Stepwise methods can limit power for hypothesis tests of cross-level interactions

Kevin M. King, Ph.D., Dale S. Kim, Connor J. McCabe, Sean P. Lane
University of Washington, University of California, Los Angeles, University of
Washington, Purdue University

Author Note

Correspondence concerning this article should be addressed to Kevin M. King, Department of Psychology, University of Washington, Box 351525, Seattle WA 98195-1525, (206) 685-3571 (f), kingkm@uw.edu.

Author Contributions: KMK generated the idea for the study. DSK wrote the analysis code and analyzed the data. SPL verified the simulation code and confirmed computations and formulas. KMK wrote the first draft of the manuscript. CJM and DSK generated the tables and figures. All authors critically edited the manuscript and approved the final version of the manuscript.

Abstract

In multilevel models, stepwise methods are commonly used to test cross-level interactions, where a cluster level variable explains differences in the effect of an observation level variable on the outcome. Researchers often wish to establish that there is between cluster variance in slopes before testing whether an observed cluster level variable explains between cluster variance in slopes. In the stepwise method, between cluster slope variance (i.e. random slopes) is required before a cross-level interaction is tested. We argue that this requirement unnecessarily reduces the power to detect true cross-level interactions, because it imposes an unnecessary constraint on the power to detect valid interactions. In short, the stepwise approach would only be valid if, in the same data, the stepwise approach always identifies true interactions that are also identified by a direct test of the interaction model. Using Monte Carlo simulations, we demonstrate that this is not the case. The power to detect a true interaction was especially low when the residual slope variance (i.e. variance unexplained by the interaction), the variance of the moderator, the number of observations per cluster, or the number of clusters was small. We recommend that researchers directly test interactions that are of interest, regardless of the presence of random slope variance.

Keywords: multilevel modeling, moderation, cross-level interaction, power analysis

Stepwise methods can limit power for hypothesis tests of cross-level interactions

Introduction

Researchers are often interested in explaining individual differences in associations between a focal predictor and an outcome. These types of hypotheses can be conceptualized as a fixed effect, where a third variable (a moderator) accounts for individual differences in the association between the focal predictor and the outcome (i.e., moderation). Moderation refers to the effect of a independent variable (the focal predictor) on an outcome that is conditional on the level of another independent variable (the moderator), and is tested in linear models by including the product of the two variables, described as the interaction term, in a statistical model (McCabe et al., 2018; Aiken et al., 1991). Hypotheses about individual differences in associations can also be tested using random effects, which estimate the variance of a distribution of differences in associations across some unit of clustering (such as organizations or individuals). In order to estimate a distribution of random effects, multiple observations are required to estimate the association for each unit of clustering, such as when individuals are measured repeatedly over time, or when students are sampled within classrooms. Finally, researchers can use cross-level interactions to explain variation in a random slope effect, where a cluster level variable is used to explain differences in the effect of an observation level variable across clusters. For example, Chen & Vazsonyi (2013) showed that the relation between adolescents' future orientation and problem behavior was stronger among students attending schools with more positive future orientation climates. In a longitudinal sample of 259 adolescents studied at up to four time points, we recently showed that the effects of age on the probability of alcohol use in the past year was strongest for those who reported more stressful life events over time (King et al., 2017). Given the importance of cross-level interactions in psychological theory (Aguinis & Culpepper, 2015), the goal of the current study was to test what factors influence the power to detect cross-level interactions in clustered data.

Cross-level interactions are commonly tested using multilevel models (MLMs), which are an extension of linear regression that allow hypothesis tests in clustered data

(Raudenbush & Bryk, 2002). Equation 1 provides an example MLM, where y_{ij} is an outcome collected for multiple observations i clustered within units j that is predicted by a within cluster variable (x_{ij}) and a between cluster moderator (z_j).

$$\begin{aligned}
 y_{ij} &= \beta_{0j} + \beta_{1j}x_{ij} + \epsilon_{ij} \\
 \beta_{0j} &= \gamma_{00} + \gamma_{01}z_j + u_{0j} \\
 \beta_{1j} &= \gamma_{10} + \gamma_{11}z_j + u_{1j} \\
 \epsilon_{ij} &\sim \mathcal{N}(0, \sigma^2) \\
 \begin{bmatrix} u_{0j} \\ u_{1j} \end{bmatrix} &\sim \mathcal{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \tau = \begin{bmatrix} \tau_{00} & \tau_{01} \\ \tau_{01} & \tau_{11} \end{bmatrix} \right)
 \end{aligned} \tag{1}$$

In MLMs, the cross-level interaction (γ_{11}) is estimated as a fixed effect, with a focal predictor measured at the observation level (x_{ij}), and a moderator (z_j) measured at the cluster level. The coefficient for the cross-level interaction (γ_{11}) estimates how the average between cluster association of x_{ij} and y_{ij} (β_{1j}) differs as a linear function of z_j (Bauer & Curran, 2005). In other words, a cross-level interaction means that a focal variable's effect varies as a function of an observed cluster level variable, and may be understood by modeling (or plotting) how the effect of the focal variable changes as the cluster level variable changes, in the same way that interactions in models with only fixed effects (such as OLS regression) are probed (Bauer & Curran, 2005; McCabe et al., 2018). Unexplained variance at the observation level (ϵ_{ij}) is assumed to be normally distributed with a mean of 0 and a variance of σ^2 . Unexplained variance at the cluster level is assumed to be uncorrelated with ϵ_{ij} , and normally distributed with variance τ_{00} . The random slope (i.e., between cluster differences in β_{1j}) is estimated in MLMs by including u_{1j} , with expectation 0 and variance τ_{11} for the association between the within-cluster variable x_{ij} and the outcome (Raudenbush & Bryk, 2002), which estimates unexplained variance in β_{1j} .

Much prior research has examined the power to detect interactions in multiple regression (Cohen et al., 2003), highlighting the importance of factors such as sample size, multi-collinearity between the interaction term and the focal predictor and moderator, and moderation effect size (among others; Aguinis 1995). Interactions in

MLMs are complicated by the presence of random slopes which affect the power to detect cross-level interactions because the focal effect varies across units of clustering (Bauer & Curran, 2005). When random slopes are not involved, such as in interactions between observation-level variables with fixed effects, or interactions among cluster-level variables, the same factors that influence the power to detect interactions in multiple regression would be expected to generalize. When random slopes are estimated, a few studies have shown that the power to detect cross-level interactions depends on cluster size and the number of clusters (i.e. sample size), the magnitude of the interaction γ_{11} , the variance of the residual ϵ_{ij} , and the variance of the random slope τ_{11} (LaHuis & Ferguson, 2009; Mathieu et al., 2012). Power was primarily affected by the effect size of the interaction, moderated by the random slope variance, where variability in slopes strengthened the effects of the interaction effect size on power.

However, this research was limited because it did not reflect common practices in testing cross-level interactions. Specifically, researchers often wish to establish that there is *any* between cluster variance in slopes before testing whether an observed cluster level variable explains between cluster variance in slopes. For example, while introducing a novel method for characterizing slope variability, Aguinis & Culpepper (2015), argued that understanding the amount of variance in the outcome attributable to cluster differences in slopes was critical before researchers searched for predictors of those slopes. Thus, a recommendation for testing cross-level interactions is to follow a stepwise procedure, requiring two statistical tests comparing a sequence of three models to support an interaction hypothesis. First, a model with a fixed slope is estimated (Model 1), and compared to a random slope model (Model 2) with a likelihood ratio test (because significance tests of the variance of random effects may be unreliable; Pawitan 2001). The cross-level interaction (Model 3, i.e. Equation 1) is only tested when the focal variable has a random slope via the significance of the likelihood ratio test.

This recommendation is common in statistical textbooks on MLM, as well as manuscripts that serve as MLM tutorials. Aguinis outlined the logic for this approach, arguing that “there is no point in examining which particular moderators might explain

slope variance...if such variance is nonexistent” (Aguinis et al., 2013, p. 1501). Here, cross-level interactions are described as a “slopes as outcomes” model, because cluster level predictors are accounting for between cluster variation in slopes. For example, researchers have proposed using the estimate and standard error of the variance of the random slope to conduct a significance test (or derive a confidence interval) to establish whether the variance differed from zero before proceeding to test for a cross-level interaction (Bliese, 2002; Singer & Willett, 2003; Hofmann, 1997). Other researchers have recommended the use of a likelihood ratio test to qualify the inclusion of the random slope parameter (Bliese, 2002). In either case, statistical significance for the random slope test has been considered necessary prior to testing the cross-level interaction effect. On the other hand, a few authors have recommended against the stepwise method, and suggested that a direct test of a cross-level interaction (where the interaction is tested without respect to the presence of a random slope) is justified. Snijders & Bosker (1999) advocated that *a-priori* hypotheses should be tested, noting that testing the cross-level interaction is justified given the researcher has an *a-priori* hypothesis that the interaction exists. Aguinis et al. (2013) also argued that even if a random slope is not found, it is worth testing cross-level interactions when there is a strong *a-priori* hypothesis for such interactions, particularly because the power to detect a random slope may be low in many studies.

We argue that requiring a random slope model to be significant before testing for a cross-level interaction before testing for the latter imposes an unnecessary constraint on power. Because the stepwise approach relies on a likelihood ratio test to determine whether a random effect is present, it depends on the following logic: First, in Model 1, σ^2 and τ_{00} absorbs all the unmodeled variance attributable to the cross-level interaction term, τ_{11} and τ_{01} . Second, in Model 2, modeling τ_{11} and τ_{01} is expected to reduce σ^2 and τ_{00} by the amount of variance attributable to the cross-level interaction term and the variance from τ_{11} and τ_{01} (which was previously unmodeled). Finally, this results in a lower likelihood for Model 2 relative to Model 1. The degree to which changes in the parameterization of variance across models influences the likelihood ratio test is a

function of the number of clusters, the number of observations, and the magnitudes of the estimates. Thus it is difficult to apply a general rule-of-thumb that will be effective in all cases. The goal of the current study was to vary each of these parameters in order to understand how likelihood ratio tests, and thus the stepwise interaction approach, influences the power to detect effects.

Methods

We designed a simulation analysis to compare the power to detect a true cross-level interaction with a stepwise approach relative to a direct test of the interaction. The true interaction model is Equation 1, above.

The change in likelihoods across models is influenced by the number of clusters, the number of observations, and the magnitude of unmodeled variances in the model. Because prior research had established how the residual ϵ_{ij} and interaction effect influenced power (LaHuis & Ferguson, 2009; Mathieu et al., 2012), we focused instead on unmodeled variance attributable to the random slope and interaction term, i.e. the variance of the random slope τ_{11} and the variance of the interaction term $\text{Var}(z_j)$.

The simulations proceeded as follows. We generated data from a cross-level interaction model as specified in Equation 1, for a range of clusters (n) with (m) observations per cluster for a range of two representative types of studies. First, many longitudinal studies have a few observations of many individuals. To represent scenarios common to these studies, we conducted simulations where $m = 3, 5$ and 10 observations were obtained from $n = 150, 300$ and 500 clusters, for 9 total sample sizes ranging from 450 to 5000. Second, many organizational studies of people clustered within groups have many observations of relatively fewer individuals. For these, we conducted simulations where $m = 10, 20$ and 30 observations were obtained from $n = 50, 100$ and 150 clusters. Thus we had 9 total sample sizes ranging from 500 to 4500.

Data were generated with the following parameters. All fixed effects (including the interaction) were specified at 1, with $\tau_{00} = 15$, and $\sigma^2 = 300$, and $\tau_{01} = 0$. The observation level predictor x_{ij} was sampled from a standard normal distribution with a

mean of 0 and variance of 1, then was centered within cluster by subtracting the cluster mean. z_j was sampled independently from x_{ij} , from a normal distribution with a mean of 0 and variances of 0.5, 10, 20, and 30. We varied τ_{11} with values of 0, 5, 10, and 15. We simulated 1,000 datasets for each of 2 factors ($\text{Var}(z_j)$ and τ_{11}) with 4 levels each, for a 16 factor design.

To aid in the interpretation of the data generating models, Table 1 illustrates $\text{Var}(y_{ij})$, the variance of the outcome, and r_{xz} , the effect size of the interaction, across conditions. It also includes ICC_β (Aguinis & Culpepper, 2015), which represents an estimate of the expected random slope variance with and without the interaction for all conditions, which can be used as an indicator of how much variance in the slope of x_{ij} is explained by z_j .

To conduct the simulations, we fit three models to each dataset. Model 1 was the random intercept only model, Model 2 was the random intercept and slope model, and Model 3 was the fully specified cross-level interaction model with both the random slope and the cross-level interaction. In Models 2 and 3, we did not estimate a correlation between the random intercept and slope. To study the power to detect the interaction, we compared the difference in the deviance for three model comparisons. Power for the stepwise approach was estimated by two comparisons, Model 2 vs. Model 1, and Model 3 vs. Model 2. An interaction would be supported in the stepwise approach when both comparisons were significant. The direct test was represented by one comparison, Model 3 vs 2, and was supported if it was significant. Thus the direct test was nested in the stepwise approach. For all comparisons, we estimated the frequency with which the deviance exceeded the critical value, which estimates the power to detect the cross-level interaction. We then computed the “proportion of power lost”, which was the proportion of datasets in which the direct approach was significant but the stepwise approach was not, divided by the total number of significant direct interaction tests. Because both the stepwise and direct tests compared Model 3 vs. 2, this effectively computed the percent of tests where a detectable, true interaction was missed using the stepwise method because the random slope was not detected (Model 2

vs. Model 1). If the stepwise approach does not reduce the power to detect a true interaction, we would expect the proportion of power lost to be close to zero, and the stepwise and direct approaches to have identical power.

Results

As would be expected, the proportion of power lost, where the true interaction was detected by the direct test but not the stepwise method depended on all four variables in the simulations: $\text{Var}(z_j)$, τ_{11} , m and n . Figures 1 and 2 illustrate our findings for the longitudinal and organizational cases, respectively, reporting the % of the time the true interaction would be supported by both the stepwise and direct tests (“Both”) and by the direct test but not the stepwise test (“Direct”). First, as would be expected, increasing both m and n increased the probability that the stepwise method and the direct method would provide identical results. In line with previous work (Lane & Hennes, 2018), increasing the number of clusters provided better relative power compared to increasing the sample size in both the longitudinal and organizational cases. For example, when $m = 10$ and $n = 50$, the direct test was more likely to detect interactions than the stepwise method, especially when both $\text{Var}(z_j)$ and τ_{11} were relatively low. Figure 3 illustrates this with scatterplots of the deviance values for all combinations of $\text{Var}(z_j)$ and τ_{11} for $m = 10$ and $n = 50$. Each point in the scatterplot illustrates the deviance for the comparison of Model 2 with Model 1 (X axis) by the deviance of the comparison of Model 3 with Model 2 (Y axis) in the same dataset. The horizontal dotted lines indicate the threshold where the deviance test was significant for Model 3 vs. Model 2, and the vertical dotted lines indicate the threshold where the deviance test was significant for Model 2 vs. Model 1. In other words, points below the horizontal dotted line had a non-significant deviance comparing Model 3 vs. Model 2, and those to the left of the vertical dotted line were non-significant comparing Model 2 vs. Model 1. It illustrates that even for large values of $\text{Var}(z_j)$ and τ_{11} , the direct test still detected some interactions that the stepwise test would have missed.

When $m = 3$ and $n = 150$, the direct test was almost always more likely to detect

interactions, but when $m > 10$ and $n > 100$, the stepwise and direct methods performed equally well. Figure 4 illustrates this again with scatterplots of the deviance values for $m = 3$ and $n = 150$, while Figure 5 illustrates this for $m = 10$ and $n = 500$. When either $\text{Var}(z_j)$ or τ_{11} were very small (i.e. $\text{Var}(z_j) = 0.5$ or $\tau_{11} = 0.5$), there was a higher probability that neither the direct or the stepwise method would successfully detect the interaction. Increasing values of $\text{Var}(z_j)$ or τ_{11} decreased the advantage of the direct test, such that when either were relatively large, the stepwise and direct methods were relatively equivalent.

Figure 6 illustrates the proportion of power lost by using the stepwise method across all conditions for the organizational samples, and Figure 7 illustrates this for longitudinal samples. As discussed above, the amount of power lost was highest at lower values of $\text{Var}(z_j)$, τ_{11} , m and n .

Discussion

Although making a test of the cross-level interaction conditional on the presence of a random slope is commonly recommended and practiced, this stepwise approach led to substantial loss of power when the variance of the moderator $\text{Var}(z_j)$, residual slope variance τ_{11} , and the number of observations (m) and clusters (n) were small, generally speaking. Power loss in the stepwise approach was especially problematic when the variance of the moderator or the random slope were smaller, regardless of sample and cluster size. As the residual random slope variance, or the variance of the moderator approached zero, a stepwise approach would be more likely to reject a true cross-level interaction, even though both terms are, by definition, unrelated to the magnitude of the interaction coefficient. This finding highlights the misfit between the stepwise approach and the hypothesis that researchers desire to test. This problem was amplified when the number of observations or clusters were smaller, but it was especially driven by cluster size. Prior studies had shown that increasing the size of clusters provided greater power for within cluster effects (Lane & Hennes, 2018), and our findings support this work. For example, with 10 observations per cluster and 500 clusters, the stepwise

approach would provide misleading results only for very small levels of $\text{Var}(z_j)$ and small to moderate levels of τ_{11} . On the other hand, a similar loss of power was observed with 30 observations per cluster and only 100 clusters. This suggests that studies that relied on the stepwise method to detect interactions have been especially vulnerable to power loss when the number of clusters were low (such as in longitudinal studies). In general, our research suggests that many hypothesized cross-level interactions have not been tested (or reported) in prior studies because the stepwise method incorrectly suggested that the cross-level interaction was null.

Researchers commonly test for random slopes using the likelihood ratio test, which provides evidence for the presence/absence of the random slope. Aguinis & Culpepper (2015) recently explicated a method to characterize cluster level differences in slopes in terms of an intraclass correlation (ICC_β). They illustrated that traditional ICCs are largely influenced by cluster differences in levels of the outcome, and are only partially influenced by cluster differences in slopes. Thus they argued that computing ICC_β would aid inference by describing the magnitude of cluster level differences in slopes, helping researchers determine whether there was sufficient variance to detect a cross-level interaction. Our findings build upon this work, suggesting that for particularly small values of ICC_β , researchers should be cautious about using a stepwise method to make decisions about whether or not to test for cross-level interactions.

Other authors have recently argued that most research using multilevel models has hypotheses about fixed, but not random effects (McNeish et al., 2017). Our findings are in line with this argument, in that we show that testing the random slopes only model limits these models' power to detect true effects that researchers care about. A few authors have noted this argument before, arguing that strong *a-priori* hypotheses for interactions justify testing cross-level interactions (Snijders & Bosker, 1999; Aguinis et al., 2013). However, no research directly tested whether the power to detect a true interaction is different for a direct test of an interaction vs. a stepwise approach, or explicated precisely why a direct approach may be preferred. Our findings suggest that statistical power depends on many factors, and a general decision making rule like the

stepwise approach will be difficult to apply without some loss of statistical power. As such, we recommend that authors avoid stepwise approaches in MLMs, and directly test interactions that are of interest (regardless of whether or not they model residual random slopes).

References

- Aguinis, H. (1995). Statistical Power with Moderated Multiple Regression in Management Research. *Journal of Management*, 21(6), 1141–1158. doi: 10.1177/014920639502100607
- Aguinis, H., & Culpepper, S. A. (2015). An expanded decision-making procedure for examining cross-level interaction effects with multilevel modeling. *Organizational Research Methods*, 18(2), 155–176.
- Aguinis, H., Gottfredson, R. K., & Culpepper, S. A. (2013). Best-practice recommendations for estimating cross-level interaction effects using multilevel modeling. *Journal of Management*, 39(6), 1490–1528.
- Aiken, L. S., West, S. G., & Reno, R. R. (1991). *Multiple regression: Testing and interpreting interactions*. Sage.
- Bauer, D. J., & Curran, P. J. (2005). Probing Interactions in Fixed and Multilevel Regression: Inferential and Graphical Techniques. *Multivariate Behavioral Research*, 40(3), 373–400.
- Bliese, P. D. (2002). Multilevel random coefficient modeling in organizational research: Examples using SAS and S-PLUS. , 401-445.
- Chen, P., & Vazsonyi, A. T. (2013). Future orientation, school contexts, and problem behaviors: A multilevel study. *Journal of Youth and Adolescence*, 42(1), 67–81.
- Cohen, J., Cohen, P., West, S. G., & Aiken, L. (2003). *Applied Multiple Regression / Correlation Analysis for the Behavioral Sciences* (1st ed.). Hillsdale NJ: Lawrence Erlbaum Associates. doi: 10.2307/2064799
- Hofmann, D. A. (1997). An overview of the logic and rationale of hierarchical linear models. *Journal of management*, 23(6), 723–744.

- King, K. M., Pedersen, S. L., Louie, K. T., Pelham Jr, W. E., & Molina, B. S. (2017). Between-and within-person associations between negative life events and alcohol outcomes in adolescents with ADHD. *Psychology of Addictive Behaviors*, 31(6), 699.
- LaHuis, D. M., & Ferguson, M. W. (2009). The accuracy of significance tests for slope variance components in multilevel random coefficient models. *Organizational Research Methods*, 12(3), 418–435.
- Lane, S. P., & Hennes, E. P. (2018). Power struggles: Estimating sample size for multilevel relationships research. *Journal of Social and Personal Relationships*, 35(1), 7–31.
- Mathieu, J. E., Aguinis, H., Culpepper, S. A., & Chen, G. (2012). Understanding and estimating the power to detect cross-level interaction effects in multilevel modeling. *Journal of Applied Psychology*, 97(5), 951–966.
- McCabe, C. J., Kim, D. S., & King, K. M. (2018). Improving present practices in the visual display of interactions. *Advances in Methods and Practice of Psychological Science*, 1(2).
- McNeish, D., Stapleton, L. M., & Silverman, R. D. (2017). On the unnecessary ubiquity of hierarchical linear modeling. *Psychological Methods*, 22(1), 114.
- Pawitan, Y. (2001). *In all likelihood: statistical modelling and inference using likelihood*. Oxford University Press.
- Raudenbush, S. W., & Bryk, A. S. (2002). *Hierarchical linear models: Applications and data analysis methods* (Vol. 1).
- Singer, J. D., & Willett, J. B. (2003). *Applied longitudinal data analysis: Modeling change and event occurrence*. Oxford university press.
- Snijders, T., & Bosker, R. (1999). *Multilevel analysis: An introduction to basic and applied multilevel analysis*. London: Sage.

Table 1

Table 1. Interaction Effects and Intraclass Correlation Coefficients Across Study Conditions.

XZ_{ij} and Z_j Effect Size		Large				Medium				Small				Near-Zero			
t_{11} Effect Size		Large	Med.	Small	Zero	Large	Med.	Small	Zero	Large	Med.	Small	Zero	Large	Med.	Small	Zero
	$Var(y_{ij})$	391	386	381	376	371	366	361	356	351	346	341	336	332	327	322	317
	r_{xz}	0.277	0.279	0.281	0.282	0.232	0.234	0.235	0.237	0.169	0.170	0.171	0.173	0.039	0.039	0.039	0.040
	ICC $_{\beta}$, no interaction	0.150	0.133	0.117	0.100	0.117	0.100	0.083	0.067	0.083	0.067	0.050	0.033	0.052	0.035	0.018	0.002
	ICC $_{\beta}$, with interaction	0.050	0.033	0.017	0.000	0.050	0.033	0.017	0.000	0.050	0.033	0.017	0.000	0.050	0.033	0.017	0.000
	% reduction in ICC $_{\beta}$	0.667	0.750	0.857	1.000	0.571	0.667	0.800	1.000	0.400	0.500	0.667	1.000	0.032	0.048	0.091	1.000

Note. ICC $_{\beta}$ = Intraclass β correlation, Med. = Medium. Intercept variances were set to 15 and X variances set to 1 across all conditions.

Corresponding variances for large, medium, small, and near-zero effect sizes for XZ_{ij} and Z_j were 30, 20, 10, and .5, respectively.

Corresponding variances for large, medium, small, and zero effect sizes for t_{11} were 15, 10, 5, and 0, respectively.

Figure 1. Success of direct and stepwise cross-level interaction tests by number of observations per cluster (n , 150, 300, 500) and number of clusters (m , 3, 5, 10) and $\text{Var}(z_j)$, and τ_{11} for longitudinal samples

Figure 2. Success of direct and stepwise cross-level interaction tests by number of observations per cluster (n , 50, 100, 150) and number of clusters (m , 10, 20, 30) and $\text{Var}(z_j)$ and τ_{11} for organizational samples

Figure 3. Deviance values for all simulations with 50 clusters of 10 observations each.

Note: Deviance values for Model 2 vs. Model 1 are plotted on the X axis, and those for the Model 3 vs Model 2 are plotted on the Y axis. Dotted lines denote critical values of a χ^2 test.

Figure 4. Deviance values for all simulations with $m = 3$, $n = 150$.

Note: Deviance values for Model 2 vs. Model 1 are plotted on the X axis, and those for the Model 3 vs Model 2 are plotted on the Y axis. Dotted lines denote critical values of a χ^2 test.

Figure 5. Deviance values for all simulations with $m = 10$, $n = 500$.

Note: Deviance values for Model 2 vs. Model 1 are plotted on the X axis, and those for the Model 3 vs Model 2 are plotted on the Y axis. Dotted lines denote critical values of a χ^2 test.

Figure 6. Proportion of power lost for organizational type samples (high m , low n).

Note: This is the proportion of datasets which were significant for the interaction test but not the random slope test, divided by the total significant number interaction tests. That is, of the properly detected interaction models, these are the proportion of models lost due to the stepwise testing method.

Figure 7. Proportion of power lost for longitudinal type samples (low m , high n).

Note: This is the proportion of datasets which were significant for the interaction test but not the random slope test, divided by the total significant number interaction tests. That is, of the properly detected interaction models, these are the proportion of models lost due to the stepwise testing method.