6.00)	6.79 (6.64)	14.14 (9.00)	10.79 (7.35)	69.50 (76.33)	58.67 (75.83)	56.00 (75.17)	49.00 (56.83)	46.83 (59.50)	42.50 (53.83)
40%	11.92 (8.17)	25.17 (12.42)	20.25 (9.00)	11.00 (9.58)	20.92 (17.08)	14.25 (10.25)	56.50 (68.25)	45.88 (64.25)	43.25 (62.38)	36.75 (45.25)	33.25 (38.75)	30.75 (39.75)

Note. Values in the parentheses are obtained with the purification procedures. Inflated FPRs ($\geq 7.5\%$) are marked in bold. FPR = false positive rate; TPR = true positive rate; OR = odds ratio; LR = logistic regression; MH = Mantel-Haenszel; DIF = differential item functioning.

Table 4. Averaged FPR and TPR (%) in Study 2.

	FPR						TPR					
	Common items			Unique items			Common items			Unique items		
	OR	LR	MH	OR	LR	MH	OR	LR	MH	OR	LR	MH
Impact = 0												
0%	3.70 (4.10)	5.80 (5.30)	4.60 (4.90)	4.65 (4.85)	5.45 (5.50)	0.00 (0.00)	—	—	—	—	—	—
10%	3.78 (4.22)	6.67 (5.44)	5.56 (4.78)	4.00 (4.11)	5.72 (5.28)	0.00 (0.00)	83.00 (84.00)	86.00 (87.00)	86.00 (85.00)	62.00 (63.00)	65.00 (64.50)	3.00 (3.00)
20%	4.63 (4.63)	12.75 (5.00)	11.50 (6.63)	6.06 (5.50)	9.19 (5.38)	0.00 (0.00)	87.50 (90.50)	87.00 (91.50)	86.50 (87.50)	61.00 (65.50)	61.00 (71.25)	0.25 (0.75)
30%	7.71 (8.14)	21.00 (6.71)	19.51 (12.29)	5.79 (5.57)	12.36 (5.36)	0.00 (0.00)	81.00 (86.00)	77.33 (89.67)	73.33 (78.67)	49.00 (57.33)	48.00 (64.83)	0.33 (0.50)
40%	15.17 (11.17)	31.17 (13.17)	27.50 (20.00)	9.92 (7.92)	17.58 (10.33)	0.00 (0.00)	63.75 (73.75)	58.50 (76.50)	56.25 (60.00)	34.25 (49.13)	35.00 (51.38)	0.13 (0.38)
Impact = 1												
0%	3.90 (4.20)	5.20 (4.80)	4.20 (4.10)	3.45 (3.50)	5.45 (5.50)	0.00 (0.00)	—	—	—	—	—	—
10%	5.33 (5.22)	7.56 (6.56)	6.56 (5.67)	4.50 (4.33)	5.67 (4.89)	0.00 (0.00)	93.00 (95.00)	88.00 (89.00)	87.00 (85.00)	57.00 (59.50)	51.50 (53.50)	2.50 (2.00)
20%	8.00 (7.25)	13.25 (7.00)	10.63 (8.88)	5.94 (5.69)	8.06 (6.94)	0.06 (0.31)	84.50 (86.50)	80.00 (86.00)	78.50 (78.50)	53.25 (57.75)	49.75 (59.00)	0.75 (1.50)
30%	10.59 (7.57)	16.56 (7.00)	13.29 (8.71)	7.50 (5.79)	11.86 (5.79)	0.50 (0.21)	78.00 (85.67)	72.33 (89.67)	71.33 (74.67)	47.67 (56.83)	41.67 (59.17)	1.83 (3.17)
40%	20.00 (17.67)	30.00 (14.00)	25.00 (20.50)	10.33 (10.00)	16.33 (9.33)	0.08 (0.00)	56.75 (64.50)	50.75 (72.25)	47.25 (49.00)	34.63 (43.13)	28.88 (43.25)	0.00 (0.13)

Note. Values in the parentheses are obtained with the purification procedure. Inflated FPRs ($\geq 7.5\%$) are marked in bold. FPR = false positive rate; TPR = true positive rate; OR = odds ratio; LR = logistic regression; MH = Mantel-Haenszel.

examinees answering the unique items was only half the number of examinees answering the common items. The LR method yielded significantly inflated FPRs for both the common and unique items, but the purification procedure reduced the inflation. The MH method yielded FPRs close to zero under all conditions, which might be because the *difMH* function did not handle missing data properly. Thus, using the MH method was not recommended when data are missing by design. In sum, the OR method is quite robust to missing data in yielding well-controlled FPRs and high TPRs.

The test score of the common items (ranging from 0-10) was also used as the matching variable in the LR method, which resulted in slightly better control of FPRs than using the booklet test score as the matching variable. Nevertheless, this matching method may not always be applicable, as common items across booklets are not always available in practice. For instance, large-scale assessments often adopt a matrix sampling design of booklets, so that examinees answer a subset of items and there is no common item across all examinees (please see “Simulation Study 3” section for details).

Simulation Study 3

In large-scale tests, such as the Program for International Student Assessment (PISA) and Trends in International Mathematics and Science Study (TIMSS), the matrix sampling design is often adopted to assemble multiple booklets. Simulation Study 3 mimicked such a design and evaluated how the three methods would perform. A total of 30 items were grouped into three item blocks as follows: Items 1 to 8 in Block 1, Items 9 to 18 in Block 2, and Items 19 to 30 in Block 3. The first booklet was composed of Blocks 1 and 2, the second booklet was composed of Blocks 2 and 3, and the third booklet was composed of Blocks 1 and 3. This matrix sampling had two features: (a) there was no item common to all three booklets, and (b) the numbers of items were different for different booklets. Both the reference group and the focal group had 600 examinees, in which one third of the examinees took each of the three booklets. The procedure for data generation and the three manipulated factors, including percentage of DIF items, impact, and purification procedure, were identical to those in Simulation Study 2. Each item block contained zero to four DIF items. The DIF magnitude was set at 0.5 logits, and all DIF items favored the reference group. Because no items common to all three booklets were available, an examinee’s test score on his or her booklet was used as the matching variable in the MH and LR methods. A total of 100 replications were carried out under each condition. FPRs and TPRs were the outcome variables.

Results

Table 5 summarizes the averaged FPRs and TPRs for all conditions in Simulation Study 3. As in the findings for Simulation Study 2, the OR method yielded acceptable FPRs and satisfactory TPRs when tests had 30% DIF items or less, and the scale purification procedure was helpful in reducing FPRs when tests had 40% or more DIF items. The LR method yielded inflated FPRs when tests had 20% or more DIF items, and the scale purification procedure could not bring FPRs back to the 5% nominal level when impact = 1. Compared with the OR and LR methods, the MH method performed extremely poor, which may have been due to the problem in handling missing data in the *difMH* function. Overall, the findings were in alignment with those of Simulation Study 2, showing that the OR method was efficient in detecting DIF when data were missing by design.

Table 5. Averaged FPR and TPR (%) in Study 3.

	FPR			TPR		
	OR	LR	MH	OR	LR	MH
Impact = 0						
0%	3.70 (3.90)	4.77 (4.97)	0.00 (0.00)	— —	— —	— —
10%	4.48 (4.67)	6.93 (5.48)	0.11 (0.04)	77.67 (80.00)	82.67 (85.00)	21.67 (23.67)
20%	4.54 (4.17)	8.25 (5.04)	0.13 (0.08)	71.83 (77.17)	74.67 (84.83)	15.50 (16.83)
30%	6.71 (4.90)	13.10 (5.14)	0.19 (0.19)	67.00 (77.22)	63.11 (79.78)	8.44 (9.56)
40%	12.59 (9.94)	21.83 (8.88)	0.61 (0.50)	46.67 (59.33)	47.25 (67.58)	4.67 (4.92)
Impact = 1						
0%	3.93 (3.93)	6.13 (8.93)	3.10 (2.60)	— —	— —	— —
10%	4.07 (4.11)	4.89 (7.11)	4.96 (3.74)	70.33 (72.00)	79.67 (83.00)	28.67 (33.00)
20%	5.04 (4.46)	7.71 (8.38)	3.67 (2.33)	70.83 (75.67)	71.50 (82.50)	18.83 (21.50)
30%	7.19 (6.38)	9.43 (8.09)	1.90 (1.52)	57.78 (66.33)	49.11 (74.22)	10.44 (13.00)
40%	12.00 (9.94)	10.78 (49.67)	15.44 (9.00)	45.67 (54.25)	55.00 (84.92)	27.25 (28.08)

Note. The values in parentheses were obtained with the purification procedure. Inflated FPRs ($\geq 7.5\%$) are marked in bold. FPR = false positive rate; TPR = true positive rate; OR = odds ratio; LR = logistic regression; MH = Mantel-Haenszel.

An Empirical Example of PISA 2015

The PISA is a large-scale assessment of students' performance in science (the major domain in 2015), reading, and mathematics (the two minor domains). Items in the three domains are assembled into several blocks, and each booklet is composed of four item blocks. This study used data from a sample of 14,530 Australian students (50.7% boys and 49.3% girls) taking the computer-based assessment (CBA) from PISA 2015. In the CBA design, a student is assigned to one of 66 booklets. The data matrix is rather sparse, as on average, 2,000 scores were observed for each item.

Only responses to the dichotomous items were analyzed. There were 46, 50, and 121 dichotomous items in the reading, mathematics, and science tests, respectively. For boys, the average rates of accuracy were 57.3%, 51.5%, and 56.2% for the three subjects, respectively; for girls, they were 60.9%, 49.7%, and 54.7%, respectively. At face value, the results suggest that boys outperformed girls in mathematics and science, while girls outperformed boys in reading.

The MH, LR, and OR methods with scale purification procedures were utilized to investigate gender DIF in the three subjects. Boys and girls were treated as the focal and reference groups, respectively, in the DIF analysis. Note that the number of items varied among the booklets. For example, in the reading test, the number of items in the 66 booklets was between 0 and 20. This implied that using the test score as the matching variable in the MH or LR method would be problematic.

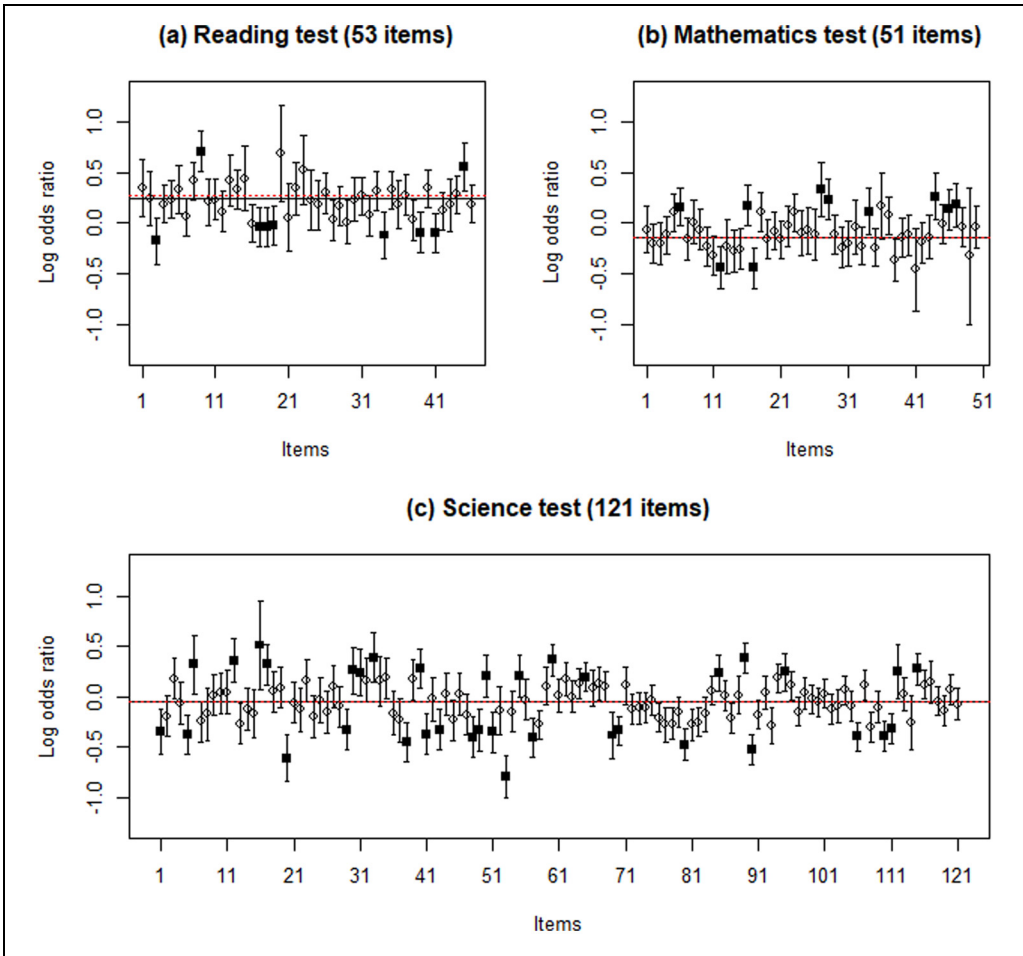


Figure 1. Log odds ratios and 95% confidence intervals for three tests in the PISA 2015.

Note. Circles and squares refer to non-DIF and DIF items, respectively; solid and dotted lines are the medians and means across non-DIF items, respectively. PISA = Program for International Student Assessment; DIF = differential item functioning.

The MH method yielded different results from those found using the LR and OR approaches. The MH, LR, and OR methods identified zero, 13, and 13 DIF items, respectively, in the reading test; zero, 16, and 15 DIF items, respectively, in the mathematics test; and zero, 52, and 59 DIF items, respectively, in the science test. In accordance with the findings in Simulation Studies 2 and 3, the MH method performed extremely poorly when there were missing data. Comparing the results from the LR and OR methods, coincidentally, both methods detected gender DIF in five, seven, and 45 items in the reading, mathematics, and science tests, respectively. Below, the results of the OR method were focused.

Because of the large sample size, a trivial DIF could be statistically significant. Le's (2009) criteria were adopted in this analysis, in which an item is ultimately flagged as having substantial DIF when two conditions are jointly satisfied: (a) $\hat{\lambda}_i$ is significant at the .05 nominal level, and (b) the DIF size (i.e., $|\hat{\lambda}_i - \tilde{\lambda}|$) is larger than 0.25 logits. There were nine,

10, and 36 items exhibiting substantial gender DIF in the reading, mathematics, and science tests, respectively. The results of DIF detection with scale purification are presented in Figure 1. Extreme $\hat{\lambda}$ values were bilaterally distributed in all tests, and the means and medians of $\hat{\lambda}$ for non-DIF items approximately overlapped. Specifically, seven, two, and 19 items in the reading, mathematics, and science tests favored boys, whereas two, eight, and 17 items favored girls, respectively.

Discussion and Conclusion

The MH and LR methods use test scores to match examinees from different groups for DIF detection. However, test scores cannot represent examinees' ability levels properly when tests have DIF items or the impact is large. Solutions to this predicament include scale purification procedures and identification of a set of DIF-free items to serve as a matching variable (Kopf et al., 2015). Scale purification and identification of DIF-free items, however, may be rather inconvenient and challenging for practitioners. In this study, a simple solution that does not rely on a matching variable was provided. In the newly proposed OR method, the log OR of two groups of examinees on their responses to each item is computed, and an item is deemed as having DIF when its log OR is an outlier against the sample median. If a practitioner knows how to compute the log OR (Equation 6), its standard error (Equation 7), and the sample median, the OR method can be easily implemented. The scale purification is also straightforward. All we need to do is recalculate the sample median from those items identified as DIF-free in the previous step.

The findings from a series of simulation studies show that the OR method yielded satisfactory FPRs and high TPRs when tests had 30% or fewer DIF items, and the scale purification procedure could reduce the inflated FPRs when tests had 40% or more DIF items. In contrast, the MH and LR methods were vulnerable to high percentages of DIF items and missingness. Thus, the OR method is recommended because of its high effectiveness and feasibility in DIF detection.

A distinct advantage of the OR method is its robustness to missing data. Missing responses weaken the relationship between an examinee's test score and latent ability level, such that standard DIF detection methods become thorny (Robitzsch & Rupp, 2009). Bank (2015) reviewed nine studies that examined DIF detection in missing data and recommended the use of deletion (including analysis-wise and listwise) and imputation. Although deleting cases with missing responses is simple, it may lead to a substantial loss of samples and, consequently, a lower power of DIF assessment (Finch, 2011). In addition, deletion is not always feasible when missing data are formulated by design, such as in booklets in large-scale assessments (e.g., the PISA). Although imputation, an advanced approach for treating missing data, may be more effective than deletion (Finch, 2011; Robitzsch & Rupp, 2009), selecting an imputation algorithm to add unsubstantial information into observed data is somewhat arbitrary and may fabricate the result of DIF assessment. Seemingly, deletion inevitably sacrifices some information, while imputation exaggerates the limited information.

Unlike deletion and imputation, the OR method utilizes information by employing observed data and considering the incompleteness of item responses to identify DIF items. As in Equation 7, a small number of valid responses would result in a large error variance of $\hat{\lambda}_i$. Conditional on the value of the log ORs, the more the missing data, the larger the error variance. This feature ensures that the OR method is not vulnerable to missing data, as the error variance for statistical testing will not be overestimated due to deletion or underestimated due to imputation.

The robust OR method is rather robust to missing data, and it might be applicable in detecting DIF in computerized adaptive testing (CAT). Because many responses are missing at

random in CAT, existing DIF detection methods require a tedious procedure to obtain the ability estimate (or expected score over the entire item pool) as a matching variable to detect DIF in new items (Lei, Chen, & Yu, 2006). The robust OR method can be an alternative in such cases, but its performance needs further investigation.

Although the OR method was developed to detect DIF between two groups, it can be used to detect DIF in more than two groups. For example, once the reference group is identified, each focal group can be compared with the reference group individually (Finch, 2016). Like the MH method, the OR approach is effective in detecting uniform DIF when the data follow the Rasch model. It was suspected that the OR method will not perform well when data follow the two- or three-parameter logistic model or the DIF is nonuniform, which should be verified in future research. In recent years, polytomous items have been widely used in educational and psychological tests. It is important for future study to extend the OR method to detect DIF in polytomous items.


Declaration of Conflicting Interests


The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.


Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: Hui-Fang Chen received the City University of Hong Kong Strategic Research Grant (Project No. CityU 21615416) and the Research Grants Council Early Career Scheme, the Hong Kong Special Administrative Region (Project No. 9048095) for this study.

ORCID iD

Kuan-Yu Jin  <https://orcid.org/0000-0002-0327-7529>

Hui-Fang Chen  <https://orcid.org/0000-0002-5037-4559>

Wen-Chung Wang  <https://orcid.org/0000-0001-6022-1567>

References

- Agresti, A. (2002). *Categorical data analysis* (2nd ed.). New York, NY: John Wiley.
- Angoff, W. H., & Ford, S. F. (1973). Item-race interaction on a test of scholastic aptitude. *Journal of Educational Measurement, 10*, 95-106. doi:10.1111/j.1745-3984.1973.tb00787.x
- Bank, K. (2015). An introduction to missing data in the context of differential item functioning. *Practical Assessment, Research & Evaluation, 20*(12). Retrieved from <http://pareonline.net/getvn.asp?v=20&n=12>
- Clauser, B., Mazor, K., & Hambleton, R. K. (1993). The effects of purification of matching criterion on the identification of DIF using the Mantel-Haenszel procedure. *Applied Measurement in Education, 6*, 269-279. doi:10.1207/s15324818ame0604_2
- Cohen, A. S., Kim, S. H., & Wollack, J. A. (1996). An investigation of the likelihood ratio test for detection of differential item functioning. *Applied Psychological Measurement, 20*, 15-26. doi:10.1177/014662169602000102
- Dorans, N. J., & Kulick, E. (1986). Demonstrating the utility of the standardization approach to assessing unexpected differential item performance on the Scholastic Aptitude Test. *Journal of Educational Measurement, 23*, 355-368. doi:10.1111/j.1745-3984.1986.tb00255.x
- Finch, W. H. (2011). The impact of missing data on the detection of nonuniform differential item functioning. *Educational and Psychological Measurement, 71*, 663-683. doi:10.1177/0013164410385226

- Finch, W. H. (2016). Detection of differential item functioning for more than two groups: A Monte Carlo comparison of methods. *Applied Measurement in Education, 29*, 30-45. doi:10.1080/08957347.2015.1102916
- French, B. F., & Maller, S. J. (2007). Iterative purification and effect size use with logistic regression for differential item functioning detection. *Educational and Psychological Measurement, 67*, 373-393. doi: 10.1177/0013164406294781
- Hidalgo, M. D., & López-Pina, J. A. (2004). Differential item functioning detection and effect size: A comparison between logistic regression and Mantel–Haenszel procedures. *Educational and Psychological Measurement, 64*, 903-915. doi:10.1177/0013164403261769
- Holland, P. W., & Thayer, D. T. (1988). Differential item performance and the Mantel–Haenszel procedure. In H. Wainer & H. I. Braun (Eds.), *Test validity* (pp. 129-145). Hillsdale, NJ: Lawrence Erlbaum.
- Kopf, J., Zeileis, A., & Strobl, C. (2015). Anchor selection strategies for DIF analysis: Review, assessment, and new approaches. *Educational and Psychological Measurement, 75*, 22-56. doi: 10.1177/0013164414529792
- Le, L. T. (2009). Investigating gender differential item functioning across countries and test languages for PISA science items. *International Journal of Testing, 9*, 122-133. doi:10.1080/15305050902880769
- Lei, P.-W., Chen, S.-Y., & Yu, L. (2006). Comparing methods of assessing differential item functioning in a computerized adaptive testing environment. *Journal of Educational Measurement, 43*, 245-264. doi: 10.1111/j.1745-3984.2006.00015.x
- Lord, F. M. (1980). *Application of item response theory to practical testing problems*. Hillsdale, NJ: Lawrence Erlbaum.
- Magis, D., Béland, S., Tuerlinckx, F., & De Boeck, P. (2010). A general framework and an R package for the detection of dichotomous differential item functioning. *Behavior Research Methods, 42*, 847-862. doi:10.3758/brm.42.3.847
- Magis, D., & De Boeck, P. (2012). A robust outlier approach to prevent Type I error inflation in differential item functioning. *Educational and Psychological Measurement, 72*, 291-311. doi: 10.1177/0013164411416975
- Magis, D., & De Boeck, P. (2014). Type I error inflation in DIF identification with Mantel–Haenszel: An explanation and a solution. *Educational and Psychological Measurement, 74*, 713-728. doi: 10.1177/0013164413516855
- Narayanan, P., & Swaminathan, H. (1996). Identification of items that show nonuniform DIF. *Applied Psychological Measurement, 20*, 257-274. doi:10.1177/014662169602000306
- Raju, N. S. (1988). The area between two item characteristic curves. *Psychometrika, 53*, 495-502. doi: 10.1007/bf02294403
- Rasch, G. (1960). *Probabilistic models for some intelligence and attainment tests*. Copenhagen, Denmark: Institute of Education Research.
- R Core Team. (2016). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Robitzsch, A., & Rupp, A. A. (2009). Impact of missing data on the detection of differential item functioning: The case of Mantel–Haenszel and logistic regression analysis. *Educational and Psychological Measurement, 69*, 18-34. doi:10.1177/0013164408318756
- Rogers, H. J., & Swaminathan, H. (1993). A comparison of logistic regression and Mantel–Haenszel procedures for detecting differential item functioning. *Applied Psychological Measurement, 17*, 105-116. doi:10.1177/014662169301700201
- Shealy, R., & Stout, W. (1993). A model-based standardization approach that separates true bias/DIF from group ability differences and detects test bias/DTF as well as item bias/DIF. *Psychometrika, 58*, 159-194. doi:10.1007/bf02294572
- Wang, W.-C., & Su, Y.-H. (2004). Effects of average signed area between two item characteristic curves and test purification procedures on the DIF detection via the Mantel–Haenszel method. *Applied Measurement in Education, 17*, 113-144. doi:10.1207/s15324818ame1702_2