Running Head: Interpretation of Spurious Differential Item Functioning

Interpretation of Spurious Differential Item Functioning

Yongnam Kim¹ & Seo Young Lee²

¹ Department of Education, Seoul National University, 1 Gwanak-ro, Gwanak-gu, Seoul 08826, Republic of Korea. Email: ykims@snu.ac.kr

² Prometric LLC, 1501 South Clinton Street Suite 1400, Baltimore, MD 21224, U.S.A. Email: seoyoung.lee@prometric.com

January 2022

ACKNOWLEDGEMENTS

The authors are grateful to the students in the item response theory class, taught in the 2020 Fall term in the Department of Education at Seoul National University, for discussions that helped to develop the ideas presented in this article. This work was partially supported by the New Faculty Startup Fund from Seoul National University.

ABSTRACT

This article introduces a counterfactual discussion on the difficulty in transforming statistical evidence of differential item functioning (DIF) into the substantive interpretation of item fairness or bias. Presented causal graphs that describe how the data were generated and the accompanying graphical explanation provide an intuitive explanation to understand which DIF should be considered genuine or spurious. In particular, the article presents three scenarios where spurious DIF is structurally embedded and will be detected by various statistical methods. The first scenario considers a case where an unknown common cause produces a spurious association while the other two scenarios consider cases where a collider variable plays a role in inducing an unexpected type of spurious DIF. Findings about colliders, especially due to sample selection, can be useful for resolving some confusion in DIF applications when analyzing a selected sample such as military data.

KEYWORDS: Differential item functioning, DIF, item bias, fairness, do-operator, causal graph, collider, sample selection, counterfactual

INTRODUCTION

The concept of *item fairness* or *item bias* was formulated within the measurement context in the 1960s and 1970s (e.g., Angoff & Sharon, 1974; Cleary & Hilton, 1968), and psychometricians have developed various statistical techniques to empirically verify it. Since then, the original concept was somewhat decontextualized, and a purely statistical term, *differential item functioning* (DIF), has been popularized (Zumbo, 2007). In the practice of test development and evaluation, it is now common that researchers and practitioners use logistic regression and other statistical methods in order to investigate whether an item functions differently across groups of examinees. In the recent literature, it is usually admitted that statistical evidence of DIF does not always indicate that the item is truly *biased* against a group (AERA, APA, & NCME, 2014; de Ayala, 2009; Finch & French, 2019; Martinková et al., 2017; Penfield & Camilli, 2007; Shepard, Camilli, & Williams, 1985; Zumbo & Gelin, 2005). This notion has been referred to as *spurious DIF* and the related challenge is well expressed by Angoff (1993): “[Test developers] are often confronted by DIF results that they cannot understand; and no amount of deliberation seems to help explain why some perfectly reasonable items have large DIF values” (Angoff, 1993, p. 19).

Relatively few studies have explicitly focused on the question of why DIF occurs and whether it can be considered spurious or genuine in the sense of whether or not the item under investigation is unfair to a group of examinees (Zumbo, 2009). Cohen and Bolt (2005) discussed a mixture modeling approach that might help to understand the reason of DIF by latent classes. Andrich and Hagquist (2012) showed that a specific DIF-detecting method such as the Mantel-Haenszel (MH) procedure can turn a truly DIF-free item to a seemingly DIF item. They referred to it as *artificial* DIF because the resulting spurious DIF is due to an artifact of the specific MH

procedure. While the previous attempts have focused on statistical techniques, this article aims to provide a conceptual basis for DIF and item bias. Using recent advances in causal inference (e.g., Imbens & Rubin, 2014; Pearl, 2009), the article introduces a formal counterfactual discussion that helps to resolve the difficulty in transforming statistical evidence into the substantive interpretation of item fairness or bias. It provides a counterfactual definition of item bias, which is clearly distinguished from the definition of DIF. In particular, causal graphical models or directed acyclic graphs (Pearl, 2009; Spirtes, Glymour, & Scheines, 2000) are used to describe how the item response data were generated and how each causally interacts with each other. It is demonstrated that much confusion as to spurious DIF can be resolved by understanding the data-generating process. A related, but not identical, graphical approach to DIF can be found in Kreiner and Christensen (2002). The discussion here, however, is not restricted to Rasch models and includes a novel explanation for spurious DIF due to *collider* variables. Further, the article presents three scenarios one might encounter in practice, where a spurious DIF is structurally embedded and will be detected by various statistical methods, including but not limited to the MH procedure, and also across 1PL to 3PL item response models.

The remainder of this article is organized as follows. First, the basic setup and preliminary concepts, including the formal definitions of DIF and item bias, are introduced. The next section then presents causal graphical models for DIF and item bias and provides a visual explanation of the difference between genuine and spurious DIFs. The corresponding scenarios for different types of spurious DIF are also presented. The following section illustrates those scenarios using simulation. The article concludes with a summary of findings and implications for future research on DIF.

PRELIMINARY CONCEPTS

Setup and Notation

Suppose that examinees' latent trait θ , aimed to be measured through a test, causes their item response Y . Although it is not necessary for the main findings of this article, one may think of θ as continuous and Y as binary, as is typical in the literature. For example, latent math ability θ affects examinees' math item responses such that they are correct ($Y = 1$) or incorrect ($Y = 0$). The test consists of multiple items, but a single item that produces the response Y is under investigation. One's interest is whether the item is biased and unfair to a group of examinees defined by G , which is an observed group characteristic. G may be an exogenous and fixed group factor such as gender or race, which is popular in the DIF literature, but may also be a potentially endogenous factor such as language or income group, or whether a student attends a special education program. In addition to the observed group characteristics, there might be an unobserved or hidden variable, denoted by U , of which one is *not* aware but that interacts with other variables in the item response process. U may be another latent trait, which would then be naturally conceptualized by *multidimensionality* (in terms of θ and U), or just another group characteristic that is unknown and not recorded in the observed data. For all the variables, no distributional restrictions are required although the regular normal and binomial distributions will be used in the later simulation for the demonstration purposes.

$P(Y|G = g)$, or simply $P(Y|g)$, stands for a regular conditional probability distribution. It indicates the probability of item response Y conditional on a specific group characteristic $G = g$. According to Pearl (2009), it has a semantic of "seeing," that is, one simply stratifies data by G and looks at the distribution of Y in each stratum of g . This article shall use another distinctive notation, referred to as the *do*-operator (Pearl, 2009), to represent a counterfactual or

interventional idea.¹ $P(Y|do(G = g))$, or simply $P(Y|do(g))$, stands for an *intervention* distribution of Y if the variable G were intervened and all units were set to take the value of g . In contrast to $P(Y|g)$, $P(Y|do(g))$ has a semantic of “doing.” It is a counterfactual probability of item response (Y) that would have appeared if group characteristic G were set to g . For example, it captures a hypothetical idea such as “what if the examinee were a female, instead of a male?” which is beyond the simple observation by stratifying data on G .² The difference between $P(Y|g)$ and $P(Y|do(g))$ is crucial and, in fact, it is the foundation of modern causal inference concepts (Pearl & Mackenzie, 2018). Throughout this article the *do*-operator suggested by Pearl (2009) will be used, but a similar definition would also be possible with Rubin’s potential outcomes notation (Imbens & Rubin, 2015). In short, $P(Y|g)$ represents an association between G and Y , while $P(Y|do(g))$ represents a causation between G and Y , that is, the causal influence of G on Y . As will be clearer later, the distinction is the key to understand the difference between DIF and item fairness.

Graphical Representations and *D*-Separation

The idea that the trait θ *causes* the item response Y is represented as $\theta \rightarrow Y$ in causal graphs. The arrow does not specify a particular functional form of the relationship such as linear

¹ Strictly speaking, Pearl (2009) distinguishes intervention from counterfactuals, which belong to rung 2 and rung 3, respectively, in his causal hierarchy (observation belongs to rung 1). For the main findings of this article, such a distinction is not necessary, and, in fact, it needs another in-depth discussion. For ease of exposition, this article shall use counterfactual and interventional, interchangeably.

² Again, strictly speaking, such an individual counterfactual statement belongs to rung 3, counterfactuals, in Pearl’s causal hierarchy, and it cannot be captured by the *do*-operator. The authors believe that the distinction might not be helpful to many psychometricians who are not familiar with the formal causal theory and may even prevent readers to find the benefit of the findings of this article. The authors admit that it is not accurate but shall slightly abuse Pearl’s concepts in order to emphasize and convey the fundamental difference between $P(Y|g)$ and $P(Y|do(g))$, which is the essence of the difference between DIF and item bias.

or logistic but only states that the value of Y is somehow affected by the value of θ .³ Causal graphs encode (un)conditional (in)dependences among variables, and the concept of d -separation (Pearl, 1988, 2009) helps to intuitively extract such information. In a causal graph where the causal Markov and faithfulness hold,⁴ statistical association between a pair of variables is transmitted only through the open paths between them. If there exist no open paths, they are statistically independent. A path is a sequence of adjacent variables in causal graphs that appear only once, regardless of the direction of arrows. It is not open, or blocked, if (i) any non-collider middle variables in the path are conditioned on or (ii) the path contains a collider that is not conditioned on (and no variable that is affected by the collider is conditioned on). A collider is a variable that is affected by two adjacent variables in the path. For example, θ serves as a collider in the path $G \rightarrow \theta \leftarrow U$, and thus the path remains blocked. It implies that G and U are unconditionally independent of each other. However, the path becomes open if θ is conditioned on because it is a collider. G and U are then dependent of each other conditional on θ . See Pearl (2009) for a more formal and detailed explanation of d -separation and the related graphical concepts.

Definitions

This article provides formal definitions of DIF and item bias.⁵ In the literature, DIF is generally defined as a purely statistical concept that describes a systematic difference in item

³ It is thus a nonparametric structural equation model (Pearl, 2012). Therefore, the graphical representation will embrace many different kinds of item response models, including not only Rasch models but also 3PL (even higher) models.

⁴ The two conditions establish the equivalence of (un)conditional (in)dependences between graphical and probabilistic statements. For more details, see Pearl (2009) or Spirtes et al. (2000).

⁵ This article does not distinguish ‘item bias’ and ‘item impact’ as Ackerman (1992) and some others did. Both terms would belong to ‘item bias,’ defined in *definition 2* in this article. A causal effect of a group indicator on an item response, that is, an item bias defined here, might be

responses between groups at the same level of the latent trait. Among many others, Penfield and Camilli (2005) explicitly defined DIF as a conditional association as below:

Definition 1 (Differential Item Functioning; DIF) *Suppose that examinees' latent trait θ causes their item response Y . Let G denote an observed group indicator of the examinees. We say that the item functions differentially across G , or it has DIF in terms of G , if G and Y are associated conditional on θ : $P(Y|g, \theta) \neq P(Y|g', \theta)$.*

The distinct values of the random variable G are denoted by g and g' . For example, $g = \text{male}$ and $g' = \text{female}$ when G denotes gender, or it could be $g = \text{white}$ and $g' = \text{nonwhite}$ when G denotes race. It is a general one such that it embraces both uniform and non-uniform DIFs and it is not restricted on binary variables for item responses and group indicators. The most notable thing in this definition is that it shows that DIF is a statistical concept that concerns association instead of causation.

In contrast, item bias is a substantive concept that is grounded on the causal context. Such an idea, although it is mostly implicit, can be found in the DIF literature. Angoff (1993) considered item bias as “[a] *cause* of the large disparity in scores” (Angoff, 1993, p. 4, emphasis added). Finch and French (2019) wrote, “an item demonstrates bias [...] [when] there exists a systematic process, other than the construct being measured, that *influences* the response to the item for some examinees but not for others” (Finch & French, 2019, p. 279, emphasis added). Beyond the observed group differences, the concept of item bias should also be able to deal with

viewed as relevant or irrelevant depending on the purpose of the measure (Zumbo & Gelin, 2005). Such a distinction might be necessary when judging construct validity but is not the focus of this article.

“an *explanation* as to the source of group differences on an item” (Finch & French, 2019, p. 280, emphasis added; see also AERA et al., 2014, p. 51, for a similar statement).

As those ideas invoke the concept of *intervention* instead of observation, it is necessary to use a different notation that is capable to represent the concept of causation between a group characteristic on item response. Using the *do*-operator, item bias can be formally defined as below:

Definition 2 (Item Bias or Unfairness) *An item is said to be biased and unfair to a group of examinees defined by G if G influences their item response Y holding θ constant:*

$$P(Y|do(g), do(\theta)) \neq P(Y|do(g'), do(\theta)).^6$$

$P(Y|do(g), do(\theta))$ is a counterfactual probability of Y that would have appeared if G were set to g and the latent trait were set to θ . It is crucial to note that the expression of examinees’ latent trait θ has been changed. Besides G , *definition 2* also has $do(\theta)$, which is reflected by the statement, “holding θ constant,” instead of “conditioning on θ .” That is, item bias is defined as an item response difference by a group characteristic (i.e., impact of G on Y) when latent trait is *intervened* to have a specific value. In fact, *definition 2* corresponds the controlled direct effect, $CDE(\theta)$, defined by Pearl (2014) in mediation analysis. It is tempted to define item bias as $P(Y|do(g), \theta) \neq P(Y|do(g'), \theta)$, where the latent trait is expressed as *do*-free, but this will

⁶ One might be reluctant to use the *do*-operator on G and thus to view G as a “cause.” This is probably because popular group indicators in the DIF and item bias literature are gender and race, which are usually considered non-manipulable (see Holland, 1986). Although it is debatable, this article considers that labeling gender or race as a “cause” is a useful way to understand the concept of item bias or unfairness. Using the well-known Berkeley admission example (Bickel, Hammel, & O’Connell, 1975), Pearl and Mackenzie (2018) defined gender bias or discrimination as a direct causal effect of gender on admission. For a more in-depth discussion about this issue, see VanderWeele and Hernán (2012). Also, note that some DIF applications consider a clearly manipulable group characteristic such as language or income group.

bring a problem known as the “mediation fallacy” (Pearl & Mackenzie, 2014). That is, conditioning on a mediator does not allow one to identify the direct causal effect if there is confounding between mediator and outcome. Thus, the definition of $CDE(\theta)$ requires that the mediator should be intervened such that the mediator value would be constant. Such a case will be presented in Figure 2B later. Simply speaking, the crude definition, $P(Y|do(g), \theta) \neq P(Y|do(g'), \theta)$, will fail to distinguish spurious DIF from item bias in a causal structure described in Figure 2B.

From the two definitions above, it is clear that DIF may indicate the presence of item bias, but does not imply it (AERA et al., 2014; de Ayala, 2009). The notion follows the well-known maxim, “correlation does not imply causation.” If a DIF (association) is, at least partially, due to a true item bias (causation), then it can be referred to as a genuine DIF. If a DIF (association) is not accompanied by a true item bias (causation), then it is referred to as a spurious DIF:

Definition 3 (Genuine DIF) *An observed DIF for an item is referred to as genuine if there is item bias: $P(Y|g, \theta) \neq P(Y|g', \theta)$ and $P(Y|do(g), do(\theta)) \neq P(Y|do(g'), do(\theta))$.*

Definition 4 (Spurious DIF) *An observed DIF for an item is referred to as spurious if there is no item bias: $P(Y|g, \theta) \neq P(Y|g', \theta)$ and $P(Y|do(g), do(\theta)) = P(Y|do(g'), do(\theta))$.*

CAUSAL DIAGRAMS FOR DIFFERENTIAL ITEM FUNCTIONING

Graphical Understanding of Differential Item Functioning and Item Bias

The basic premise of the item response process is that examinees’ latent trait θ causes their item response Y (i.e., $\theta \rightarrow Y$; see Figure 1). When investigating item bias and/or DIF, the

FIGURE 1. *Causal graphs for item responses where θ denotes examinees' latent ability or trait, Y their item response, G an observed group indicator of interest, and U an unobserved factor. (A) and (B) describe both item bias and DIF in terms of G . (C) shows neither item bias nor DIF in terms of G .*

main interest is how a group indicator of interest, G , is related to θ and Y . According to the definition of item bias (*definition 2*), item bias occurs when G directly affects Y , holding θ constant (by another intervention on the latent trait). The idea of “doing $G = g$ ” stated in *definition 2* is expressed by $G \rightarrow Y$ in casual graphs as in Figures 1A and 1B. The subtle difference between the two graphs is an additional arrow in Figure 1B, $G \rightarrow \theta$, which indicates that the group membership affects the latent traits. This additional structure, however, does not change our understanding and interpretation of the item bias because, by definition, we consider only the direct effect of G on Y , not via θ , as item bias. This is also aligned with the fact that the group impact (i.e, the effect of G on Y) should not be confused with item bias (Camilli & Shepard, 1994). Graphically speaking, the causal path $G \rightarrow \theta \rightarrow Y$ in Figure 1B is disabled when θ is intervened. This is because when a variable is intervened, all incoming arrows on the variable should be removed in causal graphs (this means that the variable changes only by the intervention and no other factors can change it; see Pearl, 2009). After intervening on θ , both

Figures 1A and 1B show $G \rightarrow Y$, meaning that a systematic difference in item response Y was because of group characteristic G .

According to *definition 1*, the concept of DIF is expressed as an association between G and Y conditional on θ . Applying d -separation, one can see that in both Figures 1A and 1B, G and Y are connected after conditioning on θ , and the conditional association is only due to the causal path $G \rightarrow Y$; there are no other open paths that connect the two after conditioning on θ . Note that the path $G \rightarrow \theta \rightarrow Y$ in Figure 1B is blocked by conditioning on the middle variable θ . That is, in Figures 1A and 1B, the observed conditional association between G and Y , which is DIF, is only due to the direct causal impact of G on Y , which is also item bias. Thus, we refer to the DIF embedded in the both causal graphs as a genuine DIF. Running DIF analyses with the item response data described in Figures 1A and 1B, regardless of the model specification from 1PL to 3PL and the chosen DIF-detecting methods, will return clear evidence of DIF that can be validly interpreted as item bias.

On the other hand, it is not uncommon in practice that no evidence of DIF is found even when it seems like G is connected to Y . Figure 1C describes such a case. The described item is not biased in terms of G because there is no direct effect of G on Y when holding θ constant. The sole path between G and Y , $G \rightarrow \theta \rightarrow Y$, is disabled by intervening on (or doing) θ , and thus G and Y are conditionally independent of each other, which implies that DIF analyses are supposed to detect no DIF. Analogous to the statement of “no causation without association,” one may say “no item bias without DIF” because the presence of DIF is a necessary cause of item bias (Zumbo, 1999). This interpretation generally holds unless multiple open paths perfectly cancel

each other out.⁷ This article shall not consider such *unfaithful* cases in the belief that although it is possible, such a perfect offsetting would be infrequent in practice.

Spurious Differential Item Functioning Due to Hidden Factors

In reality, in addition to observed group indicators, G , there could also be unknown or hidden factors, U , that affect item responses. In a general DIF analysis, one would try to understand the results in terms of G , not in terms of U , because U remains unknown or simply because U is not of interest for the analysis. In the following, it is shown that it is highly likely that such results would not be interpretable unless the entire causal structure of item response process including U is well understood. Figure 2 presents three representative examples of such a case. Note that all the three causal graphs do not show item bias in terms of G because G does not directly impact Y when intervening on θ (i.e., no arrow of $G \rightarrow Y$). Nonetheless, one would detect a DIF in terms of G from all the data generated by those models, that is, spurious DIF. Note that in all the causal graphs in Figure 2 and the corresponding scenarios, which will be presented, it is valid to say that the item is biased *in terms of* U because U directly causes Y . However, in this article U represents an unknown hidden factor. It is through G , not U , that one aims to investigate DIF and item bias. All the presented scenarios below demonstrate how one's interpretation of DIF and the judgement of item bias in terms of G , not in terms of U , could be distorted and misleading when the unknown factor U plays a role in the item response process.

1. Spurious DIF due to a hidden common cause. The first case described in Figure 2A is an example of a spurious DIF due to an unknown common cause variable. In the structure, G is associated with Y not because of a direct effect of G on Y but because of U , which serves as

⁷ Multiple paths between G and Y , after conditioning on θ , can be found in Figure 2.

FIGURE 2. Causal graphs for spurious differential item functioning (DIF) when an unknown or hidden factor U plays a role in the item responses. Although neither causal graph implies the item bias in terms of G because there is no arrow $G \rightarrow Y$, all graphs imply the spurious DIF in terms of G because G and Y are connected conditional on θ via (A) $G \leftarrow U \rightarrow Y$, (B) $G \rightarrow \leftarrow U \rightarrow Y$, and (C) $G \rightarrow S \leftarrow U \rightarrow Y$ due to the sample selection based on S .

confounding between G and Y . Some corresponding classic examples are how ice cream sales are associated with the number of drowning deaths (due to temperature) or how at the country level chocolate consumption is associated with the number of Nobel Laureates (due to per capita income). No one says that drowning deaths are *because of* ice cream sales and that we should consume more chocolate to have more Nobel Laureates. As such, although there is a systematic association between G and Y conditional on θ , in Figure 2A, the group characteristic is not the cause of the item response difference, therefore, it is not item bias but DIF in terms of G .

Note that because in the causal graph in Figure 2A, the group factor of interest, G , is affected by the unknown factor U , some popular invariant group characteristics that cannot be affected by other factors such as race or gender would be implausible in this case. However, for the cases where group characteristics are determined by some other factors, they can be an example described in Figure 2A. In DIF studies, such group characteristics have been widely used. For example, Angoff (1993) mentioned a case where whether student's participation in a

special education program or not was a group factor for DIF analysis, and Crane et al. (2006) investigated whether survey measures of shoulder function differed by patients' surgical history. For many cross-cultural studies, different language groups, which could be affected by race, have also been considered as a group factor for DIF analysis (e.g., Gibbons et al., 2011). As addressed above for Figures 1A and 1B, even after an additional arrow $G \rightarrow \theta$ is added to Figure 2A, the whole point still holds. This is because all the additional connections between G and Y via θ are disconnected or blocked by intervening on or conditioning on θ . Below is a hypothetical scenario for a spurious DIF structurally embedded in the causal graph in Figure 2A:

Scenario 1 (Figure 2A) *Suppose that θ denotes math ability and U reading ability, and they are independent of each other. A math item response Y is affected not only by θ but also by U because one could not solve the math problem without reading and understanding the required task (i.e., multidimensionality; $\theta \rightarrow Y \leftarrow U$). Further suppose that the information of whether or not students watched a new Star Wars movie was collected (G), and more students who have higher reading ability, U , tend to have watched the movie because they have read the Star Wars books and are interested in the story ($U \rightarrow G$). A teacher, or an analyst, conducted a DIF analysis in terms of G was confused by an apparent DIF result of a math item because he or she believed that the math question itself does not contain or require any knowledge of the movie.*

Unless the researcher or teacher realizes the underlying connection between the movie-watching (G) and math item response (Y) through reading ability (U), some confusion would arise. It does not make intuitive sense that students had lower math scores because they did not watch the movie. It shows that a hidden common cause can make the interpretation of DIF evidence puzzling and misleading.

2. *Spurious DIF due to the latent trait as a collider.* The previous scenario is not the only one that generates a spurious DIF despite the absence of item bias. More counterintuitive ones due to conditioning on a collider are presented in Figures 2B and 2C. In the causal inference literature, such examples are referred to as Berkson’s paradox (Berkson, 1946; see also Pearl, 2009) or an endogenous selection bias (Elwert & Winship, 2014).

First, note that the causal graph in Figure 2B differs from the graph in Figure 1C (where neither item bias nor DIF was embedded) in that it adds a confounding structure due to U between θ and Y , $\theta \leftarrow U \rightarrow Y$. This added structure makes a completely new type of a spurious DIF. By *definition 1*, DIF is an association between G and Y conditional on θ . Unlike Figure 1C, the main latent trait θ in Figure 2B serves as a collider in the path $G \rightarrow \theta \leftarrow U \rightarrow Y$. It is a collider because two arrow heads are colliding at θ . Thus, G and Y become associated conditional on the collider θ (see the previous section “Preliminary Concepts”). The resulting spurious association is less intuitive and hard to interpret. For example, in the original example by Berkson (1946), he showed that two diseases such as diabetes and cholecystitis, which are likely independent of each other in the general public, were negatively associated when only the hospital in-patient group (i.e., collider) was considered, which seems paradoxical. A spurious DIF scenario based on the causal graph in Figure 2B is presented below:

Scenario 2 (Figure 2B) *Let G and U denote gender and race, respectively. According to an artificial theory, suppose that females have a higher reading ability level, on average, than males ($G \rightarrow \theta$). Also, suppose that white students have a higher reading ability level than nonwhite students ($U \rightarrow \theta$). If a reading item consisted of many words and expressions common among white students, the item is biased against nonwhite students ($U \rightarrow Y$). Suppose that gender itself does not directly influence Y , that is, there is no item bias in terms of gender (G). A*

researcher correctly believed that the reading item has no item bias in terms of gender (G) but, nonetheless, performed a DIF analysis in terms of it. The researcher was confused by an apparent DIF result that seems to suggest that male students have higher reading scores than their counterpart female students who have the same level of reading ability.

In scenario 2, there is no item bias in terms of gender, but the DIF analysis shows a clear DIF result. This is because of the conditional association via the collider path $G \rightarrow \theta \leftarrow U \rightarrow Y$, which becomes open by conditioning on the collider θ of this path. Unless the researcher realizes the role of race (U), as a collider, the DIF result would be hard to properly interpret.

3. Spurious DIF due to sample selection. A more intriguing example due to sample selection is presented in Figure 2C. The causal graph shows a case where the group factor G is not even associated at all with either θ or Y , but nonetheless a spurious DIF in terms of G is embedded and will be detected. In the graph, S is a sample selection indicator such that some cases are selected ($S = 1$) or not ($S = 0$). As Berkson (1946) considered only a group of hospitalized people, one might look at a selected sample data ($S = 1$) instead of the entire data in the DIF analysis. For example, it is rather common in the literature that military samples are analyzed (e.g., Woods, Oltmanns, & Turkheimer, 2009). In the graph, the group factor G only affects this selection process but does not affect θ or Y . Then, the sample selection node S serves as a collider in the path $G \rightarrow S \leftarrow U \rightarrow Y$ and connects G and Y if only the selected sample is considered and analyzed. If researchers do not carefully think about the sample selection process, the resulting spurious DIF is very confusing and hard to interpret because the group factor G is theoretically unrelated to both θ and Y . Such a hypothetical scenario is encapsulated below:

Scenario 3 (Figure 2C) *A DIF analysis was performed for a test that aims to measure spatial ability (θ) in a sample of Air Force recruits ($S = 1$) in terms of race (G). The information of gender (U) was not considered. In an artificial theory, it is known that gender (U) has an impact on spatial ability such that males have higher spatial ability than females ($U \rightarrow \theta$). Also, gender itself can directly impact peoples' responses because the test uses a form of 3D gun shooting, which are more familiar to males through video games ($U \rightarrow Y$). In contrast, race (G) has no impact on the spatial ability θ and the response Y . But race affects being selected into the sample because the recruits have more nonwhites than whites ($G \rightarrow S$). The recruits also include more males ($U \rightarrow S$). A researcher who analyzed the DIF with the Air Force sample ($S = 1$) is confused by clear evidence of DIF in terms of race (G) because he or she thought that theoretically race itself has nothing to do with both special ability and the response of the test.*

Military samples are commonly collected and analyzed in DIF analysis. Thus, scenario 3 could be more frequent in reality than one might expect.

ILLUSTRATION

Data-Generating Models

A series of simulation studies were designed for the empirical demonstration of the previous scenarios of spurious DIFs. For all the simulation conditions, item response data for 25-items with 10,000 cases were generated under three IRT models (i.e., 1PL, 2PL, and 3PL). Both G and U were considered binary and, when they are exogeneous, G and U were randomly generated under a Bernoulli distribution with $p = .5$. In scenario 1 (Figure 2A), G is affected by U and thus it was simulated as $G = 1$ if $U + \varepsilon_G > 0.5$, and $G = 0$ otherwise. The random noise ε_G followed a standard normal distribution, $N(0, 1)$. The sample selection indicator S in scenario

3 (Figure 2C) was generated as $S = 1$ if $G + U + \varepsilon_S > \tau$, and $G = 0$ otherwise, and ε_S followed $N(0, 1)$. The threshold τ was chosen to have 10,000 cases of $S = 1$ from 20,000 original cases so as to maintain the same sample size in DIF analyses as other simulation scenarios. Among the 25 items, only one item was manipulated to induce a spurious DIF in favor of $U=1$. For the illustration purpose, a large DIF that produces 0.8 difference in area between two item response curves was generated (Lee, 2017). Item discrimination and pseudo-guessing parameters were fixed across both the reference and focal groups; therefore, uniform DIF was simulated. For other items, item parameters were randomly generated, and details are given in Appendix. The latent trait θ was first randomly generated from $N(0, 1)$, and then increased by one if $G = 1$ or $U = 1$ in the corresponding scenario. Data-generation process in each scenario was replicated 1,000 times.

DIF Analysis and Evaluation

Four popular methods were used to detect DIF for all the generated data. Using the *difR* package (Magis, Beland, Tuerlinckx, & De Boeck, 2010) in R, logistic regression, MH procedure, SIBTEST, and Raju's procedure were applied, and no item purification was considered.⁸ When using IRT based methods, all IRT models were correctly chosen such that no model misspecification was made. Applying Raju's method with 3PL models, pseudo-guessing parameters were fixed to the true parameters in data-generating process. For simulation evaluation, a DIF result that is significant at the 5% alpha level was considered evidence of DIF and the average detection rates across replications in each condition were computed and evaluated.

⁸ The consequences of item purification would require a more complicated graphical discussion, which is beyond the scope of this article.

TABLE 1.

Average detection rates of spurious differential item functioning in terms of G based on the data-generating models in Figure 2

	1PL	2PL	3PL
Scenario 1 (Figure 2A)			
Logistic regression	1.000 [.046, .082]	1.000 [.050, .100]	1.000 [.048, .085]
MH	1.000 [.044, .076]	1.000 [.044, .095]	1.000 [.049, .076]
SIBTEST	1.000 [.051, .078]	1.000 [.051, .097]	1.000 [.046, .084]
Raju	1.000 [.023, .046]	1.000 [.012, .145]	1.000 [.019, .138]
Scenario 2 (Figure 2B)			
Logistic regression	.868 [.041, .105]	.798 [.044, .270]	.670 [.050, .352]
MH	.900 [.031, .064]	.818 [.043, .197]	.772 [.043, .275]
SIBTEST	.941 [.047, .091]	.950 [.056, .091]	.947 [.076, .160]
Raju	.359 [.013, .034]	.873 [.015, .208]	.895 [.016, .175]
Scenario 3 (Figure 2C)			
Logistic regression	.659 [.042, .077]	.647 [.039, .069]	.684 [.039, .066]
MH	.644 [.037, .082]	.641 [.036, .065]	.678 [.033, .064]
SIBTEST	.651 [.038, .075]	.629 [.041, .066]	.666 [.039, .070]
Raju	.604 [.018, .109]	.525 [.011, .141]	.530 [.014, .147]

Note. Numbers without brackets indicate the average DIF detection rates of the DIF-induced item and numbers with brackets indicate the range of the average DIF detection rates of 24 DIF-free items.

Results

The simulation study shows that spurious DIF is rooted in the causal structure where the data were generated. Regardless of the item response models and the chosen DIF-detecting methods, statistical evidence of DIF was found with a higher probability. Table 1 summarizes the

results. Across all simulation conditions, the average rates of the DIF-induced item (numbers without brackets) were far higher, compared to the average detection rates of all non-DIF items (numbers in brackets). For example, in scenario 3 where a sample selection process was assumed, the average detection rate for the DIF-induced item by logistic regression was 68.4% while the average rate for the DIF-free items ranged from 3.9% to 6.6%. This pattern holds in general although some differences between DIF-detecting methods and item response models seem to exist. Thus, one very likely finds evidence of DIF for the item in DIF analysis using the dataset generated by the structure as in Figure 2C. However, the item is not biased or unfair to a group of examinees defined by G because in the causal graph G does not directly cause Y .

DISCUSSION

Zumbo (2007) wrote that “[t]he matter of wanting know *why* DIF occurs” (Zumbo, 2007, p. 228, emphasis added) is the *key issue of the new generation of DIF research* that follows prior research on developing sophisticated DIF-detecting methods. This article argues that the difficulty of interpreting spurious DIF—that is, the question of why DIF occurs—comes from the fundamental difference between association and causation and the inability to distinguish them in DIF analysis. Rigorous mathematical frameworks that are capable of dealing with causation have been developed (Pearl, 1988; Rubin, 1974), and recently many statistical applications in various disciplines have started embracing such formal approaches. This article is an example of how a traditional measurement research question can benefit from recent advances in causal inference. Using causal graphical models, the difference between genuine and spurious DIFs was visually explained. In particular, this article provided three scenarios that produce spurious DIF. Spurious DIF due to conditioning on colliders has never been discussed in the DIF literature before but

analyzing a selected sample frequently happens in practice. The discussion and the presented scenarios would be helpful for researchers and practitioners who may encounter “DIF results that they cannot understand” (Angoff, 1993, p. 19).

The proposed causal graphs and scenarios need to be viewed as a prototype. They cannot be applied to all situations, at least not in their original forms. However, the basic principles articulated in the article can be easily modified and extended in accordance with any given specific causal story. Recently, some researchers have started investigating the use of various covariates for DIF analysis (e.g., Gibbons et al., 2011; Tay, Huang, Vermunt, 2016). The graphical models approach can also help one to figure out which covariates are necessary to control for and which covariate should *not* be controlled for (i.e., backdoor or adjustment criteria). Additionally, the discussion of DIF in the mediation or moderation context (Zumbo & Gelin, 2005) might be further firmly framed with recent causal studies that are bringing new insights into traditional mediation and moderation analysis (e.g., Pearl, 2014; VanderWeele, 2009). Although this article discusses a rather simple case where local independence is assumed, cases where it is violated can also be easily formalized and investigated using causal graphs (e.g., Kreiner & Christensen, 2002).

APPENDIX

Following item parameters were used in the simulation. For 1PL models across scenarios, the discrimination parameter was held at 1.2 for all 25 items while the difficulty parameters were $-.63, .18, -.84, 1.60, .33, -.82, .49, .74, .58, -.31, 1.51, .39, -.62, -2.22, 1.13, -.05, -.02, .94, .82, .59, .92, .78, .08$, and -1.99 for 24 non-DIF items. For the DIF-induced item (i.e., 25th item), $.40$ ($U = 0$) and $-.40$ ($U = 1$) were used. The difficulty parameters were generated

from a standard normal distribution and were rounded to two decimal places. For 2PL models, the same difficulty parameters were used, and the discrimination parameters were generated from a random uniform distribution ranging from .7 to 1.5 and again were rounded to two decimal places. The chosen discrimination parameters were .86, 1.25, 1.43, .93, .78, 1.26, 1.12, 1.35, 1.47, .79, .92, 1.09, .96, 1.15, .91, .86, 1.01, 1.41, 1.14, 1.37, 1.41, 1.28, .87, .88, and 1.20, respectively. For 3PL models, other than the DIF-induced item, the same difficulty parameters were used for 24 non-DIF items. For the DIF-induced item, the item difficulty parameter was set to .5 for $U = 0$ and $-.5$ for $U = 1$ to produce .8 difference in area between the two item response curves. The item discrimination parameters were the same with 2PL models for all 25 items. The pseudo-guessing parameters were fixed as .2 across all the items.

REFERENCES

- Ackerman, T. A. (1992). A didactic explanation of item bias, item impact, and item validity from a multidimensional perspective. *Journal of Educational Measurement*, 29(1), 67-91.
- American Educational Research Association, American Psychological Association, & National Council on Measurement in Education (Eds.). (2014). *Standards for educational and psychological testing*. American Educational Research Association.
- Andrich, D., & Hagquist, C. (2012). Real and artificial differential item functioning. *Journal of Educational and Behavioral Statistics*, 37(3), 387-416.
- Angoff, W. H. (1993). Perspectives on differential item functioning methodology. In P. W. Holland, & H. Wainer (Eds.), *Differential item functioning* (pp. 3-24). New York, NY: Routledge.

- Angoff, W. H., & Sharon, A. T. (1974). The evaluation of differences in test performance of two or more groups. *Educational and Psychological Measurement*, 34, 807-816.
- Berkson, J. (1946). Limitations of the Application of Fourfold Table Analysis to Hospital Data. *Biometrics Bulletin*, 2(3), 47-53.
- Bickel, P. J., Hammel, E. A., & O'Connell, J. W. (1975). Sex bias in graduate admissions: Data from Berkeley. *Science*, 187(4175), 398-404.
- Cleary, T. A., & Hilton, T. L. (1968). An investigation of item bias. *Educational and Psychological Measurement*, 28(1), 61-75.
- de Ayala, R. J. (2009). *The theory and practice of item response theory*. New York, NY: The Guilford Press.
- Elwert, F., & Winship, C. (2014). Endogenous selection bias: The problem of conditioning on a collider variable. *Annual Review of Sociology*, 40, 31-53.
- Finch, W. H., & Frensh, B. F. (2019). *Education and Psychological Measurement*. New York, NY: Routledge.
- Gibbons, L. E., Crane, P. K., Mehta, K. M., Pedraza, O., Tang, Y., Manly, J. J., ... & Mungas, D. (2011). Multiple, correlated covariates associated with differential item functioning (DIF): Accounting for language DIF when education levels differ across languages. *Ageing Research*, 2(1), 19-25.
- Holland, P. W. (1986). Statistics and causal inference. *Journal of the American Statistical Association*, 81(396), 945-960.
- Imbens, G., & Rubin, D. (2015). *Causal inference for statistics, social, and biomedical sciences: An introduction*. New York, NY: Cambridge University Press.

- Kreiner, S., & Christensen, K. B. (2002). Graphical Rasch models. In M. Mesbah, B. F. Cole, & M.-L. T. Lee (Eds.) *Statistical methods for quality of life studies* (pp. 187-203). Boston, MA: Springer.
- Lee, S. (2017). Detecting differential item functioning using the logistic regression procedure in small samples. *Applied Psychological Measurement*, 41, 30-43.
- Magis, D., Beland, 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.
- Martinková, P., Drabinová, A., Liaw, Y. L., Sanders, E. A., McFarland, J. L., & Price, R. M. (2017). Checking equity: Why differential item functioning analysis should be a routine part of developing conceptual assessments. *CBE—Life Sciences Education*, 16(2), rm2.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems*. San Mateo: Morgan Kaufman.
- Pearl, J. (2009). *Causality: Models, reasoning, and inference* (2nd ed.). New York, NY: Cambridge University Press.
- Pearl, J. (2012). The causal foundations of structural equation modeling. In R. H. Hoyle (Ed.), *Handbook of structural equation modeling* (pp. 68–91). New York, NY: The Guilford Press.
- Pearl J. (2014). Interpretation and identification of causal mediation. *Psychological Methods*, 19(4), 459–481.
- Pearl, J., & Mackenzie, D. (2018). *The book of why: The new science of cause and effect*. New York, NY: Basic Books.
- Penfield, R. D., & Camilli, G. (2007). Differential item functioning and item bias. In C. R. Rao, & S. Sinharay (Eds.), *Handbook of Statistics*, 26: *Psychometrics* (pp. 125-167).

- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5), 688-701.
- Shepard, L. A., Camilli, G., & Williams, D. M. (1985). Validity of approximation techniques for detecting item bias. *Journal of Educational Measurement*, 22(2), 77-105.
- Spirtes, P., Glymour, C. N., & Scheines, R. (2001). *Causation, prediction, and search* (2nd ed.). Cambridge, MA: MIT Press.
- Tay, L., Huang, Q., & Vermunt, J. K. (2016). Item response theory with covariates (IRT-C) assessing item recovery and differential item functioning for the three-parameter logistic model. *Educational and Psychological Measurement*, 76(1), 22-42.
- VanderWeele, T. J. (2009). On the distinction between interaction and effect modification. *Epidemiology*, 20(6), 863–871.
- VanderWeele, T. J., & Hernán, M. (2012). Causal effects and natural laws: Towards a conceptualization of causal counterfactuals for nonmanipulable exposures, with application to the effects of race and sex. In C. Berzuini, P. Dawid, & L. Bernardinelli (Eds.), *Causality: Statistical perspectives and applications* (pp. 101-113). Chichester, John Wiley and Sons.
- Woods, C. M., Oltmanns, T. F., & Turkheimer, E. (2009). Illustration of MIMIC-model DIF testing with the schedule for nonadaptive and adaptive personality. *Journal of Psychopathology and Behavioral Assessment*, 31(4), 320–330.
- Zumbo, B. D. (1999). *A handbook on the theory and methods of differential item functioning (DIF)*. Ottawa, Ontario, Canada: Directorate of Human Resources Research and Evaluation, Department of National Defense.

Zumbo, B. D. (2007). Three generations of DIF analyses: Considering where it has been, where it is now, and where it is going. *Language Assessment Quarterly*, 4(2), 223-233.

Zumbo, B. D., & Gelin, M. N. (2005). A matter of test bias in educational policy research: Bringing the context into picture by investigating sociological/community moderated (or mediated) test and item bias. *Journal of Educational Research & Policy Studies*, 5(1), 1-23.