Computational Tools for Probing Interactions in Multiple Linear Regression, Multilevel Modeling, and Latent Curve Analysis

Kristopher J. Preacher
Patrick J. Curran
Daniel J. Bauer
University of North Carolina at Chapel Hill

Simple slopes, regions of significance, and confidence bands are commonly used to evaluate interactions in multiple linear regression (MLR) models, and the use of these techniques has recently been extended to multilevel or hierarchical linear modeling (HLM) and latent curve analysis (LCA). However, conducting these tests and plotting the conditional relations is often a tedious and error-prone task. This article provides an overview of methods used to probe interaction effects and describes a unified collection of freely available online resources that researchers can use to obtain significance tests for simple slopes, compute regions of significance, and obtain confidence bands for simple slopes across the range of the moderator in the MLR, HLM, and LCA contexts. Plotting capabilities are also provided.

Hypotheses involving multiplicative interaction or moderation effects are common in the social sciences. Interaction effects are typically evaluated by testing the significance of a multiplicative term consisting of the product between two or more predictor variables controlling for associated lower order main effects (e.g., Cohen, 1978). When a significant interaction is found, it is common to further decompose or "probe" this conditional effect to better understand the structure of the relation (e.g., Aiken & West, 1991).

Although interactions arise in many seemingly disparate analytic frameworks, interactions within these frameworks share a common computational form. In a series of articles, we have explored these computational linkages within three major analytic frameworks: the multiple linear regression (MLR) model (Bauer &

437

Curran, 2005), the random effects multilevel/hierarchical linear model (HLM) (Bauer & Curran, 2005; Curran, Bauer, & Willoughby, 2006), and the structural equation-based latent curve analysis (LCA) model (Curran, Bauer, & Willoughby, 2004). In these articles we also describe a variety of specific tests that we believe are helpful in the probing of complex interaction terms in MLR, HLM, and LCA. Although we view these tests as potentially powerful and widely applicable, many of the required values are computed by hand and these calculations can be quite cumbersome and consequently error-prone. The complexity of many of these tests also may substantially reduce the likelihood that these methods will be used in practice.

Our goal here is to integrate across our three prior articles and capitalize on the shared computational linkages in testing interactions in order to develop a set of freely available online tools that automate the testing and plotting of these complex effects. We begin by presenting a single shared notational system to define the general estimation of two-way interactions. This general expression can equivalently define higher order interactions stemming from MLR, HLM, and LCA. We briefly review some methods available for the further probing and plotting of these effects. We present online calculators designed to be easily accessible and to automate the calculation of a variety of tests for probing interactions described in Bauer and Curran (in press) and Curran et al. (2004, in press). Finally, we demonstrate our calculators by probing an interaction in MLR.

Interactions

An interaction implies that the magnitude of the relation between one predictor and the criterion varies as a function of at least one other predictor. It is often convenient to think of one predictor as a *focal predictor* and all other predictors involved in product terms with the focal predictor as *moderators* hypothesized to affect the relationship between the focal predictor and the criterion (although this distinction is arbitrary given the symmetry of the interaction). For simplicity, we will focus primarily on the case where there is a single moderator of interest. We denote the criterion y, the focal predictor x, and the moderator z. A general expression for a regression equation involving a multiplicative interaction effect is:

$$y_i = \gamma_0 + \gamma_1 x_i + \gamma_2 z_i + \gamma_3 x_i z_i + \varepsilon_i, \tag{1}$$

where the γs represent path coefficients or regression weights and x and z are typically centered about their respective means. The *prediction equation* is defined as the expected value of y conditioned on specific values for x and z:

$$E\big[y\big|(x,z)\big] = \hat{\gamma}_0 + \hat{\gamma}_1 x + \hat{\gamma}_2 z + \hat{\gamma}_3 xz, \tag{2}$$

where the carat symbol (^) represents the sample estimate of the corresponding parameter. The test of the coefficient $\hat{\gamma}_3$ in Equation 2 is an omnibus test of the interaction effect (Aiken & West, 1991; Saunders, 1956).

Simple Intercepts and Simple Slopes

Equation 2 can be rearranged in terms of a *simple intercept* and *simple slope,* as follows:

$$E\left[y\,|\,(x,z)\right]=\left(\hat{\gamma}_0+\hat{\gamma}_2 z\right)+\left(\hat{\gamma}_1+\hat{\gamma}_3 z\right)x. \tag{3}$$

This highlights the fact that the simple intercept (first bracketed term) and simple slope (second bracketed term) for the regression of y on x vary as a function of z (Aiken & West, 1991). The simple intercept and simple slope are *compound coefficients* that can be compactly expressed as:

$$\begin{aligned}\hat{\omega}_0 &= \hat{\gamma}_0+\hat{\gamma}_2 z \\ \hat{\omega}_1 &= \hat{\gamma}_1+\hat{\gamma}_3 z.\end{aligned} \tag{4}$$

Because $\hat{\omega}_0$ and $\hat{\omega}_1$ are linear combinations of parameters and specific values of z, it is useful to evaluate how the estimates change as the value of the moderator changes—a topic we now review.

Probing Interactions in MLR, HLM, and LCA

Two approaches have been outlined in the literature for evaluating the form of interactive effects: the *simple slopes technique* and the *Johnson–Neyman technique.*

Simple Slopes Technique

The traditional approach to probing significant interaction effects is to choose several conditional values of z at which to evaluate the significance of the simple slope for the regression of y on x (Aiken & West, 1991; Rogosa, 1980, 1981). The significance of simple intercepts is rarely of interest, and thus is not covered here. For any conditional value of the moderator(s), the significance of $\hat{\omega}_1$ may be found by first deriving the variance of the simple slope. For example, the calculation of the variance of $\hat{\omega}_1$ in Equation 4 is a simple function of the variances and covariances of the parameter estimates:

$$\mathrm{var}\left(\hat{\omega}_1\,|\,z\right)=\mathrm{var}\left(\hat{\gamma}_1\right)+2z\,\mathrm{cov}\left(\hat{\gamma}_1,\hat{\gamma}_3\right)+z^2\,\mathrm{var}\left(\hat{\gamma}_3\right), \tag{5}$$

where var($\cdot$) and cov($\cdot$) are the corresponding elements from the sample estimate of the asymptotic covariance (ACOV) matrix. The standard error of $\hat{\omega}_1$, $SE_{\hat{\omega}_1}$, is simply the square root of this quantity. Finally, we can form the critical ratio to perform a significance test for the difference between $\hat{\omega}_1$ and zero:

$$t = \frac{\hat{\omega}_1}{SE_{\hat{\omega}_1}}, \tag{6}$$

439

the significance of which is determined by comparing the obtained t to a t distribution at the desired α level and degrees of freedom $(df) = N - p - 1$, where N is the sample size and p is the number of predictors.

In order to employ the simple slopes method described above, conditional values of the moderator, denoted cv_z, must be chosen (what Rogosa, 1980, refers to as the "pick-a-point" approach). For dichotomous moderators, cv_z assumes values of the dichotomy (usually 0 and 1). For continuous moderators, the specific choices for cv_z are less obvious and may be any value of scientific interest. In the absence of theoretically meaningful values, Cohen and Cohen (1983) recommend choosing values at the mean of z and at 1 SD above and below the mean of z.

The Johnson–Neyman Technique

Despite its broad usefulness, the simple slopes method has an important limitation: the choices of cv_z are ultimately arbitrary. An alternative is the Johnson–Neyman (J-N) technique (Johnson & Neyman, 1936). The J-N technique essentially works backwards from the critical ratio defined in Equation 6. Instead of calculating t as a function of $\hat{\omega}_1$ and $SE_{\hat{\omega}_1}$ and a chosen value cv_z, the J-N technique calculates the cv_z that yields a specific t (usually the critical t value to obtain a p value of .05 at a given df). The conditional values that are returned by the J-N technique define the *regions of significance* on z, and represent the range of z within which the simple slope of y on x is significantly different from zero at the chosen α. Assuming solutions exist, the result will be two values: upper and lower boundaries of the region of significance. In many cases, the regression of y on the focal predictor is significant at values of the moderator that are *less* than the lower bound and *greater* than the upper bound, while the regression is nonsignificant at values of the moderator falling *within* the region. However, there are some cases in which the opposite holds (i.e., the significant slopes fall within the region boundaries).

Confidence Bands

Both the simple slopes and J-N techniques rely on traditional null hypothesis testing logic. It is well known that confidence intervals (CIs) provide more information than hypothesis tests, and methodologists increasingly recommend the use of CIs in addition to, or in place of, hypothesis tests whenever possible (Wilkinson & APA Task Force, 1999). The general formula for a $100 \times (1 - \alpha)\%$ CI for a simple slope (Cohen, Cohen, West, & Aiken, 2003) is:

$$CI_{\hat{\omega}_1} = \hat{\omega}_1 \pm t_{crit} SE_{\hat{\omega}_1}. \tag{7}$$

Because the formula for $SE_{\hat{\omega}_1}$ relies on particular choices for z, $CI_{\hat{\omega}_1}$ varies as a function of the moderator. When $CI_{\hat{\omega}_1}$ is plotted across all relevant values of z, the result is a pair of *confidence bands* (Bauer & Curran, in press; Rogosa, 1980, 1981; Tate, 2004).

The same methods used to represent interactions in MLR in terms of simple intercepts and simple slopes (i.e., Equation 4) can also be used to represent interactions in HLM and LCA. In each modeling context, the dependency of the criterion on the focal predictor can be represented as a function of the moderator(s) by defining two compound coefficients, $\hat{\omega}_0$ and $\hat{\omega}_1$.

In HLM with two predictors, interactions may occur between two Level 1 predictors (Case 1), between two Level 2 predictors (Case 2), or between Level 1 and Level 2 predictors (Case 3, or cross-level interaction). A cross-level (Case 3) interaction occurs when the random slope of a Level 1 predictor is predicted by a Level 2 predictor. This last type of interaction is probably most commonly encountered in HLM, so we will focus on it here. In cross-level interactions, either the Level 1 or Level 2 predictor may be chosen as the focal predictor. Typically, however, it is the Level 1 predictor that is chosen to be the focal predictor. The Level 1 equation is:

$$y_{ij} = \beta_{0j} + \beta_{1j}x_{ij} + r_{ij}, \tag{8}$$

where x_{ij} represents the Level 1 predictor for the ith individual nested within the jth group. Because the intercept and slope in HLM are viewed as random variables, they can be expressed in the Level 2 equations:

$$\beta_{0j} = \gamma_{00} + \gamma_{01}w_j + u_{0j}$$
$$\beta_{1j} = \gamma_{10} + \gamma_{11}w_j + u_{1j}, \tag{9}$$

where w_j represents the Level 2 predictor for the jth group. Substituting the Level 2 equations into Equation 8 results in the reduced form expression:

$$y_{ij} = \left(\gamma_{00} + \gamma_{10}x_{ij} + \gamma_{01}w_j + \gamma_{11}x_{ij}w_j\right) + \left(u_{0j} + u_{1j}x_{ij} + r_{ij}\right). \tag{10}$$

The prediction equation derived from Equation 10 is:

$$E\left[y\,|\,(x,w)\right] = \hat{\gamma}_{00} + \hat{\gamma}_{10}x + \hat{\gamma}_{01}w + \hat{\gamma}_{11}xw. \tag{11}$$

In Equation 11, the test of the coefficient $\hat{\gamma}_{11}$ is an omnibus test of the interaction effect. Defining x as the focal predictor and w as the moderator results in the prediction equation:

$$E\left[y\,|\,(x,w)\right] = \left(\hat{\gamma}_{00} + \hat{\gamma}_{01}w\right) + \left(\hat{\gamma}_{10} + \hat{\gamma}_{11}w\right)x.$$

This rearrangement again highlights the simple intercept and slope of the regression of y on x at specific values of w. Using the same notation as before, these are defined as

$$\hat{\omega}_0 = \hat{\gamma}_{00} + \hat{\gamma}_{01}w$$
$$\hat{\omega}_1 = \hat{\gamma}_{10} + \hat{\gamma}_{11}w. \tag{12}$$

The simple slope, $\hat{\omega}_1$, may be interpreted and evaluated using the simple slopes, regions of significance, and confidence band strategies described earlier.

Interactions in LCA

The same methods used to represent interaction effects in MLR and HLM can also be used to represent interactions in LCA, an application of structural equation modeling (SEM). LCA represents the repeated measures of a dependent variable y as a function of latent factors that capture different aspects of change, typically an intercept factor (η_α) and one or more slope factors (η_β). The equation for the repeated measures of y for individual i at time t is:

$$y_{it} = \eta_{\alpha i} + \lambda_t \eta_{\beta i} + \varepsilon_{it}, \tag{13}$$

where λ_t represents a fixed factor loading on the slope factor corresponding to a value of the time metric (see Curran et al., 2004, for details).

The factors representing the intercept and slope then often serve as endogenous (dependent) variables in other model equations. For instance, the latent curve factors could be regressed on the exogenous predictor x:

$$\eta_{\alpha i} = \mu_\alpha + \gamma_1 x_i + \zeta_{\alpha i}$$
$$\eta_{\beta i} = \mu_\beta + \gamma_2 x_i + \zeta_{\beta i}. \tag{14}$$

The reduced form equation is then:

$$y_{it} = \left(\mu_\alpha + \mu_\beta \lambda_t + \gamma_1 x_i + \gamma_2 \lambda_t x_i\right) + \left(\zeta_{\alpha i} + \zeta_{\beta i} \lambda_t + \varepsilon_{it}\right). \tag{15}$$

The fixed component (first part) of Equation 15 contains an intercept term (i.e., μ_α), conditional main effects for *time* (i.e., μ_β) and the exogenous predictor x (i.e., γ_1), and the interaction of *time* and x (i.e., γ_2). Thus, the regression of y on *time* depends in part on the level of x, so x can be considered the moderator and *time* the focal predictor. Taking the expectation of Equation 15 and rearranging yields a prediction equation that highlights the role of x as a moderator of the *time* effect (the slope of the individual trajectories):

$$E\left[y \mid (\lambda_t, x)\right] = \left(\hat{\mu}_\alpha + \hat{\gamma}_1 x\right) + \left(\hat{\mu}_\beta + \hat{\gamma}_2 x\right)\lambda_t. \tag{16}$$

This expression shows that LCA interactions may also be cast in terms of a simple intercept and simple slope, $\hat{\omega}_0$ and $\hat{\omega}_1$, respectively:

$$\hat{\omega}_0 = \hat{\mu}_\alpha + \hat{\gamma}_1 x$$
$$\hat{\omega}_1 = \hat{\mu}_\beta + \hat{\gamma}_2 x. \tag{17}$$

The techniques for calculating simple slopes, regions of significance, and confidence bands described earlier all apply here as well.

Online Tools

Because these methods for evaluating interactions are not fully incorporated into commonly used statistical software, we have developed a unified collection of online resources that implement these techniques at: http://www.quantpsy.org. From a link on this page, researchers may access a primer on regression techniques and any of six interactive Web pages designed to facilitate the probing two-way and three-way interactions in MLR, HLM, and LCA. We now briefly describe the contents of these pages, followed by a worked example of a cross-level interaction in a multilevel model.

At the bottom of the two-way interaction pages are tables where researchers may enter information from analysis output. These tables require the user to input all four estimated regression coefficients, the variances of these coefficients and two of the covariances from the ACOV matrix, *df*, and α. By default, α is set equal to .05, but this may be changed by the user. The tables contain optional sections called "Points to Plot." Users may input values of the focal predictor here to obtain coordinates useful in producing visual representations of interaction effects. These sections assume that the *x*-axis on such a plot represents values of the focal predictor, the *y*-axis represents the dependent variable, and that separate lines are desired for each conditional value of the moderator.

Once all the necessary information is entered, clicking on the "Calculate" button will produce output in the "Output Window." If there is a problem, it will be noted in the "Status" cell toward the middle of the table. The first two sections of the output simply repeat the information entered into the table. The region of significance, and simple intercepts and simple slopes calculated at the boundaries of this region, are provided by default. If conditional values of the moderator are entered into the table, the utility will also provide simple intercepts and simple slopes at each desired value. If simple intercepts and simple slopes are required for more than three conditional values of the moderator, the tables can be reused any number of times by changing the desired conditional values and clicking on the "Calculate" button again. Each line of output includes a point estimate of the simple intercept or simple slope of interest, the standard error, a critical ratio, and a *p* value. Finally, if the "Points to Plot" option is selected, coordinates for lines will be reported at the end of the output. Contents of the output window are in plain text and may be copied and pasted into any word processing program for further editing.

The tables for three-way interactions are considerably larger than the two-way interaction tables, primarily because many more elements of the ACOV matrix must be entered, but they are entirely analogous to the two-way tables. The only significant difference is that the three-way interaction tables enable the user to request simple intercepts and simple slopes of y on x at conditional values of both moderators z and w.

The HLM two-way interaction tool has three separate tables for the Case 1, Case 2, and Case 3 interactions described above. In addition, the user has the option of entering custom *df* for tests of simple intercepts, tests of simple slopes, or both simple intercepts and simple slopes. If either of these boxes is left blank, asymptotic z tests will be conducted instead of t tests. The HLM three-way interaction tool is analogous to the MLR three-way interaction tool. However, it is limited to the case in which there is a cross-level interaction between a single Level 1 (focal) predictor and two Level 2 moderators.

The LCA interaction tools are similar to the MLR and HLM tools. The LCA two-way interaction tool contains two tables: one table is used for situations involving *time* as the focal predictor and the other is used for situations involving an exogenous predictor of slopes as the focal predictor. The LCA three-way interaction tool treats *time* as the focal predictor, and allows users to request simple intercepts and simple slopes at conditional values of two exogenous predictors, x_1 and x_2.

If conditional values are entered for the moderator(s), these tools also produce graphs of the simple regression lines for the focal predictor at those values. Specifically, they generate R syntax that can be submitted to an Rweb server (or to a local R application for better resolution) via the click of a button to generate the plot. R syntax is also generated for the plotting of confidence bands, which again can be submitted to R or Rweb with the click of a button to produce the plot.

An Example

We now provide an illustration of the MLR two-way interaction tool. For this example we rely on data from the National Longitudinal Survey of Youth (NLSY). Data consist of scores on measures of math ability (*math*; the Peabody Individual Achievement Test math subsection) and associated predictor variables from 956 students (52% female) ranging in age from 59 to 156 months. Of interest in this analysis was the finding of a significant interaction between the predictors antisocial behavior (*antisocial*) and *hyperactivity* (both centered) in predicting math test scores. The model is of the same form as Equations 1–3, where *antisocial* takes the place of x and *hyperactivity* takes the place of z, except that age, sex, grade, and minority status were added as covariates. The results from fitting the model are reported in Table 1. The main effect of *antisocial* is positive but nonsignificant, and the main effect of *hyperactive* is negative and significant. In addition, the significant negative coefficient associated with the interaction term (-0.3977) indicates that the relationship between *antisocial* and *math* tends to be more strongly negative for individuals with higher overall *hyperactivity*.

TABLE 1

Output for Antisocial × Hyperactivity Interaction With Asymptotic (Co)Variances

Effect	Coefficient	*SE*	Critical Ratio
Intercept (γ_0)	38.0690	0.3224	118.08
Antisocial (γ_1)	0.0373	0.2681	0.14
Hyperactivity (γ_2)	−0.7995	0.2148	−3.72
Anti. × hyper. (γ_3)	−0.3977	0.1429	−2.78

Sample Asymptotic Covariance Matrix of the Regression Coefficients				
Effect	Intercept	Antisocial	Hyperactivity	Anti. × Hyper.
Intercept	0.10394			
Antisocial	0.01295	0.07186		
Hyperactivity	−0.00032	−0.02680	0.04614	
Anti. × hyper.	−0.02136	−0.01239	0.00031	0.02042

Descriptive Statistics				
Variable	Mean	*SD*	Minimum	Maximum
Antisocial	0	1.342	−1.344	4.656
Hyperactivity	0	1.546	−1.940	3.060

Note. Results pertaining to covariates were obtained, but were omitted from this table because they are not relevant for our calculations.

Probing the interaction further using our MLR two-way interaction tool provides additional information to aid in the interpretation of this conditional effect. We entered specific values of (centered) *hyperactivity* into the utility to assess the effect of *antisocial* on *math* at specific conditional values of the moderator, in this case at the mean and at 1 *SD* above and below the mean (i.e., where centered *hyperactivity* = −1.546, 0, and +1.546). As *hyperactivity* increases, the slope relating *math* to *antisocial* becomes more strongly negative. The simple slope is 0.652 at −1 *SD* ($p = .102$, not significant), 0.037 at the mean of *hyperactivity* ($p = .889$, not significant), and −0.577 at +1 *SD* ($p = .044$, significant). Submitting the R syntax for the simple regression lines to the Rweb server produces the plot of these effects shown in Figure 1.

Probing shows that the simple slope of *math* regressed on *antisocial* is significant at only one of the chosen conditional values of *hyperactivity*; this finding adds substantially to our understanding of the interaction effect. However, the choice of ±1 *SD* for the moderator was essentially arbitrary. Of more interest are the values of *hyperactivity* for which the simple slope is statistically significant. The region of significance on the moderator (*hyperactivity*) ranges from −2.320 to 1.493, indicating that any given simple slope outside this range is statistically significant. Given that centered *hyperactivity* ranges from about −1.94 to about +3.06, this indicates that the effect of *antisocial* on *math* is significant only for relatively high observed values of *hyperactive*. Submitting the second R program produces a plot of the confidence
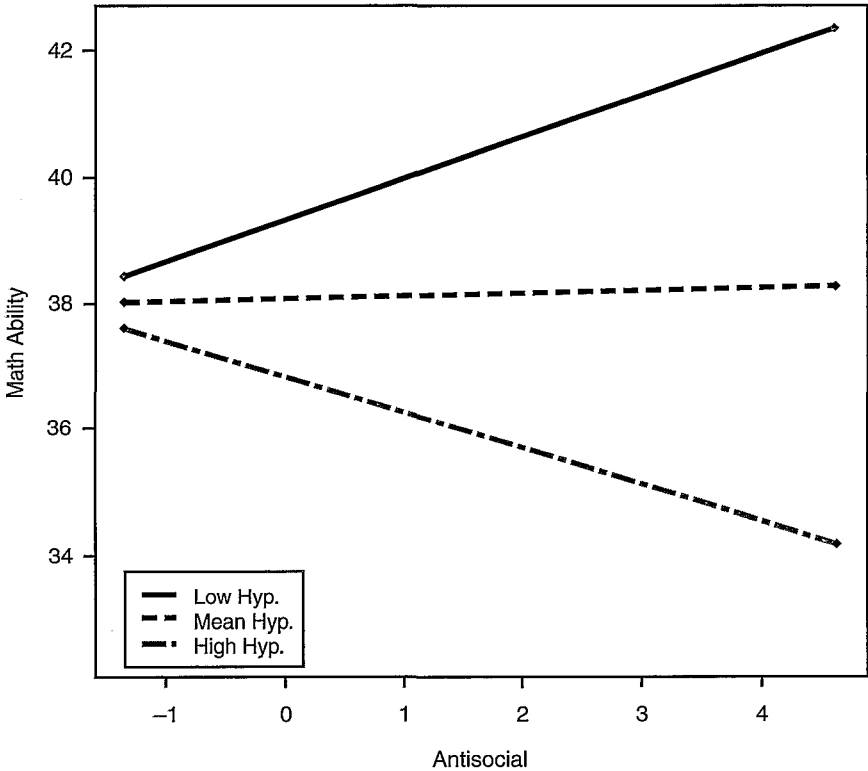
445

FIGURE 1. *Mean plot illustrating the interaction of antisocial behavior and hyperactivity. More negative slopes correspond to higher levels of hyperactivity.*

bands (see Figure 2). The default values of the moderator (−10 and +10) were changed to −3 and +3 in Figure 2 to better match the range of observed data. Because the confidence bands do not include simple slopes of zero for values of *hyperactivity* above 1.493, it can be concluded that the simple slope of *math* regressed on *antisocial* is significantly different from zero for values of *hyperactivity* above this point.

Discussion

When a significant interaction is detected in MLR, HLM, or LCA, the interpretation of the interaction becomes critically important. Aside from the point estimate itself, much important information about the interaction can be obtained by evaluating the simple slopes, regions of significance, and confidence bands. Indeed, it is difficult to imagine a situation where the interpretation of an interaction would not be improved by using these methods. Yet two issues have impeded the application of these methods in practice: uncertainty that methods developed for MLR also
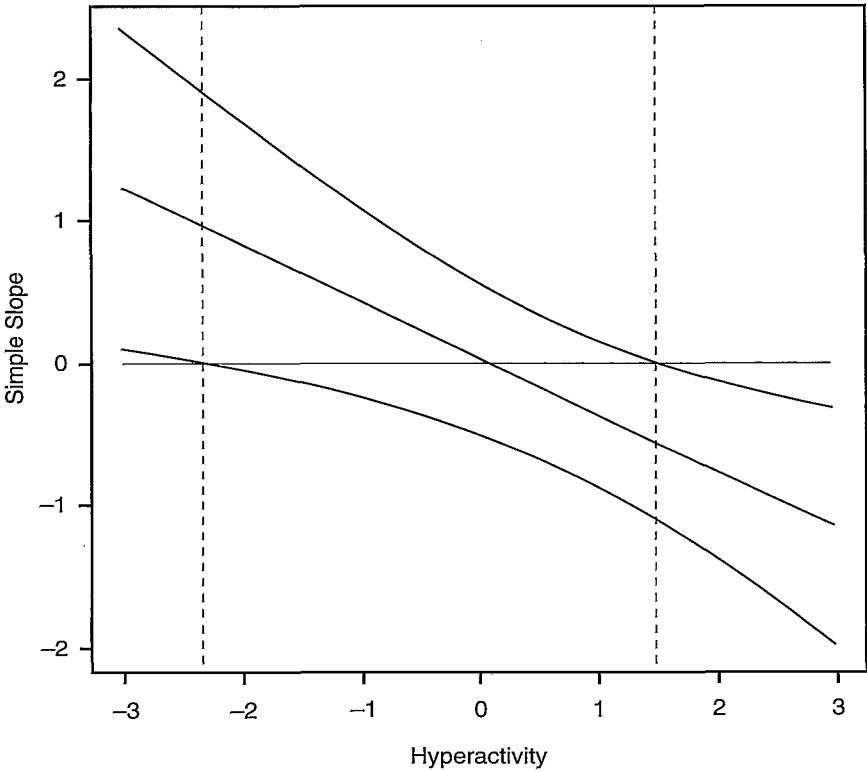
446

FIGURE 2. *Plot illustrating confidence bands for observed sample values of hyperactivity.*

apply to HLM and LCA and a lack of implementation in software. As reviewed here, however, these methods are equally valid and useful for all three statistical models. Furthermore, the online tools described above now make these methods freely available to any researcher, regardless of their statistical software of choice. Together, we believe that these developments will enable applied researchers to better interpret and communicate complex conditional effects in MLR, HLM, and LCA, including both two-way and three-way interactions.

References

Aiken, L. S., & West, S. G. (1991). *Multiple regression: Testing and interpreting interactions.* Newbury Park, CA: Sage.

Bauer, D. J., & Curran, P. J. (2005). Probing interactions in fixed and multilevel regression: Inferential and graphical techniques. *Multivariate Behavioral Research, 40,* 373–400.

Cohen, J. (1978). Partialed products are interactions; partialed powers are curve components. *Psychological Bulletin, 85,* 858–866.

Cohen, J., & Cohen, P. (1983). *Applied multiple regression/correlation analyses for the behavioral sciences* (2nd ed.). Hillsdale, NJ: Erlbaum.

Cohen, J., Cohen, P., West, S. G., & Aiken, L. S. (2003). *Applied multiple regression/correlation analysis for the behavioral sciences* (3rd ed.). Hillsdale, NJ: Erlbaum.

Curran, P. J., Bauer, D. J., & Willoughby, M. T. (2004). Testing main effects and interactions in latent curve analysis. *Psychological Methods, 9*(2), 220–237.

Curran, P. J., Bauer, D. J., & Willoughby, M. T. (2006). Testing and probing interactions in hierarchical linear growth models. In C. S. Bergeman & S. M. Boker (Eds.), *The Notre Dame Series on Quantitative Methodology: Vol. 1. Methodological issues in aging research* (pp. 99–129). Mahwah, NJ: Erlbaum.

Johnson, P. O., & Neyman, J. (1936). Tests of certain linear hypotheses and their applications to some educational problems. *Statistical Research Memoirs, 1*, 57–93.

Rogosa, D. (1980). Comparing nonparallel regression lines. *Psychological Bulletin, 88*, 307–321.

Rogosa, D. (1981). On the relationship between the Johnson-Neyman region of significance and statistical tests of parallel within group regressions. *Educational and Psychological Measurement, 41*, 73–84.

Saunders, D. R. (1956). Moderator variables in prediction. *Educational and Psychological Measurement, 16*, 209–222.

Tate, R. (2004). Interpreting hierarchical linear and hierarchical generalized linear models with slopes as outcomes. *The Journal of Experimental Education, 73*, 71–95.

Wilkinson, L., & the APA Task Force on Statistical Inference. (1999). Statistical methods in psychology journals: Guidelines and explanations. *American Psychologist, 54*(8), 594–604.

Authors

KRISTOPHER J. PREACHER is now Assistant Professor of Quantitative Psychology at the University of Kansas, 1415 Jayhawk Blvd., Room 426, Lawrence, KS 66045-7556; preacher@ku.edu. His areas of specialization include factor analysis, structural equation modeling, multilevel modeling, growth curve analysis, model fit, and the assessment of mediation and moderation effects.

PATRICK J. CURRAN is a Professor in the L. L. Thurstone Psychometric Laboratory in the Department of Psychology at the University of North Carolina at Chapel Hill, CB #3270 Davie Hall, Chapel Hill, NC 27599-3270; curran@unc.edu. His areas of specialization are structural equation modeling and multilevel modeling, particularly as applied to longitudinal data.

DANIEL J. BAUER is Assistant Professor of Psychology, Department of Psychology, University of North Carolina at Chapel Hill, CB #3270 Davie Hall, Chapel Hill, NC 27599-3270; dbauer@email.unc.edu. His areas of specialization are mixed-effects models, structural equation models, and finite mixture models, with an emphasis on applications to longitudinal data on adolescent problem behavior.