Interpretation and Visualization of Moderation Effects and Random Slopes in Multilevel Models

Julie A. Lorah [a] ✉ ⓘ

[a]Counseling and Educational Psychology, Indiana University

Abstract ■ Interpretation of complex effects and models can be one of the most challenging and important aspects of quantitative data analysis. The present study tackles this issue for moderation effects, including random slope effects, for multilevel models. To demonstrate the generalization of these procedures beyond the basic multilevel model, the multilevel logistic regression model is used. A moderation effect may be useful when a researcher would like to assess how a particular relationship differs for different groups or different levels of a moderator variable. When the moderator under consideration is a random effect, a random slope model arises. The random slope model has various applications; for example, when observations are nested within individuals comprising a longitudinal design, a random slopes model can be used to assess individual growth trajectories for the subjects in the study. However, these useful effects may be particularly difficult to interpret substantively. Therefore, the present study suggests a method combining the traditional aspects of plotting moderation effects with quantities of interest (QI) computation. Specific suggestions and examples, including R syntax, for associated data visualizations are provided.

Keywords ■ moderation, interaction, random slopes, multilevel logistic regression, interpretation, data visualization. **Tools** ■ R.

✉ Jlorah@iu.edu

Acting Editor ■ Denis Cousineau (Université d'Ottawa)

Introduction

Interpreting results and creating appropriate data visualizations is an important, but often overlooked, part of the quantitative data analysis process. Although guidance is available for applied researchers interpreting many of the more common quantitative models, it is often unclear how to do so for more complex effects and models. The present contribution addresses this gap by suggesting a novel combination of procedures and data visualizations to interpret moderation effects and random slopes in multilevel models. Although these procedures are first demonstrated through a multilevel regression model, they generalize easily to more complex models; to demonstrate this point, the procedure is then shown with the multilevel logistic regression model for binary outcomes. The present study contributes to the literature by offering concrete steps and examples for interpretation of moderation in complex multilevel models; in addition, there is very little literature available to aid researchers in interpreting

random slopes effects and the present study fills this gap.

Multilevel Models

It is common for data to be nested or clustered, such as when subjects are nested within neighborhoods or observations are nested within subjects (i.e. longitudinal data). Ignoring this type of nesting can result in severely inflated Type I error rates (Snijders & Bosker, 2012) and is not recommended. Instead, this type of nested data is ideal for analysis with the multilevel model because predictors can be included at all levels within the model. For example, if students are nested within schools, both student-level and school-level predictors can be added. The multilevel model is similar to the regression model, but with the addition of a random effect which allows the outcome to vary randomly by group membership. The model is described in more detail below, but for readers unfamiliar with the multilevel model, there are several helpful introductory sources including tutorials (see, for example: Kahn, 2011; Reise & Duan, 1999; Woltman, Feldstain, MacKay, & Rocchi,

2012) and texts (see, for example: Snijders & Bosker, 2012).

Moderation and Random Slope Effects

Moderation effects exist when a given relationship varies as a function of a third, moderator variable. Consider a hypothetical example: perhaps the relationship between studying habits and academic achievement is moderated by a student's level of anxiety such that students with high anxiety benefit less from studying than students with low anxiety. In other words, the relationship between studying and achievement is stronger for students with low anxiety. Although the scenario just described is hypothetical, this description serves to demonstrate the nuanced and helpful findings possible by examining moderation effects. The multilevel model with a level-one moderation effect is (Snijders & Bosker, 2012):

$$
\begin{aligned}
Y_{ij} &= \beta_0 + \beta_1 \times Xij + \beta_2 \times Mij + \beta_3 \times X_{ij}M_{ij} + u_{0j} + e_{ij} \\
Var(u_{0j}) &= \tau_0^2 \\
Var(e_{ij}) &= \sigma^2 \\
Cov(u_{0j}, e_{ij}) &= 0
\end{aligned}
\tag{1}
$$

where i represents the level-one unit which is the individual; j represents the level-two unit which is the group membership (e.g., school, hospital, etc.); Y represents the continuous outcome measure; β_0 is the intercept, all other β are slope coefficients, X is the independent variable, M is the moderator variable, u_{0j} is the level-two random effect distributed with mean of zero and estimated variance τ_0^2; and e_{ij} is the level-one random effect distributed with mean of zero and estimated variance σ^2. The interaction effect is represented by the product term X times M, and although both variables are measured at level-one in the present specification, more generally, the product term may be constructed as the product of two level-one variables, two level-two variables or the combination of a level-one and level-two variable.

Continuing the example just provided, the outcome measure, academic achievement would be represented by Y in Equation 1; the independent variable, studying habits, would be represented by X; and the moderator variable, anxiety, would be represented by M.

The significance of the moderation effect can be assessed with a t-test on β_3 (Aiken & West, 1991). If the moderation is significant, the result should be plotted since there is no other consistent way to ascertain the nature of the effect (Aiken & West, 1991); however, if the moderation effect is not significant, researchers must carefully assess the situation since a non-significant effect can be the result of no true effect or low power (i.e. Type II error; Lorah, 2020). Research recommends conducting power analyses for the interaction effect prior to analysis, mean-centering predictors, considering a main-effects-only model, and considering information criteria approaches to guard against the possibility of misleading interpretation when the moderation effect is not significant (Lorah, 2020).

In addition to fixed-effect interaction effects, the multilevel regression model also allows for an interaction between a fixed and random effect; this is also referred to as the random slopes model (Lorah, 2018). For example, imagine an investigation of a sample of high school students nested within schools. The categorical variable, school membership, represents the random effect. Perhaps this investigation considers a moderation effect whereby the relationship between a student's studying habits and their academic achievement is moderated by school membership. Notice that this simply represents a moderation effect between a fixed effect (studying habits) and a random effect (school membership), which is also known as a random slope effect. This hypothetical random slope effect could be used to demonstrate whether or not the relationship between studying habits and academic achievement is different in different schools. The multilevel model with a random slope effect is (Snijders & Bosker, 2012):

$$Y_{ij} = \beta_0 + \beta_1 \times X_{ij} + u_{0j} + u_{1j} \times X_{ij} + e_{ij}$$
$$Var(u_{0j}) = \tau_0^2$$
$$Var(u_{1j}) = \tau_1^2$$
$$Cov(u_{0j}, u_{1j}) = \tau_{01} \tag{2}$$
$$Var(e_{ij}) = \sigma^2$$
$$Cov(u_{0j}, e_{ij}) = 0$$
$$Cov(u_{1j}, e_{ij}) = 0$$

where all notation is consistent with Equation 1 and u_{1j} represent level-two residuals.

Multilevel Logistic Regression Models

Procedures suggested in the present study generalize beyond the multilevel model to more complex models: this will be demonstrated with the multilevel logistic regression model, which is appropriate for nested data with a binary outcome, such as whether or not a student graduated from high school. This model uses the logit, or log-odds, of the probability of the binary outcome taking on a value of one versus zero (Snijders & Bosker, 2012). Although many researchers interpret the results of logistic

regression models by computing the odds and interpreting this quantity, methodologists point to the difficulty in truly conceptualizing the scope of an effect through presentation of odds (Long, 1997). For this reason, researchers are recommended to use predicted probabilities (Long, 1997) to interpret model results, and that is the method used in the current study.

Multilevel Logistic Regression with Moderation and Random Slopes

The multilevel logistic regression model with a fixed effect interaction at level-one can be specified as follows (Snijders & Bosker, 2012):

$$logit(p_{ij}) = \beta_0 + \beta_1 \times X_{ij} + \beta_2 \times M_{ij} + \beta_3 \times X_{ij}M_{ij} + u_{0j}$$
$$\text{where } logit(p) = \ln\left(\frac{p}{1-p}\right) \tag{3}$$
$$Var(u_{0j}) = \tau_0^2$$

where all notation is consistent with Equation 1 and p represents the probability of the binary outcome variable taking on a value of one (versus zero, assuming dummy-coding). If the outcome being modeled was a binary indicator of whether a student graduated high school or not, then p in Equation 3 would represent the probability that the given student graduates high school. Note that the parameter e_{ij} (individual-level residual) is no longer needed in

Equation 3 (or in Equation 4, shown below) since the probability (p) of an outcome is being modeled, represented as an aggregate across participants, rather than the outcome value itself.

The multilevel logistic regression model can also be estimated with a random slope effect and this model is specified as follows (Snijders & Bosker, 2012):

$$logit(p_{ij}) = \beta_0 + \beta_1 \times X_{ij} + u_{0j} + u_{1j} \times X_{ij}$$

$$\text{where } logit\,(p) = \ln\left(\frac{p}{1-p}\right)$$

$$Var(u_{0j}) = \tau_0^2$$

$$Var(u_{1j}) = \tau_1^2$$

$$Cov(u_{0j}, u_{1j}) = \tau_{01}$$

(4)

where all notation is consistent with Equations 1 and 2.

Interpretation of Moderation and Random Slopes Effects

Interpreting moderation and random slopes effects can be difficult. Typically, to interpret moderation effects, the simple slopes approach is recommended (Aiken & West, 1991) and involves plotting the relationship between the independent and dependent variables for certain discrete values of the moderator variable. This is essentially a process of computing quantities of interest (QI; Carsey & Harden, 2014) since the procedure involves computing model-based predictions of a relevant outcome or QI. The choice of specific values for the moderator variable is arbitrary, but for a categorical moderator, each category may be plotted and for a continuous moderator, often the mean (or median) and one standard deviation above and below the mean are used. The literature provides detailed descriptions of implementing this approach in the multiple regression model (for example, see Aiken & West, 1991; Jaccard, Turrisi, & Wan, 1990; Jose, 2013; Lorah & Wong, 2018; Lorah & Miksza, 2019) including R code for doing so (Caron, Valois, & Gellen-Kamel, 2020), but it may be unclear how to generalize this approach for more complicated models, such as the multilevel model or a logistic regression model for binary outcomes or to more complicated effects, such as the random slope effect.

In particular, there is very little literature available to aid the researcher in interpreting random slopes effects. The available literature indicates that the effect may be interpreted by using τ_1^2 to indicate the hypothetical value for specific random slopes, for example, the range in which 95% of the slopes are expected to fall (Lorah, 2018). Although this may be helpful, it may still be difficult for the

reader to gain a sense of the scope of the effect with this interpretation; therefore, the present study builds on this by suggesting a method for plotting these effects by generalizing the simple slopes procedure to additional models and effects.

In order to implement the simple slopes procedure generally, the researcher will need a way to choose specific, relevant values of the moderator variable and the model prediction equation. For example, based on the multilevel model provided in Equation 1, the prediction equation can be used to compute predicted values for the outcome, Y:

$$Y_{ij} = \beta_0 + \beta_1 \times X_{ij} + \beta_2 \times M_{ij} + \beta_3 \times X_{ij}M_{ij} + u_{0j} \quad (5)$$

where variables are consistent with those defined in Equation 1 and predicted values of Y represent the QI for this model. The regression coefficients (the β terms) will be obtained after estimation, and the values for the predictor variables (X and M, in this case) can be chosen substantively. The random effect (u_{0j}) is centered around zero with estimated variance (τ_0^2) and so a value of zero would represent the mean level two unit; or a relevant value from the distribution can be chosen. If additional control variables are included in the model, a relevant value, such as the mean, can be chosen. In this way, a predicted value for Y can be computed for any relevant combination of predictor variables. This prediction equation can also be used to create a plot for interpretation of moderation by substituting a specific value for M and plotting the resulting curves (see below for a demonstration), which is also referred to as the simple slopes procedure.

For the multilevel logistic regression, the prediction equation is a bit more complicated and can be derived by solving for p in Equation 3 (Snijders & Bosker, 2012):

$$p = \frac{e^{\beta_0 + \beta_1 \times X_{ij} + \beta_2 \times M_{ij} + \beta_3 \times X_{ij} \times M_{ij} + u_{0j}}}{1 + e^{\beta_0 + \beta_1 \times X_{ij} + \beta_2 \times M_{ij} + \beta_3 \times X_{ij} \times M_{ij} + u_{0j}}}$$

(6)

where variables are consistent with those defined in Equation 1 and the value of p, or the predicted probably of an event occurring, represents the QI for this model. Similarly, after substituting back in the values of the parameter estimates based on the model results ($\beta_0 - \beta_3$), this equation can be used to compute predicted probabilities (values of p) for model interpretation. In other words, the predicted probability represents the probability of success (i.e. $Y = 1$) on the outcome variable, given specific values for the predictor variables (i.e. X and M). Note that this idea of providing model-predicted values is a very general and useful tool for model interpretation that can be applied to many models.

Research Questions

1. How can the moderation effects within a multilevel model and multilevel logistic regression model be interpreted substantively through appropriate data visualization?
2. How can the results from research question 1 be generalized more broadly to more complex effects, such as the random slopes effect?

Methods

In order to demonstrate interpretation of interaction and random slopes effects in the multilevel regression model, hypothetical data are generated and analyzed. Hypothetical data are used rather than real data so that the reader can easily replicate the analysis and demonstration and because the demonstration is not intended to provide any sort of substantive conclusions, simulated data precludes any attempt to do so.

Data are simulated for four models: the multilevel regression with moderation (Equation 1); the multilevel model with random slopes (Equation 2); the multilevel logistic regression with moderation (Equation 3); and the multilevel logistic regression with random slopes (Equation 4).

For moderation models, data are simulated according to Equations 1 and 3 with the following properties: total sample size = 200; number of groups (level-2 units) = 20; $\beta_0 = \beta_1 = \beta_2 = .3$; $\beta_3 = .2$; $\tau_0 = \sigma = .5$; X is normally distributed with mean 0 and standard deviation 1; and M is binary with $p = .5$. In the example interpretations that follow, X will represent studying habits, where higher values indicate better studying habits; M will represent anxiety measured with two categories: low anxiety and high anxiety; and the outcome measure, Y, will represent academic achievement where higher values indicate higher achievement. Note that although X is simulated according to the normal distribution, real data may not approximate the normal distribution as precisely, which can

impact results. Data were generated using R (R Core Team, 2021) and models were estimated with the `glmer()` function from the lme4 package (Bates, Maechler, Bolker, & Walker, 2015). All R syntax is included in the appendix. For the random slopes effect, data are simulated according to Equations 2 and 4. All quantities are the same as for the interaction plots, with a few differences. For the multilevel model with random slopes, number of groups (level-2 units) = 10; $\beta_1 = 5$; and $\tau_0 = \tau_1 = \sigma = 2$. For the multilevel logistic regression with random slopes, number of groups (level-2 units) = 10; $\beta_0 = \beta_1 = 1$; and $\tau_0 = \tau_1 = 1.5$. Otherwise, the process is identical.

For interpretation of model results, providing a measure of effect size, in addition to the visualizations, is helpful. For this purpose, R^2 is computed for each of the four models using the MuMIn package (Barton, 2019) in R (R Core Team, 2021) with the `r.squaredGLMM()` function which computes variance explained in a multilevel model by partitioning variance (Johnson, 2014). This function provides both marginal (variance explained by the fixed effects) and conditional (variance explained by the fixed and random effects) measures (Barton, 2019) and conditional measures are reported in the present analysis since interest lies in both the fixed as well as random effects.

Results and Discussion

Results from the simulated data are presented in the context of the hypothetical research questions and variables introduced previously.

Figure 1

Figure 1 presents the results of a moderation effect in the context of a multilevel model. Beginning with overall model fit, the R^2 value for the model is 0.67, indicating about 67% of the variance in achievement is attributable to the fixed and random effects in the model. A significant moderation effect in this example indicates that the relationship between studying habits and achievement differs depending on the student's anxiety levels. The plot further clarifies this relationship by indicating that the relationship between studying habits and achievement is positive, but stronger for those students with lower anxiety. Specifically, for a student with high anxiety, each standard deviation increase in studying habit behavior is related to about 0.3 standard deviation increase in achievement, while for students with low anxiety, each standard deviation increase in studying habit behavior is related to about 0.5 standard deviation increase in achievement. These specific relationships can be found by substituting the specific coded values for anxiety (i.e. 0 or 1) into the model equation (Equation 5) through the simple slopes procedure (see procedure described above, or refer to additional sources

Figure 1 ■ Moderation plot for multilevel model
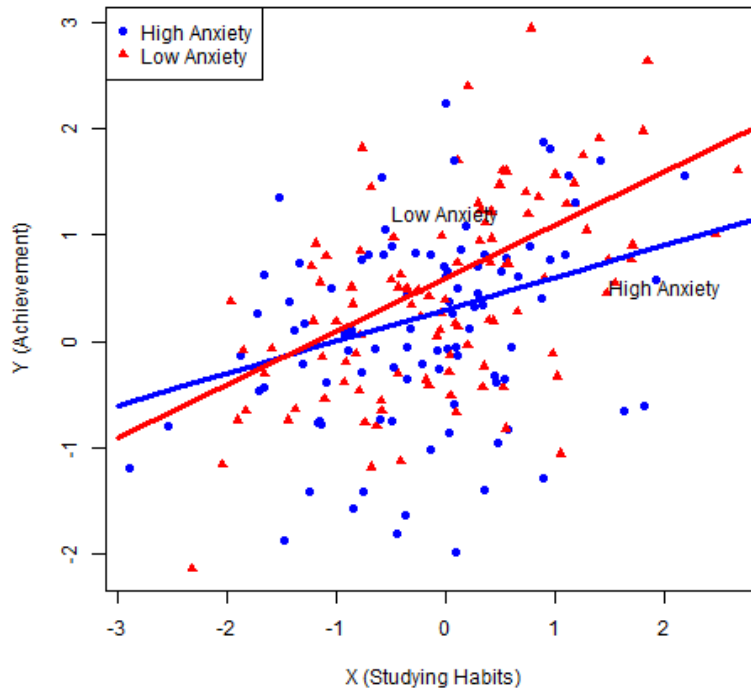
such as Aiken & West, 1991). These results may indicate that studying is most beneficial when the student is not overwhelmed with anxiety.

Figure 1 also includes observed data where red triangles are employed to represent students with low anxiety and blue circles are employed to represent students with high anxiety. The inclusion of observed data in this graphic can offer various benefits, including a more nuanced picture of the data; a better sense of the empirical distribution for each group; and clarification of how much, if at all, the regression lines representing the model under consideration are extrapolating past observed data. In the present example, it is clear that there is still quite a bit of variation within each group, and that the model-based regression lines do not extrapolate beyond observed values, which is desirable.

Some of the values included in Figure 1 (as well as Figures 2-4) are included in Table 1. This is provided to help clarify what is actually being plotted in the data visualization. In addition, the table is provided to demonstrate the advantages of plotting these results versus just providing a table for these results. Inspecting the table results in very little insight into the nature of the interaction effect; particularly compared with the result displayed in Figure 1.

Figure 2

Generalizing this procedure to the multilevel logistic regression provides the results in Figure 2. Beginning with overall model fit, $R^2 = 0.09$, indicating about 9% of the variance in achievement is attributable to the fixed and random effects in this model. The outcome of interest is now a binary indicator of graduating versus not graduating. Note that the logistic regression model proceeds by modeling the probability of the event (i.e. graduation) rather than the binary event itself. Accordingly, the results that are plotted are predicted probabilities of graduating. Figure 2 indicates that students with better studying habits are generally more likely to graduate, but that this relationship is moderated by the student's anxiety such that students with low anxiety are more likely to benefit from studying.

Note that the relationship between studying habits and probability of graduation is non-linear for each group (for both low and high anxiety students). This is a property of the model under examination, which employs a logit link function for describing the relationship between predictors and the probability of a binary outcome (see Equation 3). Substantively, this seems reasonable as well. Figure 2 indicates that for low anxiety students, there are some-

Figure 2 ■ Moderation plot for multilevel logistic regression
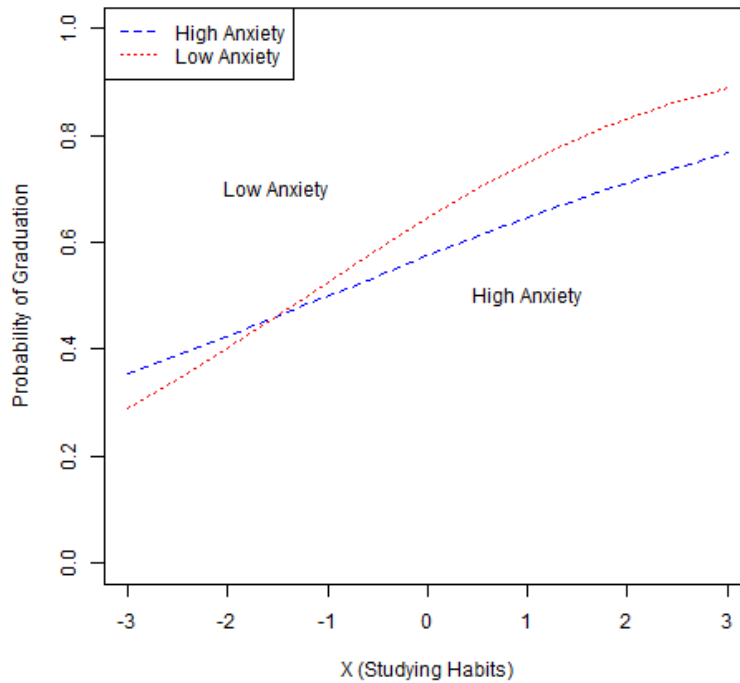
what of diminishing returns as the student approaches the highest values for studying habits, as the curve flattens a bit at the highest values of X. This is substantively reasonable, as one might expect bigger changes in a relevant outcome for improving poor studying habits compared with fine-tuning already fairly good studying habits.

Figure 3

This procedure is also applicable to random slope effects, as demonstrated in Figure 3. The R^2 for this model is 0.90, indicating a large percentage of the variance in achievement is explained by the fixed and random effects. Random slopes essentially imply that the moderator is the random effect. Accordingly, this model is examining whether the relationship between studying habits and achievement is moderated by school membership. A significant random slope effect indicates that this relationship is different in different schools. Figure 3 shows that depending on the school in which the student is enrolled, studying could have a stronger or weaker relationship, but that in general, this relationship is positive. Note that similar to Figure 3, the observed data is included where each different colored point represents an observation from a different school (the schools are identified by their number, from one to ten, as shown on the legend). Specific schools of substantive interest can thereby be identified in this visualization.

Figure 4

Finally, this method for interpretation of random slopes may be generalized to the multilevel logistic regression model. For this model, $R^2 = 0.68$, indicating the 68% of the variance in achievement is explained by the model. Figure 4 demonstrates that the relationship between studying habits and probability of graduation differs by school membership. In most schools this relationship is generally positive, in some schools it is slightly negative, and in one school there is not a strong relationship at all (the curve is fairly flat). Note that similar to Figure 2, this visualization clarifies the non-linear relationship being modeled between studying habits and the probability of graduation. Also, similar to Figure 3, a different color for each school is used to allow investigation of specific schools. For example, it appears that the slightly negative relationship between studying habits and probability of graduation is observed in schools 4 and 5, and so these schools might be interesting to further investigate, since they may be outliers in this context.

Further Generalizations

Note that there are a wide variety of ways in which these procedures could be generalized to slightly different models. For example, the demonstrations were provided with

Table 1 ■ First 6 rows of data used to create Figures 1-4, respectively

Figure 1 data			Figure 2 data		
M	X	QI	M	X	QI
0	−3	−0.6	0	−3	0.354344
1	−3	−0.9	1	−3	0.289051
0	−2.4	−0.42	0	−2.4	0.396517
1	−2.4	−0.6	1	−2.4	0.354344
0	−1.8	−0.24	0	−1.8	0.440286
1	−1.8	−0.3	1	−1.8	0.425558

Figure 3 data			
X	u_{0j}	u_{1j}	QI
-3	2.383864	−1.33442	−8.31289
-2.4	2.383864	−1.33442	−6.11354
-1.8	2.383864	−1.33442	−3.91419
-1.2	2.383864	−1.33442	−1.71484
-0.6	2.383864	−1.33442	0.484513
0	2.383864	−1.33442	2.683864

Figure 4 data			
X	u_{0j}	u_{1j}	QI
-3	1.62702	−2.10365	0.99737
-2.4	1.62702	−2.10365	0.994912
-1.8	1.62702	−2.10365	0.990181
-1.2	1.62702	−2.10365	0.981135
-0.6	1.62702	−2.10365	0.964057
0	1.62702	−2.10365	0.93258

Note. X represents x-axis; QI represents y-axis; different values of M are used to create different lines for each moderator (i.e. low anxiety or high anxiety); different values of u_{0j}/u_{1j} are used to create different lines for each group (i.e. school).

a binary moderator variable. If the moderator were continuous, relevant values could be chosen, such as the first, second, and third quartile; or the mean, one standard deviation above the mean, and one standard deviation below the mean. Plotting could then proceed analogously. Alternatively, although the independent variable is continuous in the present demonstration, it is also possible to model a categorical independent variable, in which case the discrete levels of the variable could each be used in the plot.

Another common scenario in real data might involve a large number of groups (level-two units). For example, if the present demonstration included 200 schools, instead of 10 schools, the plot might be too busy to be truly informative. In this case, a smaller number of groups could be randomly selected. Alternatively, specific relevant values for the slopes could be chosen and then plotted. For example, the literature indicates that one way to interpret random slopes effects is to report the average slope, and a "high" and "low" value such as the first and third quartile; one standard deviation above and below the mean; or the values within which 95% of slopes are expected to fall (Lorah, 2018). Accordingly, rather than plotting the ex-

pected regression lines for specific observed groups (for example, schools), the plot could be created by plotting the expected slopes of an explicitly defined high, low, and average group.

Lastly, the present study describes computation of predicted outcome values (for the multilevel model) and predicted probabilities (for the multilevel logistic regression model). These two outcomes could more generally be considered quantities of interest (QI) which describes any model-based quantity that is relevant for model interpretation; use of simulation methods can further provide confidence intervals around these quantities of interest (see Carsey & Harden, 2014, for detailed instructions for doing so).

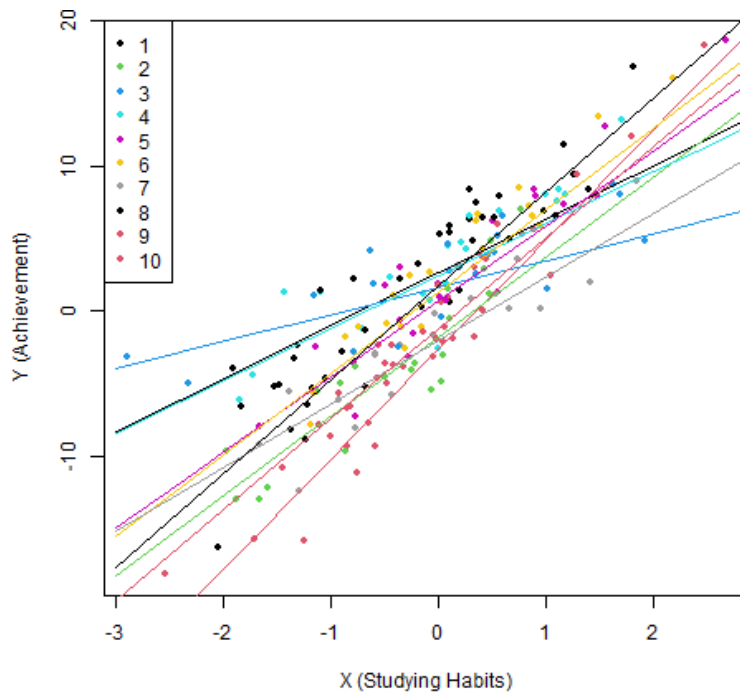### *Summary of Analysis Steps*

In summary, the present study suggests the following steps for analyzing and interpreting model results when a moderation or random slope effect is included in a multilevel model:

1. Specify and estimate model of interest (see R code, part 1b, 2b, 3b, or 4b).

Figure 3 ■ Random slopes plot for multilevel model

2. Choose specific values for all predictor variables included in the model (see R code, part 1c, 2c, 3c, or 4c, section "values for varying conditions").

   (a) First, choose specific values of substantive interest for the moderator variable. The simple slopes literature provides guidance for how to do this, but typically the mean and a high and low value are chosen for continuous moderators and each possible value is chosen for categorical moderators. For example, in Figure 1, low anxiety and high anxiety were chosen for the moderator values. Or, in the case of a random slopes model, each group (i.e. level-two unit) can be used. In the example analysis, all 10 schools are plotted.

   (b) Second, choose specific values of substantive interest for the independent variable. For example, consider Figure 1. The independent variable, X (studying habits), is displayed on the x-axis, and this continuous variable ranges from -3 to +3. When choosing appropriate values for the independent variable, be careful not to extrapolate beyond values present in the observed data.

   (c) Third, choose specific values of substantive interest for any control variables (any additional predictors). In the example models, there are no additional control variables, so this step may be skipped. However, in many applied examples, there will be additional predictors beyond the moderator (M) and the independent variable (X). For these remaining predictors, choose a baseline value such as a reference group (for categorical predictors), or the mean value (for continuous or categorical predictors). Since these control variables may be less substantively interesting than the moderator or independent variables, a single value for each control variable should be sufficient for the purposes of computing the appropriate quantity of interest.

3. Next, compute the quantity of interest (i.e. predicted values or predicted probabilities) for each of these sets of predictors from step 2 (see R code, part 1c, 2c, 3c, or 4c). In the present study, the quantity of interest is either the predicted outcome value (Equation 5) or the predicted probability (Equation 6).

4. Create appropriate tables and/or plots of the results (see R code, par 1d, 2d, 3d, or 4d).

Conclusions

The present study has suggested and demonstrated a relatively straight-forward and accessible method for interpreting interaction and random slopes effects in multilevel models using the concepts of quantities of interest, and

Figure 4 ■ Random slopes plot for multilevel logistic regression
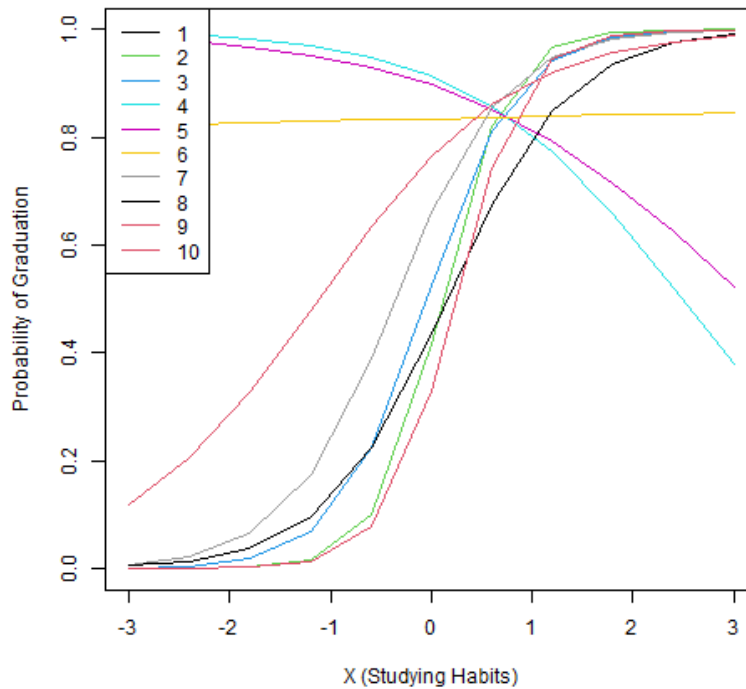
moderation (or random slopes) plots. The prior demonstration has shown that creating these plots results in far more interpretable results compared with tables or simply interpreting significance or direction of effects and that this procedure can be generalized to many other, similar models and effects. Future research should continue to pursue these graphical methods of interpretation, perhaps extending the findings to other more complicated models and effects. In addition, future applied studies can use these techniques and instructions to further elucidate the study results in various ways. In particular, applied studies that include nesting, such as children nested within families or students nested within schools may find multilevel models and random slopes effects particularly relevant. Moderation effects could be used to examine demographic predictors such as the intersection of race and gender or psychological constructs, such as protective factors for various outcomes.

References

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

Barton, K. (2019). Mumin: Multi-model inference (Version 1.43.15). Retrieved from https://CRAN.R-project.org/package=MuMIn

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1), 1–48. doi:10.18637/jss.v067.i01

Caron, P., Valois, P., & Gellen-Kamel, A. (2020). Some computational description of moderation analysis. *The Quantitative Methods for Psychology*, *16*(1), 9–20.

Carsey, T. M., & Harden, J. J. (2014). *Monte carlo simulation and resampling methods for social science*. Sage: Los Angeles.

Jaccard, J., Turrisi, R., & Wan, C. K. (1990). *Interaction effects in multiple regression*. Newbury Park: Sage Publications.

Johnson, P. C. (2014). Extension of nakagawa & schielzeth's rglmm2 to random slopes models. *Methods in Ecology and Evolution*, *5*, 944–946. doi:10.1111/2041-210X.12225

Jose, P. E. (2013). *Doing statistical mediation & moderation*. New York: The Guilford Press.

Kahn, J. H. (2011). Multilevel modeling: Overview and applications to research in counseling psychology. *Journal of Counseling Psychology*, *58*(2), 257–271. doi:10.1037/a0022680

Long, J. S. (1997). *Advanced quantitative techniques in the social sciences series, vol 7. : Regression models for*

categorical and limited dependent variables. Raccoon City: Sage Publications, Inc.

Lorah, J. A. (2018). Effect size measures for multilevel models: Definition, interpretation, and timss example. *Large-scale Assessments in Education*, *6*(8), 1–9. doi:10.1186/s40536-018-0061-2

Lorah, J. A. (2020). Interpretation of main effects in the presence of non-significant interaction effects. *The Quantitative Methods for Psychology*, *16*(1), 33–45.

Lorah, J. A., & Miksza, P. (2019). Applications of moderation analysis for music education research. *Bulletin of the Council for Research in Music Education*, *220*, 21–41. doi:10.5406/bulcouresmusedu.220.0021

Lorah, J. A., & Wong, Y. J. (2018). Contemporary applications of moderation analysis in counseling psychology. *Journal of Counseling Psychology*, *65*, 629–640. doi:10.1037/cou0000290

R Core Team. (2021). *R: A language and environment for statistical computing [computer software]*. Vienna, Austria. URL: R Foundation for Statistical Computing. Retrieved from https://www.R-project.org/

Reise, S. P., & Duan, N. (1999). Multilevel modeling and its application in counseling psychology research. *The Counseling Psychologist*, *27*(4), 528–551. doi:10.1177/0011000099274003

Snijders, T. A. B., & Bosker, R. J. (2012). *Multilevel analysis: An introduction to basic and advanced multilevel modeling*. Los Angeles: Sage.

Woltman, H., Feldstain, A., MacKay, J. C., & Rocchi, R. (2012). An introduction to hierarchical linear modeling. *Tutorials in Quantitative Methods for Psychology*, *8*(1), 52–69.

Appendix: R Syntax

R syntax to produce Figure 1: Moderation plot for multilevel model

```
library(lme4) #lmer function
library(MuMIn)#r.squaredGLMM function

##1a. Simulate data##
#####################

#set constants#
set.seed(50)
b0<- .3 #intercept
b1<- .3 #slope for X
b2<- .3 #slope for M
b3<- .2 #slope for X*M
L2sd<-.5 #square root of tau-squared(zero); level-two random effect
L1sd<-.5 #square root of sigma-squared; residual (for linear model only)
N<-200 #total sample size for simulated data
NumGroups<-20 #total level-2 sample size (number of groups)

#generate variables#
X<-rnorm(N) #independent variable
M<-rbinom(N,1,.5) #moderator variable
ID<-rep(seq(from=1,to=NumGroups,by=1),times=N/NumGroups) #group ID number
U0j<-rep(rnorm(NumGroups,mean=0,sd=L2sd),times=N/NumGroups) #values for the u0j random effect
eij<-rnorm(N,mean=0,sd=L1sd) #values for the residual (for linear model only)

##generate Y (continuous)
Y_continuous<-b0+b1*X+b2*M+b3*X*M+U0j+eij

#create data frame to store simulated variables
mydata1<-data.frame(ID=ID,Y=Y_continuous,X=X,M=M)

##1b. Estimate model##
#####################

M0<-lmer(Y~X*M+(1|ID),data=mydata1)
r.squaredGLMM(M0)

##1c. Compute QI values for M0 (linear model)##
#####################
```
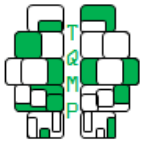
```
#This portion is a loop to find the quantities of interest
    #while varying the values of M and X

lbx<- -3 #this is the lower bound of the x-axis for the plot
    #recommend choosing a value somewhere around observed min(X)
ubx<- 3 #this is the upper bound of the x-axis for the plot
    #recommend choosing a value somewhere around observed max(X)
StepSize<-(ubx-lbx)/10

#values for the varying conditions#
values_M<-c(0,1)
values_X<-seq(from=lbx,to=ubx,by=StepSize)

#num. rows in result (the matrix where QI values will be stored)
ncond<-length(values_M)*length(values_X)

#create a matrix to store the resulting QI values
result<-matrix(ncol=3,nrow=ncond)
colnames(result)<-c("M","X","QI")
result[,1]<-rep(values_M,length.out=ncond)
result[,2]<-rep(values_X,each=length(values_M),length.out=ncond)

#loop to compute actual values of QI#
for(i in 1:(ncond)){
M<-result[i,1]
X<-result[i,2]
result[i,3]<-b0+b1*X+b2*M+b3*X*M
}

FinalResult<-as.data.frame(result)
#display first 6 rows of the data frame
head(FinalResult)

##1d. Plot QI values for M0 (linear model)##
#####################

#create data subset for each line#
subset1<-FinalResult[FinalResult$M==0,]
subset2<-FinalResult[FinalResult$M==1,]
color_values=c("blue","red")#one color per factor level
pch_values=c(16,17)#one shape per factor level

#Plot 1#
#line for M=0, then for M=1#
bmp(file="Plot1.bmp")
plot(mydata1$X,mydata1$Y,
    col=color_values[factor(mydata1$M)],
    pch=pch_values[factor(mydata1$M)],
    ylab="Y (Achievement)",xlab="X (Studying Habits)")
lines(x=subset1$X,y=subset1$QI,type="l",col=color_values[1], lwd=3)
lines(x=subset2$X,y=subset2$QI,type="l",col=color_values[2], lwd=3)
text(2,.5,"High Anxiety")  #M=0
text(0,1.2,"Low Anxiety")  #M=1
legend("topleft",
    legend=c("High Anxiety","Low Anxiety"),
    pch=pch_values,
    col=color_values)
dev.off()
```

R syntax to produce Figure 2: Moderation plot for multilevel logistic regression

```
##2a. Simulate data##
####################

#set constants#
```

```
set.seed(50)
b0<- .3 #intercept
b1<- .3 #slope for X
b2<- .3 #slope for M
b3<- .2 #slope for X*M
L2sd<-.5 #square root of tau-squared(zero); level-two random effect
N<-200 #total sample size for simulated data
NumGroups<-20 #total level-2 sample size (number of groups)

#generate variables#
X<-rnorm(N) #independent variable
M<-rbinom(N,1,.5) #moderator variable
ID<-rep(seq(from=1,to=NumGroups,by=1),times=N/NumGroups) #group ID number
U0j<-rep(rnorm(NumGroups,mean=0,sd=L2sd),times=N/NumGroups) #values for the u0j random effect

#generate prob(Y)#
ExpXB<-exp(b0+b1*X+b2*M+b3*X*M+U0j)
ProbY<-ExpXB/(1+ExpXB)

#generate Y (binary)
Y_binary<-rbinom(N,1,ProbY)

#create data frame to store simulated variables
mydata2<-data.frame(ID=ID,Y=Y_binary,X=X,M=M)

##2b. Estimate model##
######################

M1<-glmer(Y~X*M+(1|ID),family=binomial(link="logit"),data=mydata2)
r.squaredGLMM(M1)

##2c. Compute QI values for M1 (logistic model)##
######################

#This portion is a loop to find the quantities of interest
   #while varying the values of M and X

lbx<- -3 #this is the lower bound of the x-axis for the plot
   #recommend choosing a value somewhere around observed min(X)
ubx<- 3 #this is the upper bound of the x-axis for the plot
   #recommend choosing a value somewhere around observed max(X)
StepSize<-(ubx-lbx)/10

#values for the varying conditions#
values_M<-c(0,1)
values_X<-seq(from=lbx,to=ubx,by=StepSize)

#num. rows in result (the matrix where QI values will be stored)
ncond<-length(values_M)*length(values_X)

#create a matrix to store the resulting QI values
result<-matrix(ncol=3,nrow=ncond)
colnames(result)<-c("M","X","QI")
result[,1]<-rep(values_M,length.out=ncond)
result[,2]<-rep(values_X,each=length(values_M),length.out=ncond)

#loop to compute actual values of QI#
for(i in 1:(ncond)){
M<-result[i,1]
X<-result[i,2]
result[i,3]<-exp(b0+b1*X+b2*M+b3*X*M)/(1+exp(b0+b1*X+b2*M+b3*X*M))
}

FinalResult<-as.data.frame(result)
#display first 6 rows of the data frame
head(FinalResult)
```
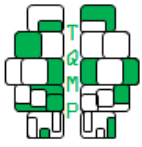
```
##2d. Plot QI values for M1 (logistic model)##
######################

#create data subset for each line#
subset1<-FinalResult[FinalResult$M==0,]
subset2<-FinalResult[FinalResult$M==1,]
color_values=c("blue","red")#one color per factor level
lty_values=c(2,3)#one line type per factor level

#Plot 2#
#curve for M=0, then for M=1#
bmp(file="Plot2.bmp")
plot(x=subset1$X,y=subset1$QI,type="l",ylim=c(0,1),col=color_values[1],
    lty=lty_values[1],
    main="Predicted Probabilities",
    ylab="Probability of Graduation",xlab="X (Studying Habits)")
lines(x=subset2$X,y=subset2$QI,type="l",col=color_values[2],
    lty=lty_values[2])
text(1,.5,"High Anxiety")
text(-1.5,.7,"Low Anxiety")
legend("topleft",
    legend=c("High Anxiety","Low Anxiety"),
    lty=lty_values,
    col=color_values)
dev.off()
```

R syntax to produce Figure 3: Random slopes plot for multilevel model

```
##3a. Simulate data##
######################

#set constants
set.seed(50)
b0<- .3
b1<- 5
L2sd<-2
L2slope<-2
L1sd<-2
N<-200
NumGroups<-10

#generate variables
X<-rnorm(N)
ID<-rep(seq(from=1,to=NumGroups,by=1),times=N/NumGroups)
U0j<-rep(rnorm(NumGroups,mean=0,sd=L2sd),times=N/NumGroups)
U1j<-rep(rnorm(NumGroups,mean=0,sd=L2slope),times=N/NumGroups)
eij<-rnorm(N,mean=0,sd=L1sd)

##generate Y (continuous)
Y_continuous<-b0+b1*X+U1j*X+U0j+eij

#create data frame to store simulated variables
mydata3<-data.frame(ID=ID,Y=Y_continuous,X=X)

##3b. Estimate model##
######################

M0<-lmer(Y~X+(X|ID),data=mydata3)
r.squaredGLMM(M0)

##3c. Compute QI for M0 (linear model)##
######################

lbx<- -3 #this is the lower bound of the x-axis for the plot
```

```
   #recommend choosing a value somewhere around observed min(X)
ubx<- 3 #this is the upper bound of the x-axis for the plot
   #recommend choosing a value somewhere around observed max(X)
StepSize<-(ubx-lbx)/10

#values for the varying conditions
values_X<-seq(from=lbx,to=ubx,by=StepSize)
#ranef() for u0j and u1j values
estimated_u0j<-ranef(M0)$ID[,1]
estimated_u1j<-ranef(M0)$ID[,2]

#num. rows in result
ncond<-length(values_X)*length(estimated_u0j)

#create a matrix to store the resulting QI values
result<-matrix(ncol=4,nrow=ncond)
colnames(result)<-c("X","u0j","u1j","QI")
result[,1]<-rep(values_X,length.out=ncond)
result[,2]<-rep(estimated_u0j,each=length(values_X),length.out=ncond)
result[,3]<-rep(estimated_u1j,each=length(values_X),length.out=ncond)

for(i in 1:(ncond)){
X<-result[i,1]
u0j<-result[i,2]
u1j<-result[i,3]
QI1<-b0+b1*X+X*u1j+u0j
result[i,4]<-QI1
}

FinalResult<-as.data.frame(result)

#display first 6 rows of the data frame
head(FinalResult)

##3d. Plot QI values for M0 (linear model)##
######################

#create subset for each line
subsets<-list()
for(i in 1:20) {
    subsets[[i]]<-
FinalResult[FinalResult$u0j==estimated_u0j[i]&FinalResult$u1j==estimated_u1j[i],]
}

########
#Plot 3#
#curve for each of the 12 groups (level 2 units)
bmp(file="Plot3.bmp")

plot(mydata3$X,mydata3$Y,pch=20,col=factor(mydata3$ID),
    ylab="Y (Achievement)",xlab="X (Studying Habits)")
for(i in 1:NumGroups){
lines(x=subsets[[i]]$X,y=subsets[[i]]$QI,type="l",col=i)
}

legend("topleft",
    legend=levels(factor(ID)),
    pch=20,
    col=factor(levels(factor(ID))),
    bg="white")

dev.off()
```
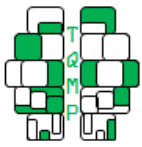
R syntax to produce Figure 4: Random slopes plot for multilevel logistic regression

```
##4a. Simulate data##
#######################

#set constants
set.seed(50)
b0<- 1
b1<- 1
L2sd<-1.5
L2slope<-1.5
N<-200
NumGroups<-10

#generate variables
X<-rnorm(N)
ID<-rep(seq(from=1,to=NumGroups,by=1),times=N/NumGroups)
U0j<-rep(rnorm(NumGroups,mean=0,sd=L2sd),times=N/NumGroups)
U1j<-rep(rnorm(NumGroups,mean=0,sd=L2slope),times=N/NumGroups)

#generate prob(Y)
ExpXB<-exp(b0+b1*X+U1j*X+U0j)
ProbY<-ExpXB/(1+ExpXB)

#generate Y
Y_binary<-rbinom(N,1,ProbY)

#create data frame to store simulated variables
mydata4<-data.frame(ID=ID,Y=Y_binary,X=X)

##4b. Estimate model##
#######################

M1<-glmer(Y~X+(X|ID),family=binomial(link="logit"),data=mydata4)
r.squaredGLMM(M1)

##4c. Compute QI for M1 (logistic model)##
#######################

lbx<- -3 #this is the lower bound of the x-axis for the plot
    #recommend choosing a value somewhere around observed min(X)
ubx<- 3 #this is the upper bound of the x-axis for the plot
    #recommend choosing a value somewhere around observed max(X)
StepSize<-(ubx-lbx)/10

#values for the varying conditions
values_X<-seq(from=lbx,to=ubx,by=StepSize)
#ranef() for u0j and u1j values
estimated_u0j<-ranef(M1)$ID[,1]
estimated_u1j<-ranef(M1)$ID[,2]

#num. rows in result
ncond<-length(values_X)*length(estimated_u0j)

#create a matrix to store the resulting QI values
result<-matrix(ncol=4,nrow=ncond)
colnames(result)<-c("X","u0j","u1j","QI")
result[,1]<-rep(values_X,length.out=ncond)
result[,2]<-rep(estimated_u0j,each=length(values_X),length.out=ncond)
result[,3]<-rep(estimated_u1j,each=length(values_X),length.out=ncond)

for(i in 1:(ncond)){
X<-result[i,1]
u0j<-result[i,2]
u1j<-result[i,3]
QI1<-exp(b0+b1*X+X*u1j+u0j)/(1+exp(b0+b1*X+X*u1j+u0j))
result[i,4]<-QI1
}
```

```
FinalResult<-as.data.frame(result)

#display first 6 rows of the data frame
head(FinalResult)

##4d. Plot QI values for M1 (logistic model)##
######################

#create subset for each line
subsets<-list()
for(i in 1:12) {
    subsets[[i]]<-
FinalResult[FinalResult$u0j==estimated_u0j[i]&FinalResult$u1j==estimated_u1j[i],]
}

########
#Plot 4#

#curve for each of the 20 groups (level 2 units)
bmp(file="Plot4.bmp")
plot(x=subsets[[1]]$X,y=subsets[[2]]$QI,type="l",ylim=c(0,1),
    col=factor(levels(factor(ID)))[1],
    main="Predicted Probabilities, Random Slopes (Schools)",
    ylab="Probability of Graduation",xlab="X (Studying Habits)")
for(i in 2:NumGroups){
lines(x=subsets[[i]]$X,y=subsets[[i]]$QI,type="l",
    col=factor(levels(factor(ID)))[i])
}
legend("topleft",
    legend=levels(factor(ID)),
    col=factor(levels(factor(ID))),
    lty=1,
    bg="white")

dev.off()
```

Citation