

Two-Part Models for Father–Child Relationship Variables: Presence in the Child’s Life and Quality of the Relationship Conditional on Some Presence

Kimberly L. Henry *Colorado State University, Colorado School of Public Health*

Thao P. Tran *Colorado State University*

Della V. Agbeke *Colorado State University*

Hyanghee Lee *Colorado State University*

Anne Williford *Colorado State University*

John J. Dziak *The Pennsylvania State University*

ABSTRACT Parent–child relationship variables are often measured using a two-part approach. For example, when assessing the warmth of the father–child relationship, a child is first asked if they have contact with their father; if so, the level of warmth they feel toward him is ascertained. In this setting, data on the warmth measure is missing for children without contact with their father, and such missing data can pose a significant methodological and substantive challenge when the variable is used as an outcome or antecedent variable in a model. In both cases, it is advantageous to use an analytic method that simultaneously models whether the child has contact with the father, and if they do, the degree to which the father–child relationship is characterized by warmth. This is particularly relevant when the two-part variable is measured over time, as contact status may change. We offer a pragmatic tutorial for using two-part variables in regression models, including a brief overview of growth modeling, an explanation of the techniques to handle two-part variables as predictors and outcomes in the context of growth modeling, examples with real data, and syntax in both R and Mplus for fitting all discussed models.

KEYWORDS: two-part models, two-part predictors, two-part growth models, father–child relationships, father–child contact

doi: 10.1086/714016

In family-based studies, researchers commonly assess variables that describe the parent–child relationship—for example, the quality of parenting, harsh discipline, or the level of attachment between a parent and child. These variables are of

common interest in social work practice, particularly in the child welfare system, where the focus is often to enhance parent–child relationships and parenting practices (Chaviano et al., 2018). As such, studies investigating these kinds of variables must use appropriate analytic techniques to capture reliable representations of the phenomena under study, especially in applied sciences such as social work, where the goal is to inform clinical practice and social policy.

In many cases, however, these sorts of variables are irrelevant—or are at least qualitatively different—if a child and parent do not have contact with one another. Therefore, the collection of this type of variable is often carried out using a two-step measurement approach. For example, when assessing the warmth of the relationship between a child and their father, a child may be first asked if they have contact with their father (e.g., “How often do you spend time with your father?”), and if some minimal threshold of contact is met (e.g., at least some contact), the child is then asked a series of questions to ascertain the level of warmth. As a result, a child without contact with their father will have a missing score by design for the items that measure warmth. Furthermore, this is an unusual kind of missing data in which the true value is not just simply unknown, but the meaning of the variable changes or becomes unclear (for a more extreme example, consider a job satisfaction question for an unemployed person). The presence of missing data in this case can pose a significant challenge when the variable is used as either an outcome of an antecedent variable (e.g., Does paternal depression affect the warmth of the relationship between child and father?) or as a predictor of some later outcome (e.g., Is more warmth in the father–child relationship associated with healthy child development?). Although one might be tempted to treat the covariate of interest (e.g., warmth) as missing during times when the child and father do not have contact, or to record the lowest possible value, it is advantageous to use an analytic method that simultaneously models whether the child has contact with the father, and if they do, the degree to which the father–child relationship is characterized by warmth. In the case of no contact, warmth is not simply missing and is certainly not missing at random, an assumption of most modern missing-data methods (Collins et al., 2001). We cannot assume that no contact between father and child is the same as the lowest possible score for warmth. If the father and child do not have contact, certain father–child relationship characteristics (e.g., warmth, supervision, harsh parenting) are not relevant or are at least qualitatively different than for dyads with contact. These issues become even more relevant when the two-part variable is measured over time, as contact status tends to change for a portion of the population. This is common when studying families over time, particularly in the child welfare system, where single-parent families have been historically overrepresented (Whittaker & Tracy, 1990) and contact with fathers is often limited, or at least assumed to be (Brewsaugh & Strozier, 2016).

Models that consider two-part variables are readily available but are not often used in applied research. In this paper, we offer a pragmatic tutorial for using two-part variables in regression models, including a brief overview of growth modeling, explanation of the techniques to handle two-part variables as predictors and outcomes in the context of growth modeling, examples with real data, and syntax in both R (R Core Team, 2014) and Mplus (Muthén & Muthén, 2017) for fitting all discussed models. Our goal is to provide social work researchers with an easy-to-use framework for capitalizing on this sort of data that is commonly encountered in studies of families.

An Example From the Rochester Intergenerational Study

In this paper, we provide an example derived from the Rochester Intergenerational Study (RIGS). We briefly summarize RIGS here (see Thornberry et al., 2018, for more detailed information). The genesis of RIGS is the Rochester Youth Development Study (RYDS), a birth cohort of 1,000 adolescents (referred to as Generation 2 [G2]; G2's primary caregiver is referred to as Generation 1 in the Rochester Studies), representative of the 7th- and 8th-grade public school population in Rochester, NY, in 1988. Adolescents who were at high risk for antisocial behavior were oversampled (by oversampling males and adolescents residing in high-crime areas of the city). The RYDS adolescents (G2s) were followed from 1988 to 2006 across 14 interviews. RIGS began in 1999 as the oldest biological child of G2 participants was identified (if born previously; $n = 370$ in Year 1); these children are referred to as Generation 3 (G3). New firstborn G3s were added to the sample in each subsequent year when G3 was 2 years of age. Annual interviews of G2s were completed each year after the family's entry into RIGS and continued until the child (G3) turned 18. G3s completed annual interviews beginning at age 8. Over the course of RIGS, data were collected from 539 parent–child (G2–G3) dyads. All procedures for RYDS and RIGS were overseen by the University at Albany Institutional Review Board. The examples used in this tutorial consider all G2 fathers and their G3 child (273 father–child dyads).

For this tutorial, we use the child's (G3) report of several variables at ages 14, 15, and 16. Our focus is on the relationship between the child (G3) and their father (G2). Most children in the study had continuous contact with their mother, but continual contact with fathers was less common. Details and descriptive statistics of the variables we consider in this tutorial are reported in Table 1. The two-part variable considered in this tutorial is a measure of the child's sense of warmth toward their father. At each interview, children were asked how often they see their father; response options included 0 (*never*), 1 (*almost never*), 2 (*sometimes*), and 3 (*often*). Across all children and years considered, children did not have contact with their father on 18% of the measurement occasions. Children who reported a score of 1 or higher for contact were administered an 11-item measure of warmth toward their

Table 1*Descriptive Statistics for Considered Variables*

Variable (Variable Name)	Mean/ Percentage	SD	Description
Child's ID variable (NEWID)	–	–	–
Child's sex (G3MALE)	49%	–	1 = male, 0 = female
Biological parents live in one household at G3 age 10 (INTACT)	30%	–	1 = G3 lives with both biological parents in one household; 0 = G3 does not live with both biological parents in one household
Child's contact with father (G3CON)	84%	–	Based on the question "How often do you see your father?" Response choices included 0 (<i>never</i>), 1 (<i>almost never</i>), 2 (<i>sometimes</i>), and 3 (<i>often</i>). G3CON = 0 if the child's response = 0 and G3CON = 1 if the child's response ≥ 1 . A score of 0 on G3CON was also assigned if the child's father was deceased.
Age 14	81%	–	
Age 15	80%	–	
Age 16	80%	–	
Child's warmth toward father (G3WARM)			An 11-item scale that assessed the child's feelings of warmth toward their father. The question stem stated, "How often would you say that . . ."; example items include "get along well with your father" and "you think your father is terrific." Response choices included 1 (<i>never</i>), 2 (<i>almost never</i>), 3 (<i>sometimes</i>), and 4 (<i>often</i>). The mean of the 11 items formed the scale. Cronbach's alpha was .86–.87 across ages.
Age 14	3.34	0.50	
Age 15	3.32	0.50	
Age 16	3.26	0.51	

Child's report of stressful life events (G3EVENT)			An 18-item inventory of whether a series of potentially stressful life events had occurred since the date of last interview (approximately 1 year). Example items included: "failed a course at school," "someone in your family died," and "your parents split up." Each item was scored 0 (<i>no</i>) or 1 (<i>yes</i>). The sum of the number of stressful events was calculated at each interview.
Age 14	4.42	2.82	
Age 15	4.46	2.82	
Age 16	4.61	2.64	
Child's report of substance use (G3SUBS)			1 = use alcohol or cannabis at least once per month; 0 = do not use alcohol or cannabis or use less frequently than once per month
Age 14	3%	–	
Age 15	6%	–	
Age 16	12%	–	
Child's depressive symptoms (G3DEP)			An 11-item scale that assessed depressive symptoms during the prior year. The question stem stated, "Since your last interview, how often did you . . ."; example items included "feel lonely" and "feel depressed or very sad." Response choices included 0 (<i>never</i>), 1 (<i>almost never</i>), 2 (<i>sometimes</i>), and 3 (<i>often</i>). The mean of the 11 items formed the scale. Cronbach's alpha was .89–.92.
Age 14	1.07	0.64	
Age 15	1.09	0.69	
Age 16	1.16	0.66	

Note. N = 273.

father (Hudson, 1996). Example items, scale anchors, and psychometric information are reported in Table 1. Relevant items were reverse coded, and the average of the 11 items was taken for each child at each study year to create a scale score of warmth toward father. Thus, at each interview year, we have a binary indicator of whether the child saw their father ($G3CON = 0$ if they never saw their father, and $G3CON = 1$ if they saw their father [even if only minimally]). In addition, conditional on some contact, we have a score for warmth toward father at each measurement occasion ($G3WARM$). This two-part variable is at the center of the demonstrations presented in this tutorial.

In order to consider a wide variety of models that a researcher may encounter, we present an assortment of types of predictors and outcomes. In the first part of the demonstration, we consider the two-part variable (contact/warmth) as an outcome. In this context, we use one time-invariant variable to predict the two-part outcome: whether the child ($G3$), and both the biological mother and father, lived in the same household at age 10 ($INTACT$). We also use one time-varying variable to predict the two-part outcome: child's report of stressful life events in the life of the family during the year preceding the interview ($G3EVENT$). Both variables are described in Table 1.

In the second part of the demonstration, we consider the two-part variable (contact/warmth) as a predictor of the child's depressive symptoms (Radloff, 1977), which were measured from age 14 to age 16 ($G3DEP$). The depressive-symptoms scale is described in Table 1.

A Brief Introduction to Growth Curve Analysis

The modeling of change over time is common in family research. Here, we provide a brief primer to set the stage for the forthcoming discussion of growth curve analysis with two-part variables. We focus on the analysis of longitudinal data as a standard two-level multilevel regression model in which repeated measures (Level 1 of the multilevel model: within persons) are nested in persons (Level 2 of the multilevel model: between persons; Laird & Ware, 1982). We first consider a growth model for a continuous measure. The linear growth model for a repeated measure y_{it} , where t denotes the measurement occasion (e.g., time or age of the child) and i denotes the individual, is written as

$$y_{it} = \pi_{0i} + \pi_{1i} \times time_{it} + e_{it} \quad (1.1)$$

$$\pi_{0i} = \beta_{00} + r_{0i} \quad (1.2)$$

$$\pi_{1i} = \beta_{10} + r_{1i} \quad (1.3)$$

$$e_{it} \sim N(0, \sigma_e^2) \quad (1.4)$$

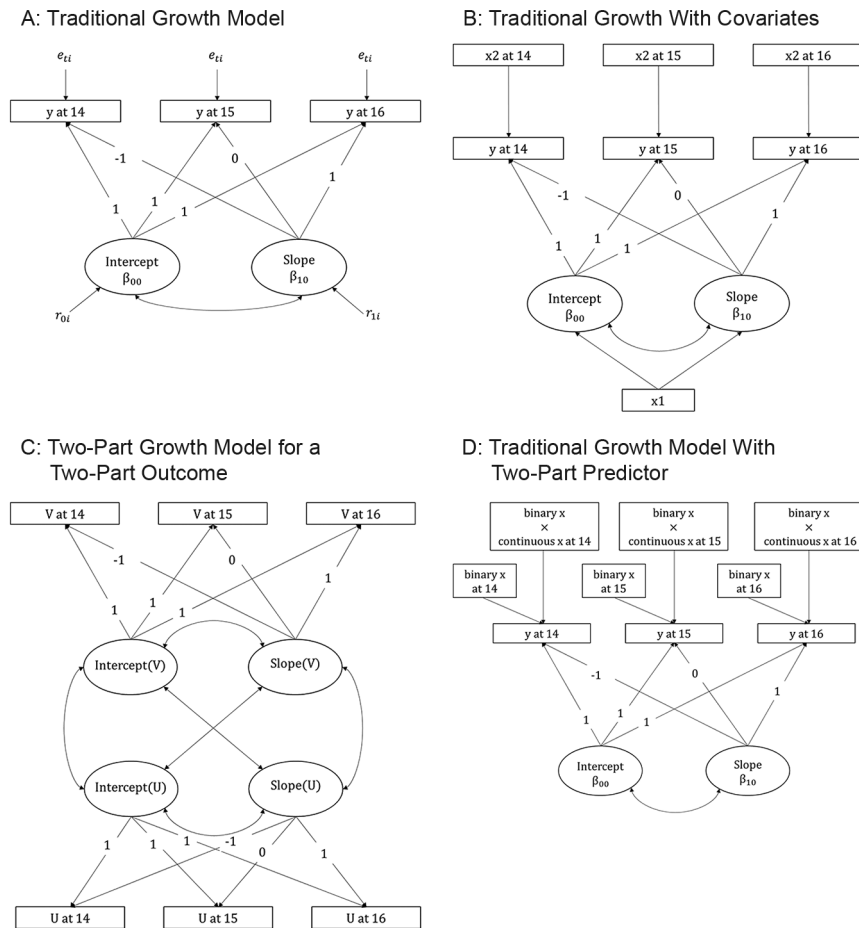
$$r_{0i}, r_{1i} \sim \text{MVN} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_0^2 & \\ & \sigma_1^2 \end{bmatrix}, \sigma_{01} \right). \quad (1.5)$$

Equation 1.1 represents the Level 1 (within-persons) model and indicates that variable y for individual i at time t is defined by an intercept (π_{0i}), a slope (π_{1i}), and a residual (e_{it}). The intercept and slope are subscripted with an i to denote that each individual in the sample has an intercept and a slope to describe their trajectory of change in y over time; these are commonly referred to as *growth parameters*. The intercept represents the individual's predicted score on y when time = 0, and the slope represents the individual's predicted change in y for a 1-unit increase in time. Residuals represent the difference between the observed y score and the predicted y score for each person at each timepoint.

Equations 1.2 and 1.3 represent the Level 2 (between-persons) models. Equation 1.2 denotes that the *random intercepts* (i.e., intercepts vary across individuals) can be described by a fixed mean (β_{00}), the average predicted score for y when time = 0, and a residual (r_{0i}) for each individual, which captures their deviation from β_{00} . Equation 1.3 denotes that the *random slopes* (i.e., slopes vary across individuals) can be described by a fixed effect of time (β_{10}), the average predicted rate of change in y , and a residual (r_{1i}) for each individual, which captures their deviation from β_{10} . Equation 1.4 delineates the Level 1 residuals (e_{it}), which are assumed to be *normally and conditionally independently distributed* (i.e., conditional on the individual random effects (σ_0^2, σ_1^2) with a mean of 0 and a common variance (σ_e^2). Last, the variances of the Level 2 random intercepts and slopes (σ_0^2, σ_1^2), shown in Equation 1.5, are assumed to be multivariate normally distributed and freely covary (σ_{01}), indicating that one's level at time = 0 may be related to their rate of change over time. We may also represent the growth model graphically. Figure 1A translates the equations just described to a graphical depiction of the model.

Consider an example in which change in children's depressive symptoms is modeled from ages 14 to 16. We begin with an *unconditional growth model*: a model that specifies change in depressive symptoms but does not include predictors of change in depressive symptoms over time. The syntax for this example is presented in the online Appendix (Script Set 1). We fit the model using the *nlme* package in R Version 3.1 (Pinheiro et al., 2020) and Mplus Version 8.2 (Muthén & Muthén, 2017). Throughout the text and tables of this manuscript, we present the results from R (and R packages). In some cases, the results differ slightly from the results produced by Mplus due to different estimation methods. We regressed depressive symptoms across time on a vector of ones (to define the intercept, the default model setup in both *nlme* and Mplus) and a time variable (G3AGE15) that is coded $-1, 0,$ and 1 corresponding to child age 14, 15, and 16, respectively. Thus, the intercept represents the child's depressive symptoms at age 15 (when time = 0), and the slope captures linear change in depressive symptoms for each year from age 14 to 16. We defined

Figure 1. Graphical Depiction of Described Models



Note. Residuals are depicted in Figure 1A only for parsimony.

the center timepoint (age 15) as the intercept in this example, but any age may be selected. The fixed-effect estimate of the intercept (β_{00}) is 1.104, denoting the average predicted depressive symptoms at age 15. There is substantial variability in the intercept estimates—variance(r_{0i})—across children ($\sigma_0^2 = .290$ for variance, $\sigma_0 = .539$ for standard deviation), indicating that children differ in their level of depressive symptoms at age 15. Most R packages print out the standard deviation and correlation for the random effects by default, whereas Mplus prints out the variance and covariance for the random effects by default. For most functions in R, the variance of the random effects can be printed via the *VarCorr* function (see the example at the bottom of the R syntax for Script Set 1 in the online Appendix). The fixed-effect estimate of the linear slope (β_{10}) is .036, denoting the average predicted rate of

change in depressive symptoms each year from age 14 to age 16. This small relative value suggests a minor yearly increase in depressive symptoms over this period. The variability in the linear slope estimates across children ($\sigma_2^2 = .020$, $\sigma_1 = .142$) captures differences in the rate of change of depressive symptoms across children (i.e., some children may increase depressive symptoms rapidly, some may not change, and some may decrease depressive symptoms over the 3-year period). The correlation of the random intercept and slope (standardized σ_{01}) is .179, indicating a small positive correlation. Specifically, greater depressive symptoms at age 15 are associated with a greater increase in depressive symptoms across time. The variance of the Level 1 time-specific residuals [$v(e_{it})$]—which were constrained to be equivalent over time—is .133. This is indicative of substantial within-person variability in depressive symptoms not accounted for by time.

An unconditional growth model can be extended to consider predictors of the outcome modeled. Consider a model that adds one time-invariant predictor (G3 sex; the variable is called G3MALE) and a time-varying predictor (G3 stressful life events measured at ages 14, 15, and 16; the variable is called G3EVENT). To evaluate the effect of a time-invariant variable (e.g., G3MALE), we added the variable as a predictor of the intercept (e.g., depressive symptoms at age 15) and slope for age (represented by an interaction between G3MALE and G3AGE15). To evaluate the effect of a time-varying variable (e.g., G3EVENT), we simply added the predictor as a main effect, allowing stressful life events at time t to predict depressive symptoms at time t . Figure 1B provides a graphical depiction of the model, where x_1 represents a time-invariant covariate (e.g., G3MALE) and x_2 represents a time-varying covariate (e.g., G3EVENT). The results of the model are presented in Table 2. We estimated that at age 15, holding stressful life events constant, G3 males had an average depressive-symptoms score 0.367 units lower than the average of G3 females; however, the change in depressive symptoms across time was quite similar for G3 males and G3 females (the estimate for the sex by age interaction is $-.039$). We also estimated that, holding G3 sex and G3 age constant, each additional stressful life event increased the predicted depressive-symptoms score by 0.048 units.

Growth Model for a Binary Outcome

This framework for modeling growth can be easily extended to consider categorical outcomes, for example, to capture level and change in a binary y variable. For a binary outcome, a growth model can be fit using either a logistic or probit regression model (Hedeker & Gibbons, 2006). We focus here on the logistic model. In the case of a binary y , it is assumed that y is a discretized form of an underlying continuous latent outcome (y^* ; the asterisk denotes that y is latent). As an example, consider the child's regular use of substances (0 = no regular use of substances, 1 = use alcohol or cannabis at least once per month) at ages 14, 15, and 16. We can imagine an underlying continuum of regular substance use (G3SUBS*). When a child's value of

Table 2*Results of a Conditional Traditional Growth Model for Continuous and Binary Outcomes*

Parameter	Continuous Outcome: Depressive Symptoms		Binary Outcome: Substance Use	
	Estimate	SE	Estimate	SE
Fixed effects				
Intercept	1.071	0.055	-6.846	1.140
Age (centered at age 15)	0.054	0.025	0.895	0.361
Male	-0.367	0.062	-0.488	0.640
Age (centered at age 15) × Male	-0.039	0.036	0.709	0.539
Stressful life events	0.048	0.008	0.393	0.100
Random effects				
sd(intercept)	0.461	-	-	-
sd(age [centered at age 15])	0.117	-	2.662	-
corr(intercept, age)	0.145	-	-	-
sd(residual)	0.372	-	-	-

Note. Sd = standard deviation; corr = correlation. Although not reported here, the output (shown in the online Appendix) of the models includes *p*-values for testing statistical significance of the fixed effects.

G3SUBS* exceeds a certain level, also referred to as a *threshold*, their observed score of regular substance use is 1, and 0 otherwise. Similar to a continuous outcome, each individual's trajectory of change in y^* can be described by an intercept (the predicted log odds of y^* when time = 0) and a slope (the predicted change in the log odds of regular substance use for a 1-unit increase in time). We expect variability in the log odds of the outcome at time 0 (i.e., variability in the intercepts across children) and variability in the rate of change over time (i.e., variability in the slopes across time). Moreover, covariance in the intercepts and slopes across children can be estimated. In these ways, growth models for continuous and binary outcomes are similar.

It is important to note one key difference between a growth model for a continuous outcome and a binary outcome. For a continuous outcome, the mean growth trajectory for the population is simply defined as the linear function of the intercept (β_{00}) and the slope (β_{10}). However, whereas this is also the case for the latent response variable when y is categorical (y^*), it is not the case for the observed categorical y . Hedeker and colleagues (2018) explained that the estimates derived from generalized linear mixed models (GLMM, including growth models for categorical outcomes) have what Neuhaus and colleagues (1991) referred to as a *subject-specific interpretation*. That is, the subject-specific regression estimates from a resultant GLMM represent

the expected change in the outcome for a 1-unit increase in the predictor, holding constant other predictors and all random effects in the model. The *population-average regression estimates* represent the expected change in the outcome for a 1-unit increase in the predictor, holding constant other predictors but not random effects. Thus, the subject-specific parameter estimates from a growth model of a categorical outcome differ in interpretation from the corresponding population-average estimates. Although both the subject-specific and population parameter estimates are useful, the former are applicable for drawing inferences for individuals, and the latter are applicable for drawing population-based inferences (Zeger et al., 1988). Population estimates are needed to calculate population-average predicted probabilities of the categorical outcome based on covariate levels. That is, to calculate the response probability for the population at a particular point in time (e.g., probability of regular substance use at age 15), it is necessary to calculate and then average the individual response probabilities for all individuals in the population. For binary growth models fit as a GLMM, obtaining the population-average estimates is a postestimation (i.e., after fitting the GLMM and obtaining the subject-specific estimates) endeavor. Hedeker et al. (2018) published a technique for carrying this out in SAS. Their recommendations are also implemented in the *GLMMadaptive* package (Rizopoulos, 2020) for R, and we demonstrate its use in the following example.

Consider an example to model the log odds of the child using substances at least once per month at each measurement occasion (G3 age 14, 15, 16). The syntax for this example is presented in Script Set 2 in the Appendix (online). Time is specified using the same technique as the model for depressive symptoms (centered at age 15). A logit link function and a maximum likelihood estimator are used. The fixed-effect estimate of the intercept was -5.152 , which represents the predicted log odds of regular substance use at age 15 (a subject-specific effect) for a participant with a random effect of 0. There was some variability in the intercept estimates across children, indicating that children differed in the log odds of regular substance use at age 15. The fixed-effect estimate of the slope was 1.176 , denoting the predicted rate of change in the log odds of substance use from age 14 to age 16 (also a subject-specific effect). This presents the expected change in the log odds of regular substance use for each year the child grows older; the positive value indicates that the log odds of regular substance use increase over time. The variability in the slope (σ_1^2) was very small, and in fact, when fit using the *GLMMadaptive* package in R, the slope variability cannot be estimated. This is often the case with growth models for binary outcomes (Long et al., 2009). Therefore, the random effect for the slope is constrained to zero in this example.

The *GLMMadaptive* package allows for estimation of the population-average coefficients (referred to as *marginal coefficients* in the output) as well as the subject-specific effects just presented. The population-average intercept was estimated to be -2.708 , and the population average slope was estimated to be $.719$. These log

odds estimates can then be used to calculate the predicted probability of regular substance use at any given age. For example, at age 15, the estimated probability of regular substance use is calculated as follows: $\exp(-2.708) \div (1 + \exp(-2.708)) = .06$. At age 16 the estimated probability of regular substance use is calculated as follows: $\exp(-2.708 + (.719 \times 1)) \div (1 + \exp(-2.708 + (.719 \times 1))) = .12$.

Following the same protocol for the continuous growth model example, the binary growth model can easily be extended to include predictors. We added G3 sex as a predictor of both the intercept and slope of the binary growth model and added G3 stressful life events as a time-varying predictor. Results are presented in Table 2. Holding stressful life events constant, we did not find strong evidence that the intercept (log odds of regular substance use at age 15) and the slope (rate of change in the log odds of substance use across time) differed between males and females. (Although the estimates are not close to 0 for the effect of G3MALE on the intercept and slope, the standard errors are very large with respect to the estimates, and therefore, there is a great deal of uncertainty about the effect of sex in this instance.) However, we did find evidence that G3 stressful life events were associated with greater log odds of regular substance use (holding sex and age constant). The marginal coefficient estimate for stressful life events was .259; by exponentiating this value— $\exp(.259) = 1.29$ —we estimated that in the population, a 1-unit increase in stressful life events was associated with about a 30% increase in the odds of regular substance use during adolescence.

With a simple explanation of fitting growth models for continuous and binary outcomes complete, we are now ready to turn our attention to consideration of two-part variables and our example of father–child contact and warmth in the father–child relationship conditional on some contact.

When a Two-Part Variable is an Outcome

Olsen and Schafer (2001) introduced a two-part growth model for a longitudinal semicontinuous variable (y_{it}), defining a *semicontinuous variable* as one that arises from two processes: one that determines if the response is zero, and a second that determines the level of y_{it} if nonzero. This type of model has been commonly used to model substance use by individuals over time (Brown et al., 2005). For example, if y_{it} was a measure of alcohol problems, a score of 0 would represent people who reported no alcohol use. Among those who do drink, we would observe their score for frequency/severity of alcohol problems. Olsen and Schafer (2001) demonstrated that a variable of this nature can be modeled using two correlated latent growth curves: one to model the likelihood of being a drinker versus a nondrinker (i.e., the binary part of a two-part model as it contrasts nondrinkers [0] to drinkers [1]) and a second to model alcohol problems (i.e., the continuous part of the two-part model as it measures the frequency/severity of alcohol problems) among those who reported alcohol use. To specify this model, a growth curve is fit to both parts (i.e.,

one growth model for the binary part and one growth model for the continuous part), and the random growth parameters (e.g., the intercepts and slopes) are allowed to covary. Xu and colleagues (2014) presented a more recent examination of two-part growth models in a multivariate context.

In extending our motivating example, the two-part model described by Olsen and Schafer may be applied (with some change in interpretation) to situations in which the two parts represent whether or not the child has contact with their father at each observed age (which may be fitted as a binary growth model), and conditional on having some contact, the level of warmth that the child feels toward their father at each observed age (which may be fitted as a continuous growth model). In the RIGS, children were first asked how often they had contact with their father during the past year, and only those with some contact were asked about the warmth they felt toward their father. Thus, the outcome variable in this scenario is broken into two parts, U and V for the binary and continuous parts, respectively:

$$U_{ii} = \begin{cases} 1 & \text{if contact with father} = \text{yes} \\ 0 & \text{if contact with father} = \text{no} \end{cases}$$

and

$$V_{ii} = \begin{cases} g(\text{G3 warmth}) & \text{if } U_{ii} = 1 \\ \text{irrelevant} & \text{if } U_{ii} = 0 \end{cases}$$

For V_{ii} , the continuous part of the two-part outcome, g , is a monotone increasing function that allows V_{ii} to be normally distributed—that is, approximately Gaussian (Olsen & Schafer, 2001). For skewed variables, meeting this assumption may necessitate a transformation (i.e., log, exponential, square root). Because two-part models are often used to accommodate outcomes that are skewed to the right, a log transformation is commonly applied, but any transformation that normalizes a non-normal distribution can be accommodated. Figure 1C provides a graphical depiction of an unconditional two-part regression model.

In the RIGS example, at each of the three observation periods, we have a binary indicator of whether the child had some contact with their father (G3CON = 1 for at least some contact, G3CON = 0 for no contact); for those with some contact, we have a measure of the warmth that they felt toward their father (G3WARM). As such, G3WARM is irrelevant (or at least qualitatively different) and was not assessed during measurement occasions when G3CON = 0. In the distribution of G3WARM, the variable is skewed to the left; that is, there was a tendency toward higher warmth scores. Therefore, we exponentiated the G3WARM scores (G3WRM2PRT), which moved the distribution of the continuous measure of our two-part outcome toward normality ($M = 30.4$, $SD = 12.7$, skew = .2).

Using the longitudinal variables G3CON and G3WRM2PRT, two correlated growth curve models (i.e., correlated random-effects models) may be fit to the data—one for the logit probability that G3CON = 1 (i.e., $U_{ii} = 1$) and one for the mean conditional

response of G3WRM2PRT [i.e., $E(V_{it}|U_{it} = 1)$]. Each is defined by an intercept (i.e., the predicted score for the modeled variable when age equals 15 [the centering point for G3AGE15]) and a slope (i.e., the predicted rate of change for the modeled variable for each year the child grows older), akin to the unconditional continuous and binary growth models described in the introduction. The intercept and slope within each part, and the intercepts and slopes across parts, are allowed to freely covary. The syntax for the unconditional two-part growth model (i.e., a model without predictors of the growth parameters) is presented in the online Appendix (Script Set 3) for *GLMMadaptive* in R and in Mplus.

For the binary part of the model, the intercept was 3.921. Note that when fitting a two-part growth model, *GLMMadaptive* models the log odds of a 0 score (e.g., no contact) rather than a score of 1 (e.g., contact); therefore, it is necessary to take the negative of the printed value for an interpretation that is consistent with the traditional binary growth model fit earlier. This value represents the average predicted log odds that a child will have some contact with their father at age 15 for a participant with a random effect of zero. Note that this is a subject-specific effect. The slope for time (centered at age 15) was $-.224$ (taking the negative of the printed value is also necessary for the slope coefficients), indicating that the log odds of contact tend to decline over time, though the standard error for this effect is large. For the continuous part of the model, the intercept was 27.975, providing the predicted level of warmth at age 15 conditional on having some contact with the father. Recall that to arrive at a normal distribution for our warmth score, we exponentiated the original warmth score. Thus, to return to the original scale, we took the natural log [e.g., $\ln(27.975) = 3.3$]. The slope for time was -0.983 , indicating that the warmth toward father tended to decline over time. Specifically, each 1-unit increase in age was associated with a 0.983-unit decrease in G3WRM2PRT. In our example, random effects for the intercept and slope were estimated for the continuous growth model, allowing for variation across children in warmth at age 15 as well as the rate of change in warmth across time. A random effect for the intercept of the binary growth model was included (i.e., allowing for variation in the log odds of contact at age 15 across children); however, there was not substantial variability in the binary slope and thus its variance was constrained to zero. Covariances of the continuous random intercept and slope, and the binary intercept, were all modeled. The covariation of the random intercept of the binary part of the growth model with the random intercept and slope of the continuous part of the growth model allowed the presence or absence of father contact across time to provide information about the trajectory of warmth over the observation period. In this way, children who had no contact across the observation period contributed little to the estimation of the growth model for warmth (Olsen & Schafer, 2001).

For a two-part growth model, *GLMMadaptive* will also provide marginal (i.e., population-averaged) coefficients. These coefficients average over random effects,

and they also average between participants who have and do not have observed values on the continuous part of the model. Unfortunately, the marginal coefficients must not be used in our example because they rely on the assumption that $V_{ti} = 0$ if $U_{ti} = 0$. This assumption is reasonable for substance use variables; for example, it is reasonable to assume that cigarettes per week for nonsmokers is zero. However, as discussed earlier, it does not apply well to the current situation. It is not reasonable to assume that all children with no contact with their father should have zeroes on the warmth scale. However, it is not reasonable to ignore them either, as these cases are not missing at random. Thus, in this situation there is no readily interpretable marginal model for the entire sample because it is not reasonable to average well-defined and poorly defined variables together.

Predictors of the growth parameters for each part are easily incorporated by regressing the growth parameters on a time-invariant covariate and regressing the individual scores at each measurement occasion on a time-varying covariate. We augmented our example two-part model in this way, and the results of the model with a time-invariant predictor (i.e., whether the child lived with both biological parents in the same household at age 10 [INTACT]) and a time-varying predictor (i.e., stressful life events as reported by the child [G3EVENT]) are presented in Table 3. First, consider the results for the continuous growth model for level of warmth toward father. Holding stressful life events constant, and conditional on some contact, our model predicted that children living in a household with both the biological mother and father at age 10 would, on average, have a G3WRM2PRT score [i.e., $\exp(\text{G3WARM})$] 4.740 units higher at age 15 (the centering point for time) than children who did not live with both biological parents at age 10. This effect does not show evidence of systematically changing over time (i.e., the interaction with time is not significantly different than 0). Holding constant the child's age and living situation at age 10 (i.e., INTACT), each 1-unit increase in stressful life events was associated with a 0.544-unit decrease in G3WRM2PRT—that is, $\exp(\text{G3WARM})$. Next, consider the results for the binary growth model. Living situation appears to be the only variable reliably related to contact. At age 15 (the centering point for time), the log odds of contact were 3.884 units higher if the child lived with both biological parents at age 10 (a subject-specific effect), and this effect appears to be constant across time (i.e., the interaction with time is near 0).

When a Two-Part Variable is a Predictor

Dziak and Henry (2017) outlined a method for examination of a two-part variable as a predictor (rather than as an outcome). In this setting, some variable (y_{ti}) is measured over time, and the desire is to determine if a two-part variable ($x_{\text{binary}_{ti}}$ and $x_{\text{numeric}_{ti}}$), also measured over time, is predictive of y_{ti} . For example, consider the unconditional growth model we presented earlier for children's depressive symptoms (G3DEP) from age 14 to age 16. We may be interested in determining if contact

Table 3

Results of a Conditional Growth Model for a Two-Part Outcome (Contact With Father and Warmth Toward Father)

Parameter	Estimate	SE
Fixed effects for continuous growth model (child's warmth toward father) ^a		
Intercept	29.114	1.281
Age (centered at age 15)	-0.882	0.599
Intact family at age 10	4.740	1.625
Age (centered at age 15) × Intact	-0.173	0.788
Stressful life events	-0.544	0.163
Fixed effects for binary growth model (contact with father) ^b		
Intercept	2.579	0.587
Age (centered at age 15)	-0.241	0.192
Intact family at age 10	3.884	0.986
Age (centered at age 15) × Intact	0.028	0.633
Stressful life events	0.079	0.075
Random effects		
sd(continuous intercept)	10.921	-
sd(continuous slope for age [centered at age 15])	0.867	-
sd(binary intercept)	3.808	-
corr(continuous intercept, continuous slope for age)	-0.321	-
corr(continuous intercept, binary intercept)	0.679	-
corr(binary intercept, continuous slope for age)	-0.311	-
sd(residual)	6.900	-

Notes. Sd = standard deviation; corr = correlation. Although not reported here, the output (shown in the online Appendix) of the models includes *p*-values for testing statistical significance of the fixed effects.

^aEstimates are in the metric of exp(warmth).

^bEstimates present the log odds of contact equal to 1 (i.e., at least some contact with father).

with father (G3CON, the binary part of the two-part variable $x_{\text{binary},i}$) is predictive of depressive symptoms, and conditional on having some contact, if the level of warmth that the child reports toward their father (G3WARM, the numeric or continuous part of the two-part variable $x_{\text{numeric},i}$) is predictive of the child's depressive symptoms.

In this setting in which the child's warmth (G3WARM) is an *exogenous variable* (i.e., a variable not influenced by other variables in the model), missingness of child's warmth during periods of time when they do not have contact with the father poses a major challenge. Without proper handling, all measurement occasions when the child does not have contact with the father (G3CON = 0) will be deleted in a listwise fashion. This may have serious ramifications for power and the ability to

obtain unbiased estimates. Although an analyst may be inclined to impute the missing values—for example, assigning the lowest score for warmth or using multiple imputation to model what the missing value might have been had the child been in contact with the father—Dziak and Henry (2017) described the reasons why these approaches are undesirable. The former assumes that not having contact with one’s father is the same as having very low feelings of warmth toward him, which is clearly problematic. The latter assumes that the data are missing at random (an unlikely and untestable assumption), which would lead to biased estimates. Moreover, these methods also miss the opportunity to examine important substantive questions about the potential differential effects of the two parts. For example, a two-part approach to modeling the predictor in the current example would allow for the determination of whether both the presence of contact with the father—and conditional on some contact, the level of warmth that the child has toward their father—are associated with the child’s depressive symptoms. Such findings could have important implications for informing the delivery of targeted supports and services for children and adolescents, particularly for youth in the child welfare system, whose risk for mental health challenges is heightened (McNeil et al., 2020).

Dziak and Henry outlined a simple approach for modeling this type of two-part predictor in a longitudinal model. The approach involves a simple recoding of the variables. First, one centers the continuous part of the variable ($G3WARM_{it}$) at the mean ($G3WARM_{ci} = G3WARM_{it} - \text{mean}[G3WARM_{it}]$) and then assigns all cases in which the binary part of the variable ($G3CON_{it}$) equals 0 to have a score of 0 on the centered version of the new continuous variable (e.g., if $G3CON = 0$ then $G3WARMX_{it} = 0$; if $G3CON = 1$ then $G3WARMX_{it} = G3WARM_{ci}$). Then, when the growth model is fit, the outcome ($G3DEP_{it}$) is regressed on the relevant time metric (e.g., $G3AGE$ centered at age 15), $G3CON_{it}$, and the product term of $G3CON_{it}$ and $G3WARMX_{it}$.

In this way, the regression coefficient for $G3CON_{it}$ represented the expected difference in $G3DEP_{it}$ for cases in which $G3CON_{it} = 1$ (e.g., contact with father) and $G3WARM_{ci} = 0$ (e.g., average warmth in the sample for the currently described example) and cases in which $G3CON_{it} = 0$ (e.g., no contact with father). In other words, we arrived at a comparison of father–child dyads in contact who had a relationship characterized by average warmth to father–child dyads with no contact. It is important to note that any other relevant comparison (i.e., other than when the continuous part of the variable is at the mean) is possible. This simply requires centering $G3WARM_{ci}$ at the desired comparison point. For example, one could subtract the score that is one standard deviation below the mean if it is desired to compare cases in which $G3CON_{it} = 1$ (e.g., contact with father) and warmth is one standard deviation below the mean with cases in which $G3CON_{it} = 0$ (e.g., no contact with father). The regression coefficient for $G3WARMX_{it}$ represented the expected difference in $G3DEP_{it}$ for a 1-unit increase in $G3WARM_{ci}$ when $G3CON_{it} = 1$ (e.g., child is in

contact with father). As usual, the regression coefficient associated with both parts of the variable is adjusted for time and any other covariates in the model. Figure 1D provides a graphical depiction of the growth model with a two-part predictor.

Table 4 presents results of the model just described, in which contact with father and (conditional on some contact) the child's perception of the warmth of the relationship were treated as time-varying predictors of the child's depressive symptoms from age 14 to age 16. Script Set 4 of the online Appendix presents the syntax needed to fit this model in R (the *nlme* package) and Mplus. In the first model (presented in Table 4), G3WARM is centered at a score of 2—a very low score for warmth in the sample. Here, we find that when warmth is at this low level, having contact with the father was associated with a child's depressive-symptoms score that was 0.421 units higher than for children with no father contact. In addition, conditional on having some contact, each 1-unit increase in warmth was associated with a 0.282-unit decrease in depressive symptoms.

If we instead center warmth at a score of 4 (see the second example under Script Set 4 in the online Appendix)—the highest possible score for warmth—we find that when warmth is very high, having contact with father is associated with a child depressive-symptoms score 0.143-units lower than would be expected if the child had no contact with their father. The effect of warmth, conditional on some contact, remains the same as in the previous example because centering the

Table 4

Results of a Conditional Growth Model of G3 Depressive Symptoms for a Two-Part Predictor (Contact With Father and Warmth Toward Father)

Parameter	Estimate	SE
Fixed effects		
Intercept	1.062	0.058
Age (centered at age 15)	0.030	0.018
Contact with father	0.421	0.081
Warmth toward father (centered at low score for warmth)	-0.282	0.050
Random effects		
sd(intercept)	0.500	—
sd(age [centered at age 15])	0.138	—
corr(intercept, age)	0.133	—
sd(residual)	0.366	—

Note. Sd = standard deviation; corr = correlation. Although not reported here, the output (shown in the online Appendix) of the models includes *p*-values for testing statistical significance of the fixed effects.

continuous part of the two-part predictor only shifts the interpretation of the effect of the binary part of the two-part predictor.

In summary, this model suggests that when contrasting depressive symptoms as a function of father contact, the level of warmth of the relationship is important. When warmth is low, contact with the father is associated with worse depressive symptoms for the child than no contact at all, but when warmth is high, contact with the father is associated with fewer depressive symptoms for the child. Moreover, when children have contact with their father, we observed a negative relationship between warmth and depressive symptoms: Greater warmth toward the father was associated with fewer depressive symptoms. Again, these distinctions are important for understanding effective prevention and intervention practices when serving youth who have no to little or intermittent contact with their fathers.

Discussion

In this paper, we presented a solution to a common measurement challenge faced by family researchers: the specification and consideration of two-part variables in which assessment of certain parenting or relationship variables depends on the level of contact between family members. For example, in considering the expression of warmth between father and child, care must be taken to account for periods when the father and child do not have contact. Beyond the warmth example presented in this tutorial, many other types of variables common in family research suffer from the same situation. For example, characterization of parental supervision, discipline tactics, or harsh parenting during times when a parent and child do not have contact is indeterminate. In these cases, the characterizations of parenting should not simply be considered missing and handled via conventional methods (e.g., listwise deletion, multiple imputation) because the scores are clearly not missing at random, an assumption of these methods (Rubin, 1987). Nor should researchers be forced to plug in a value (e.g., the worst possible score for the parenting characteristic). Rather, we argue that during times of no parent–child contact, variables of this nature should not be treated as missing data in the ordinary sense but rather as a two-part variable that can be represented as a pair of interrelated variables. We presented techniques and syntax to account for these sorts of two-part variables as both predictors and as outcomes in longitudinal research. The illustrative examples that we offered provide context to the types of questions that can be answered in family-based research and highlight the flexibility and predictive performance of models to accommodate two-part variables.

The methods discussed in this paper not only offer a statistical solution to the challenge posed by two-part variables in family research but also allow for more nuanced assessment of the questions at hand. In the case of two-part outcomes, dual-trajectory longitudinal growth models can simultaneously model change in contact over time, and conditional on having some contact, change in the relationship

variable of interest (e.g., expression of warmth between father and child) over time. Once specified, one can study how covariates may differentially impact growth in each part of the model.

In the case of two-part predictors, the two-part representation of a predictor can offer unique insight over an analog approach that considers a covariate as exclusively unidimensional (Dziak & Henry, 2017). A two-part specification of a predictor allows one to determine if the absence of contact is different depending on the level of the continuous part of the relationship construct of interest. In our example, we demonstrated that the effect of having no father contact on a child's depressive symptoms depends on whether the comparison is with father-child dyads who are in contact but have low warmth or father-child dyads who are in contact and have high warmth. The beneficial effect of contact was only observed when comparing no-contact dyads to in-contact dyads with high warmth. This is a nuanced finding that would be lost if father-child dyads without contact were excluded from the analysis or if contact and warmth were not studied in tandem using a two-part approach. It should be noted that the method described by Dziak and Henry (2017) is only appropriate when the missing data are meaningless (e.g., depressive symptoms for deceased individuals) or qualitatively different in meaning (e.g., our warmth example for absent vs. present fathers), not when it is simply unknown. If the missing values are meaningful but unobserved, then the method described by Dziak and Henry is inappropriate (Greenland & Finkle, 1995; Jones, 1996) and full-information maximum likelihood or multiple imputation is the appropriate approach for handling the missing data.

The techniques illustrated in this paper for a two-part outcome and a two-part predictor allow researchers to use all available data and avoid the need to exclude cases or timepoints where data are missing because the father-child pair do not have contact, or to artificially impute a score (i.e., assume a father-child dyad without contact would have the worst possible score for warmth). Thus, the approaches we suggest in this paper maximize power, minimize bias, and maintain the generalizability of findings to the identified and sampled population.

In sum, family researchers who use the types of two-part variables described here may more thoroughly answer their questions related to these variables by adopting the approaches discussed in this paper. The real data examples, and syntax for specifying these models in both Mplus and R, provide researchers with a clear roadmap for applying these methods in their own work.

Author Notes

Kimberly L. Henry, PhD, is a professor in the Colorado State University Psychology Department and at the Colorado School of Public Health.

Thao P. Tran, MS, is a doctoral student in the Colorado State University Psychology Department.

Della V. Agbeke, MS, is a doctoral student in the Colorado State University Psychology Department.

Hyanghee Lee, PhD, is a postdoctoral fellow in the Colorado State University Psychology Department.

Anne Williford, PhD, is an associate professor at the Colorado State University School of Social Work.

John J. Dziak, PhD, is a research associate professor at the Edna Bennett Pierce Prevention Research Center at The Pennsylvania State University.

Correspondence concerning this article should be directed to Kimberly L. Henry, Department of Psychology, Colorado State University, Fort Collins, CO 80523-1876, or via e-mail to kim.henry@colostate.edu.

Acknowledgments

Support for the Rochester Intergenerational Study is provided by the National Institute on Drug Abuse (R01DA020195). We thank the participants of the Rochester Intergenerational Study for their support, and we thank Drs. Adrienne Freeman-Gallant and Becky Chu for their assistance in collecting and curating the data for the project. The research in this paper was also supported by award P50 DA039838 from the National Institute on Drug Abuse.

References

- Brewsough, K., & Strozier, A. (2016). Fathers in child welfare: What do social work textbooks teach our students? *Children and Youth Services Review, 60*, 34–41. <https://doi.org/10.1016/j.childyouth.2015.11.015>
- Brown, E. C., Catalano, R. F., Fleming, C. B., Haggerty, K. P., & Abbott, R. D. (2005). Adolescent substance use outcomes in the Raising Healthy Children project: A two-part latent growth curve analysis. *Journal of Consulting and Clinical Psychology, 73*(4), 699–710. <https://doi.org/10.1037/0022-006X.73.4.699>
- Chaviano, C. L., McWey, L. M., Lettenberger-Klein, C. G., Claridge, A. M., Wojciak, A. S., & Pettigrew, H. V. (2018). Promoting change among parents involved in the child welfare system: Parents' reflections on their motivations to change parenting behaviors. *Journal of Social Work, 18*(4), 394–409. <https://doi.org/10.1177/1468017316654340>
- Collins, L. M., Schafer, J. L., & Kam, C. M. (2001). A comparison of inclusive and restrictive strategies in modern missing data procedures. *Psychological Methods, 6*(3), 330–351. <https://doi.org/10.1037/1082-989x.6.4.330>
- Dziak, J. J., & Henry, K. L. (2017). Two-part predictors in regression models. *Multivariate Behavioral Research, 52*(5), 551–561. <https://doi.org/10.1080/00273171.2017.1333404>
- Greenland, S., & Finkle, W. D. (1995). A critical look at methods for handling missing covariates in epidemiologic regression analyses. *American Journal of Epidemiology, 142*(12), 1255–1264. <https://doi.org/10.1093/oxfordjournals.aje.a117592>
- Hedeker, D., du Toit, S. H. C., Demirtas, H., & Gibbons, R. D. (2018). A note on marginalization of regression parameters from mixed models of binary outcomes. *Biometrics, 74*(1), 354–361. <https://doi.org/10.1111/biom.12707>
- Hedeker, D., & Gibbons, R. D. (2006). *Longitudinal data analysis*. John Wiley & Sons.
- Hudson, W. H. (1996). *WALMYR assessment scales scoring manual*. WALMYR Publishing Company.
- Jones, M. P. (1996). Indicator and stratification methods for missing explanatory variables in multiple linear regression. *Journal of the American Statistical Association, 91*(433), 222–230. <https://doi.org/10.1080/01621459.1996.10476680>

- Laird, N. M., & Ware, J. H. (1982). Random-effects models for longitudinal data. *Biometrics*, 38(4), 963–974. <https://doi.org/10.2307/2529876>
- Long, J. D., Loeber, R., & Farrington, D. P. (2009). Marginal and random intercepts models for longitudinal binary data with examples from criminology. *Multivariate Behavioral Research*, 44(1), 28–58. <https://doi.org/10.1080/00273170802620071>
- McNeil, S. L., Andrews, A. R., & Cohen, J. R. (2020). Emotional maltreatment and adolescent depression: Mediating mechanisms and demographic considerations in a child welfare sample. *Child Development*, 91(5), 1681–1697. <https://doi.org/10.1111/cdev.13366>
- Muthén, L. K., & Muthén, B. O. (2017). *Mplus user's guide* (8th ed.).
- Neuhauser, J. M., Kalbfleisch, J. D., & Hauck, W. W. (1991). A comparison of cluster-specific and population-averaged approaches for analyzing correlated binary data. *International Statistical Review/Revue Internationale de Statistique*, 59(1), 25–35. <https://www.jstor.org/stable/1403572>
- Olsen, M. K., & Schafer, J. L. (2001). A two-part random-effects model for semicontinuous longitudinal data. *Journal of the American Statistical Association*, 96(454), 730–745. <https://doi.org/10.2307/2670310>
- Pinheiro, J., Bates, D., DebRoy, S., Deepayan, S., EISPACK authors, Heisterkamp, S., Willigen, B. Van, & R-Core. (2020). *Package "nlme"* <https://cran.r-project.org/web/packages/nlme/nlme.pdf>
- R Core Team. (2014). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. <https://www.r-project.org/>
- Radloff, L. S. (1977). The CES-D Scale: A self-report depression scale for research in the general population. *Applied Psychological Measurement*, 1(3), 385–401. <https://doi.org/10.1177/014662167700100306>
- Rizopoulos, D. (2020). *GLMMadaptive: Generalized linear mixed models using adaptive Gaussian quadrature*. <https://cran.r-project.org/package=GLMMadaptive>
- Rubin, D. B. (1987). *Multiple imputation for nonresponse in surveys*. John Wiley & Sons. <https://doi.org/10.1002/9780470316696>
- Thornberry, T. P., Henry, K. L., Krohn, M. D., Lizotte, A. J., & Nadel, E. L. (2018). Key findings from the Rochester Intergenerational Study. In V. I. Eichelsheim & S. G. A. van de Weijer (Eds.), *Intergenerational continuity of criminal and antisocial behavior: An international overview of current studies*. Routledge.
- Whittaker, J. K., & Tracy, E. M. (1990). Family preservation services and education for social work practice: Stimulus and response. In J. K. Whittaker, J. Kinney, E. M. Tracy, & C. Booth (Eds.), *Reaching high-risk families: Intensive family preservation in human services* (pp. 1–12). <https://doi.org/10.4324/9781315128047>
- Xu, S., Blozis, S. A., & Vandewater, E. A. (2014). On fitting a multivariate two-part latent growth model. *Structural equation modeling*, 21(1), 131–148. <https://doi.org/10.1080/10705511.2014.856699>
- Zeger, S. L., Liang, K.-Y., & Albert, P. S. (1988). Models for longitudinal data: A generalized estimating equation approach. *Biometrics*, 44(4), 1049–1060. <https://doi.org/10.2307/2531734>

Manuscript submitted: September 2, 2020

First revision submitted: December 4, 2020

Second revision submitted: January 11, 2021

Accepted: January 12, 2021

Electronically published: June 3, 2022