Ubiquitous bias & false discovery due to model misspecification in analysis of statistical interactions: The role of the outcome's distribution and metric properties

Benjamin W. Domingue^{1,†}, Klint Kanopka², Sam Trejo³, Mijke Rhemtulla⁴, and Elliot M. Tucker-Drob⁵

¹Graduate School of Education, Stanford University & Center for Population Health Sciences, Stanford Medicine

²Graduate School of Education, Stanford University

³La Follette School of Public Affairs & Department of Sociology, University of Wisconsin–Madison

⁴Department of Psychology, University of California, Davis

⁵Department of Psychology & Population Research Center, University of Texas at Austin

[†]Correspondence about the paper should be sent to ben.domingue@gmail.com.

Abstract

Lived experience suggests substantial heterogeneity in how people react to a common stimuli. Previous work has considered several potential sources of bias and confusion in studying interactions but less attention has been devoted to the nature of the outcome variable in such studies. Here, we consider bias and false discovery associated with estimates of interaction parameters as a function of the distributional and metric properties of the outcome. Focusing on a variety of models for non-continuously distributed outcomes (binary and count outcomes), we show that attempts to use the linear model for recovery lead to catastrophic levels of bias and false discovery. Focusing on transformations of normally distributed variables (i.e., censoring and departures from interval scaling), we show that linear models again produce spurious interaction effects when such interactions are absent from the generating model. We provide explanations offering geometric and algebraic intuition as to why interactions can be a challenge for these incorrectly specified models. In light of these findings, we make two specific suggestions. First, a careful consideration of the outcome’s distributional properties should be a standard component of interaction studies. Second, researchers should approach research focusing on interactions with heightened levels of scrutiny.

1 Introduction

Lived experience suggests substantial heterogeneity in how people react to a common stimuli.¹ This suggests the hypothesis that some feature (of the stimuli, environment, or person) may explain this heterogeneity. For example, which psychological factors account for variation in the success of smoking cessation programs (Halpern et al., 2016)? Does the level of public regard moderate the effect of racial discrimination on depressive symptoms (Seaton & Iida, 2019)? Are effects of growth mindset interventions moderated by environmental features (Claro et al., 2016; Yeager et al., 2019)? These few examples are a small portion of the larger literature probing for such heterogeneities in empirical settings. We believe, along with others (Bryan et al., 2021), that such questions are crucial.

The standard tool for analysis of such heterogeneity is the inclusion of interaction terms in statistical models. However, analysis of such models is fraught. We are not the first to note this problem; this issue has been remarked upon regularly over the last several decades (e.g., Loftus, 1978; Lubinski and Humphreys, 1990; Rohrer and Arslan, 2021; Tabery, 2008); we briefly discuss previously highlighted methodological issues below. Here, we focus specifically on issues arising due to measurement properties of the outcome variables. Statistical models for interactions rely upon strong assumptions regarding the distribution and measurement

¹This is a paraphrase of what we have heard Jeremy Freese call the First Law of Sociology—“Some do, some don’t”—but psychologists also use it (Haaf & Rouder, 2019).

properties of the outcome variable. When these assumptions are not met, conventional regression models may produce biased interaction effect estimates and high rates of type 1 error; that is, such models tend to produce spurious positive findings due to misspecification. Our work builds specifically on previous studies emphasizing the role of outcome scale in subsequent interpretation (Loftus, 1978) but with a focus on bias and false discovery.

1.1 Previously expressed concerns regarding analysis of statistical interactions

We assume that interest is in the interaction of two predictors (x and z) in the study of some outcome (y). The outcome is a function of the predictor variables and model parameters: $y = f(\beta_0 + \beta_1x + \beta_2z + \beta_3xz)$ for some function f (e.g., we may consider $f(t) = t + \epsilon$ where $\epsilon \sim \text{Normal}(0, \sigma^2)$). Interest is in an estimate of β_3 . Previous studies have emphasized several potential problems:

1. Studies examining interaction may yield false positives if interactions between additional covariates (w) are not also included in the model (Keller, 2014). If w is a covariate of interest, then analysis of the interaction xz should be based on models that also include xw and zw .
2. If the true data generating model is based on x^2 , interaction studies focusing on xz may lead to false positives if x and z are correlated (Lubinski & Humphreys, 1990; MacCallum & Mar, 1995).
3. The substantive implications of findings of interactions may be highly contingent on the nature of x and z . In particular, endogeneity may lead to multiple viable interpretations (Fletcher & Conley, 2013).
4. Under certain configurations, statistical power of interaction studies will be substantially lower than in studies of main effects (see Section 16.4 of Gelman et al., 2020).

We highlight these previous concerns for two purposes. First, they help to emphasize that there are numerous challenges to the statistical analysis of interactions. Second, much of the previous scholarship on the estimation of interactions has focused specifically on the “right hand side” of the model; relatively less focus has been paid to the outcomes. We focus explicitly on the “left hand side” and show that characteristics of the outcome can have major implications for our ability to recover the relevant parameters.

1.2 Outcome types

We consider a range of outcomes that collectively span a broad range of outcomes of interest in psychology. Below we briefly describe the outcomes, where they may be used in psychological research, and the typical analytic approaches used to study them.

- Binary outcomes are common in analysis of, for example, mental health diagnoses (Ancelin et al., 2017; Culverhouse et al., 2018; Stringa et al., 2020). Such outcomes may be analyzed via logistic regression, a specific version of the generalized linear model (Nelder & Wedderburn, 1972).
- Count outcomes occur when interest is in the number of times something occurred; they are used to quantify, for example, drug/alcohol use (Angosta et al., 2019; Meyers et al., 2019) or frequency of medical care (Minkovitz et al., 2005; Richardson & Ratner, 2005). Such outcomes may be analyzed via Poisson regression, a specific version of the generalized linear model (Nelder & Wedderburn, 1972).
- Censored outcomes occur when a scale is limited and values above or below a threshold are effectively capped. In many cases, censoring can be readily observed (i.e., by looking at a histogram). Test scores, for example, commonly have ceilings or floors (e.g., Koedel and Betts, 2010). However, many scales used in psychological research may similarly exhibit censoring; e.g., depression scales (Djukanovic et al., 2017; Page et al., 2007). Censored data may be modeled via the Tobit model (Tobin, 1958).
- Non-interval outcomes: Many psychological constructs are measured via aggregating scores across multiple tests or item responses to produce composite indices. For example, the number of correct responses on an achievement test or the number of symptoms indicated on a depression scale can be summed to produce a composite score, or aggregated using more advanced methods (e.g. empirical Bayes) for estimating factor scores in the context of factor analytic or item response theory models (van der Linden, 2016). However, neither the raw sum scores nor the outcomes produced by latent variable models necessarily have interval scales. For instance, when a test has an unbalanced distribution of easy and difficult items relative to the ability distribution of the sample, the relation between the latent ability and the composite index may become nonlinear (see Figure 3 of Tucker-Drob, 2009). The interval interpretation allows us to suppose that a one unit change across the scale consistently has the same meaning. This interpretation is valid for many physical measures (e.g., a one meter change in length always means a change by a standard amount) but does not necessarily hold for psychological measures (Michell, 2008). Data may have structure necessary for interval interpretations (Domingue, 2014), but in many cases we do not have strong evidence of such structure. A failure to have this equal interval property can have implications for a variety of scale uses (Ballou, 2009). Given that the nature of the non-interval scale is typically unknown, appropriate analysis is typically impossible.

Alongside the analytic approaches tailored for the specific features of the various outcomes, it is also common for the linear model to be used for analyzing outcomes of all types. For example, the linear probability model is frequently used in the case of binary outcomes (Gomila, 2020). We similarly consider analysis for

all outcomes via both tailored approaches and a (misspecified) linear alternative. The linear model is a valuable tool and robust for many purposes but we show that it can lead to highly misleading results when used to study interactions as a function of the geometry of the outcome variable.

1.3 Key contributions

We focus on describing bias and type 1 error across the four outcomes when divided into two categories. We first examine non-continuous outcomes (binary and count outcomes) We show that application of linear modeling techniques in such a scenario results in disastrous levels of type 1 error. Second, we consider continuously measured outcome variables that do not meet classical assumptions of the linear model (censored and non-interval outcomes) We show that such outcomes—when viewed as transformed versions of classic linear outcomes—similarly produce undesirable levels of type 1 error when analyzed with the linear model. We also provide geometric and algebraic guidance for why type 1 error rates are so high. Evidence from both settings suggest a need for heightened scrutiny of the characteristics of the outcome in interaction studies.

2 Methods

2.1 A running example

We consider a running example so as to describe results in a common way. We suppose that we observe some outcome y and interest is in predictors x and z as well as their interaction. We'll assume that $(x_i, z_i) \sim \text{MVN}(\mu, \Sigma)$ where $\mu = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$ and $\Sigma = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$. We focus on $\rho = 0$ so as to clarify that the issue is about the product term even if x and z are not associated. We define a quantity, Γ , that plays a key role throughout. This quantity is defined as

$$\Gamma_i = \beta_0 + \beta_1 x_i + \beta_2 z_i + \beta_3 x_i z_i. \quad (1)$$

In interaction studies, interest is typically in estimates of β_3 . Specific choices for the relevant parameters (i.e., $\beta_0, \beta_1, \beta_2, \sigma^2$) and sample size (N) are documented throughout but we view them as suggestive; they are chosen to help us illustrate the pronounced problems of bias and false positives *irrespective* of the specific configuration of parameters.

If $\mathbb{E}(y_i|x_i, z_i) = \Gamma_i$, then the linear model (LM) is appropriate for estimation of the β parameters. We focus here on cases wherein $\mathbb{E}(y_i|x_i, z_i) \neq \Gamma_i$. In Section 3.1, we consider transformations of Γ that yield discrete outcomes; we do this by supposing that $\mathbb{E}(y_i|x_i, z_i) = g(\Gamma_i)$ for different g . Suppose, for example, that we are interested in binary outcomes; in that case, we may assume that $\mathbb{E}(y_i) = \sigma(\Gamma_i)$ where σ is

the logistic sigmoid, $\sigma(x) = 1/(1 + e^{-x})$. We consider outcomes belonging to the family of generalized linear models (GLM; Nelder and Wedderburn, 1972) with non-identity link (binary and count outcomes). In Section 3.2, we suppose that $\mathbb{E}(y_i^*|x, z) = \Gamma_i$ but instead of observing y^* we observe some transformed version y (i.e., $g(y_i^*) = y_i$ for some transformation g). We focus on transformations g that induce metric limitations in y_i .

2.2 Four approaches for illustrating bias and false discovery

We take four approaches to illustrating bias and false discovery.

1. The first two approaches are based on visualizations for select values of the relevant parameters. We first consider figures that plot y or $\mathbb{E}(y)$ as a function of xz . Such figures capture the breakdowns in symmetry that induce false positives.
2. We then consider plots showing y or $\mathbb{E}(y)$ for select values of z . These plots illustrate the fact that the best fit lines for different values of z are not parallel (thus indicating an interaction effect) as we vary the key parameters.
3. We next consider Taylor Series expansions of f . Taylor series expansions allow for representation of a function f as a function of an infinite sum of f 's derivatives. Use of this technique requires certain smoothness conditions on f (and thus won't be applicable in all cases), but, where applicable, allows us to offer additional insight into the algebra leading to type 1 errors.
4. Finally, to demonstrate the ubiquity of the problem, we calculate bias and type 1 error rates via simulation studies as a function of the relevant model parameters and sample size.

This sequence will allow us to see the correlations between xz and the outcome, illustrate the differences in linear fits that suggest the existence of an interaction, offer an algebraic rationale for why this comes about, and give some guidance on how large a problem this is under different conditions.

2.3 Simulation Details

In all cases, interest will focus on estimates of β_3 . We denote such estimates as $\widehat{\beta}_3$ but emphasize that $\beta_3 = 0$ in fact. We consider sample sizes of 250 to 1000.² Note that the problems identified for samples of 1000 would only be magnified in larger samples. For any given configuration of conditions, we consider

²To offer context for this choice of sample size, we surveyed 150 studies published between January 1, 2016 and October 20, 2020 focused on depression (based on the Pubmed query: `(depression[Title]) AND (heterogeneity[Title] OR moderation[Title] OR interaction[Title])`). Of the ascertained studies, 30% were empirical studies along the lines of what we consider here. The sample sizes ranged widely (min=22, median=576, IQR=292-2279, max=134357); our choice of 250-1000 is meant to reflect the center of this distribution.

results from analysis of 1000 simulated datasets. All analyses were conducted in R. Estimation of linear, logistic, and Poisson models is done via base R functions; estimation of Tobit models via tailored packages (Henningsen, 2010). For each specific combination of parameters, we simulated/analyzed 1,000 datasets. Code to replicate all analyses is available on GitHub.³

3 Results

3.1 Transformations of Γ

3.1.1 Binary outcomes

Suppose y_i is a binary outcome; we suppose that y_i is a Bernoulli random variable with $\Pr(y_i = 1) = (1 + \exp(-\Gamma))^{-1}$. That is, we assume that data are generated via the logistic model but the key observation regarding false discovery and the LM is not sensitive to this fact (e.g., the Probit model could be used). We suppose that the LM is used for estimation; that is, we ignore the fact that y_i is binary and assume that $y_i \sim \text{Normal}(\Gamma_i, \sigma_y^2)$. Application of the LM may be useful given that parameter estimates are straightforward to interpret and there may also be computational reasons for its use (e.g., estimation in the context of relatively large datasets with hierarchical structure can become computationally expensive and the LM may help to offset this fact). However, we emphasize one key point: Interaction analysis with the LM may be a flawed approach depending on the prevalence of the outcome. We have discussed this issue in a specific context elsewhere (Domingue et al., 2020; Trejo et al., 2018) but here show it in its general form.

We focus here on the problem induced by variation in the intercept (β_0), which is related to the prevalence of the outcome (i.e., $\bar{y}$). Figure 1 considers $\mathbb{E}(y)$ as a function of xz . When $\beta_0 = 0$, there is a clear structure to the attached plot but we emphasize that there is no linear association between $\mathbb{E}(y)$ and xz (note the small correlation). However, as β_0 increases, the symmetry in the scatterplot breaks down; the resulting asymmetry leads to a large correlation and thus false positives.

To illustrate the underlying geometry driving this correlation, Figure 2 shows $\mathbb{E}(y)$ as a function of x along with the implied logistic (solid lines) and linear (dashed lines) fits for two different values of z ($z = \pm 1$). When $\beta_0 = 0$, the geometry of the $\mathbb{E}(y)$ values is such that both the logistic and linear fits for $z = \pm 1$ are clearly parallel (thus indicating $\beta_3 = 0$). However, as β_0 increases, both the slope and intercept of the linear fit exhibit a dependence on z , and the geometry of the $\mathbb{E}(y)$ values is such that there is an implied interaction. This is due to the fact that the $z = 1$ values in blue are observed at a location where the true underlying sigmoid is quite flat thus leading to an inference that $\beta_3 < 0$. For sufficiently large β_0 , the problem begins to

³See <https://github.com/ben-domingue/interaction-problems>.

Figure 1: Scatterplot of $\mathbb{E}(y)$ as a function of xz for different values of β_0 ($\beta_1 = \beta_2 = 1, \beta_3 = 0, N = 1000$) when y is a binary outcome.

resolve given that both $z = -1$ and $z = 1$ values are observed wherein the underlying sigmoid is relatively flat.

Figure 2: Illustration of the geometry driving false discovery due to variation in β_0 when the linear model is used ($\beta_1 = \beta_2 = 1, \beta_3 = 0, N = 1000$) for analysis of binary outcomes. Blue and red dots represent those data points within half a unit of their respective z values (i.e., z such that $|z - 1| < 0.5$ are in blue and z such that $|z + 1| < 0.5$ are in red. Fitted lines for ± 1 are similarly shaded). Solid lines are fits from logistic regression model while dashed lines are fits from linear model.

As an algebraic illustration of the problem, suppose that $\mathbb{E}(y|x, z) = \sigma(\beta_0 + \beta_1 x + \beta_2 z)$, where $\sigma(\cdot)$ is the standard logistic sigmoid. Using a Taylor series expansion about zero, we can write

$$\mathbb{E}(y|x, z) = \frac{1}{2} + \frac{1}{4}(\beta_0 + \beta_1 x + \beta_2 z) + \frac{1}{48}(\beta_0 + \beta_1 x + \beta_2 z)^3 + \dots \quad (2)$$

When expanding the cubic term, the highest weighted term is an interaction term including x and z : $\frac{1}{8}\beta_0\beta_1\beta_2 xz$. By modeling this expectation as $\mathbb{E}(y|x, z) = b_0 + b_1 x + b_2 z + b_3 xz$, we should expect false

discovery of an interaction effect due to this term. Balanced classes (or $\beta_0 = 0$) improve the situation but do not fully solve this problem, as the cubic term in the Taylor series still expands to contain the higher order interaction terms $\frac{1}{16}\beta_1^2\beta_2x^2z$ and $\frac{1}{16}\beta_1\beta_2^2xz^2$ that the LM may attribute to an interaction between x and z .

With both this geometric and algebraic intuition at hand, we focus on the associated bias in estimates $\widehat{\beta}_3$ from the LM as a function of β_0 where $\beta_3 = 0$. Given that x and z are centered at 0, expected prevalence of the outcome is $\frac{e^{\beta_0}}{1+e^{\beta_0}}$ (e.g., expected prevalence is 0.88 for $\beta_0 = 2$). In the left panel of Figure 3, we can see that LM estimates of β_3 are unbiased when $\beta_0 = 0$ but heavily biased when $\beta_0 \approx 2$; in our simulated datasets, the 10th through 90th percentile of estimates is far from zero. This bias decreases, for reasons shown in Figure 2, for large values of β_0 . Note that the degree of bias is independent of sample size. The right panel of Figure 3 shows the implications in terms of type 1 error for two selected values of β_0 . Here we also consider GLM estimates; note that they have uniformly correct type 1 error (given $\alpha = 0.05$). When $\beta_0 = 0$, levels of false discovery are appropriate (as expected given the right panel). However, for $\beta_0 = 2$, false discovery is uniformly problematic and increasing sample size only increases the level of false discoveries. With binary outcomes, analysis of statistical interactions needs to carefully consider the level of prevalence if the LM is to be considered.

Figure 3: Variation in bias (left) and Type 1 error rate (right) associated with $\widehat{\beta}_3$ when $\beta_3 = 0$ for binary outcomes (for $\beta_0 = 0, \beta_1 = \beta_2 = 1$). Left: Estimates of β_3 based on the LM for two sample sizes as a function of β_0 . Shaded regions capture span of estimates between 10th and 90th percentile while the solid line shows median estimates. Right: Levels of type 1 error as a function of sample size for two values of β_0 .

3.1.2 Count outcomes

We now pivot to another outcome in the GLM framework and focus on the scenario where y_i is a count outcome. We suppose that data are generated via the Poisson model; thus, $\mathbb{E}(y_i) = \exp(\Gamma_i)$. False discovery is again a catastrophic problem if the linear model is used in this setting. We begin to illustrate this fact

via Figure 4. We first show $\mathbb{E}(y)$ as a function of xz but this time we introduce variation in β_2 . Note that there is no linear correlation apparent in the scatterplot when β_2 is relatively small. However, increase in β_2 induces an association despite the fact that $\beta_3 = 0$. We illustrate the geometry leading to this false discovery in Figure 5. We continue to vary β_2 and hold $\beta_1 = 0.2$. As β_2 increases, seeming variation in $\mathbb{E}(y)$ as a function of z is introduced. This would imply $\beta_3 \neq 0$ in the incorrectly specified linear approach.

Figure 4: Scatterplot of $\mathbb{E}(y)$ as a function of xz for different values of β_2 ($\beta_0 = 1, \beta_1 = 0.2, N = 1000$) when y_i is a count outcome.

As an algebraic illustration, we can again consider a Taylor series expansion. Suppose that $\mathbb{E}(y|x, z) = \exp(\beta_1 x + \beta_2 z) = \exp(\beta_1 x) \exp(\beta_2 z)$. We can use Taylor series expansion to write this as

$$\mathbb{E}(y|x, z) = (1 + \beta_1 x + (\beta_1 x)^2/2 + \dots)(1 + \beta_2 z + (\beta_2 z)^2/2 + \dots) = 1 + \beta_1 x + \beta_2 z + \beta_1 \beta_2 xz + \dots \quad (3)$$

If we mistakenly assume that $y = b_1 x + b_2 z + b_3 xz + \epsilon$, we would anticipate $b_3 = \beta_1 \beta_2$ if we omit higher order terms. Thus, if either β_1 or β_2 does not equal zero, false discovery will result when the LM is deployed for analysis.

Figure 5: Illustration of the geometry driving false discovery due to variation in β_2 when the linear model is used ($\beta_0 = 1, \beta_1 = 0.2, N = 1000$) for analysis of count outcomes. Blue and red dots represent those data points within half a unit of their respective z values (i.e., z such that $|z - 1| < 0.5$ are in blue and z such that $|z + 1| < 0.5$ are in red. Fitted lines for ± 1 are similarly shaded). Solid lines are fits from Poisson regression model while dashed lines are fits from linear model.

Implications of this problem for bias and false discovery are shown in Figure 6. At left, bias in $\widehat{\beta}_3$

increases as a function of β_2 ; the performance of the LM for analysis of count outcomes will be sensitive to the magnitude of the main effects. At right, we observe that the bias introduced by β_2 induces type 1 errors. Analysis of larger samples will be plagued by false positives in such an instance. However, analysis of the appropriately specified GLM yields the correct type 1 error rate.

Figure 6: Variation in bias (left) and Type 1 error rate (right) associated with $\widehat{\beta}_3$ when $\beta_3 = 0$ for count outcomes (for $\beta_0 = 0, \beta_1 = 0.5$). Left: Estimates of β_3 based on the LM for two sample sizes as a function of β_2 . Shaded regions capture span of estimates between 10th and 90th percentile while the solid line shows median estimates. Right: Levels of type 1 error as a function of sample size for two values of β_2 .

3.2 Transformations of y^*

3.2.1 Censored Outcomes

We first consider transformations g of y^* that lead to scales with a ceiling or floor (Garin, 2014). We focus on a floor but the same concerns would apply to ceilings. To simulate data, we utilize the Tobit model (Tobin, 1958). We suppose that

$$y_i^* \sim N(\Gamma_i, \sigma_{y^*}^2). \quad (4)$$

However, we do not observe y_i^* ; rather, we observe

$$y_i = g(y_i^*) = \begin{cases} y_i^* & y_i^* > c \\ c & y_i^* \leq c \end{cases} \quad (5)$$

for some constant c .⁴

We begin to illustrate the reason for type 1 errors when the LM is deployed due to the existence of floors in Figure 7. We show y as a function of xz in gray versus y^* in red. When the floor is sufficiently low, few

⁴Note that the censoring violates the smoothness conditions required to conduct the Taylor series expansion.

points are affected and recovery with the censored data yields the correct inference. This is due to the fact that we observe the full shape of (xz, y) . However, when the floor begins to censor a significant number of cases, we no longer observe the bottom half of the shape. When these points are raised towards the y mean, they are downweighted in the sum leading to the correlation coefficient and thus lead to the resulting (spurious) positive correlation between xz and y .

Figure 7: Scatterplot of y and y^* as a function of xz for different values of c ($\beta_0 = 0, \beta_1 = \beta_2 = 1, \sigma_y^2 = 0.25, N = 1000$) when y_i is a censored outcome.

The implications for fitted trajectories for different values of z are shown in Figure 8. Trajectories for relatively large values of z are unaffected (i.e., the solid and dashed blue lines are similar in all cases). However, as c increases, there is an increasing effect on the red line; when c is at its largest, the trajectory based on observed values is flatter than the trajectory based on uncensored values thus suggesting an interaction under the naive LM.

Figure 8: Illustration of the geometry driving false discovery due to variation in c when the linear model is used ($\beta_0 = 0, \beta_1 = \beta_2 = 1, \sigma_y^2 = 0.25, N = 1000$) for analysis of censored outcomes. Blue and red dots represent those data points within half a unit of their respective z values (i.e., z such that $|z - 1| < 0.5$ are in blue and z such that $|z + 1| < 0.5$ are in red. Fitted lines for ± 1 are similarly shaded). Solid lines are fits from Tobit regression model while dashed lines are fits from linear model.

Figure 9 extends this geometric logic to an analysis of bias and type 1 error as a function of c . On the left, we first consider LM-based estimates of $\hat{\beta}_3$. As c increases, so does bias in $\hat{\beta}_3$. This translates into an increase in type 1 error which we consider on the right. Bias leads to highly elevated levels of false discovery even for modestly sized samples given the configurations considered here even for a modest value of c . The

bias leads to catastrophic type 1 error as the floor affects a larger proportion of outcomes (i.e., $c = -1$), here, increasing sample size serves only to increase the salience of the problem. In contrast, Tobit estimates produce type 1 errors at the expected rate for both values of c irrespective of sample size.

Figure 9: Variation in bias (left) and Type 1 error rate (right) associated with $\widehat{\beta}_3$ when $\beta_3 = 0$ for outcomes with a floor ($\beta_0 = 0, \beta_1 = \beta_2 = 1, \sigma_y^2 = 1$). Left: Estimates of β_3 based on the LM for two sample sizes as a function of c . Shaded regions capture span of estimates between 10th and 90th percentile while the solid line shows median estimates. Right: Levels of type 1 error as a function of sample size for two values of c .

3.2.2 Non-interval Outcomes

As a final example, we consider the impact of distortions that lead to non-interval scales on our ability to make accurate inferences. To do this, we consider a monotonic but non-linear transformation of y_i^* . Such a transformation is motivated by, for example, the work of Michell (1986) which surfaces the issue of measurement scales in the context of psychological constructs. Scales are frequently assumed to be interval but need not be. In particular, we consider “Lord’s transformation” (Briggs & Betebenner, 2009) which stretches the scale for larger values of y_i^* . That is, we suppose that $\mathbb{E}(y_i^*|x, z) = \Gamma_i$ but that we observe

$$y_i = g(y_i^*) = 1.05^{(y_i^* - \alpha)/\lambda}. \quad (6)$$

Here we consider $\alpha = 0$ and allow λ to vary. This transformation has the effect of making a one unit difference in y have differential meaning in the y^* metric, see illustration in Figure A.1. When λ is small, distances between smaller values of y^* are compressed while distances between larger values are inflated. Differences between the distances quickly dissipate as λ increases. The transformation’s effect has similarities to the effect of the a floor in that, in both cases, differences between small values of y^* are being minimized (i.e., false positives are generated here for the same reason as in Figure 7); given these similarities, we provide Figures A.2 and A.3 but do not further discuss them.

Next we look to a Taylor series expansion of Equation 6 about zero. Supposing that $\alpha = 0$, we have $\mathbb{E}(y|x, z) = 1.05^{\frac{1}{\lambda}\Gamma}$. Looking at the first three terms, we see:

$$\mathbb{E}(y|x, z) = 1.05^{\frac{1}{\lambda}\Gamma} = 1 + \ln 1.05 \frac{1}{\lambda}\Gamma + \frac{1}{2} \ln^2 1.05 \left(\frac{1}{\lambda}\Gamma\right)^2 + \dots$$

Expanding the third term, we produce a term of $\frac{\ln^2 1.05}{\lambda^2} \beta_1 \beta_2 xz$. If both β_1 and β_2 are not equal to zero, this term will thus lead to false discovery of an interaction coefficient when the LM is used. Given that λ is in the denominator, this problem is exacerbated at smaller values of λ .

Figure 10 focuses on the bias and type 1 error induced by this transformation. Note first that we do not consider estimates that are adjusted for the outcome's specific distributional properties; the fact that the outcome is not measured on an interval scale is typically not something that is directly known (Domingue, 2014). This increases the salience of the following observation about the LM estimates since they would likely be used in place of the correct alternative. In particular, we can see a clear sensitivity in estimates of β_3 to the value of λ . For small values of λ , there are larger deviations from the true interval scale which leads to an increase in bias. This then translates into a higher level of type 1 error rate when λ is relatively small.

Figure 10: Variation in bias (left) and Type 1 error rate (right) associated with $\widehat{\beta_3}$ when $\beta_3 = 0$ for outcomes observed following Lord's transformation ($\beta_0 = 0, \beta_1 = \beta_2 = 1, \sigma_y^2 = 1, \alpha = 0$). Left: Estimates of β_3 based on the LM for two sample sizes as a function of λ . Shaded regions capture span of estimates between 10th and 90th percentile while the solid line shows median estimates. Right: Levels of type 1 error as a function of sample size for two values of λ .

4 Discussion

The study of interactions is of clear interest in many settings. However, the specific properties of the outcome variables in question may make accurate inference in such settings challenging. To illustrate this point, we considered two types of outcomes: outcomes based on a transformation of Γ and outcomes based on a transformation of some intermediate variable y^* where $\mathbb{E}(y_i^*|x, z) = \Gamma_i$. In both cases, there are a variety of problems that may threaten inference; some of these issues are relatively easy to diagnose and resolve, others less so. We separately discuss implications for analysis of both types of outcomes before concluding with general thoughts on analysis of interactions.

For outcomes based on transformations of Γ , there are clear gains to using the correctly specified alternative. In all cases, the correctly specified approach led to correct levels of type 1 error whereas the LM-based approach led to highly elevated levels of type 1 error. Although there are other considerations that may lead one away from the GLM approaches we consider here (Battey et al., 2019; Gomila, 2020; Ryan et al., 2018), we note that the problems we emphasize are unique to identification of interaction coefficients; analysis of main effects via the linear model is not necessarily compromised in the same way. We would urge heightened scrutiny when attempting to model interaction effects using alternatives that may be incorrectly specified.

For outcomes based on transformations of y^* , our results suggest that common problems will serve to increase type 1 error. Some of these problems are perhaps fairly easy to identify (i.e., the existence of a floor) and can frequently be addressed using standard solutions (i.e., models for censored data). Other problems (i.e., non-interval scales) may be hard to identify and nevertheless have deleterious effects on statistical inference. We view this as a rationale for a form of statistical humility—in particular, a single study should rarely be viewed as a dispositive piece of evidence regarding heterogeneity absent very large samples⁵ and very high-quality measurement—especially when studying outcomes with challenging measurement properties of the kind that abound in social science.

Much of what we discuss is not novel in the sense that these findings are anticipated by other studies focusing on power and false discovery in the context of misspecified models or other issues (e.g., censored outcomes). We view the issues of interactions as warranting special attention given interest in such studies. Issues of power are frequently emphasized in, for example, research proposals. Our findings suggest the need for carefully constructed power analyses if the goal is to derive appropriate guidance about, for example, sample sizes; such studies may be constructed using recently developed techniques (Jaccard & Brinberg, 2021). We think power calculations that are tailored to the situation at hand are important given that conclusions about power may be sensitive to conditions such as the correlation between various predictors.

⁵Although, we would note that increased sample size only serves to exacerbate some problems.

One potential response to these findings is that the problem of bias is easily avoided if correctly specified models are used. On the one hand, we agree and think this an underappreciated issue in the context of interaction studies. On the other hand, we would also note that in a great many cases it is not trivial to identify the correct model (given incomplete information about difficult-to-measure outcomes). In our view, a degree of humility is imperative.

Turning to the issue of false discovery, we make two key points. First, the issues of false discovery raised here cannot be dealt with by utilizing robust standard errors. The problems illustrated in Figures 2, 5, and 7 show that it is the parameter estimates themselves—not merely the standard errors—that are flawed (for a related argument about the limitations of robust standard errors, see King and Roberts, 2015). We have focused on type 1 error given that it directly calls attention to the likelihood of high rates of false positives in many interaction studies focused on outcomes that present measurement challenges. The second (closely related) point is that biased parameter estimates of the kind observed here may lead to a large quantity of spurious findings in research focusing on interactions. This is a concern and should lead to increased attention to the nature of the outcome variable in studies of statistical interactions.

We acknowledge limitations. We don't consider analysis using techniques such as structural equation modeling or mixed models. There are also outcome types that we do not consider. Despite the fact that our simulations only speak conclusively to some cases, we would still argue that these demonstrations evidence need for increased caution in studies of interactions.

If the “some do, some don't” formulation holds true—and, we believe that it does in many cases—interaction studies will clearly be of interest. Yet, poorly-designed interaction studies can lead fields into dead-ends: spurious interactions will be detected and true interactions will be obscured, reversed, or overestimated. Consider, for example, the era of candidate gene-by-environment studies (Duncan & Keller, 2011) which is now viewed as a vast literature consisting of almost entirely false positives. A failure to steer clear of dead-ends can lead to wasting large quantities of resources—both in terms of finite research dollars and even scarcer researcher time. Thus, we encourage researchers pursuing questions focused on studies of heterogeneity of the kind that invoke statistical interactions to take a more realistic perspective on the quality of inference likely to result from their particular data context.

References

- Ancelin, M.-L., Scali, J., Norton, J., Ritchie, K., Dupuy, A.-M., Chaudieu, I., & Ryan, J. (2017). Heterogeneity in hpa axis dysregulation and serotonergic vulnerability to depression. *Psychoneuroendocrinology*, 77, 90–94.

- Angosta, J., Steers, M.-L. N., Steers, K., Riggs, J. L., & Neighbors, C. (2019). Who cares if college and drinking are synonymous? identification with typical students moderates the relationship between college life alcohol salience and drinking outcomes. *Addictive behaviors*, 98, 106046.
- Ballou, D. (2009). Test scaling and value-added measurement. *Education finance and Policy*, 4(4), 351–383.
- Battey, H., Cox, D., & Jackson, M. (2019). On the linear in probability model for binary data. *Royal Society open science*, 6(5), 190067.
- Briggs, D., & Betebenner, D. (2009). Is growth in student achievement scale dependent. *Unpublished manuscript*.
- Bryan, C. J., Tipton, E., & Yeager, D. S. (2021). Behavioural science is unlikely to change the world without a heterogeneity revolution. *Nature human behaviour*, 5(8), 980–989.
- Claro, S., Paunesku, D., & Dweck, C. S. (2016). Growth mindset tempers the effects of poverty on academic achievement. *Proceedings of the National Academy of Sciences*, 113(31), 8664–8668.
- Culverhouse, R. C., Saccone, N. L., Horton, A. C., Ma, Y., Anstey, K. J., Banaschewski, T., Burmeister, M., Cohen-Woods, S., Etain, B., Fisher, H. L., et al. (2018). Collaborative meta-analysis finds no evidence of a strong interaction between stress and 5-HTTLPR genotype contributing to the development of depression. *Molecular psychiatry*, 23(1), 133–142.
- Djukanovic, I., Carlsson, J., & Årestedt, K. (2017). Is the hospital anxiety and depression scale (hads) a valid measure in a general population 65–80 years old? a psychometric evaluation study. *Health and quality of life outcomes*, 15(1), 1–10.
- Domingue, B. (2014). Evaluating the equal-interval hypothesis with test score scales. *Psychometrika*, 79(1), 1–19.
- Domingue, B., Trejo, S., Armstrong-Carter, E., & Tucker-Drob, E. (2020). Interactions between polygenic scores and environments: Methodological and conceptual challenges. *Sociological Science*, 7, 465–486.
- Duncan, L. E., & Keller, M. C. (2011). A critical review of the first 10 years of candidate gene-by-environment interaction research in psychiatry. *American Journal of Psychiatry*, 168(10), 1041–1049.
- Fletcher, J. M., & Conley, D. (2013). The challenge of causal inference in gene–environment interaction research: Leveraging research designs from the social sciences. *American journal of public health*, 103(S1), S42–S45.
- Garin, O. (2014). Ceiling effect. *Encyclopedia of quality of life and well-being research*. Dordrecht: Springer Netherlands, 631–633.
- Gelman, A., Hill, J., & Vehtari, A. (2020). *Regression and other stories*. Cambridge University Press.
- Gomila, R. (2020). Logistic or linear? estimating causal effects of experimental treatments on binary outcomes using regression analysis. *Journal of Experimental Psychology: General*, 1–27.

- Haaf, J. M., & Rouder, J. N. (2019). Some do and some don't? accounting for variability of individual difference structures. *Psychonomic Bulletin & Review*, 26(3), 772–789.
- Halpern, S. D., French, B., Small, D. S., Saulsgiver, K., Harhay, M. O., Audrain-McGovern, J., Loewenstein, G., Asch, D. A., & Volpp, K. G. (2016). Heterogeneity in the effects of reward-and deposit-based financial incentives on smoking cessation. *American journal of respiratory and critical care medicine*, 194(8), 981–988.
- Henningsen, A. (2010). Estimating censored regression models in r using the censreg package. *R package vignettes collection*, 5(2), 12.
- Jaccard, J., & Brinberg, M. (2021). Monte carlo simulations using extant data to mimic populations: Applications to the modified linear probability model and logistic regression. *Psychological Methods*.
- Keller, M. C. (2014). Gene $\times$ environment interaction studies have not properly controlled for potential confounders: The problem and the (simple) solution. *Biological psychiatry*, 75(1), 18–24.
- King, G., & Roberts, M. E. (2015). How robust standard errors expose methodological problems they do not fix, and what to do about it. *Political Analysis*, 159–179.
- Koedel, C., & Betts, J. (2010). Value added to what? how a ceiling in the testing instrument influences value-added estimation. *Education Finance and Policy*, 5(1), 54–81.
- Loftus, G. R. (1978). On interpretation of interactions. *Memory & Cognition*, 6(3), 312–319.
- Lubinski, D., & Humphreys, L. G. (1990). Assessing spurious “moderator effects”: Illustrated substantively with the hypothesized (“synergistic”) relation between spatial and mathematical ability. *Psychological bulletin*, 107(3), 385.
- MacCallum, R. C., & Mar, C. M. (1995). Distinguishing between moderator and quadratic effects in multiple regression. *Psychological Bulletin*, 118(3), 405.
- Meyers, J. L., Salvatore, J. E., Aliev, F., Johnson, E. C., McCutcheon, V. V., Su, J., Sally, I., Kuo, C., Lai, D., Wetherill, L., et al. (2019). Psychosocial moderation of polygenic risk for cannabis involvement: The role of trauma exposure and frequency of religious service attendance. *Translational psychiatry*, 9(1), 1–12.
- Michell, J. (1986). Measurement scales and statistics: A clash of paradigms. *Psychological bulletin*, 100(3), 398.
- Michell, J. (2008). Is psychometrics pathological science? *Measurement*, 6(1-2), 7–24.
- Minkovitz, C. S., Strobino, D., Scharfstein, D., Hou, W., Miller, T., Mistry, K. B., & Swartz, K. (2005). Maternal depressive symptoms and children's receipt of health care in the first 3 years of life. *Pediatrics*, 115(2), 306–314.

- Nelder, J. A., & Wedderburn, R. W. (1972). Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)*, 135(3), 370–384.
- Page, A. C., Hooke, G. R., & Morrison, D. L. (2007). Psychometric properties of the depression anxiety stress scales (dass) in depressed clinical samples. *British Journal of Clinical Psychology*, 46(3), 283–297.
- Richardson, C. G., & Ratner, P. A. (2005). Sense of coherence as a moderator of the effects of stressful life events on health. *Journal of Epidemiology & Community Health*, 59(11), 979–984.
- Rohrer, J. M., & Arslan, R. C. (2021). Precise answers to vague questions: Issues with interactions. *Advances in Methods and Practices in Psychological Science*.
- Ryan, W., Evers, E., & Moore, D. A. (2018). False positive poisson. *Available at SSRN 3270063*.
- Seaton, E. K., & Iida, M. (2019). Racial discrimination and racial identity: Daily moderation among black youth. *American Psychologist*, 74(1), 117.
- Stringa, N., Milaneschi, Y., van Schoor, N. M., Suanet, B., van der Lee, S., Holstege, H., Reinders, M. J., Beekman, A. T., & Huisman, M. (2020). Genetic liability for depression, social factors and their interaction effect in depressive symptoms and depression over time in older adults. *The American Journal of Geriatric Psychiatry*.
- Tabery, J. (2008). Ra fisher, lancelet hogben, and the origin (s) of genotype–environment interaction. *Journal of the History of Biology*, 41(4), 717–761.
- Tobin, J. (1958). Estimation of relationships for limited dependent variables. *Econometrica: journal of the Econometric Society*, 24–36.
- Trejo, S., Belsky, D. W., Boardman, J. D., Freese, J., Harris, K. M., Herd, P., Sicinski, K., & Domingue, B. W. (2018). Schools as moderators of genetic associations with life course attainments: Evidence from the wls and add heath. *Sociological science*, 5, 513–540.
- Tucker-Drob, E. M. (2009). Differentiation of cognitive abilities across the life span. *Developmental psychology*, 45(4), 1097.
- van der Linden, W. J. (2016). *Handbook of item response theory: Volume 1: Models*. CRC Press.
- Yeager, D. S., Hanselman, P., Walton, G. M., Murray, J. S., Crosnoe, R., Muller, C., Tipton, E., Schneider, B., Hulleman, C. S., Hinojosa, C. P., et al. (2019). A national experiment reveals where a growth mindset improves achievement. *Nature*, 573(7774), 364–369.

A Supplemental Figures

Figure A.1: Effect of transformation when $\alpha = 0$ for various values of λ .

Figure A.2: Scatterplot of y (y^* suppressed given the difference in scales) as a function of xz for different values of λ ($\beta_0 = 0, \beta_1 = \beta_2 = 1, \sigma_y^2 = 0.25, \alpha = 0, N = 1000$) when y_i is a transformed outcome.

Figure A.3: Illustration of the geometry driving false discovery due to variation in λ when the linear model is used ($\beta_0 = 0, \beta_1 = \beta_2 = 1, \sigma_y^2 = 0.25, \alpha = 0, N = 1000$) for analysis of transformed outcomes. Blue and red dots represent those data points within half a unit of their respective z values (i.e., z such that $|z - 1| < 0.5$ are in blue and z such that $|z + 1| < 0.5$ are in red. Fitted lines for ± 1 are similarly shaded). Dashed lines are fits from linear model.