CHAPTER 4

MODELING GROWTH USING MULTILEVEL AND ALTERNATIVE APPROACHES

Janet K. Holt

In recent years, the array of linear growth modeling techniques has expanded greatly, from traditional methods of univariate or multivariate repeated measures analysis to more flexible random coefficients models that include multilevel growth models, latent growth curve models, and growth mixture modeling. This chapter provides general descriptions of and contrasts methods across three classes of growth models: multilevel growth modeling, including linear, quadratic, piecewise, and shift models; latent growth curve models, including models for linear, quadratic, and estimated growth curves; and growth mixture models, including estimation and prediction of latent growth classes as well as prediction of distal outcomes. While not an exhaustive list of growth modeling methods, this is a comparison of some of the more popular, cutting-edge methods in the literature. In this chapter, I demonstrate how analysts can use these classes of growth models in longitudinal or developmental studies to model change across time, describe how to use each of these models to address substantive questions about change, and summarize the strengths/weaknesses, assumptions, and data

Multilevel Modeling of Educational Data, pages 111–159

requirements for each of these classes of models. I begin with a general discussion of each approach and provide examples from the research literature. For selected models, an applied demonstration of the technique is also included.

MULTILEVEL MODELING

Reconceptualizing Longitudinal Models As Multilevel Models

Traditional approaches for analyzing longitudinal data utilize repeated measures ANOVA or MANOVA techniques; however, these methods place severe constraints on the form of the data. The two most problematic constraints in repeated measures analyses are that all subjects must have an equal number of data points and that the data collection schedule needs to be *time-structured,* such that the planned schedule of data collection must be at the same times for all individuals. By default, these traditional longitudinal analyses use listwise deletion to discard participants without full data for all time points. This often results in a much-reduced data set that does not accurately represent the originally sampled population and that is likely to be biased.

In contrast, in multilevel modeling (MLM), the data are a series of observations nested within the individual; therefore, the structure of the data can be person-specific and much more flexible. This approach allows for data that are collected at unequally-spaced waves of data collection and that are *time-unstructured* (i.e., different data collection schedules for different individuals) and *unbalanced* (i.e., different number of observations for each individual). This flexibility also can translate to person-specific growth trajectories. Analysts can estimate the variation in growth patterns and investigate relationships with covariates to model both the intra- and inter-person variability. This reconceptualization of growth modeling results in a flexible modeling approach that more aptly captures the inherent complexity in growth processes.

Multilevel Linear Growth Models

Analysts can use the basic multilevel linear growth model to assess both initial status and linear change over time. Equations 4.1 and 4.2 describe this model with random coefficients:

$$y_{ti} = \pi_{0i} + \pi_i (time - time_1)_{ti} + e_{ti} \tag{4.1}$$

$$\pi_{0i} = \beta_{00} + r_{0i}$$
$$\pi_{1i} = \beta_{10} + r_{1i} \tag{4.2}$$

for $i = 1,\ldots, n$ subjects across $t = 1,\ldots, T$ waves. The growth parameters, π_{0i} and π_{1i}, represent the intercept and linear rate of change, respectively, for person i, and e_{ti} is the within-person residual not accounted for by the specified growth parameters. If time 1 is the initial time point assessed in the data, then the intercept represents the initial value on the dependent variable. The level-one equation (see Equation 4.1) is the individual growth model and specifically describes the outcome at time 1, the intercept and the rate of change for person i, and random fluctuations around the linear growth trajectory. The level-two equations (see Equation 4.2) describe the between-person variability in the growth parameters: the intercepts, π_{0i}, and the linear slopes, π_{1i}. The level-two residuals, r_{0i} and r_{1i}, represent the random, between-person differences in the growth parameters, π_{0i} and π_{1i}, respectively; and the fixed effects in this model, β_{00} and β_{10}, represent the average intercept and the average rate of growth, respectively. The level-two equations allow us to model the variability in the growth parameters across persons. Together, Equations 4.1 and 4.2 represent the unconditional linear growth model with random slopes and intercepts.

Data Requirements and Assumptions

Data Requirements

MLM allows the analysis of incomplete data as long as data are missing at random (MAR; i.e., the missingness pattern can be related to observed values of other variables in the data set; Little & Rubin, 2002). As previously mentioned, multilevel growth models do not require the data to be time-structured or balanced. However, these models do require one more wave of data than the number of growth parameters in the level-one growth model (see Equation 4.1). Therefore, a linear model with two growth parameters in the level-one equation, π_{0i} and π_{1i}, would require at least three waves of data. This is the minimum requirement, but one can estimate the parameters with greater precision with additional waves of data.

Assumptions

MLM makes assumptions regarding both the random components of the level-one and level-two models as well as specification assumptions about the relationship of the variables to the random components. For a detailed explanation of assumptions, see Raudenbush and Bryk (2002, pp. 255–256). One important assumption is that both level-one and level-two residuals are independently and normally distributed. Nonnormality introduced at level one will bias the standard errors at both levels one and two. Examination of the residuals with normal probability plots is an accepted procedure for checking whether the data meet this assumption. Analysts need to construct separate normal probability plots for each level of residuals (e.g., e_{il}, r_{0i}, r_{1i}). To normalize the data and resolve many non-normality problems, analysts can use common data transformation procedures. (See Judd & McClelland, 1989, for an excellent discussion of data transformations.)

Form of the Data

General data analysis software typically lays out data on repeated observations in a multivariate, or *wide*, format, where each observation is represented by a different variable in a separate column within the database. For example, Figure 4.1 presents an SPSS screen shot of data on reading achievement for a sample of children across five waves of data collection: 1, 2, 3, 4, and 8. In this dataset, described in more detail in a later section,

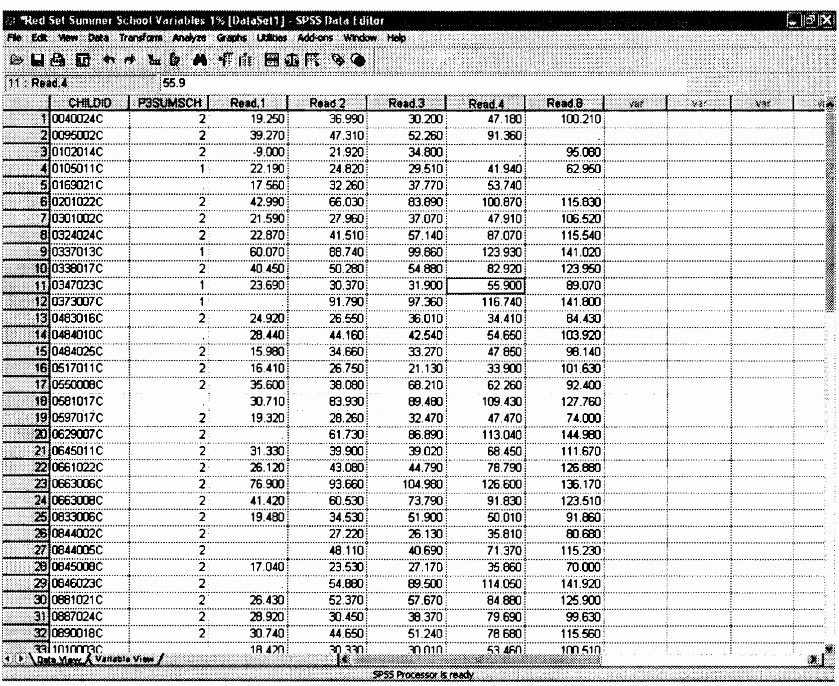

	CHILDID	P3SUMSCH	Read.1	Read.2	Read.3	Read.4	Read.8	var	var	var
1	0040024C	2	19.250	36.990	30.200	47.180	100.210			
2	0095002C	2	39.270	47.310	52.260	91.360	.			
3	0102014C	2	-9.000	21.920	34.800	.	95.080			
4	0105011C	1	22.190	24.820	29.510	41.940	62.950			
5	0169021C	.	17.560	32.260	37.770	53.740	.			
6	0201022C	2	42.990	66.030	83.890	100.870	115.830			
7	0301002C	2	21.590	27.960	37.070	47.910	106.520			
8	0324024C	2	22.870	41.510	57.140	87.070	115.540			
9	0337013C	1	60.070	88.740	99.860	123.930	141.020			
10	0338017C	2	40.450	50.280	54.880	82.920	123.950			
11	0347023C	1	23.690	30.370	31.900	55.900	89.070			
12	0373007C	1	.	91.790	97.360	116.740	141.800			
13	0483016C	2	24.920	26.550	36.010	34.410	84.430			
14	0484010C	.	28.440	44.160	42.540	54.650	103.920			
15	0484025C	2	15.980	34.660	33.270	47.850	98.140			
16	0517011C	2	16.410	26.750	21.130	33.900	101.630			
17	0550008C	2	35.600	38.080	68.210	62.260	92.400			
18	0581017C	.	30.710	83.930	89.480	109.430	127.760			
19	0597017C	2	19.320	28.260	32.470	47.470	74.000			
20	0629007C	2	.	61.730	86.890	113.040	144.980			
21	0645011C	2	31.330	39.900	39.020	68.450	111.670			
22	0661022C	2	26.120	43.080	44.790	78.790	126.880			
23	0663006C	2	76.900	93.660	104.980	126.600	136.170			
24	0663008C	2	41.420	60.530	73.790	91.830	123.510			
25	0833006C	2	19.480	34.530	51.900	50.010	91.860			
26	0844002C	2	.	27.220	26.130	35.810	80.680			
27	0844005C	2	.	48.110	40.690	71.370	115.230			
28	0845008C	2	17.040	23.530	27.170	35.860	70.000			
29	0846023C	2	.	54.880	89.500	114.050	141.920			
30	0881021C	2	26.430	52.370	57.670	84.880	125.900			
31	0887024C	2	28.920	30.450	38.370	79.690	99.630			
32	0890018C	2	30.740	44.650	51.240	78.680	115.560			
33	1010003C		18.420	30.330	30.010	53.460	100.510			

Figure 4.1 Multivariate, or *wide*, data layout.

data were collected at five "terms," that is, fall and spring of kindergarten, fall and spring of first grade, and spring of third grade. Note that the analyst formatted each measure of reading achievement as a separate column in this data layout and that there is only one row per person or child ID. However, when conducting growth analyses, it is helpful to represent each time point as a separate case. Additionally, if the data are not time-structured and different persons have different data structures, it is cumbersome to structure the data in a multivariate format. For these reasons, analysts organizing repeated measures data for MLM often use a person-period, or *long*, data set. A person-period data set has multiple rows per person, one for each time point at which the person has a measurement. The analyst enters a variable coding the passing of time (e.g., wave, time, age, term, grade) as a separate variable in the data set. In this data layout, the number of rows equals the number of observations, whereas in a multivariate format, the number of rows equals the number of persons. Figure 4.2 compiles the time-varying measures of reading achievement into a single column called "read," and in this format, the variable "term" indicates the wave in which data collection occurred. Also note that there now are five rows of data per child, one for each time period of data collection.

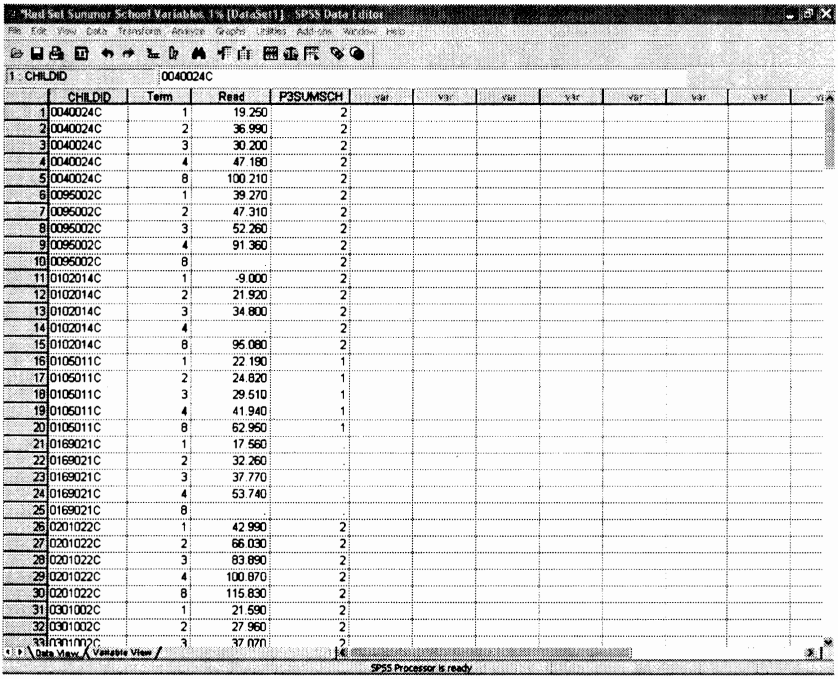

	CHILDID	Term	Read	P3SUMSCH
1	0040024C	1	19.250	2
2	0040024C	2	36.990	2
3	0040024C	3	30.200	2
4	0040024C	4	47.180	2
5	0040024C	8	100.210	2
6	0095002C	1	39.270	2
7	0095002C	2	47.310	2
8	0095002C	3	52.260	2
9	0095002C	4	91.360	2
10	0095002C	8	.	2
11	0102014C	1	-9.000	2
12	0102014C	2	21.920	2
13	0102014C	3	34.800	2
14	0102014C	4	.	2
15	0102014C	8	95.080	2
16	0105011C	1	22.190	1
17	0105011C	2	24.820	1
18	0105011C	3	29.510	1
19	0105011C	4	41.940	1
20	0105011C	8	62.950	1
21	0169021C	1	17.560	.
22	0169021C	2	32.260	.
23	0169021C	3	37.770	.
24	0169021C	4	53.740	.
25	0169021C	8	.	.
26	0201022C	1	42.990	2
27	0201022C	2	66.030	2
28	0201022C	3	83.890	2
29	0201022C	4	100.870	2
30	0201022C	8	115.830	2
31	0301002C	1	21.590	2
32	0301002C	2	27.960	2
33	0301002C	3	37.070	2

Figure 4.2 Person-period, or *long*, data layout.

Treatment of Time

To measure the passing of time, analysts typically enter a variable for time into the level-one equation, as in Equation 4.1. However, depending on the research scenario, the analyst may use an alternate variable for time that more closely corresponds to the research design and occasions of measurement. For instance, the analyst may use age, denoted in years, months, or even days, for developmental research studies; on the other hand, in school-based studies, analysts frequently use grade level to describe the passage of time.

No matter the variable chosen to denote the passage of time, the analyst should give careful consideration to centering this variable in multilevel modeling because the interpretation of the intercept, β_{00}, and all other lower-order growth parameters depends on the centering point. For instance, if the time series ranges from 24 months of age to 36 months of age and has data available at 3-month intervals, then depending on the centering of age, β_{00} could represent the average outcome at the onset of the data collection period (i.e., 24 months), at the midpoint (i.e., 30 months), or at the end of the data collection period (i.e., 36 months). Biesanz, Deeb-Sossa, Papakakis, Bollen, and Curran (2004) suggest that the coding of time in multilevel modeling should facilitate interpretability and should focus on the main period of interest in the study.

To illustrate the effects of centering, consider the situation in which a researcher is studying the effects of test coaching by monitoring test scores just prior to the coaching sessions, and at the quarterly periods during and immediately following the 12-week coaching session. Centering at 0 weeks (i.e., initial status), 6 weeks (i.e., midpoint), and 12 weeks (i.e., final status) will result in different interpretations of β_{00}. If the researcher coded the first week as 0, β_{00} would represent the expected value of coaching at initial status, just prior to the onset of coaching. This may not be the most desired interpretation for β_{00} because it is before the treatment actually occurs. In contrast, centering at 6 weeks would require recoding the origin to (*week* – 6) and would result in an intercept that is the average outcome midway through coaching, an interesting time point if the growth process is of interest. Centering at final status would require recoding the origin to (*week* – 12) and would allow an interpretation of the intercept after coaching is complete, which may be most appropriate if the final outcome is of greatest interest.

Centering the time variable does not affect β_{10} in a linear growth model because β_{10} is a constant linear rate of growth across time—in this example, from 0 to 12 weeks. Likewise, centering time does not affect the within-person variance or the residual variance in rate of change; it can, however, drastically affect both the variance of the intercept and the covariation between the intercept and rate of change. Therefore, if coaching reduces the variability in test scores, then the random variance in the intercept and the correlation between the intercept and growth rate would decrease as the centering point moves from initial status to midpoint to final status. Thus, decisions regarding centering should take into consideration the impact of the centering point on the interpretation of the intercept, the residual variance of the intercept, and the correlation between the intercept and rate of change.

Variance-Covariance Structures

The unconditional growth model in Equations 4.1 and 4.2 has an implied variance of e_{ti} equal to σ^2, the within-person residual variance, and a variance/covariance structure of the unique person effects, r_{0i} and r_{1i}, equal to:

$$\mathbf{T} = \begin{bmatrix} \tau_{00} & \tau_{01} \\ \tau_{10} & \tau_{11} \end{bmatrix} \tag{4.3}$$

(Raudenbush & Bryk, 2002). **T**, referred to as the tau matrix, is the level-two covariance structure for the model that Equations 4.1 and 4.2 describe. Assuming time is centered at initial status, the variance terms, τ_{00} and τ_{11}, describe the variance in initial status and rate of growth across individuals, respectively, whereas the covariance term, τ_{01} or τ_{10}, describes the covariation between initial status and rate of growth across individuals, which also can be expressed as a correlation. With data from three or more time points (multiwave data), the correlation between initial status and rate of growth provides an estimate of the correlation between true initial status and growth rate (Raudenbush & Bryk, 2002), although floor or ceiling effects at particular time points can influence the estimate. As noted previously, this correlation can vary with a change in centering point. Depending on the software package used, various options are available for modeling alternative covariance structures within a multilevel framework. These are more fully described under software options.

Modeling Change with Covariates

The structure of the multilevel model allows the incorporation of covariates at levels one and two. Consequently, analysts can use time-varying covariates at level one to account for variation in observations within individuals, and time-invariant covariates at level two to account for variation in growth parameters across individuals. Combined, both types of covariates allow for the formulation of rich models of change.

Time-Varying Covariates

Analysts can incorporate covariates into level one of the model to account for within-person changes that occur across observations. The level-two equations will model each of the level-one parameters. For example, Equation 4.4 (with *time* centered at $time_1$) adds a time-varying covariate, a_{ti}, to level one, and Equation 4.5 models the associated parameter, π_{2i}, in a new level-two equation. Depending on whether r_{2i} has significant variation across individuals, the analyst is able to fix the equation (i.e., remove this residual from the level-two equation) or treat it as random (i.e., keep the residual in the level-two equation). Analysts can use a significance test of the variation of the r_{2i} residuals to inform this decision. However, for models with more than two or three random effects, estimation difficulties may occur as **T** becomes increasingly complex. The number of elements in the tau matrix is r*(r+1)/2, where r is the number of random effects at level two, because we normally estimate the variances for each random effect as well as the covariances between all pairs of random effects. For this reason, parsimony is an important consideration when determining which level-one predictors to include as random and which effects to fix at level two.

$$y_{ti} = \pi_{0i} + \pi_{1i}(time - time_1) + \pi_{2i}a_{ti} + e_{ti} \tag{4.4}$$

$$\begin{aligned} \pi_{0i} &= \beta_{00} + r_{0i} \\ \pi_{1i} &= \beta_{10} + r_{1i} \\ \pi_{2i} &= \beta_{20} + r_{2i} \end{aligned} \tag{4.5}$$

Time-Invariant Covariates

One of the strengths of multilevel modeling is the ability to model cross-level effects, or interactions between variables measured at different levels of analysis. Within a growth modeling framework, this allows for modeling the relationships between effects that are repeated measures (i.e., measured within-persons) and individual-level effects (i.e., measured at the person level). Covariates assessed at the person level are termed time-invariant covariates, and analysts easily can incorporate them into the level-two equations of the multilevel growth model. Examples of time-invariant covariates include gender, or race/ethnicity. As an alternative to Equation 4.5, Equation 4.6 is a series of level-two equations that analysts could formulate to model the growth parameters in Equation 4.4. The addition of X_i, a time-invariant covariate, to the level-two equations allows the introduction of a person-level variable to study the effects of that variable on the intercept via

β_{01}; the effects of X_i on the linear growth rate via β_{11}, and the effects of X_i on the time-varying covariate via β_{21}. These latter two parameters represent the cross-level interactions of a person-level variable, X_i, with the within-person effects of Equation 4.4. Note that Equation 4.6 includes the same predictor in all level-two equations; however, the analyst may choose to include different time-invariant predictors within any of the level-two models.

$$\begin{aligned} \pi_{0i} &= \beta_{00} + \beta_{01}X_i + r_{0i} \\ \pi_{1i} &= \beta_{10} + \beta_{11}X_i + r_{1i} \\ \pi_{2i} &= \beta_{20} + \beta_{21}X_i + r_{2i} \end{aligned} \tag{4.6}$$

Modeling Change in Growth Rates

Although a simple linear model of change is appropriate for many growth scenarios, there are instances in which the linear model is not the best fit, and the analyst should examine other alternatives. Consider the situation in which subjects grow in a linear trajectory but then growth slows and the rate of change lessens (i.e., decelerates) or, alternatively, the growth increases (i.e., accelerates) over time. As this description illustrates, more complex growth curves may involve changes in the growth rate. Alternately, the change in growth rate may be abrupt, and thus represent separate phases of growth. Deviations from a constant linear rate of growth may be due to physical, cognitive, or other internal developmental processes or can be due to transitions that occur in the society or institutions to which the subjects belong. Whatever the cause, analysts often can examine resultant changes in growth rates by using growth models that extend beyond the simple linear form. The following sections on polynomial multilevel growth models and multiphase multilevel growth models will consider these alternatives to simple linear growth.

Polynomial Multilevel Growth Models

The addition of terms that include higher-order time variables (e.g., time-squared, time-cubed) can be used to account for changes in growth rates. A quadratic growth curve includes the square of the time variable, and the corresponding coefficient represents the degree of acceleration or deceleration in growth that occurs over time; that is, whether or not the curve is tapering off (decelerating) or rapidly increasing (accelerating) as the time variable increases. Typically, analysts test the quadratic model with a likelihood ratio test to determine if it provides a better fit than the linear model. This is done by constructing an hypothesis test comparing the restricted model (e.g., linear model) to the more complex alternative model (e.g., quadratic model). The likelihood ratio test compares deviances and *df* for these two nested models using the χ^2 difference test. A statistically significant χ^2 test indicates that the more complex model is warranted. (See McCoach & Black, 2008, Chapter 7 this volume for more details about the chi-square difference test). Equation 4.7 describes the unconditional quadratic growth model:

$$
\begin{aligned}
y_{ti} &= \pi_{0i} + \pi_{1i}(time - time_1) + \pi_{2i}(time - time_1)^2 + e_{ti} \\
\pi_{0i} &= \beta_{00} + r_{0i} \\
\pi_{1i} &= \beta_{10} + r_{1i} \\
\pi_{2i} &= \beta_{20} + r_{2i}
\end{aligned}
\tag{4.7}
$$

In this model, the interpretation of the linear coefficient changes somewhat from the linear growth model. Recalling that time is centered at the first time-point of data collection, the intercept, π_{0i}, remains the initial status for person *i*. However, π_{1i} now represents the rate of change *at initial status* (i.e., when $time - time_1 = 0$) for person *i*. We refer to this effect as the *instantaneous* rate of change at initial status. This change in interpretation arises because the quadratic model no longer has a single linear rate of change; instead, there is a different rate of change at every time point. Rates of change in a quadratic model are estimated by the slopes of lines tangent to the growth curve at each point on the curve. These "simple-slopes" change across the entire time-span of the growth curve. In a polynomial model, the simple-slopes are equivalent to the first derivative, with respect to time, of the level-one equation evaluated at each specific value of time; for the quadratic model, the equation for these simple-slopes is: $\pi_{1i} + 2\pi_{2i}(time - time_1)$. Because time is centered at $time_1$ in this example, analysts can interpret π_{1i} as the instantaneous rate of growth *at initial status*. If the analyst recentered the data to the midpoint, then π_{1i} would represent the instantaneous rate of change *at the midpoint*. The interpretation of π_{1i} depends on the placement of the origin for time (Biesanz, et al., 2004).

The new level-one parameter, π_{2i}, represents the acceleration/deceleration apparent in the growth curve for person *i* across time. When π_{2i} is positive, acceleration is occurring, and the growth curve is convex to the time axis (i.e., the instantaneous linear growth rates are increasing). In contrast, if π_{2i} is negative, then deceleration is occurring, and the growth curve is concave to the time axis (i.e., the instantaneous linear growth rates are slowing). This new parameterization of growth requires at least four time points because there are three growth parameters in the level-one equation of the quadratic growth model: π_{0i}, π_{1i}, and π_{2i}.

The level-two equation for π_{2i} is a model for the acceleration/deceleration parameter for person *i*, and describes its variation as a function of the grand mean for acceleration/deceleration, β_{20}, and the unique effects of the person (see Equation 4.7). Overall, the person-level parameters, β_{00}, β_{10}, and β_{20}, represent the population average intercept at initial status, the average instantaneous rate of growth at initial status, and the average acceleration/deceleration across time, respectively. The residuals, r_{0i}, r_{1i}, and r_{2i}, are the unique effects of person *i* on their initial status, instantaneous growth at initial status, and acceleration/deceleration rates, respectively.

The unconditional quadratic growth model has random slopes and intercepts, an implied variance of e_{ti} equal to σ^2, and a variance/covariance structure of r_{0i}, r_{1i}, and r_{2i} equal to:

$$\mathbf{T} = \begin{bmatrix} \tau_{00} & \tau_{01} & \tau_{02} \\ \tau_{10} & \tau_{11} & \tau_{12} \\ \tau_{20} & \tau_{21} & \tau_{22} \end{bmatrix} \tag{4.8}$$

(Raudenbush & Bryk, 2002). Analysts could expand this model further to include individual effects on the mean growth parameters (i.e., a conditional growth model) by incorporating level-two time-invariant covariates, similar in structure to Equation 4.6. As with the linear model, the level-two covariates of the different growth parameters do not need to be the same. Additionally, organizational effects on the person-level parameters can be assessed through a three-level model (e.g., time nested within individuals, nested within organizations) by incorporating level-three time-invariant covariates.

Analysts also can formulate more complex polynomial growth models. For example, a cubic model includes a parameter for time-cubed. The parameter for the cubed term represents the change in the acceleration or deceleration that is occurring over time. If examined graphically, a cubic model has one inflection point, indicating that the quadratic acceleration/deceleration pattern is changing. In multilevel polynomial growth models, analysts interpret the highest-order term (e.g., cubic parameter) across the full range of the time variable, whereas they interpret the lower-order terms (e.g., intercept, linear, and quadratic parameters) at the centering point. Researchers need to weigh the advantages of modeling a more complex model with the added data requirements. For instance, a minimum of five waves would be needed to test a cubic model. However, if this is not a trend that they would expect or be interested in, researchers may opt to forego the more complex model and analyze a less complex polynomial model that could be estimated with greater precision.

Examples From The Literature

Using data from the Longitudinal Study of American Youth (LSAY), Ma (2005) studied both student and school effects on the rate of mathematics achievement growth from 7th through 11th grade using a three-level linear growth model. Level one modeled the intra-individual variation across time, level two modeled the inter-individual variation within schools, and level three modeled the inter-school variation. The outcome variable was mathematics achievement, and the student-level predictors included gender, race, age, and other demographic variables. The school-level predictors included measures of both school context and school climate. Ma compared growth rates and effects of the student-level demographic variables, and effects of school-level context and climate variables, on mathematics achievement growth between students who experienced early acceleration in their mathematics instruction and those who did not, as well as among those in regular, honors, and gifted classes.

In a growth study of the onset of tense marking (a measure of grammatical development in young children), Hadley and Holt (2006) constructed quadratic growth models to estimate the growth rates in tense marking from 24 to 36 months of age in a sample of slow-developing language learners, assessed at 3-month intervals. The positive coefficients for both the linear and the quadratic components indicated that growth was increasing and accelerating across time. Hadley and Holt expected that children's changes over time in their mean length of utterance (MLU) also would affect their developmental trajectories in tense marking. Including the time-varying covariate, MLU, which also was assessed in 3-month intervals from 24 to 36 months, controlled for the variation in children's growth in tense marking that was not due solely to the child's developmental growth over time but also to the individual child's change in MLU. The researchers also modeled the interaction between MLU and time in order to determine if the tense marking growth rate covaried with changes in MLU.

Hadley and Holt (2006) also examined the influence of several time-invariant covariates on the growth parameters to determine if environmental factors (e.g., maternal education) or maturational factors (e.g., family history of speech, language, or learning disabilities) covaried with the average level of tense marking and its linear growth rate in slow-developing language learners. They assessed the time-invariant covariates at the initial evaluation and modeled their influence on the growth parameters as well as on the time-varying covariates.

Application—Linear and Quadratic Growth Models

To assess growth in reading achievement from kindergarten through third grade, I used a subsample of data from the kindergarten cohort of the Early Childhood Longitudinal Survey (ECLS-K; n = 16,400). The students completed cognitive reading assessments in the fall and spring of kindergarten, fall and spring of first grade, and the spring of third grade. These data were not gathered at equal intervals: students did not complete the assessments in the fall and spring of second grade or the fall of third grade. Furthermore, a subsample of only 30% of the base-year cohort completed the assessment in the fall of first grade (National Center for Education Statistics, 2004). Therefore, not all subjects had an equal number of data collection points. To simplify the presentation of this illustrative example, I did not use sampling weights in the current analyses. Full maximum likelihood was used for parameter estimation.

I formulated multilevel polynomial growth models to 1) determine the growth trajectory for reading achievement from kindergarten through third grade, 2) determine whether a time-varying covariate, changing schools between terms, affects the reading achievement, and 3) determine whether the growth parameters and the effect of changing schools vary depending on the time-invariant predictor, gender. Figure 4.3 graphs a random sample of 1% of the cases by gender. It is apparent that reading achievement

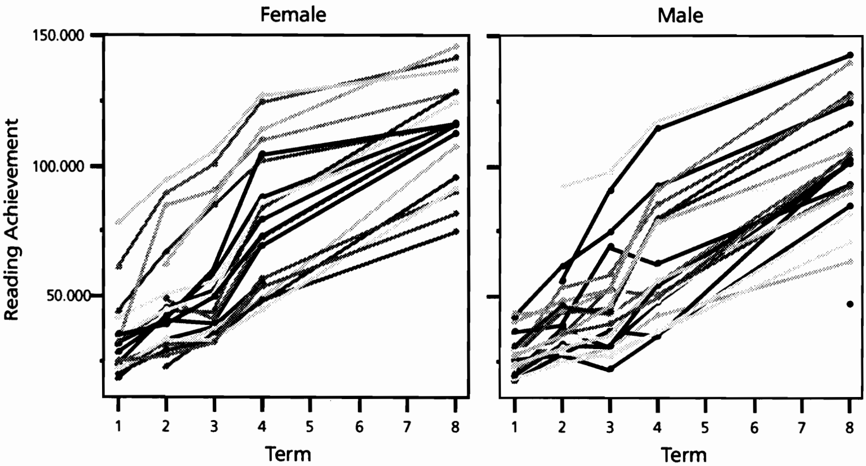

Figure 4.3 Individual growth curves of reading achievement as a function of gender for a selected sample from kindergarten through third grade.

growth is decelerating (slowing down) from kindergarten to third grade and that a linear growth model would not account for this deceleration. In these graphs, time ranges from 1 to 8, representing each term from fall and spring of kindergarten to fall and spring of third grade.

Although the graph of individual growth curves can illustrate general trends, a multilevel statistical analysis directly addresses the questions of interest. To assess reading achievement growth, I first formulated a two-level linear growth model (Model 1). Next, I formulated a quadratic growth model (Model 2) and tested the difference in the deviances to determine if the more complex quadratic model yielded a significantly better fit to the data. I then added the time-varying covariate (1 = school change, 0 = no school change) to the quadratic growth model (Model 3) to determine if there was a significant relationship between changing schools and reading achievement, controlling for growth, as well as if the growth remained statistically significant after controlling for the effect of changing schools. Finally, I entered the level-two covariate, gender (1 = male, 0 = female) as a predictor of the growth parameters and of the impact of the time-varying covariate (changing schools) on reading achievement (Model 4).

Results for the four models are presented in Table 4.1. Equation 4.9 represents the most complex model, Model 4; those parameters not included in Models 1—3 would be set to zero in this general model. Time was measured by *term*, and centered at fall of first grade. Note that Equation 4.9 contains a separate level-two equation for each time-varying effect at level one.

$$\begin{aligned}
&Reading\ Ach_{ti} = \pi_{0i} + \pi_{1i}(term - term_{F1st})_{ti} + \\
&\pi_{2i}\left(term - term_{F1st}\right)^2_{ti} + \pi_{3i}(school\ change)_{ti} + e_{ti} \\
&\pi_{0i} = \beta_{00} + \beta_{01}(gender)_i + r_{0i} \\
&\pi_{1i} = \beta_{10} + \beta_{11}(gender)_i + r_{1i} \\
&\pi_{2i} = \beta_{20} + \beta_{21}(gender)_i + r_{2i} \\
&\pi_{3i} = \beta_{30} + \beta_{31}(gender)_i + r_{3i}
\end{aligned} \qquad (4.9)$$

There was significant variation in the linear growth trajectories, as indicated by τ_{11} in Model 1 (See Table 4.1) The mean linear growth rate in reading achievement from kindergarten through third grade, as estimated by β_{10}, was positive and statistically significant. This growth rate implies that a child whose reading growth is one standard deviation above average can be expected to improve at a rate of $11.77 + \sqrt{5.09} = 11.77 + 2.26 = 14.03$ points per term (unconditional linear growth rate + 1 SD in growth rate).

In Model 2, I added the quadratic term, π_{2i}, as noted in Equation 4.9. The quadratic growth model provided a significantly better fit than the lin-

TABLE 4.1 Multilevel Growth Curve Modeling of Reading Achievement

Fixed Effects	Model 1	Model 2	Model 3	Model 4
Avgerage Achievement				
Intercept, β_{00}	49.68***	51.59***	51.63***	53.17***
Gender, β_{01}				-3.01***
Linear Effect				
Intercept, β_{10}	11.77***	13.05***	13.09***	13.45***
Gender, β_{11}				−0.71***
Quadratic Effect				
Intercept, β_{20}		−0.40***	−0.39***	−0.44***
Gender, β_{21}				0.10***
School Change				
Intercept, β_{30}			−3.15***	−3.68***
Gender, β_{31}				1.05†
Random Effects	**Variance Components and Deviance Statistics**			
Level one error	100.61	64.60	63.66	63.66
Intercept, τ_{00}	212.61***	321.20***	321.48***	319.19***
Linear growth, τ_{11}	5.09***	13.01***	13.15***	13.02***
Quadratic growth, τ_{22}		0.69***	0.68***	0.67***
School change, τ_{33}			30.45*	29.77*
Deviance	539,631.1	529,291.1	528,927.1	528,788.3
Number of parameters	6	10	15	19

***$p < .001$, **$p < .01$, *$p < .05$, †$p < .07$

ear model according to the likelihood ratio test comparing the two models, $\chi^2(4) = 10{,}340$, $p < .001$; therefore, I retained the quadratic parameter in subsequent growth models. There was significant random variation in the quadratic parameter, as indicated by τ_{22} (See Table 4.1). The linear parameter in Model 2, β_{10}, is positive, indicating a positive mean instantaneous growth rate at fall of first grade. However, the quadratic parameter, β_{20}, is negative, indicating that this growth rate is not constant but instead tends to diminish over time. Specifically, predictions for reading achievement tend to change by a factor of $(-.40)*(term - term_{F1st})^2$, evidence of a decelerating growth pattern.

In Model 3, I added the time-varying covariate, school change, as a predictor of reading achievement. The coefficient, β_{30}, was statistically significant and negative, indicating that, given the instantaneous linear growth and deceleration effects, a child who changed schools would be expected to score approximately 3 points lower in reading achievement after the switch. This effect significantly varied across students, as indicated by τ_{33}.

Finally, in Model 4 I added the time-invariant covariate, gender, to determine if gender was related to average reading achievement at the fall of first grade (for a child who did not change schools during terms); to the instantaneous rate of growth in reading achievement at the fall of first grade; to the deceleration effect; or to the effect of school change on reading achievement (i.e., β_{01}, β_{11}, β_{21}, and β_{31}, respectively; see Equation 4.9). The results indicate that, relative to females, males (*gender* = 1) who did not change schools (*schoolchange* = 0) had significantly lower reading scores at the beginning of first grade ($p < .001$), significantly lower instantaneous rates of change at the fall of first grade ($p < .001$), and significantly less deceleration in growth from kindergarten through third grade ($p < .001$). Finally, these results also showed that the negative effect of changing schools did not vary significantly between boys and girls, although there is the suggestion that females may experience the negative effect of changing schools slightly more strongly than males ($p < .07$).

In summary, the results demonstrate that the instantaneous rate of change in reading achievement is positive at the fall of first grade (i.e., indicating an increasing trend at this time-point), but across the span of kindergarten to third grade there is a deceleration in reading growth over time. Those children who changed schools had lower reading achievement scores on average, although the growth in reading was still statistically significant after accounting for school change. Finally, there were different growth trajectories for females and males and a tendency for the effect of school change to be moderated by gender, although this effect failed to reach statistical significance.

In this example, multilevel modeling allowed the estimation of achievement growth across a time period in which data were not collected at equal intervals and where data were incomplete for some individuals in the sample. Furthermore, the multilevel model allowed for variation in growth trajectories when estimating average growth, producing a growth model that better reflected the individual differences in growth patterns.

Multiphase Multilevel Growth Models

When the growth follows a continuous pattern without abrupt transitions, analysts tend to prefer the previously described multilevel linear and quadratic growth models or latent growth curve models (described in a later section). In some situations, however, growth may occur in phases and people may exhibit fairly distinct growth patterns between phases. The figures below (4.4a through 4.4c) contain examples of these patterns of