CHAPTER 2. APPLYING LGM TO EMPIRICAL DATA

Data

In the following, we demonstrate how to use growth curve models in practice. For this demonstration, we use data from the National Institute of Child Health and Human Development (NICHD) Study of Early Child Care and Youth Development (NICHD Early Child Care Research Network, 2006). Quantification and subsequent analysis of parent–child relationships offer a relatively objective way to approach familial problems. A great deal of research examines the linkages between these relationships and outcomes such as smoking or drinking (e.g., Blackson, Tarter, Loeber, Ammerman, & Windle, 1996), self-regulation ability (e.g., Wills et al., 2001), and adolescent pregnancy (see Miller, Benson, & Galbraith, 2001, for a review). Good parent–child relationships are consistently found to be associated with positive child outcomes, but the associations between parent–child relationships at any given time and later outcome measures are typically moderate in size. One of the many reasons for the lack of perfect predictive validity of parent–child relationships is that these relationships change over time. Theoretical approaches ranging from the psychodynamic (e.g., Freud, 1958) to evolutionary (e.g., Steinberg, 1989) and social-cognitive (e.g., Smetana, 1988) all predict change in parent–child relationships as children mature, albeit for different reasons. Studying change over time in parent–child relationships provides additional opportunities to increase predictive accuracy, as well as to arrive at a more complete understanding of how parent–child interactions influence children's and adolescents' undesirable behaviors.

We selected two composite measures from the Child–Parent Relationship Scale. Specifically, 15 rating-scale items from the Student–Teacher Relationship Scale (Pianta, 1993) were adapted to assess parents' report of the child's attachment to the parent. Each item was scored on a scale of 1 to 5, where 1 = *Definitely does not apply* and 5 = *Definitely applies*. These items were used to construct four composite measures, labeled Conflict With Child (CNFL) and Closeness With Child (CLSN) for mothers and for fathers. These variables were assessed during the elementary school years, Grades 1 through 6 (data were not collected in the second grade). Mothers' Closeness to Child served as the primary repeated-measures variable for most of the models to be illustrated. In conjunction with these relationship measures, a grouping variable (child gender) was selected so that we could demonstrate how groups can be included in longitudinal models. We extracted data for all children who had at least one valid CLSN score, yielding a total sample size of 1,127 children (571 boys and 556 girls). Analyses were performed on all cases with complete data on the variables of interest.[1] Descriptive information is included in Table 2.1. Mean scores for CLSN are presented in Figure 2.1. Finally, the observed covariance matrix and means for the 851 complete-data cases in our sample are given in Table 2.2.

Software

The models described below were estimated using three software packages. We used LISREL 8.8 (Jöreskog & Sörbom, 1996), Mx 1.1 (Neale et al., 2003), and Mplus 4.2 (L. K. Muthén & Muthén, 1998–2006). Other user-friendly packages are available for analyzing growth curve models, such as AMOS and EQS. Mx is especially useful for its flexibility and the fact that it is public domain software, freely available for downloading. A free student version of LISREL, capable of estimating all models in this book, is available at the SSI Web site as of this writing.[2] Ferrer, Hamagami, and McArdle (2004) provide a guide to specifying growth curve models in a variety of software applications.

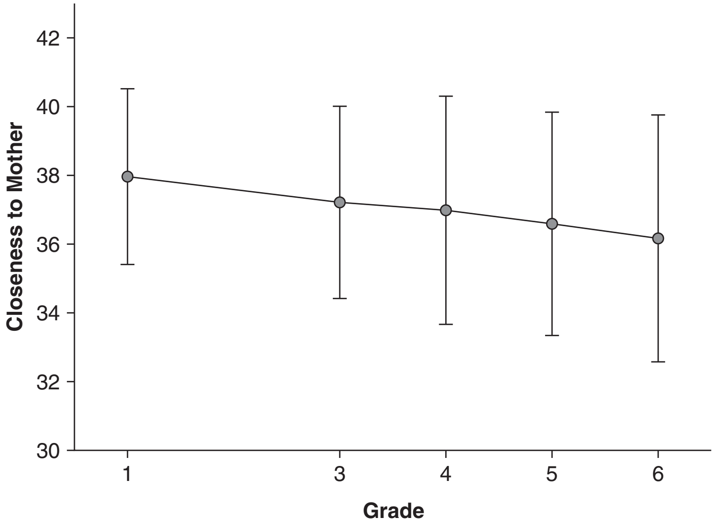

Figure 2.1 Mean Closeness Scores (Mothers), Grades 1–6. Error Bars Represent One Standard Deviation.

TABLE 2.1
Descriptive Statistics for Mother's Closeness to Child Data

	Entire Sample			*Boys*			*Girls*		
Grade	N	*Mean*	SD	N	*Mean*	SD	N	*Mean*	SD
Cases with incomplete data									
1	1,016	37.96	2.56	508	37.73	2.77	508	38.20	2.30
3	1,025	37.19	2.82	508	37.07	2.73	517	37.32	2.90
4	1,022	36.96	3.33	516	36.62	3.53	506	37.31	3.08
5	1,018	36.56	3.25	506	36.65	3.34	512	36.77	3.16
6	1,024	36.18	3.56	512	35.93	3.63	512	36.42	3.47
Cases with complete data									
1	851	37.95	2.53	417	37.76	2.66	434	38.14	2.38
3	851	37.28	2.74	417	37.20	2.62	434	37.35	2.86
4	851	37.05	3.28	417	36.74	3.45	434	37.34	3.07
5	851	36.57	3.21	417	36.34	3.31	434	36.79	3.09
6	851	36.14	3.59	417	35.84	3.73	434	36.42	3.43

TABLE 2.2
Mother–Child Closeness:
Means and Covariances of Cases With Complete Data ($N = 851$)

Grade	*Means*	*Covariances*				
1	37.9542	6.3944				
3	37.2785	3.2716	7.5282			
4	37.0463	4.1435	6.0804	10.7290		
5	36.5696	3.7058	5.1597	6.5672	10.2920	
6	36.1363	4.1286	5.7608	7.2365	7.6463	12.9085

Overview of Model-Fitting Strategy

A typical application of LGM to a problem with repeated measures contains variables measured at two levels of analysis. *Level 2 units* are the entities under study, which are usually (but not necessarily) individuals. *Level 1 units* are the repeated measurements taken on each Level 2 unit. Other variables may be measured either at Level 1 or at Level 2. Level 1 variables include the outcome (y) variable(s) and all other variables that are measured at the same occasion. Variables that are measured repeatedly and used to predict variability across repeated measures of the outcome are referred to as *time-varying covariates* (TVCs). Level 2 variables represent characteristics of the Level 2 units and thus vary *across* individuals rather than *within* individuals over time; examples might include gender, contextual variables, or stable personality traits. Level 2 predictors are often called *time-invariant covariates.* Variability in the intercept or slope factors may be explained by time-invariant covariates. For example, a researcher may be interested in the change in parents' perceptions of parent–child closeness as children age. She or he may collect data from the same parents on a number of occasions and may want to model the change in closeness over time. For this model, the researcher might choose the zero point of the time scale to lie at the first measurement occasion and then collect four subsequent waves of data from the same children. If the researcher finds that trajectories (intercepts and slopes) vary significantly across children, she or he may want to investigate whether gender differences explain some of that interindividual variability. Therefore, the researcher may introduce gender as a predictor of initial level (the intercept factor), of rate of change (the slope factor), or both. Here, gender is a time-invariant covariate because it varies over the Level 2 units (children) but not within Level 2 units.

Especially with a technique as flexible as LGM, it is helpful to specify a theoretically informed sequence of models to test before attempting to fit models to data. One then checks the fit of an a priori sequence of theoretically plausible models in the specified order. When models of increasing complexity no longer result in significant improvements in fit, one concludes than an acceptable model has been found. In the next section, we demonstrate this approach.

Model 0: The Null Model

The phrase *null model* generally is used to refer to a basis for comparison of hypothesized models. The null model in LGM is different from that in typical applications of SEM. In typical SEM applications, the null model is one in which no relationships are predicted among measured variables. Only variance parameters are estimated, and there are no latent variables. However, in the context of LGM, we define the null model to be a model in which there is no change over time and no overall variability in mean level (Widaman & Thompson, 2003). Only the mean level (intercept; α_1) and a common disturbance variance (θ_ε) are estimated:

$$\mathbf{\Lambda}_y = \begin{bmatrix} 1 \\ \vdots \\ \vdots \\ 1 \end{bmatrix}, \tag{2.1}$$

$$\mathbf{\Psi} = [0], \tag{2.2}$$

$$\boldsymbol{\alpha} = [\alpha_1], \tag{2.3}$$

$$\mathbf{\Theta}_\varepsilon = \begin{bmatrix} \theta_\varepsilon & & & \\ 0 & \ddots & & \\ & & \ddots & \\ 0 & & 0 & \theta_\varepsilon \end{bmatrix}, \tag{2.4}$$

where $\mathbf{\Lambda}_y$ is a 5 × 1 matrix representing the fixed loadings of each of our 5 occasions of measurement on the intercept factor. In this model, there is hypothesized to be no change over time, so the slope factor is omitted altogether. Matrix $\mathbf{\Psi}$ is a 1 × 1 matrix containing the variance of the intercept, which is fixed to zero in this model. $\mathbf{\Theta}_\varepsilon$ is a 5 × 5 diagonal matrix with all elements on the diagonal constrained to equality. This equality constraint represents the assumption of homoscedasticity.[3] Finally, $\boldsymbol{\alpha}$ is a 1 × 1 matrix containing the estimated population mean, α_1. If it is determined that the null model is inappropriate for the data (it usually will be), the intercept variance is usually estimated and a linear slope factor included to represent change over time. Predictors of intercept and slope can be included, following either an a priori theory-guided approach or a more exploratory approach. This simple two-parameter model, or a model adhering as closely as possible to it, will serve as the null model throughout the rest of this chapter for purposes of computing NNFI.

Model 1: Random Intercept

The random intercept model is the simplest example of a latent growth curve model. In LGM, the random intercept model is equivalent to a one-factor CFA model incorporating a mean structure, with all factor loadings fixed to 1.0 and all disturbance variances constrained to equality (see Figure 2.2). The parameter matrices are specified as in Equations 2.1, 2.2, 2.3, and 2.4, save that $\mathbf{\Psi} = [\psi_{11}]$, which corresponds to interindividual variability in overall level.

When we fit Model 1 to the data in Table 2.2 using LISREL, we obtained the parameter estimates reported in Table 2.3.[4] Results showed that there was significant unexplained intraindividual variance ($\hat{\theta}_\varepsilon$) and interindividual variance ($\hat{\psi}_{11}$). In both cases, the significance of these parameters is seen by their size relative to their standard errors, which exceeds a 2:1 ratio. This significant variation may provide a statistical rationale for fitting more

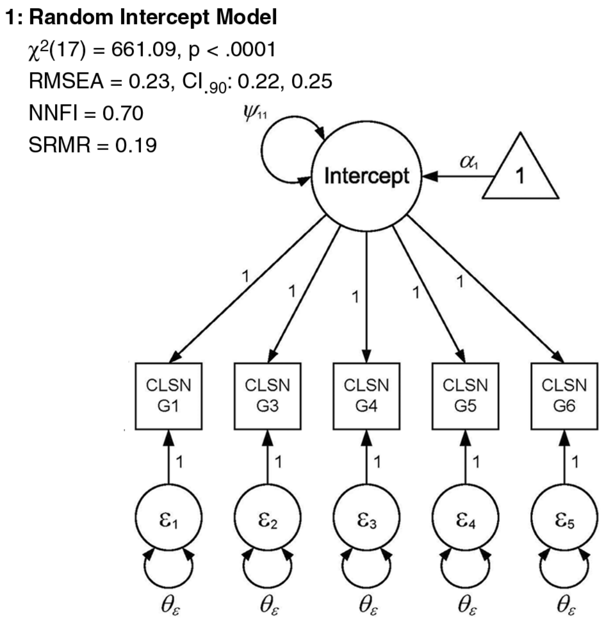

Figure 2.2 A Path Diagram Representing the Random Intercept Model.

NOTE: CLSN = Closeness with child.

complex models by, for example, including both time-varying and time-invariant covariates.

At the same time, the fit statistics show that the random intercept model does not provide an adequate fit to the data, as one might expect by examining Figure 2.1. The χ^2 test rejects the model at $p < .0001$, the RMSEA far exceeds the acceptable fit range (values less than .08 or so), and the model is characterized by large average residuals. Clearly, an intercept-only model is not appropriate for the mother–child closeness data. Next, a linear slope

TABLE 2.3
Model 1: Random Intercept Model

Parameter	*Estimate*
Mean intercept $\hat{\alpha}_1$	37.00 (0.09)
Intercept variance $\hat{\psi}_{11}$	5.27 (0.30)
Disturbance variance $\hat{\theta}_\varepsilon$	4.68 (0.11)

NOTE: Numbers in parentheses are standard errors of parameter estimates.

factor is introduced to account for the roughly linear trend observed in Figure 2.1.

Model 2: Fixed Intercept, Fixed Slope

In Model 2, we fix both the intercept and the slope, meaning that a single, average intercept parameter ($\hat{\alpha}_1$) and linear slope parameter ($\hat{\alpha}_2$) are estimated, ignoring any interindividual variation in the these aspects of change. The slope is included by adding a column to $\mathbf{\Lambda}_y$. This column now codes the slope factor and reflects two characteristics of the data. First, the loading corresponding to the first grade measure is coded 0 to place the origin of time at first grade. Second, the spacing between elements in the column reflects the fact that there was no measurement in the second grade. Thus, even though the elapsed time between the first and second measures is twice that between the second and third measures, the model still reflects linear growth because of the way in which time was coded (see Equation 2.5).

The latent growth curve specifications for Model 2 are presented in Figure 2.3. Matrix representations for the intercept and slope factors, their variances and covariances, and their means are as follows:

$$\mathbf{\Lambda}_y = \begin{bmatrix} 1 & 0 \\ 1 & 2 \\ 1 & 3 \\ 1 & 4 \\ 1 & 5 \end{bmatrix}, \tag{2.5}$$

$$\mathbf{\Psi} = \begin{bmatrix} 0 & \\ 0 & 0 \end{bmatrix}, \tag{2.6}$$

$$\boldsymbol{\alpha} = \begin{bmatrix} \alpha_1 \\ \alpha_2 \end{bmatrix}, \tag{2.7}$$

$$\mathbf{\Theta}_\varepsilon = \begin{bmatrix} \theta_\varepsilon & & \\ 0 & \ddots & \\ & & \ddots & \\ 0 & & 0 & \theta_\varepsilon \end{bmatrix}. \tag{2.8}$$

Because both the intercept and the slope are fixed, their variances and covariance are constrained to zero (hence $\mathbf{\Psi}$ contains zeroes). As in Model 1, the matrix $\mathbf{\Theta}_\varepsilon$ is constrained to represent the assumption of equal disturbance variances.

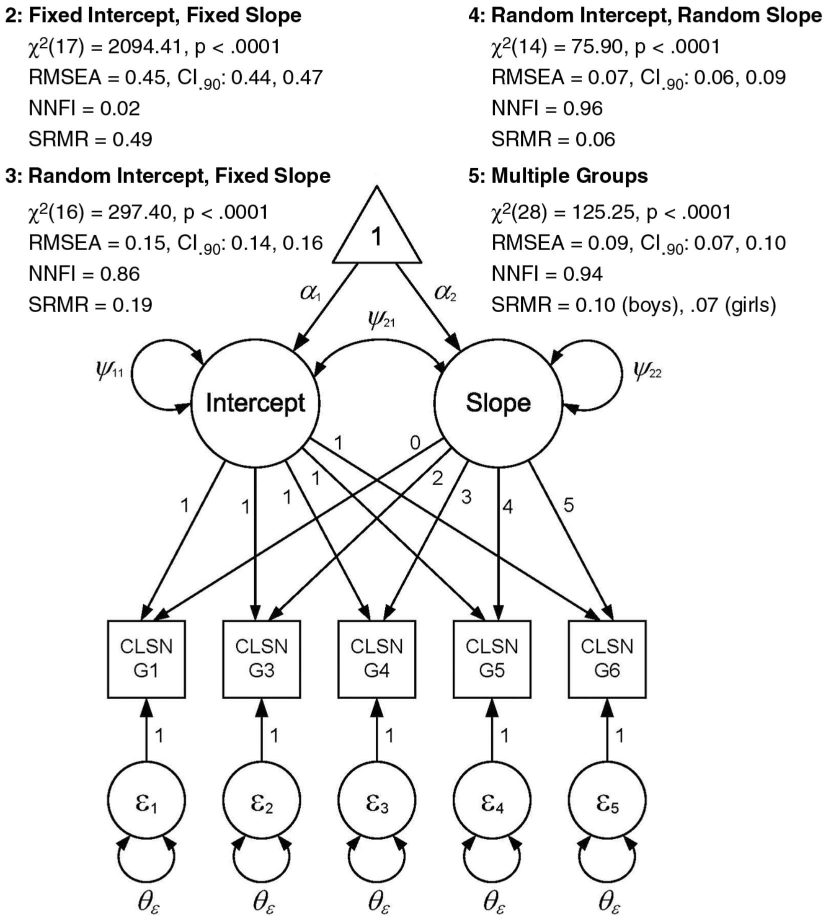

Figure 2.3 A Path Diagram Representing the General Linear Latent Growth Curve, With Random Intercept, Random Slope, Intercept–Slope Covariance, and Equal Disturbance Variances.

NOTE: CLSN = Closeness with child.

Results are reported in Table 2.4 and Figure 2.3. The new slope parameter estimate α_2 is significant and negative, reflecting the fact that closeness between mothers and children decreases over time during the elementary school years. However, the fit of Model 2 (see Figure 2.3) is much worse than that of Model 1. This poor fit comes from the way in which the intercept and slope factors are specified. Model 2 uses fixed intercept and slope factors. It constrains all mother–child pairs to share the same initial value of closeness and obliges all

TABLE 2.4
Model 2: Fixed Intercept, Fixed Slope

Parameter	*Estimate*
Mean intercept $\hat{\alpha}_1$	38.00 (0.09)
Mean slope $\hat{\alpha}_2$	–0.36 (0.03)
Disturbance variance $\hat{\theta}_\varepsilon$	9.57 (0.21)

NOTE: Numbers in parentheses are standard errors of parameter estimates.

pairs to decrease in closeness at precisely the same rate. People obviously differ, and this model does not take this into account. Another way to understand this lack of fit is to note that a fixed-intercept model implies that the between-person proportion of variance (the *intraclass correlation*, or ICC) is zero. To the extent that ICC deviates from zero in the data, the fixed-intercept model will fit more poorly. A random intercept model implies a nonzero ICC and so can account for some autocorrelation among the repeated measures. We address this issue first by relaxing the constraint on the intercept variance, or "freeing" the intercept (Model 3), and then by freeing both the intercept and slope (Model 4).

Model 3: Random Intercept, Fixed Slope

It is reasonable to suspect that not all mother–child pairs have the same level of closeness at Grade 1. However, the previous model with a fixed intercept was specified as if they did. Although we freed the intercept when we fit Model 1, we fixed it once grade was included in the form of slope factor loadings. A more realistic model would permit individuals to differ in intercept by permitting the intercept variance to be freely estimated. This slightly modified model represents interindividual variation in intercepts by estimating not only a mean intercept ($\hat{\alpha}_1$) but also an intercept variance ($\hat{\psi}_{11}$), indicating the degree to which individuals' intercepts vary about the population mean intercept. The parameter estimates from Model 3 (see Table 2.5) are similar to those from Model 2, with the average change in CLSN about –0.36 units per grade and an average intercept of approximately 38. The intercept variance ($\hat{\psi}_{11} = 5.37$) is significant, indicating that there is nontrivial variance between individuals in initial status. Allowing the intercept to be random across individuals—freeing only one parameter from Model 2—clearly improved model fit by a significant amount. As might be expected, this model fits much better than the previous fixed intercept model, with $\Delta\chi^2(1) = 1{,}797$.[5] However, there is still room for improvement, as indicated by the significant disturbance variance ($\hat{\theta}_\varepsilon = 4.21$). Just as it was reasonable to suppose that children vary at Grade 1 initial status, it is also reasonable to suppose that mother–child pairs vary in their rate of change over time.

TABLE 2.5
Model 3: Random Intercept, Fixed Slope

Parameter	*Estimate*
Mean intercept $\hat{\alpha}_1$	38.00 (0.10)
Mean slope $\hat{\alpha}_2$	–0.36 (0.02)
Intercept variance $\hat{\psi}_{11}$	5.37 (0.30)
Disturbance variance $\hat{\theta}_\varepsilon$	4.21 (0.10)

NOTE: Numbers in parentheses are standard errors of parameter estimates.

Note that the mean intercept and slope estimates in Model 3 are the same as those from Model 2; all we have done is permit the interindividual variability around the mean intercept. An interesting point is that the sum of the intercept variance and disturbance variance in Model 3 is equal to the disturbance variance from Model 2. In Model 2, the intercept variance was constrained to zero, forcing all the between-person variability to be expressed as disturbance variance. In the next model, both intercepts and slopes are permitted to vary.

Model 4: Random Intercept, Random Slope

Thus far, we have shown how to fit models with only fixed slopes. In general, however, there will be between-individual differences in both baseline level and in rate of change. Our previous models with fixed parameters each ignore some of this between-individual variability. To better reflect the nature of our data, we now specify both intercept and slope variance parameters to be random. In Model 4, every individual is allowed to have a different slope and a different intercept.[6] In addition to being more realistic in many contexts, this model allows estimation of the intercept–slope covariance (ψ_{21}). In our example, the intercept–slope covariance is interpreted as the degree to which mother–child closeness at Grade 1 is related to rate of change over time.

As in the previous models, the LGM representation of Model 4 is relatively simple. To specify this model, we free the variances of both intercept and slope factors and add an intercept–slope covariance parameter. Matrix $\mathbf{\Psi}$ therefore becomes

$$\mathbf{\Psi} = \begin{bmatrix} \psi_{11} & \\ \psi_{21} & \psi_{22} \end{bmatrix}. \qquad (2.9)$$

Estimating the parameters of Model 4 with LISREL yielded the estimates shown in Table 2.6. The intercept variance is large relative to the slope variance, and the intercept–slope covariance (0.25, corresponding to a correlation of 0.40) is significant. This indicates that children with higher intercepts have shallower negative slopes or that mother–child pairs who are closer at Grade 1 tend to decrease in closeness at a slower rate than those who are less close at Grade 1.

TABLE 2.6
Model 4: Random Intercept, Random Slope

Parameter	*Estimate*
Mean intercept $\hat{\alpha}_1$	38.00 (0.08)
Mean slope $\hat{\alpha}_2$	–0.36 (0.02)
Intercept variance $\hat{\psi}_{11}$	2.98 (0.29)
Slope variance $\hat{\psi}_{22}$	0.14 (0.02)
Intercept/slope covariance $\hat{\psi}_{21}$	0.25 (0.06)
Disturbance variance $\hat{\theta}_{\varepsilon}$	3.70 (0.10)

NOTE: Numbers in parentheses are standard errors of parameter estimates.

A chi-square difference test showed that this model with random intercept and random slope significantly improved fit over that of Model 3, which had only a random intercept, $\Delta\chi^2(2) = 221.5$, $p < .0001$. The RMSEA, NNFI, and SRMR fit indices indicate good fit as well. As noted, the significant positive covariance between intercept and slope implies that mother–child pairs who are closer in Grade 1 tend to experience less precipitous drops in closeness.

Returning to our comments about the interpretation of the intercept–slope covariance (ψ_{21}), before the reader concludes that children who have closer relationships with their mothers enjoy smaller decreases in closeness over the remaining elementary school years, we caution that the intercept–slope covariance might have been much less impressive had the zero point been defined at some other age. For example, if most children sampled had approximately identical mother–child closeness at kindergarten, then differences in slope alone would lead to the observed covariance when time is centered at Grade 1. Indeed, if the age variable is rescaled so that the intercept is defined at 2 years before Grade 1, a nonsignificant intercept/slope correlation of –.03 is obtained. This underscores the importance of the decision of where to place the origin of the time scale when fitting one's model.

Mehta and West (2000) note that if the linear LGM is an appropriate model for the data the true-score variance of the repeated measures will follow a quadratic pattern. Denoting time by t and the time origin by t^*,

$$\sigma^2_{\xi(t)} = \psi_{11} + \psi_{22}(t - t^*)^2 + 2\psi_{21}(t - t^*), \tag{2.10}$$

where $\sigma^2_{\xi(t)}$ is the true score variance at time t and ψ_{11}, ψ_{22}, and ψ_{21} are, respectively, the population intercept variance, the slope variance, and the intercept–slope covariance. In other words, the collection of individual trajectories will resemble a bow tie or fan spread. In many situations, the point at which the between-person variability is minimized (the "knot" of the bow tie) is of interest. This point, which Hancock and Choi (2006) have termed the *aperture*, can be easily calculated as the choice of time origin that minimizes the intercept variance; that is,

$$a = a^* - \frac{\hat{\psi}_{21}}{\hat{\psi}_{22}}, \tag{2.11}$$

where a^* is the occasion originally chosen for the time origin and $\hat{\psi}_{21}$ and $\hat{\psi}_{22}$ are, respectively, the estimated intercept–slope covariance and the slope variance.[7] The aperture is the point at which $\psi_{21} = 0$ (Mehta & West, 2000). Using the results of Model 4, the occasion at which children were most similar in mother–child closeness is $a = 0 - 0.2533/0.1366 = -1.85$, or nearly a year before kindergarten. Of course, caution in interpretation is warranted whenever the aperture falls outside the range of occasions for which data are observed, as it obliges the researcher to assume that a linear model is appropriate for those occasions. This may not be the case.

To summarize our progress so far, we have proceeded from estimating a null model through estimating a model with random intercept and slope. The random intercept model (Model 1) indicated that most variation occurred within individuals but that there also was a nonnegligible amount of variation between individuals. Model 2 showed what one might expect to find when between-individual variance is ignored and showed how constraining parameters to particular values can harm model fit when different individuals have different trajectories on the outcome measure. Model 3 built on Model 2 by freeing the intercept variance parameter. We found that allowing the intercept to vary across people resulted in a large improvement in model fit. The model with random intercepts and slopes (Model 4) performed much better than the previous, more constrained models (Models 2 and 3). In addition to improving overall fit, the advantages of Model 4 included discovering a significant positive covariance (ψ_{21}) between intercepts and slopes, implying that the rate at which closeness changes over time was related to closeness at Grade 1. We use Model 4 as the basis for all subsequent models. In Model 5, we examine the possibility that there are (child) gender differences in mother–child closeness trajectories.

Model 5: Multiple-Groups Analysis

The results from Model 4 imply that it is useful to think of individuals as having different intercepts and slopes. It is possible that some of this variance is systematically related to other variables of interest. For example, we can hypothesize and test group differences in mean intercept and slope for boys and girls (McArdle & Epstein, 1987). There are at least two ways to examine group differences in trajectories: (1) splitting the sample into two groups and estimating parameters in both groups simultaneously and (2) specifying the grouping variable as a predictor of both intercepts and slopes. Specifying a two-groups analysis in LISREL involves splitting the data file into two files based on gender and conducting a multisample analysis in which models are fit simultaneously to both data sets. The two models are specified in the same syntax file, and each model is as depicted in Figure 2.3. Equality constraints are imposed on key parameters in corresponding parameter matrices across groups to test hypotheses of group differences in those parameters. This multiple-groups method can be applied in principle to any number of groups, and models with different forms may be specified in the different groups. This method enables a novel approach to examining treatment effects and initial status × treatment interaction effects in longitudinal settings, as we discuss in the next chapter. Using group as a predictor variable is discussed in Model 6.

We follow the multiple-groups strategy for the closeness data, fitting separate models to boys and girls simultaneously with no cross-group constraints on model parameters. The results are reported in Table 2.7. NNFI was calculated using separately specified two-parameter null models for each group. Model fit is mediocre (RMSEA = .088; 90% confidence interval [CI] = {.072, .104}), permitting cautious interpretation of parameter estimates. All parameter estimates are significant for both boys and girls except for the intercept–slope covariance for girls ($\hat{\psi}_{21}$ = .14), which is notably lower than that for boys ($\hat{\psi}_{21}$ = .36). The correlations corresponding to these covariances are r = .22 (for girls) and r = .59 (for boys), indicating a much stronger relationship between initial status (at first grade) and change over time in mother–child closeness for boys than for girls.

Note that the intercept and the slope are lower for boys (37.85 and –0.38, respectively) than for girls (38.14 and –0.33, respectively), indicating that, on average, boys appear to start lower, and decrease at a faster pace, than girls. Given these apparent differences, if a researcher had substantive reason to test whether these differences were significant, one approach would be to constrain these parameters to equality across groups and look for a significant drop in model fit by means of a chi-square difference test. In this case, if an equality constraint is applied to intercept means (permitting slopes to vary), the result of the difference test is $\Delta\chi^2(1) = 3.11$, $p = .08$,

TABLE 2.7
Model 5: Multiple Groups Analysis

Parameter	*Estimate*
Boys	
Mean intercept $\hat{\alpha}_1$	37.85 (0.12)
Mean slope $\hat{\alpha}_2$	–0.38 (0.03)
Intercept variance $\hat{\psi}_{11}$	3.03 (0.42)
Slope variance $\hat{\psi}_{22}$	0.12 (0.03)
Intercept/slope covariance $\hat{\psi}_{21}$	0.36 (0.08)
Disturbance variance $\hat{\theta}_{\varepsilon}$	3.78 (0.15)
Girls	
Mean intercept $\hat{\alpha}_1$	38.14 (0.11)
Mean slope $\hat{\alpha}_2$	–0.33 (0.03)
Intercept variance $\hat{\psi}_{11}$	2.90 (0.39)
Slope variance $\hat{\psi}_{22}$	0.15 (0.03)
Intercept/slope covariance $\hat{\psi}_{21}$	0.14 (0.08)
Disturbance variance $\hat{\theta}_{\varepsilon}$	3.63 (0.14)

NOTE: Numbers in parentheses are standard errors of parameter estimates.

indicating that the intercepts are not significantly different. If the equality constraint is instead applied to the slope means (permitting intercepts to vary), the difference test again is nonsignificant, $\Delta\chi^2(1) = 1.47$, $p = .23$, indicating that the slopes are not significantly different. Despite appearances, there is not enough evidence to suggest that boys and girls follow different linear trajectories in mother–child closeness from Grades 1 through 6. Other cross-group constraints are possible as well, of course—theory may suggest testing for differences in disturbance variances or intercept–slope covariances (ψ_{21}). A test of the hypothesis that ψ_{21} is equal for boys and girls also is inconclusive, $\Delta\chi^2(1) = 3.83$, $p > .05$.

For examples and further discussion of multiple-groups LGM, see Curran, Harford, and Muthén (1996), Curran, Muthén, and Harford (1998), McArdle (1989), and McArdle and Epstein (1987). We close this section by noting that we assume group membership is known (observable). If groups are assumed, but membership is uncertain, the researcher may be interested in growth mixture modeling, discussed in Chapter 3. Next, we illustrate an alternative method of investigating gender differences in closeness trajectories.

Model 6: The Conditional Growth Curve Model

Instead of running multiple-groups models, the analysis of gender differences in intercept and slope may be considerably simplified by including gender as an exogenous predictor (Level 2 predictor or *time-invariant* *covariate*) of both intercept and slope in a single-group analysis. Such predictors of random variables can be introduced to account for between-individual variance in the estimates of intercept and slope. The variance parameters were freed in Models 3 and 4. Because the intercept and slope variances can be considered unexplained individual differences, they potentially can be accounted for using this technique, which can equally well use categorical or continuous Level 2 predictors. This type of model is sometimes termed a *conditional LGM* (Tisak & Meredith, 1990; Willett & Sayer, 1994), whereas Models 1 to 5 could be termed *unconditional* (Singer & Willett, 2003). Although we chose only gender as a predictor, we could choose any number of variables, for example, variables designed to measure between-individual differences in ethnicity, socioeconomic status, or parental religiosity. A significant advantage of Model 6 over Model 5 is that, because it is unnecessary to divide the sample into groups, the time-invariant covariate may be either nominal (e.g., gender in this case) or continuous. A disadvantage is that the researcher must be able to assume invariance of some model parameters across groups. For example, whereas we were free to estimate different disturbance variances for boys and girls in Model 5, we are not free to do so in Model 6 without significant additional effort.

We demonstrate the use of a grouping variable (gender) as a predictor in Model 6. In this model, we include gender as a predictor of intercept with a fixed coefficient, β_1, which is interpreted as the mean effect of gender on intercept.[8] Gender is also included as a predictor of slope, with fixed coefficient β_2. When a time-invariant covariate is included as a predictor of slopes, the effect is often called a *cross-level interaction* because time (Level 1) interacts with the covariate (Level 2) to predict the repeated measures (Cronbach & Webb, 1975; Curran, Bauer, & Willoughby, 2004; Kreft & de Leeuw, 1998; Raudenbush & Bryk, 2002). Note that if the effect of age on mother–child closeness varies across individuals, and if gender partially explains interindividual variability in that effect, then the cross-level interaction has the same interpretation as a moderation effect in traditional multiple regression analysis.[9]

Figure 2.4 contains a path diagram including gender as a predictor of both intercept and slope. The variances of intercept and slope, as well as their covariance, were reconceptualized as residual variances and a residual covariance—that is, that portion of the variance and covariance not accounted for by gender. These residual parameters are depicted in Figure 2.4 as ψ_{11}, ψ_{22}, and ψ_{21}. The parameters β_1 and β_2 represent the effects of gender on intercept and slope, respectively. Results for this model are presented in Figure 2.4 and Table 2.8. NNFI was calculated using Model 0 as a null model, augmented by estimating the mean and variance of gender.

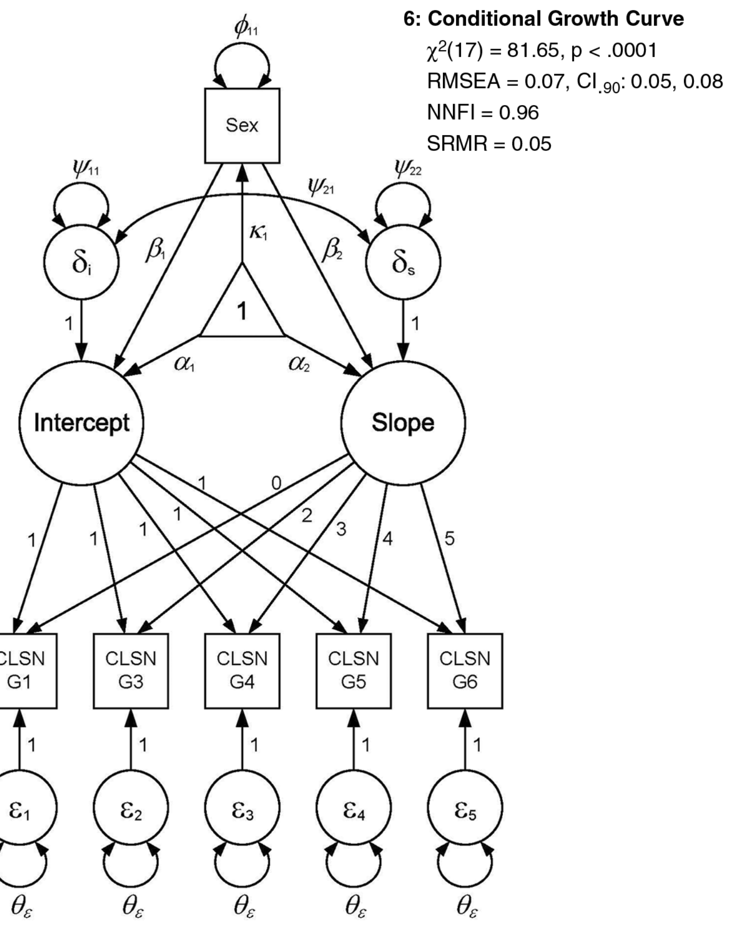

Figure 2.4 A Path Diagram Representing Gender as a Predictor of Individual Differences in Both Intercept and Slope.

NOTE: CLSN = Closeness with child.

Results for Model 6 show that the coefficient for predicting intercept from gender was –0.29; that is, mother–daughter pairs (coded 0) have higher average closeness scores than mother–son pairs (coded 1) at first grade, but not significantly higher. Mother–daughter pairs also showed less change over time than mother–son pairs, but not significantly less change. As we might expect, these results are consistent with those obtained in Model 5. For example, the difference between mean intercepts for girls and boys in Model 5 (see Table 2.7) was $37.85 - 38.14 = -0.29$, equal to the

TABLE 2.8
Model 6: Conditional Growth Curve Model

Parameter	*Estimate*
Mean intercept $\hat{\alpha}_1$	38.14 (0.11)
Mean slope $\hat{\alpha}_2$	–0.33 (0.03)
Group effect on intercept $\hat{\beta}_1$	–0.29 (0.16)
Group effect on slope $\hat{\beta}_2$	–0.05 (0.04)
Intercept variance $\hat{\psi}_{11}$	2.96 (0.28)
Slope variance $\hat{\psi}_{22}$	0.14 (0.02)
Intercept/slope covariance $\hat{\psi}_{21}$	0.25 (0.06)

NOTE: Numbers in parentheses are standard errors of parameter estimates.

estimated fixed effect of gender, $\hat{\beta}_1$, in Model 6 (see Table 2.8). Similarly, the difference in slopes in Model 5 is –0.05, equal to the fixed effect of gender on slope, $\hat{\beta}_1$, obtained in Model 6. Because Models 5 and 6 are not completely identical (the disturbance variances are different across groups), results may not always resemble each other so closely. We note that this could easily be altered by imposing a cross-group equality constraint.

Treating time-invariant covariates as predictors of growth factors may lead to difficulties. The conditional growth curve model may be understood as a mediation model in which the growth factors are hypothesized to completely mediate the effect of the time-invariant covariate on the outcome variables, but complete mediation is rarely a tenable hypothesis. In Model 6, for example, by omitting paths linking child gender directly to closeness, the direct effects are implicitly constrained to zero. If these zero constraints are inappropriate, model fit will suffer and estimated parameters likely will be biased. An alternative strategy is to relax constraints on the direct effects (therefore estimating paths linking the covariate directly to each repeated measure) and instead constrain $\hat{\beta}_1$ and $\hat{\beta}_2$ to zero (Stoel, van den Wittenboer, & Hox, 2004). However, this model may not address a question of substantive interest. In addition, it should be borne in mind that interpretation of the effects of exogenous predictors on the intercept factor will vary with the scaling of time and the location of the time origin (Stoel, 2003; Stoel & van den Wittenboer, 2003).

Model 7: Parallel Process Model

It is possible to investigate the relationship between aspects of change specific to each of two repeated-measures variables, modeling growth processes in more than one variable. This procedure permits examination of relationships among aspects of change for different variables (McArdle, 1989). For example, a researcher may be interested in modeling growth in mother–child conflict and father–child conflict simultaneously to examine the relationship between the intercept of one and the slope of the other. This kind of model, referred to variously as a *parallel process model* (Cheong, MacKinnon, & Khoo, 2003), *multivariate change model* (MacCallum et al., 1997), *cross-domain individual growth model* (Sayer & Willett, 1998; Willett & Sayer, 1994, 1995), *multiple-domain model* (Byrne & Crombie, 2003), *fully multivariate latent trajectory model* (Curran & Hussong, 2003; Curran & Willoughby, 2003), *simultaneous growth model* (Curran et al., 1996); *bivariate growth model* (Aber & McArdle, 1991), or *associative LGM* (S. C. Duncan & Duncan, 1994; T. E. Duncan, Duncan, & Strycker, 2006; T. E. Duncan et al., 1999; Tisak & Meredith, 1990), follows easily from a simple random intercept, random slope model. A parallel process model contains two sets of intercepts and slopes, one set for each repeated-measures variable. The covariances among the intercepts and slopes are estimated as well (see Figure 2.5). For this example, we used mother–child closeness and mother–child conflict as the two dependent variables (see Table 2.9). The sample consisted of 849 children (433 girls and 416 boys) having complete data for both mother–child closeness and mother–child conflict. Model fit information can be found in Figure 2.5. NNFI was calculated by specifying as a null model a fixed intercept model for both closeness and conflict, permitting occasion-specific disturbances to covary.

Table 2.9 contains some interesting results, probably few of which will come as a surprise to parents of elementary-school-age children. First, parameter estimates related to mother–child closeness are, as expected, nearly identical to those from Model 4; any discrepancies can be attributed to the fact that the sample is slightly smaller for Model 7 due to missing data. Both closeness and conflict change over time, but in opposite directions, as indicated by their mean slope estimates. The estimates for the covariances among the intercepts and slopes are reported in the "Curve covariances" section of Table 2.9, and the correlations implied by these covariances are in the section "Curve correlations." The covariance of the intercepts ($\hat{\psi}_{31}$) is significantly negative, indicating that children who were particularly close to their mothers in first grade were also those who experienced the least conflict. Similarly, the slope covariance ($\hat{\psi}_{42}$) is significantly negative, indicating that children characterized by steeper decline in closeness tended to be those who experienced accelerated conflict as they aged. Finally, conflict intercepts were negatively associated with closeness

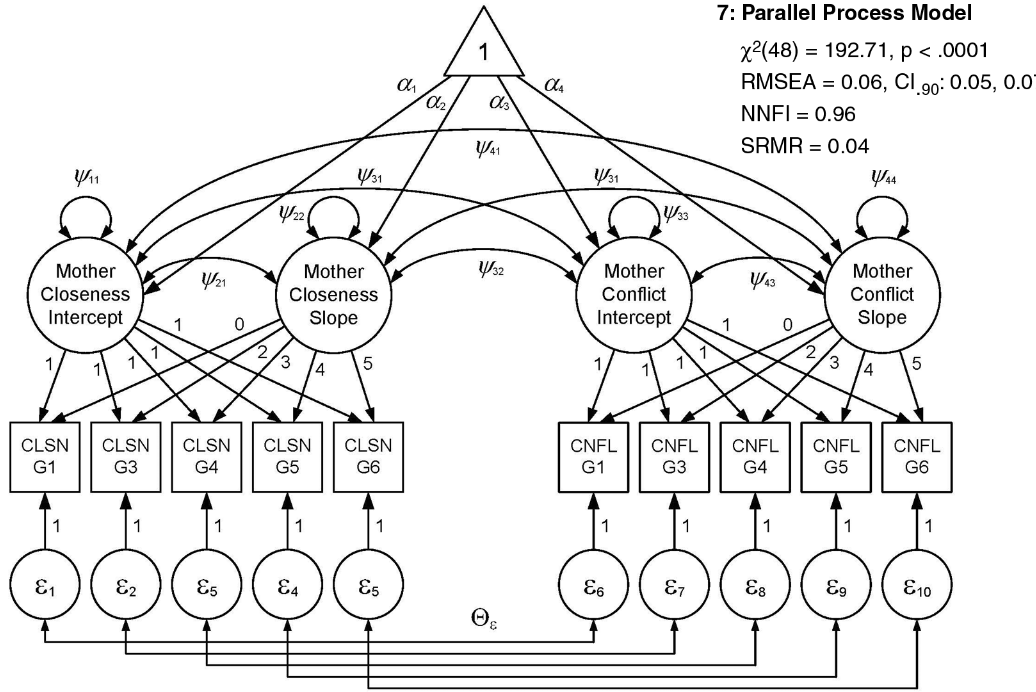

Figure 2.5 A Path Diagram Representing Simultaneous Growth in Mother–Child Closeness and Mother–Child Conflict.

NOTE: Intercepts and slopes from both variables are allowed to covary, as are disturbances for closeness and conflict measured at corresponding occasions. The disturbance variances for closeness were constrained to equality over time, as were the disturbance variances for conflict and the closeness–conflict disturbance covariances.

CLSN = Closeness with child; CNFL = Conflict with child.

TABLE 2.9
Model 7: Parallel Process Model

Parameter	*Estimate*
Mother–child closeness	
Mean intercept $\hat{\alpha}_1$	38.00 (0.08)
Mean slope $\hat{\alpha}_2$	–0.36 (0.02)
Intercept variance $\hat{\psi}_{11}$	2.99 (0.29)
Slope variance $\hat{\psi}_{22}$	0.14 (0.02)
Intercept/slope covariance $\hat{\psi}_{21}$	0.25 (0.06)
Mother–child conflict	
Mean intercept $\hat{\alpha}_3$	15.24 (0.20)
Mean slope $\hat{\alpha}_4$	0.30 (0.04)
Intercept variance $\hat{\psi}_{33}$	25.74 (1.61)
Slope variance $\hat{\psi}_{44}$	0.40 (0.06)
Intercept/slope covariance $\hat{\psi}_{43}$	–0.57 (0.23)
Curve covariances	
Intercept covariance $\hat{\psi}_{31}$	–3.88 (0.51)
CLSN intercept/CNFL slope covariance $\hat{\psi}_{41}$	0.13 (0.09)
CNFL intercept/CLSN slope covariance $\hat{\psi}_{32}$	–0.41 (0.12)
Slope covariance $\hat{\psi}_{42}$	–0.07 (0.02)
Curve correlations	
Intercept covariance	–0.44
CLSN intercept/CNFL slope covariance	0.12
CNFL intercept/CLSN slope covariance	–0.22
Slope covariance	–0.28

NOTE: Numbers in parentheses are standard errors of parameter estimates.
CLSN = Closeness with child; CNFL = Conflict with child.

slopes ($\hat{\psi}_{32}$), meaning that those first graders who demonstrated relatively more conflict with their mothers tended to experience more precipitous decreases in closeness as they got older.

Extensions to the basic parallel process model are possible. If the slopes associated with repeated measures of two variables are hypothesized to be not merely related but *causally* related, directional paths among growth factors may be specified. For example, Curran, Stice, and Chassin (1997) use a parallel process model in which both adolescent alcohol use and peer alcohol use change linearly over time. Age, gender, and parental alcoholism are used to predict aspects of change, and individual differences in intercepts from each repeated-measures variable (scaled to be at the initial measurement occasion) are used to predict variability in the other repeated-measures variable. Curran and Hussong (2002) model parallel growth in antisocial behavior and reading ability in children, predicting reading slopes with antisocial intercepts. Curran et al. (1996) specified a model in which the intercepts of alcohol use and bar patronage were hypothesized to affect one another's slopes. Raudenbush, Brennan, and Barnett (1995) used a similar approach to model predictors of simultaneous change in judgments of marital quality in husband/wife dyads, where each member of a dyad was measured at three yearly intervals. In addition, parallel processes in more than two repeated-measures variables may be specified.

Model 8: Cohort-Sequential Designs

Both the cross-sectional and longitudinal approaches potentially suffer from shortcomings when used to assess trajectories in single samples. Cross-sectional designs are sometimes prone to cohort or history effects that may mislead researchers into thinking a trend exists when one does not or masking a trend that actually exists. Longitudinal designs, on the other hand, are sometimes compromised by the threat of contamination due to repeated measurement of the same individuals. *Cohort-sequential designs* (Meredith & Tisak, 1990; Nesselroade & Baltes, 1979; Schaie, 1965, 1986; Tisak & Meredith, 1990), also called *accelerated longitudinal designs* (Miyazaki & Raudenbush, 2000; Raudenbush & Chan, 1992; Tonry, Ohlin, & Farrington, 1991) or the *method of convergence* (Bell, 1953, 1954; McArdle, 1988), have been suggested as a way to reduce the threat of these potential confounds by combining the longitudinal and cross-sectional approaches to examining developmental change. Cohort-sequential designs also greatly collapse the time needed to conduct longitudinal studies and reduce problems of attrition (Tonry et al., 1991). Consider the case in which age is the metric of time. Rather than follow the same sample of high school freshmen for 8 years through college, a researcher may instead elect to follow three cohorts (high school freshmen, high school juniors, and college freshmen) for only 4 years each. By employing multiple cohorts of subjects and measuring at only a few occasions within each cohort, a full trajectory for the entire time range of interest can be obtained. The cohort-sequential design is more appropriately thought of as an efficient data collection strategy than as a model, although this strategy leads to some interesting modeling options.

To demonstrate analysis of cohort-sequential data in our example, we created two artificial cohorts. We first randomly separated our data into two groups and then deleted data to mimic the pattern of data that might be gathered in a true cohort-sequential design. For Cohort 1, we deleted all measurements for children in Grade 6, and for Cohort 2, we deleted all measurements for children in Grade 1. This resulted in a data set in which children in Cohort 1 provided data for Grades 1, 3, 4, and 5, and children in Cohort 2 provided data for Grades 3 through 6. In randomly assigning cases to each cohort and deleting a portion of the data, our sample size was reduced to 893 (n_1 = 426, n_2 = 467). In practice, of course, data from these two cohorts would be collected concurrently over a single span of 5 years. Even less overlap is probably acceptable.

Cohort-sequential data can be examined in either of two ways. First, the researcher can aggregate the data and perform a single-group analysis. This approach involves treating uncollected data for each cohort as MCAR, which ordinarily is a safe assumption because such data are missing by design (T. E. Duncan et al., 1999; B. Muthén, 2000). Alternatively, the researcher can consider the cohorts as separate groups and perform a multiple-groups analysis (McArdle & Hamagami, 1992). This approach is similar to Model 5, with data separated by cohort and all corresponding parameters constrained to equality across cohorts. The multiple-groups option derives from one approach to dealing with missing data in which a relatively small number of "missingness" patterns are identifiable (Allison, 1987; T. E. Duncan et al., 1999; McArdle & Bell, 2000; McArdle & Hamagami, 1991; B. Muthén et al., 1987). B. Muthén (2000) described the implementation of cohort-sequential designs using these two approaches. It should be noted that treating cohorts as separate groups may lead to estimation problems in some circumstances (T. E. Duncan et al., 1999). With small groups, there may even be more measurement occasions than subjects, resulting in the necessary removal of some data, leaving some measurement occasions unrepresented. Here, our sample is quite large, so this issue is not a concern. We demonstrate both the single-group and multiple-groups approaches.

The specifications for the single-group approach are the same as for Model 4, and the path diagram is therefore that in Figure 2.3. The only difference here is that all subjects in Cohort 1 have missing values for Grade 6, and those in Cohort 2 have missing values for Grade 1. The results (see Table 2.10) are similar to those obtained for Model 4, although, as one would expect given the smaller sample size and the missing data, the standard errors are slightly larger. In spite of this, however, the results support the same conclusions drawn from analysis of the larger sample. The estimated mean closeness score at first grade is $\hat{\alpha}_1$ = 37.99, and the mean slope is $\hat{\alpha}_2$ = −0.35. Both are similar to the Model 4 estimates, as were the variances and covariance of the intercept and the slope.

The multiple-groups cohort-sequential approach is similar to the two-group analysis demonstrated earlier (Model 5), with all parameters constrained to equality across cohorts, including intercept and slope means. It is important to exercise care in specifying the equality constraints. In our analysis, each data file contains only those variables for which participants provide data. Therefore, the data file for Cohort 1 has only four variables,

TABLE 2.10
Model 8: Cohort-Sequential Design

Parameter	*Single-Group Estimate*	*Multiple-Group Estimate*
Mean intercept $\hat{\alpha}_1$	37.99 (0.10)	37.95 (0.13)
Mean slope $\hat{\alpha}_2$	–0.35 (0.03)	–0.36 (0.04)
Intercept variance $\hat{\psi}_{11}$	3.73 (0.43)	3.73 (0.43)
Slope variance $\hat{\psi}_{22}$	0.10 (0.03)	0.10 (0.03)
Intercept/slope covariance $\hat{\psi}_{21}$	0.21 (0.10)	0.21 (0.10)
Disturbance variance $\hat{\theta}_{\varepsilon}$	3.77 (0.12)	3.78 (0.12)

NOTE: Numbers in parentheses are standard errors of parameter estimates.

corresponding to measurements at Grades 1, 3, 4, and 5. Similarly, the data file for Cohort 2 contains measurements for only Grades 3 through 6. The variable for closeness at Grade 4 in Cohort 1 is the third repeated measure, whereas for Cohort 2 it is the second. The factor loading matrices for the two cohorts, then, are

$$\mathbf{\Lambda}_1 = \begin{bmatrix} 1 & 0 \\ 1 & 2 \\ 1 & 3 \\ 1 & 4 \end{bmatrix}, \tag{2.12}$$

$$\mathbf{\Lambda}_2 = \begin{bmatrix} 1 & 2 \\ 1 & 3 \\ 1 & 4 \\ 1 & 5 \end{bmatrix}. \tag{2.13}$$

The corresponding path diagram is shown in Figure 2.6.

The resulting parameter estimates on the right-hand side of Table 2.10 are highly similar to those obtained from the single-group analysis. For additional examples of the multiple-groups cohort-sequential approach, see reports by Aber and McArdle (1991), Anderson (1993), Baer and Schmitz (2000), Buist, Deković, Meeus, and van Aken (2002), T. E. Duncan, Duncan, and Hops (1993), S. C. Duncan, Duncan, and Hops (1996), and McArdle and Anderson (1990).

Note that NNFI is not provided for the single-group cohort-sequential design because some data are treated as missing. To permit the use of ML estimation with access to all ML-based fit indices, we recommend using the multiple-groups approach if complete data are available within each group, using covariances and means as input data for each cohort. If some cases

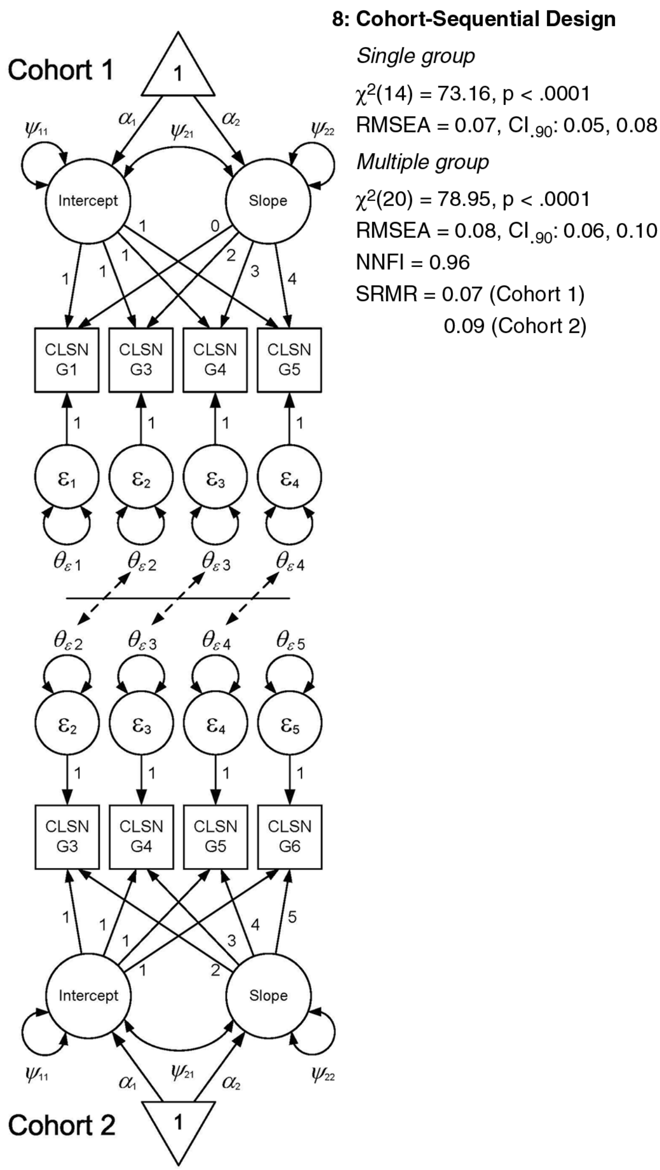

Figure 2.6 A Path Diagram Representing a Multiple-Groups Cohort-Sequential Design.

NOTE: The horizontal line represents the division between the two simultaneously estimated models. Parameters with identical labels and subscripts are constrained to equality across models.

CLSN = Closeness with child.

have missing data, we recommend using FIML to take advantage of missing data techniques that use all available data. It is rarely appropriate to discard data. In addition, if it is desired to estimate disturbance variances separately within each cohort, the multiple-groups approach is preferable.

The multiple-groups approach to testing cohort-sequential designs can be used explicitly to test for cohort effects, such as cohort differences in mean intercept or slope, by means of $\Delta\chi^2$ tests (Anderson, 1993; Meredith & Tisak, 1990). The single-group approach also may be used to test for cohort effects by including cohort as a dummy-variable predictor of intercept and slope (Raudenbush & Chan, 1992) and noting significant effects, but this model imposes the assumption of equal disturbance variances across cohorts.[10]

Model 9: Time-Varying Covariates

Earlier we defined TVCs as variables measured repeatedly and used to predict repeated measures of an outcome variable. Very little attention has been devoted to the treatment of TVCs in the LGM context, although much has been written on the subject of TVCs in the context of multilevel modeling. There are two ways to conceive of TVCs in LGM. These two methods address subtly different questions. The first, suggested by B. Muthén (1993) and illustrated or used in articles by George (2003), Bijleveld and van der Kamp (1998), and Curran and colleagues (Curran & Hussong, 2002, 2003; Curran & Willoughby, 2003; Curran et al., 1998; B. O. Muthén & Curran, 1997), is to include TVCs directly in the model as repeated exogenous predictors of the outcome, as in Figure 2.7. The β parameters in this model are interpreted as occasion-specific effects of the covariate, or as the ability of the covariate to predict occasion-specific deviations in the outcome. In this approach, the effect of a TVC may vary across time, but not across individuals. Alternatively, the β parameters could be constrained to equality to represent the hypothesis that the covariate effect remains stable over occasions. Either way, this model reflects growth in the repeated-measure variable controlling for occasion-specific effects of the TVC.

We fit the model in Figure 2.7 to the closeness data, treating mother–child conflict as a TVC measured concurrently with closeness. The results are reported in Table 2.11 and in Figure 2.7. NNFI was computed by augmenting the null model in Model 0 by estimating the means and variances of the five repeated measures of the TVC. The time-specific effect of conflict on closeness remained between –0.15 and –0.13 for all examined grades, indicating that, at each occasion, conflict tended to be negatively related to closeness after partialing out individual differences accounted for by the intercept and slope factors. To claim that this relationship is *causal* would be unjustified without first satisfying the criteria for establishing a causal relationship.

An alternative method of including TVCs in growth curve models makes use of definition variables. This approach differs somewhat from standard practice in SEM, but is equivalent to standard methods of including TVCs in multilevel modeling (e.g., Raudenbush & Chan, 1993). The variable *time* (or *age*, etc.) itself can be considered a TVC because time varies across repeated measures of the outcome. If definition variables can accommodate

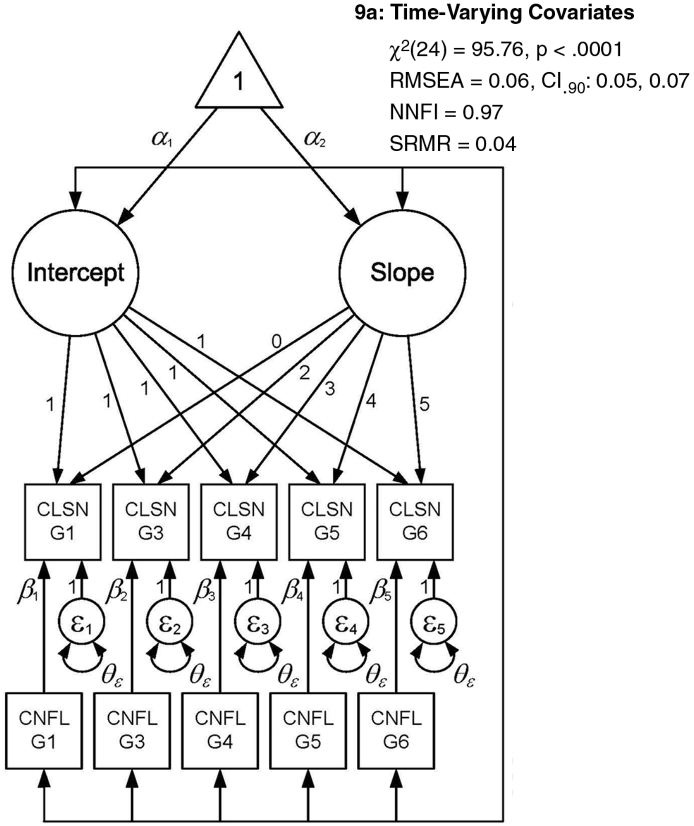

Figure 2.7 A Growth Curve Model Including a Time-Varying Covariate of the Sort Suggested by B. Muthén (1993).

NOTE: Although not individually depicted, all variances and covariances among the intercept factor, slope factor, and TVC variables are estimated.

CLSN = Closeness with child; CNFL = Conflict with child.

TABLE 2.11
Model 9a: Time-Varying Covariates

Parameter	*Estimate*
Mean intercept $\hat{\alpha}_1$	40.17 (0.41)
Mean slope $\hat{\alpha}_2$	–0.36 (0.12)
Intercept variance $\hat{\psi}_{11}$	2.41 (0.25)
Slope variance $\hat{\psi}_{22}$	0.13 (0.02)
Intercept/slope covariance $\hat{\psi}_{21}$	0.20 (0.05)
Disturbance variance $\hat{\theta}_\varepsilon$	3.52 (0.10)
Conflict effect (G1) $\hat{\beta}_1$	–0.15 (0.03)
Conflict effect (G3) $\hat{\beta}_2$	–0.13 (0.01)
Conflict effect (G4) $\hat{\beta}_3$	–0.13 (0.01)
Conflict effect (G5) $\hat{\beta}_4$	–0.13 (0.01)
Conflict effect (G6) $\hat{\beta}_5$	–0.13 (0.02)

NOTE: Numbers in parentheses are standard errors of parameter estimates.

different occasions of measurement across individuals, they can also be used to model the effects of any other TVC in a similar manner. An additional slope factor is added to the model to represent the TVC. Values of the TVC are inserted into individual data vectors, essentially giving each individual or Level 2 unit a unique $\mathbf{\Lambda}_y$ matrix. In this approach, the effect of a TVC can vary across individuals, but not across time. That is, a mean effect of the TVC across individuals is estimated and, if desired, the variance of the TVC's slope factor (and covariances with other aspects of change) also can be estimated. This approach requires raw data as input.

For example, consider the model in Figure 2.8. At each occasion of measurement (grade), each individual also provided data on mother–child conflict (the TVC). Three individual data vectors are illustrated. Numbers in diamonds represent contents of each individual data vector for the TVC. Similar notation can be found in Mehta and West (2000, p. 34). Using this method, multiple TVCs may be included by using a different slope factor for each covariate. Interactions among TVCs (e.g., the interaction between grade and mother–child conflict) may be investigated by including slope factors containing loadings equal to the products of slope loadings for the involved covariates. For example, in Figure 2.8, the loadings on the interaction factor for Person 1 would be 0, 20, 21, 72, and 80. Cross-level interactions between time-invariant and time-varying covariates may be specified by including predictors of TVC slope factors. Models specified this way are equivalent to multilevel models with Level 1 predictors.

Fitting the model in Figure 2.8 to our closeness and conflict data yielded the results reported in Table 2.12. Because raw data were used as input, fit indices such as RMSEA and NNFI are not provided. Because Mx was used

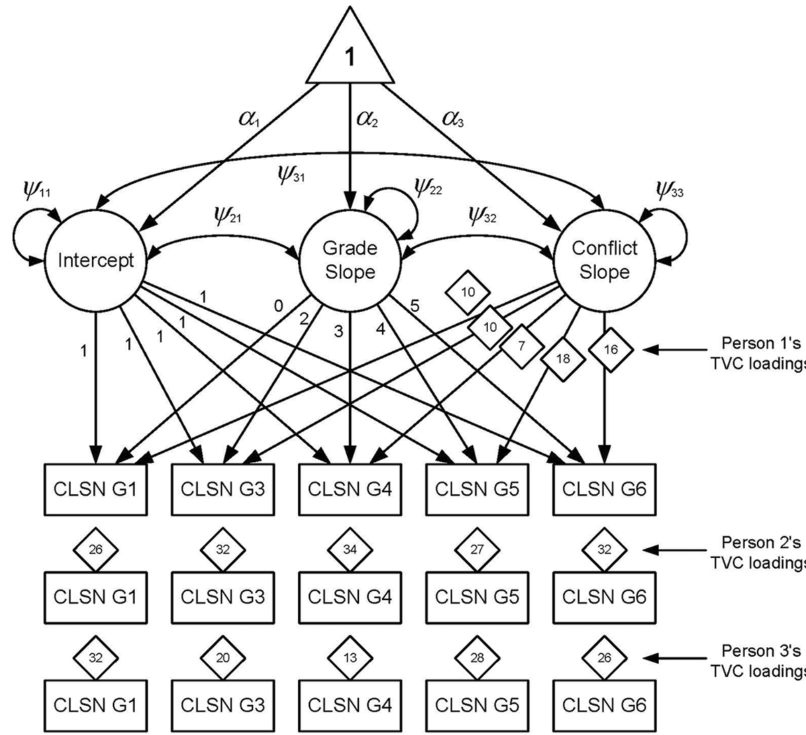

Figure 2.8 A Growth Curve Model Including a Time-Varying Covariate of the Sort Inspired by Mehta and West's (2000) Treatment of Time.

NOTE: A growth model may also be applied to the repeated measures of the covariate. Disturbance variances are present but not illustrated.

CLSN = Closeness with child; TVC = Time-varying covariate.

for this analysis, we are able to provide likelihood-based 95% confidence intervals[11] for each parameter estimate rather than the usual standard errors (Mx does not provide standard errors as a default option and, in fact, warns against their use with parameters having nonnormal sampling distributions, such as variances). All parameter estimates, save two, are statistically significant.

TABLE 2.12
Model 9b: Time-Varying Covariates

Parameter	*Estimate*	*95% Confidence Interval*
Mean intercept $\hat{\alpha}_1$	40.12	{39.82, 40.43}
Mean grade slope $\hat{\alpha}_2$	–0.31	{–0.35, –0.27}
Mean conflict slope $\hat{\alpha}_3$	–0.14	{–0.16, –0.12}
Intercept variance $\hat{\psi}_{11}$	4.93	{3.40, 6.72}
Grade slope variance $\hat{\psi}_{22}$	0.11	{0.08, 0.15}
Conflict slope variance $\hat{\psi}_{33}$	0.02	{0.02, 0.03}
Intercept/grade slope covariance $\hat{\psi}_{21}$	0.11	{–0.08, 0.29}
Intercept/conflict slope covariance $\hat{\psi}_{31}$	–0.28	{–0.39, –0.19}
Grade/conflict slope covariance $\hat{\psi}_{32}$	0.00	{–0.01, 0.01}
Disturbance variance $\hat{\theta}_\varepsilon$	3.28	{3.09, 3.48}

Model 10: Polynomial Growth Curves

The trajectories we have examined with growth curve models have been simple linear functions of *time* or *age.* The LGM user is not limited to linear functions, however. The framework presented thus far can accommodate any trajectories that are *linear in parameters* and *nonlinear in variables.* That is, basic LGM models can accommodate any trajectory in which the parameters of growth act as simple linear weights associated with transformations of the time metric. A common example is the quadratic latent growth curve (MacCallum et al., 1997; Meredith & Tisak, 1990; Stoolmiller, 1995) illustrated in Figure 2.9. In Figure 2.9, the loadings associated with the quadratic slope factor are the squares of the loadings associated with the linear slope factor. The mean of the quadratic slope (α_3) represents the degree of quadratic curvature in the trajectory.

We fit the quadratic growth curve in Figure 2.9 to the closeness data. Results are reported in Table 2.13 and Figure 2.9. The mean of the quadratic component is not significant ($\hat{\alpha}_3 = -0.019$, $p = .077$). A χ^2 difference test reveals that, compared with a purely linear model, the improvement in model fit is negligible ($\Delta\chi^2(1) = 3.12$, $p = .078$). Had there been a theoretical motive to do so, we might also have chosen to permit the quadratic

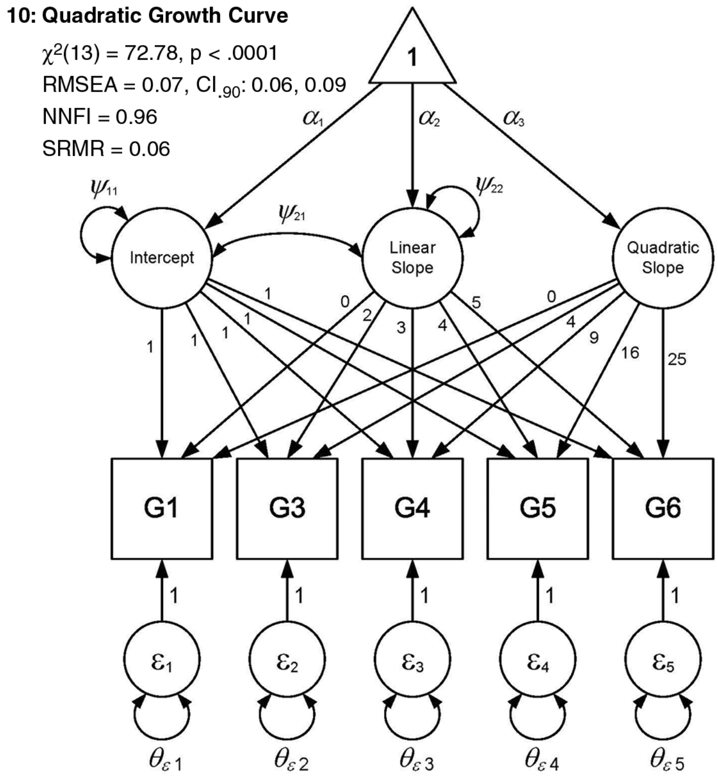

Figure 2.9 A Latent Growth Curve Model Including a Fixed Quadratic Slope Factor.

slope factor to vary randomly and to covary with the intercept and linear slope factors (Willett & Sayer, 1994) despite the lack of a mean quadratic effect. In fact, doing so results in a significant improvement in fit, but it is likely that the added complexity, in the form of three additional free parameters, overfits the data by absorbing random variability.

All too frequently, we suspect, quadratic growth curves are fit when a linear LGM does not fit adequately. We caution against this use of a quadratic model as capitalizing on possibly idiosyncratic characteristics of the particular sample under scrutiny. We suspect that there are few theories in the social sciences that naturally lead to predictions of nonlinear change that specifically imply a quadratic trend. Assuming there are enough repeated measures so that additional variance and covariance parameters will be

TABLE 2.13
Model 10: Polynomial Growth Curve

Parameter	*Estimate*
Mean intercept $\hat{\alpha}_1$	37.94 (0.09)
Mean linear slope $\hat{\alpha}_2$	–0.26 (0.06)
Mean quadratic slope $\hat{\alpha}_3$	–0.02 (0.01)
Intercept variance $\hat{\psi}_{11}$	2.98 (0.29)
Linear slope variance $\hat{\psi}_{22}$	0.14 (0.02)
Intercept/linear covariance $\hat{\psi}_{21}$	0.25 (0.06)
Disturbance variance $\hat{\theta}_\varepsilon$	3.70 (0.10)

NOTE: Numbers in parentheses are standard errors of parameter estimates.

identified, any number of polynomial growth factors may be added. However, a proper theoretical rationale must exist for adding these aspects of change.

It is possible to specify functional forms other than polynomial curves. Later we discuss structured latent curves, extensions to traditional growth curve models that can accommodate more complex functional forms that are *nonlinear in parameters*, in which the parameters of growth are no longer necessarily simple linear weights. One example is the exponential function commonly used to model population growth.

Model 11: Unspecified Trajectories

In the models described to this point, the function relating the outcome variable to time was completely defined. For example, in Figure 2.3, the paths from the slope factor to the measured variables are fixed in a linear progression, from 0 through 5, corresponding to a linear influence of grade on mother–child closeness. A creative extension of LGM involves the creation of *shape factors*, aspects of change for which the shape of the growth function (and therefore the factor loadings) are unknown and must be estimated from the data rather than specified a priori by the researcher (Meredith & Tisak, 1990). For example, we could replace the linear factor with a shape factor in Model 4, constrain the first and last loadings for this new factor to 0 and 1,[12] respectively, and estimate the remaining three loadings (λ_{22}, λ_{32}, and λ_{42}) using 0 and 1 as anchors. The estimated loadings would reveal the shape of the longitudinal trend. Although the free loadings are not proportions per se, even roughly linear growth should result in loadings that monotonically increase from 0 to 1. Alternatively, the first two loadings can be constrained to 0 and 1, in which case subsequent intervals can be interpreted by using the change occurring between the first two occasions as a benchmark (Hancock & Lawrence, 2006; Stoel, 2003). Models with shape factors are sometimes called *completely latent* (Curran & Hussong, 2002, 2003; McArdle, 1989), *fully latent* (Aber & McArdle, 1991), or *unspecified* (T. E. Duncan et al., 2006; Lawrence & Hancock, 1998; Stoolmiller, 1995; J. Wang, 2004) because the trajectory has not been specified a priori. This model is more exploratory than previously discussed models in that the researcher is not testing hypotheses about specific trajectories. Rather, the data are used to gain insight into what kind of trajectory might be appropriate. Freeing some of the loadings on a linear slope factor to create a shape factor allows direct comparison of the two models using a nested-model χ^2 difference test, essentially a test of departure from linearity. Good examples of unspecified trajectory models are provided by T. E. Duncan et al. (1993), T. E. Duncan, Tildesley, Duncan, and Hops (1995), and McArdle and Anderson (1990).

We applied an unspecified trajectory model to our mother–child closeness data, anchoring the first and fifth loadings to 0 and 5, respectively, to mirror our previous scaling of time (see Figure 2.10). A χ^2 difference test comparing Model 11 with Model 4 resulted in a nonsignificant difference ($\Delta\chi^2(3) = 2.68$, $p = .44$), indicating that a linear trend is sufficient to model the mother–child closeness data. This conclusion is bolstered by the fact that the loadings followed a nearly perfectly linear trend even without being constrained, as illustrated in Figure 2.11.

Summary

In this chapter, we described several latent growth curve models. Beginning with a basic null model, each model was applied in turn to the same data set to illustrate use of the models in practice. Beyond the basic linear LGM with random intercepts and random slopes, we showed how the model could be extended to handle multiple groups, predictors of intercept and slope factors, and growth in more than one outcome variable or more than one age cohort. We showed how TVCs may be added to a growth curve model and how polynomial or unspecified nonlinear trajectories can be modeled.

We want to emphasize that the researcher need not be restricted to investigating the progression of models presented in this chapter. If theory or past research suggests that a random intercepts, random slopes model is appropriate, then there is little reason to fit a simpler model. Likewise, if there is reason to expect a Level 2 predictor to explain individual differences in the quadratic component of a polynomial trend, it is straightforward and

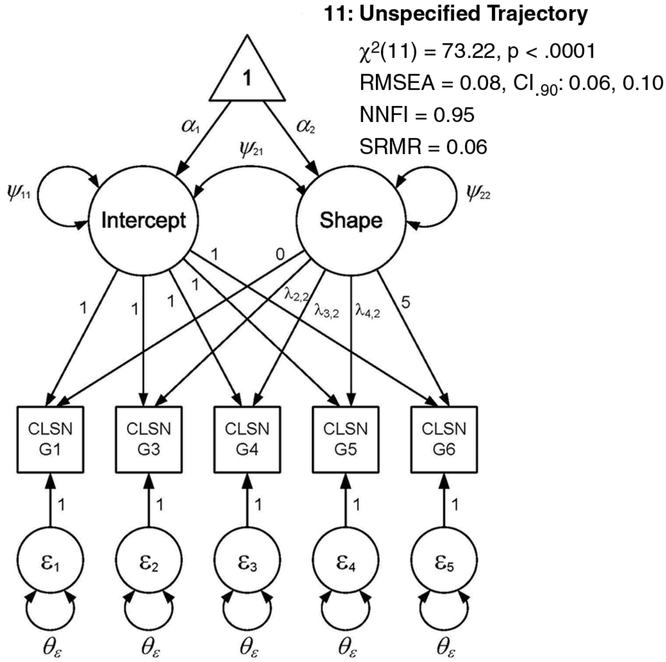

Figure 2.10 A Path Diagram Representing a Latent Growth Curve With an Unspecified Shape Factor.

NOTE: CLSN = Closeness with child.

appropriate to combine Models 6 and 10 rather than beginning with one model or the other. In short, it is important to always look to theory first. If theory is not sufficiently specific to suggest models to be investigated, then the exploratory strategy illustrated in this chapter—beginning with a null model, subsequently including a linear slope factor, and adding random effects and predictors—can be useful in helping the researcher to understand the data and appropriately model change over time.

Notes

1. We use complete data in this book for pedagogical simplicity and because the full array of SEM fit indices is available when complete data are used. The data used in subsequent analyses are provided along with syntax at http://www.quantpsy.org/.

2. LISREL is available from Scientific Software International (http://www.ssicentral.com/); Mx is available from Virginia Commonwealth University (http://www.vcu.edu/mx/); and Mplus is available from L. K. Muthén & Muthén (http://www.statmodel.com/).

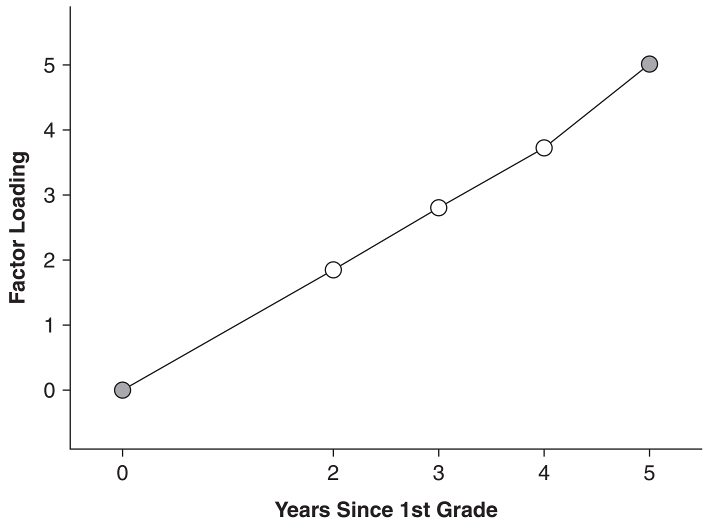

Figure 2.11 A Plot of Factor Loadings Versus Years Since the First Grade.

NOTE: Open circles represent estimated loadings. Closed circles represent fixed loadings used to anchor the estimation of the remaining loadings. Although the factor loadings follow a positive linear trend, they actually describe a *downward* trajectory because the slope mean is negative.

3. Throughout this book, we constrain the disturbance variances to equality for pedagogical simplicity and because it is reasonable to suppose that residual variability in mother–child closeness remains stable over time. This equality constraint is not required, assuming there are enough repeated measures to identify these parameters. Indeed, a strength of the LGM approach is that we can explicitly model heteroskedasticity—or any of a number of other disturbance covariance structures—by estimating different residual variances across occasions (Willett & Sayer, 1994). All things being equal, however, we recommend that preference be given to modeling homoscedasticity (equal disturbance variances) whenever possible to maximize parsimony, and because permitting residual variances to differ by occasion can sometimes mask or "soak up" nonlinearity in the trajectory, yielding deceptively good fit.

4. Syntax for this model, and for most subsequent models, is provided at our Web site (http://www.quantpsy.org/).

5. This is the nested model test we mentioned earlier, in which the difference in χ^2 values for the two models is itself treated as a χ^2 statistic with degrees of freedom equal to the difference in *df* for the two nested models.

6. To be clear, separate intercepts and slopes are not actually estimated in this procedure; rather, the model imposes a multivariate normal distribution on the latent variables and yields estimates of the means, variances, and covariances of those distributions.

7. Alternatively (and equivalently), the aperture can be directly estimated as a model parameter in some SEM programs by constraining ψ_{21} to zero and constraining the slope factors to their original fixed values minus an aperture parameter.

8. In this model, and later in Model 9, we make use of β parameters, which are elements of a matrix of path coefficients (**B**) from the full structural equation model not discussed in Chapter 1 (Bollen, 1989).

9. Because of this similarity, cross-level interactions in multilevel modeling may be decomposed, probed, and plotted according to guidelines stated by Aiken and West (1991). See Curran, Bauer, and Willoughby (2004) and Preacher, Curran, and Bauer (2006) for discussion.

10. These two approaches to testing for cohort effects are directly analogous to the two approaches for examining the effects of exogenous predictors of change discussed in Models 5 and 6.

11. Likelihood-based CIs are computed by determining the values a parameter must adopt for model fit to worsen by a given amount. For example, 95% CIs are formed by moving the parameter value away from the ML estimate in small steps—reoptimizing the model each time—until the ML fit function increases by 3.84 χ^2 units (3.84 is the critical value of χ^2 when $df = 1$) (Neale et al., 2003; Neale & Miller, 1997). This method of creating CIs can be time-consuming due to the amount of CPU-intensive reoptimization required, but likelihood-based CIs have several advantages over standard errors, which assume normality (typically untrue of variance parameters, for example), require *t* tests that are not invariant to reparameterization, and sometimes yield nonsensical results for bounded parameters. Standard errors are still de rigueur when reporting parameter estimates, but we predict that likelihood-based CIs will become more popular because of their desirable characteristics, particularly as computers improve in cost and computational efficiency.

12. Two loadings must be constrained for this model to be identified. Any two loadings will suffice, as long as they are constrained to different values.