8

Latent Growth Curve Modeling

Thus far, the examples used to motivate the utility of structural equation modeling have been based on cross-sectional data. Specifically, it has been assumed that the data have been obtained from a sample of individuals measured at one point in time. Although it may be argued that most applications of structural equation modeling are applied to cross-sectional data, it can also be argued that most social and behavioral processes under investigation are dynamic, that is, changing over time. In this case, cross-sectional data constitute only a snapshot of an ongoing dynamic process and interest might naturally center on the study of this process.

Increasingly, social scientists have access to longitudinal data that can provide insights into how outcomes of interest change over time. Indeed many important data sets now exist that are derived from panel studies (e.g., NCES, 1988; NELS:88; The National Longitudinal Study; The Longitudinal Study of American Youth; and the Early Childhood Longitudinal Study; to name a few). Access to longitudinal data allows researchers to address an important class of substantive questions—namely, the growth and development of social and behavioral outcomes over time. For example, interest may center on the development of mathematical competencies in young children (Jordan, Hanich, & Kaplan, 2003a, 2003b; Jordan, Kaplan, & Hanich, 2002). Or, interest may center on growth in science proficiency over the middle school years. Moreover, in both cases, interest may focus on predictors of individual growth that are assumed to be invariant across time (e.g., gender) or that vary across time (e.g., a student's absenteeism rate during a school year).

This chapter considers the methodology of growth curve modeling—a procedure that has been advocated for many years by researchers such as Raudenbush and Bryk (2002); Rogosa, Brandt, and Zimowski, (1982); and Willett (1988) for the study of intraindividual differences in change (see also Willett & Sayer, 1994). The chapter is organized as follows. First, we consider growth curve modeling as a general multilevel problem. This is followed by the specification of a growth curve model as a latent variable structural equation model. In this section, we consider the problem of how time is measured and incorporated into the model. The next section considers the addition of predictors into the latent growth curve model, as well as using the growth parameters as predictors of proximal and distal outcomes. This is followed by a discussion of growth curve modeling extensions that accommodate multivariate outcomes, nonlinear curve fitting, autoregressive structures.

This chapter will not consider other important issues of structural equation modeling to dynamic data. In particular, we will not consider the stationarity of factors in longitudinal factor analysis (e.g., Tisak & Meredith, 1990), nor will we consider recent developments in the merging of time-series models and structural equation models (e.g., Hershberger, Molenaar, & Corneal, 1996). For a detailed account of growth curve modeling, see Bollen and Curran (2006).

8.1 Growth Curve Modeling: A Motivating Example and Basic Ideas

To motivate the development of growth curve modeling, let us revisit the input-process-output model in Figure 1.2. A criticism of the input-process-output model, as diagrammed in Figure 1.2 is that it suggests a static educational system rather than a system that is inherently dynamic. For example, the outcomes of achievement and attitudes are, arguably, constructs that develop and change over time. Therefore, it may be of interest to adopt a dynamic perspective and ask how outcomes change over time and how those changes are influenced by time-invariant and time-varying features of the educational system. In addition to examining the change in any one of these outcomes over time, it may be of interest to examine how two or more outcomes change together over time.

For the purposes of the example that will be used throughout this chapter, we study change in science achievement and science attitudes separately and together. To set the framework for this application, Figure 8.1 shows the empirical trajectories for 50 randomly chosen students on the science achievement assessment over the five waves of LSAY. The figure shows considerable variability in both level and trend in science achievement over the waves of LSAY. Figure 8.2 shows the general trend in science attitudes over the five grade levels. Unlike achievement in science, attitudes toward science show a general linear decline over time. The advantage of growth curve modeling is that we can obtain an estimate of the initial level of science

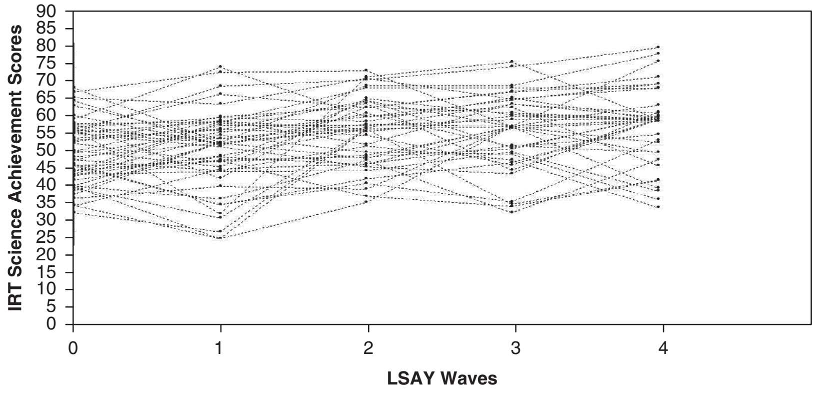

Figure 8.1 Fifty Random Science Achievement Observed Trajectories

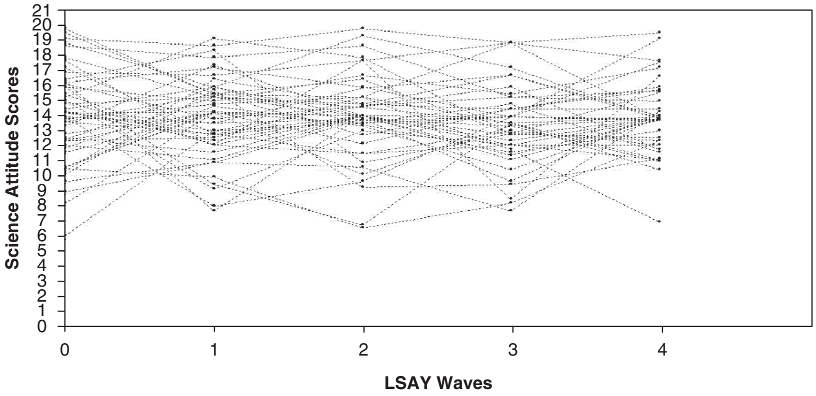

Figure 8.2 Fifty Random Science Attitude Observed Trajectories

achievement and the rate of change over time and link these parameters of growth to time-varying and time-invariant variables. In this example, such predictors will include student gender as well as teacher and parental push variables. However, in addition to applying univariate growth curve models, we also examine how these outcomes vary together in a multivariate growth curve application.

8.2 Growth Curve Modeling From the Multilevel Modeling Perspective

The specification of growth models can be viewed as falling within the class of multilevel linear models (Raudenbush & Bryk, 2002), where Level 1 represents intraindividual differences in initial status and growth, and Level-2 models individual initial status and growth parameters as a function of interindividual differences.

To fix ideas, consider a growth model for a continuous variable such science achievement. We can write a *Level-1* equation expressing outcomes over time within an individual as

$$y_{ip} = \pi_{0p} + \pi_{1p} t_i + \varepsilon_{ip}, \qquad [8.1]$$

where y_{ip} is the achievement score for person p at time i, π_{0p} represents the initial status at time $t = 0$, π_{1p} represents the growth trajectory, t_i represents a temporal dimension that here isassumed to be the same for all individuals—such as grade level, and ε_{ip} is the disturbance term. Later in this chapter, we consider more flexible alternatives to specifying time metrics.

Quadratic growth can also be incorporated into the model by extending the specification as

$$y_{ip} = \pi_{0p} + \pi_{1p} t_i + \pi_{2p} t_i^2 + \varepsilon_{ip}, \qquad [8.2]$$

where π_{2p} captures the curvilinearity of the growth trajectory. Higher-order terms can also be incorporated. In Section 8.4.2, we explore an alternative to the quadratic growth model in Equation [8.2] by allowing for general nonlinear curve fitting.

The specification of Equations [8.1] and [8.2] can be further extended to handle predictors of individual differences in the initial status and growth trajectory parameters. In the terminology of multilevel modeling, individuals would be considered *Level-2* units of analysis. In this case, two models are specified, one for the initial status parameter and one for the growth trajectory parameter. Consider, for example, a single time-invariant predictor of initial status and growth for person p, denoted as x_p. An example of such a predictor might be socioeconomic status of the student. Then, the Level-2 model can be written as

$$\pi_{0p} = \mu_{\pi_0} + \gamma_{\pi_0} x_p + \zeta_{0p} \qquad [8.3]$$

and

$$\pi_{1p} = \mu_{\pi_1} + \gamma_{\pi_1} x_p + \zeta_{1p}, \qquad [8.4]$$

where μ_{π_0} and μ_{π_1} are intercept parameters representing population true status and population growth when x_p is zero; γ_{π_0} and γ_{π_1} are slopes relating x_p to initial status and growth, respectively.

The model specified above can be further extended to allow individuals to be nested in groups such as classrooms. In this case, classrooms become a *Level-3* unit of analysis. Finally, the model can incorporate time-varying predictors of change. In the science achievement example, such a time-varying predictor might be changes in parental push or changes in attitudes toward science over time. Thus, this model can be used to study such issues as the influence of classroom-level characteristics and student-level invariant and varying characteristics on initial status and growth in reading achievement over time.

8.3 Growth Curve Modeling From the Structural Modeling Perspective

Research by B. Muthén (1991) and Willett and Sayer (1994) have shown how the general growth model described in the previous section can also be incorporated into a structural equation modeling framework. In what follows, the specification proposed by Willett and Sayer (1994) is described. The broad details of the specification are provided; however, the reader is referred to Willett and Sayer's (1994) article for more detail.

The Level-1 individual growth model can be written in the form of the factor analysis measurement model in Equation [4.24] of Chapter 4 as

$$\mathbf{y} = \boldsymbol{\tau}_y + \boldsymbol{\Lambda}_y\boldsymbol{\eta} + \boldsymbol{\varepsilon}, \qquad [8.5]$$

where $\mathbf{y}$ is a vector representing the empirical growth record for person *p*. For example, $\mathbf{y}$ could be science achievement scores for person *p* at the 7th, 8th, 9th, 10th, and 11th grades.

In this specification, $\boldsymbol{\tau}_y$ is an intercept vector with elements fixed to zero and $\boldsymbol{\Lambda}_y$ is a fixed matrix containing a column of ones and a column of constant time values. Assuming that time is centered at the seventh grade,[1] the time constants would be 0, 1, 2, 3, and 4. The matrix $\boldsymbol{\eta}$ contains the initial status and growth rate parameters denoted as π_{0p} and π_{1p}, and the vector $\boldsymbol{\varepsilon}$ contains measurement errors, where it is assumed that Cov($\boldsymbol{\varepsilon}$) is a diagonal matrix of constant measurement error variances. Because this specification results in the initial status and growth parameters being absorbed into the latent variable vector $\boldsymbol{\eta}$, which vary randomly over individuals, this model is sometimes referred to as a latent variable growth model (B. Muthén, 1991). The growth factors, as in the multilevel specification, are random variables.

Next, it is possible to use the standard structural model specification discussed in Chapter 4 to handle the Level-2 components of the growth model, corresponding. Considering the Level-2 model without the vector of predictor variables **x**, the model can be written as

$$\boldsymbol{\eta} = \boldsymbol{\alpha} + \mathbf{B}\boldsymbol{\eta} + \boldsymbol{\zeta}, \tag{8.6}$$

where $\boldsymbol{\eta}$ is specified as before, $\boldsymbol{\alpha}$ contains the population initial status and growth parameters μ_{π_0} and μ_{π_1}, $\mathbf{B}$ is a null matrix, and $\boldsymbol{\zeta}$ is a vector of deviations of the parameters from their respective population means. Again, this specification has the effect of parameterizing the true population initial status parameter and growth parameter into the structural intercept vector $\boldsymbol{\alpha}$. Finally, the covariance matrix of $\boldsymbol{\zeta}$, denoted as $\boldsymbol{\Psi}$, contains the variances and covariances of true initial status and growth.

The Level-2 model given in Equation [8.6] does not contain predictor variables. The latent variable growth model can, however, be extended to include exogenous predictors of initial status and growth. To incorporate exogenous predictors requires using the **x**-measurement model of the sort described in Chapter 4. Specifically, the model is written as

$$\mathbf{x} = \boldsymbol{\tau}_x + \boldsymbol{\Lambda}_x\boldsymbol{\xi} + \boldsymbol{\delta}, \tag{8.7}$$

where here **x** is a vector of exogenous predictors, $\boldsymbol{\tau}_x$ contains the mean vector, $\boldsymbol{\Lambda}_x$ is an identity matrix, $\boldsymbol{\xi}$ contains the exogenous predictors deviated from their means, and $\boldsymbol{\delta}$ is a null vector. This specification has the effect of placing the centered exogenous variables in $\boldsymbol{\xi}$ (Willett & Sayer, 1994, p. 374).

Finally, the full specification of the structural equation model given in Equation [4.1] can be used to model the predictors of true initial status and true growth, where, due to the centering of the exogenous predictors, it retains its interpretation as the population mean vector of the individual initial status and growth parameters (Willett & Sayer, 1994, p. 375).

An important feature of the structural equation modeling approach to growth curve modeling is its flexibility in handling structured errors. That is, the assumption of independent and homoscedastic errors can be relaxed allowing for heteroscedasticity and autocorrelation. In the former case, heteroscedasticity can be incorporated by relaxing the equality constraints among error variances in the diagonal of $\boldsymbol{\Theta}_\varepsilon$. Autocorrelation can be incorporated into growth curve models by allowing free off-diagonal elements in $\boldsymbol{\Theta}_\varepsilon$.

8.3.1 AN EXAMPLE OF UNIVARIATE GROWTH CURVE MODELING

The data for this example come from the Longitudinal Study of American Youth (LSAY; Miller, Hoffer, Sucher, Brown, & Nelson, 1992).[2] LSAY includes two sets of schools, a national probability sample of approximately 60 high schools (Cohort 1) and approximately 60 middle schools (Cohort 2). An average of 60 10th graders (Cohort 1) and 60 7th graders (Cohort 2) from each of the 60 high schools and middle schools have been followed since 1987, gathering information on students' family and school background, attitudes, and achievement.

In addition to general background information, achievement and attitude measures were obtained. Achievement tests in science and mathematics were given to the students each year. The items for the mathematics and science achievements tests were drawn from the item pool of the 1986 National Assessment of Educational Progress (NAEP) tests (NAEP, 1986).

The measure of student attitudes toward science is based on a composite which consists of an equally weighted average of four attitudinal subscales, namely interest, utility, ability, and anxiety. There are nine variables in this composite, for example, "I enjoy science"; "I enjoy my science class"; and so on. Variables were recoded so that high values indicate a positive attitude toward science. The composite is measured on a 0 to 20 metric.

For the purposes of this example, we concentrate the younger cohort, measured at grades 7, 8, 9, 10, and 11. In addition to science achievement test scores, we also include *gender* (male = 1) as a time-invariant predictor. Time-varying predictors include a measure *of parent academic push* (PAP) and *student's science teacher push* (STP). PAP is an equally weighted average of eight variables. Both student and parent responses are used in this composite. Questions asked of the students are related to parental encouragement for making good grades, doing homework, and interest in school activities. Questions asked of the parents were related to their knowledge of their child's performance, homework, and school projects. This composite is measured on a 0 to 10 metric. Although a composite measure of parent science push composite was available for Cohort 2, it was not used because the items composing this composite were not measured at all the time points.

Science teacher push is a composite based on five student response variables referring to teacher encouragement of science. Response values for this composite range from 0 to 5.

The sample size for this study was 3,116. Analyses used Mplus (L. Muthén & Muthén, 2006) under the assumption of multivariate normality of the data.

Missing data were handled by full information maximum likelihood imputation as discussed in Chapter 5. The analysis proceeds by assessing growth in science achievement and science attitudes separately, then together in a multivariate growth curve model.

Growth in Science Achievement. Column 1 of Table 8.1 presents the results of the linear growth curve model without predictors. A path diagram of this model is shown in Figure 8.3. This model is estimated allowing for heteroscedastic but non-autocorrelated disturbances. The initial status is set at

Table 8.1 Selected Results of Growth Curve Model of Science Achievement

	Model 1[a]	*Model 2*[b]	*Model 3*[c]
Effect	*Maximum Likelihood Estimates*		
Intercept	50.507*	50.632*	47.042*
Slope	2.207*	1.813*	1.810*
Var(intercept)	71.665*	68.935*	67.755*
Var(slope)	2.409*	1.563*	1.602*
r(intercept and slope)	–0.392*	–0.352*	–0.365
Intercept on gender		0.667*	0.736*
Slope on gender		–0.078	–0.083
SCIACH1 on PAP1			0.344*
SCIACH2 on PAP2			0.340*
SCIACH3 on PAP3			0.483*
SCIACH4 on PAP4			0.377*
SCIACH5 on PAP5			0.278*
SCIACH1 on STP1			0.109
SCIACH2 on STP2			0.161
SCIACH3 on STP3			0.556*
SCIACH4 on STP4			0.495*
SCIACH5 on STP5			0.293*
BIC	111359.030	116079.778	227152.878

a. Linear growth curve model—no covariates.

b. Linear growth curve model—gender as time-invariant covariate.

c. Linear growth curve model—gender as time-invariant covariate; parent academic push (PAP) and science teacher push (STP) as time-varying covariates.

*$p < .05$.

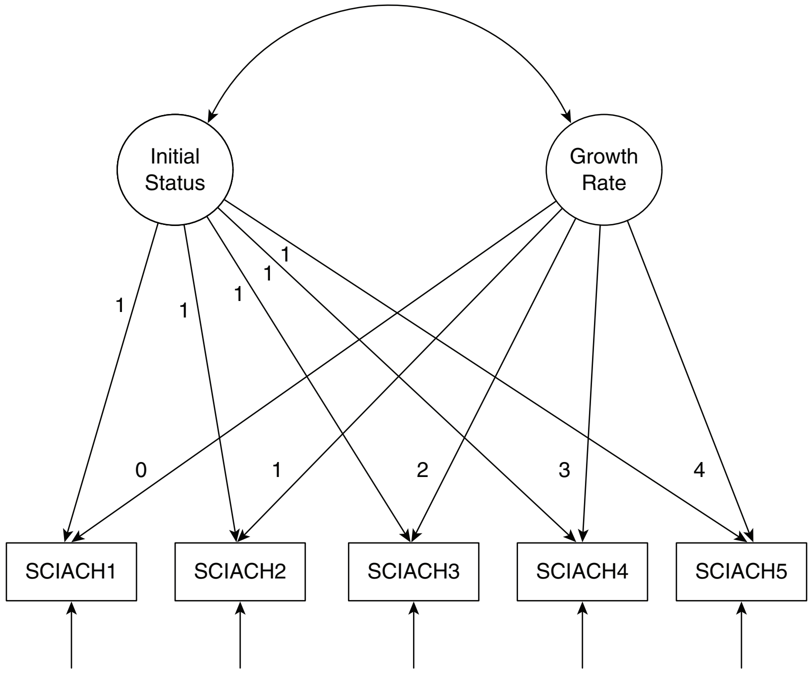

Figure 8.3 Initial Growth Curve Model of Science Achievement

seventh grade. The results indicate that the average seventh grade science achievement score is 50.51 and increases an average of 2.21 points a year. The correlation between the initial status and rate of change is negative suggesting the possibility of a ceiling effect. Figure 8.1 presents a random sample of 50 model-estimated science achievement trajectories.

Column 2 of Table 8.1 presents the results of the linear growth curve model with gender as a time-invariant predictor of initial status and growth rate. A path diagram of this model is shown in Figure 8.4. The results indicate a significant difference in favor of boys for seventh grade science achievement, but no significant difference between boys and girls in the rate of change over the five grades.

Column 3 of Table 8.1 presents the results of the linear growth curve model with the time-varying covariates of PAP and STP included. The results for gender remain the same. A path diagram of this model is shown in Figure 8.5. The results for the time-varying covariates suggest that early PAP is a stronger predictor of early science achievement compared with STP. However, the effects of both time-varying covariates balance out at the later grades.

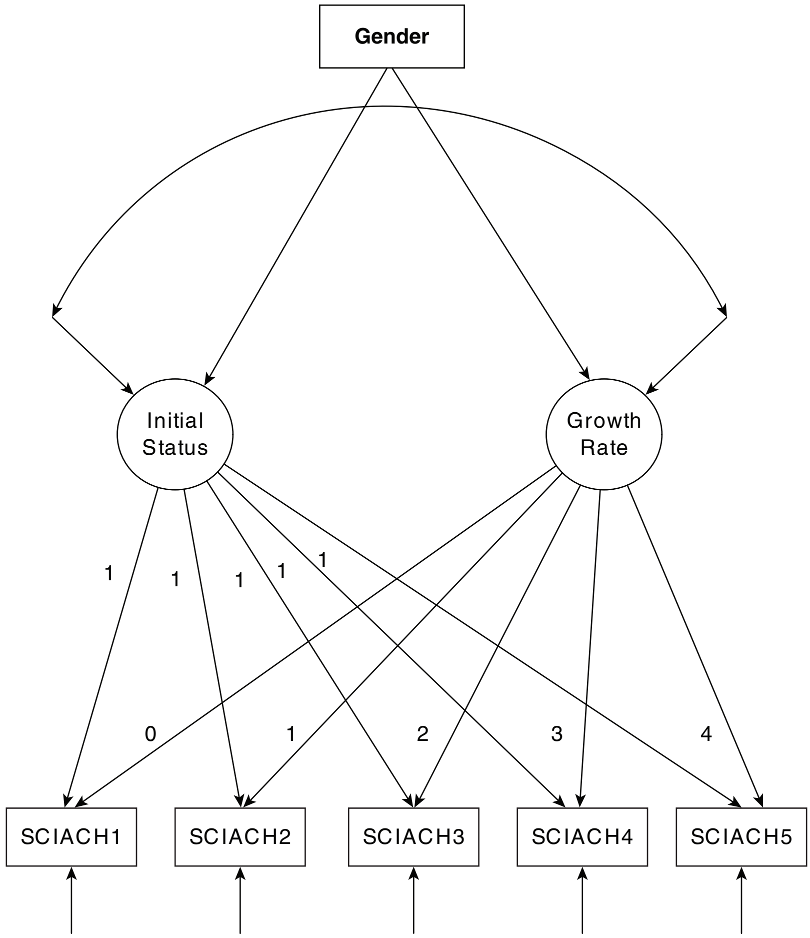

Figure 8.4 Growth Curve Model of Science Achievement With Time-Invariant Predictors

Growth in Attitudes Toward Science. Column 1 of Table 8.2 shows the results of the simple linear growth curve model applied to the science attitude data. Path diagrams for this and the remaining models are not shown. The results show a seventh grade average attitude score of 14.25 points (on a scale of 1 to 20) and a small but significant decline over time. Moreover, a strong negative correlation can be observed between initial science attitudes and the change over time. This suggests, as with achievement, that higher initial attitudes are associated with slower change in attitudes over time.

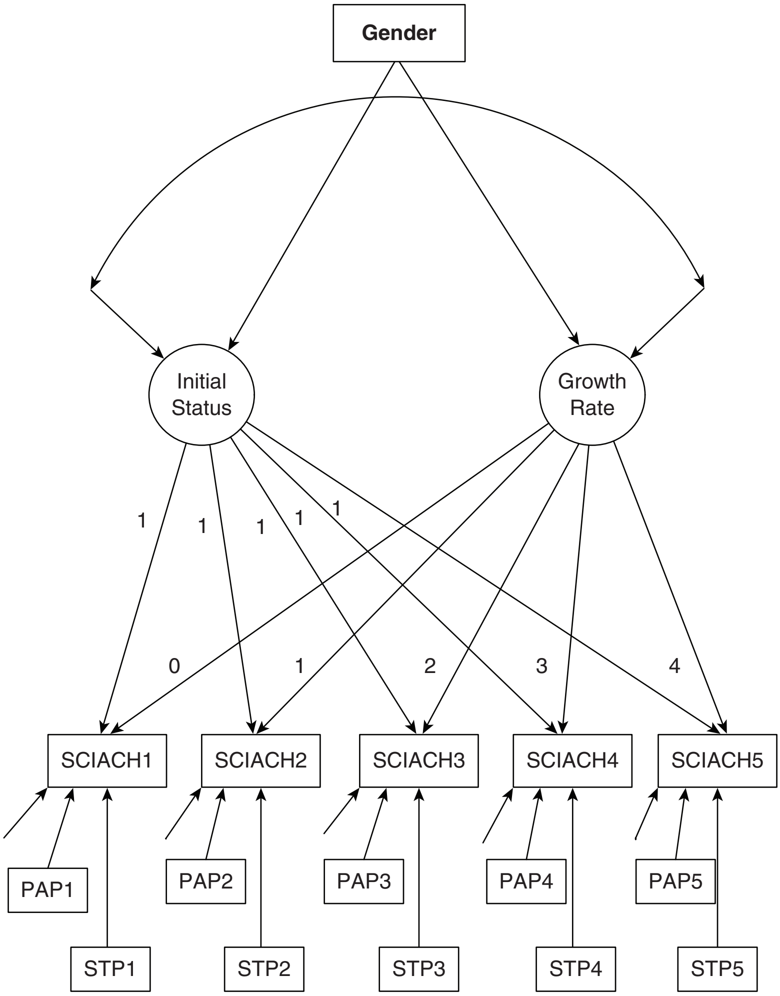

Figure 8.5 Growth Curve Model of Science Achievement With Time-Invariant and Time-Varying Predictors

Column 2 of Table 8.2 examines sex differences in initial seventh grade science attitudes and sex differences in the rate of decline over time. As with science achievement we observe initial differences in attitudes at seventh grade with boys exhibiting significantly higher positive attitudes compared with girls. However, there appears to be no sex differences in the rate of attitude change over time.

Table 8.2 Selected Results of Growth Curve Model of Attitudes Toward Science

	Model 1[a]	*Model 2*[b]	*Model 3*[c]
Effect	*Maximum Likelihood Estimates*		
Intercept	14.251*	14.058*	12.018*
Slope	–0.095*	–0.094*	0.043
Var(intercept)	3.422*	3.388*	2.889*
Var(slope)	0.121*	0.121*	0.107*
r(intercept and slope)	–0.578*	–0.578*	–0.564*
Intercept on gender		0.369*	0.413*
Slope on gender		–0.003	–0.009
ATTITUDE1 on PAP1			0.148*
ATTITUDE2 on PAP2			0.113*
ATTITUDE3 on PAP3			0.086*
ATTITUDE4 on PAP4			0.077*
ATTITUDE5 on PAP5			0.091*
ATTITUDE1 on STP1			0.284*
ATTITUDE2 on STP2			0.330*
ATTITUDE3 on STP3			0.332*
ATTITUDE4 on STP4			0.312*
ATTITUDE5 on STP5			0.218*
BIC	69089.797	73588.288	184562.117

a. Linear growth curve model—no covariates.

b. Linear growth curve model—gender as time-invariant covariate.

c. Linear growth curve model—gender as time-invariant covariate; parent academic push (PAP) and science teacher push (STP) as time-varying covariates.

*$p < .05$.

Column 3 of Table 8.2 adds the time-varying covariates to model in Column 2. The results here are somewhat different than those found for achievement. Specifically, we observe that PAP is a relatively weak predictor of science attitudes compared with STP. Moreover, an inspection of correlations between sex and each of the time-varying predictors can be interpreted as representing whether sex differences are occurring for PAP and STP. The results indicate small and mostly nonsignificant sex differences in these time-varying covariates.

8.4 Extensions of the Basic Growth Curve Model

An important feature of growth curve modeling within the structural equation modeling perspective is its tremendous flexibility in handling a variety of different kinds of questions involving growth. In this section, we consider four important extensions of growth curve modeling. First, we consider the multivariate growth curve modeling, including models for parallel and sequential processes. Second, we consider model extensions for nonlinear curve fitting. Third, we consider an extension that incorporates an autoregressive component to the model. Finally, we briefly consider some flexible alternatives to addressing the time metric. It should be noted that these three extensions do not exhaust the range of analytical possibilities with growth curve modeling. For a more comprehensive treatment of the extensions of growth curve modeling, see Bollen and Curran (2006).

8.4.1 MULTIVARIATE GROWTH CURVE MODELING

Consider the case where an investigator wishes to assess the relationship between growth in mathematics and reading proficiency. It can be argued that these achievement domains are highly related. Indeed, one may argue that because measures of mathematics proficiency require reading proficiency, reading achievement might be a causal factor for growth in mathematics proficiency. For now, however, we are only interested in assessing how these domains change together.

A relatively straightforward extension of the growth curve specification given in Equations [8.5] to [8.7] allows for the incorporation of multiple outcome measures (Willett & Sayer, 1996). Important information about growth in multiple domains arises from an inspection of the covariance matrix of $\boldsymbol{\eta}$ denoted above as $\boldsymbol{\Psi}$. Recall that in the case of univariate growth curve modeling the matrix $\boldsymbol{\Psi}$ contains the covariance (or correlation) between the initial status parameter π_0 and the growth parameter π_1. In the multivariate extension, $\boldsymbol{\Psi}$ contains the measures of association among the initial status and growth rate parameters of each outcome. Thus, for example, we can assess the degree to which initial levels of reading proficiency are correlated with initial proficiency levels in mathematics and also the extent to which initial reading proficiency is correlated with the rates of growth in mathematics. We may also ask whether rates of growth in reading are correlated with rates of growth in mathematics. As in the univariate case, the multivariate case can be easily extended to include time-invariant and time-varying predictors of all the growth curve parameters.

If both mathematics and reading proficiency are measured across the same time intervals, then we label this a *parallel growth process*. However, an interesting additional extension of multivariate growth curve modeling allows the developmental process of one domain to predict the developmental process of a later occurring outcome (see, e.g., B. Muthén & Curran, 1997). For example, one might argue that development in reading proficiency in first, second, and third grades predict the development of science achievement in fourth, fifth, and sixth grades.

For this extension, the decision where to center the level of the process is crucial. One could choose to center initial reading proficiency at first grade and initial science proficiency at fourth grade. However, it may be the case that reading proficiency at first grade shows little variation and thus may not be a useful predictor of initial science proficiency. Perhaps a more sensible strategy would be to center initial reading proficiency at the third grade and center initial science proficiency at fourth grade. One might expect more variation in reading proficiency at the third grade and this variation might be more predictive of science proficiency at the fourth grade. As in the univariate case, the issue of centering will most often be based on substantive considerations.

An Example of Multivariate Growth Curve Modeling

An inspection of Figures 8.1 and 8.2 suggests the need to study changes in science achievement and science attitudes together. In the interest of space, we fit the full time-invariant and time-varying model to the achievement and attitude data in one analysis. A path diagram for this model is not shown. The results are shown in Table 8.3. The results generally replicate those of the univariate analyses, and in the interest of space, the time-varying covariate results are not shown. However, it is important to focus on the correlations between the growth parameters for achievement and attitudes. The results indicate a positive correlation between seventh grade science achievement and seventh grade science attitudes ($r = 0.458$). Moreover, we observe that higher rates of growth in science achievement are associated with higher rates of growth in attitudes toward science ($r = 0.381$). An apparent contradiction arises when considering the negative correlation between initial science achievement and rate of change in science attitudes. Again, an explanation might be a ceiling effect, insofar as higher achievement scores are associated with higher attitudes and therefore attitudes toward science cannot change much more.

8.4.2 NONLINEAR CURVE FITTING

In practical applications of growth curve modeling, it might be the case that a nonlinear curve better fits the data. An approach to nonlinear curve fitting, suggested by Meredith and Tisak (1990), entails freeing a set of the factor loadings associated with the slope. Specifically, considering the science achievement

Table 8.3 Selected Results of Multivariate Growth Curve Model of Science Achievement and Attitudes Toward Science

Effect Estimates	*Maximum Likelihood*
Ach. intercept	46.749*
Ach. slope	2.325*
Att. intercept	12.066*
Att slope	0.064
Var(ach. intercept)	69.775*
Var(ach. slope)	2.318*
Var(att. intercept)	3.164*
Var(att. slope)	0.167*
r(ach. intercept/att. intercept)	0.458*
r(ach. intercept/att. slope)	–0.231*
r(ach. intercept/ach. slope)	–0.394*
r(att. intercept/att slope)	–0.616*
r(att. intercept/ach. slope)	–0.178*
r(ach. slope/att. slope)	0.381*
Ach. intercept on gender	0.846*
Ach slope on gender	–0.151
Att intercept on gender	0.416*
Att slope on gender	–0.012
BIC	295164.119

*$p < .05$.

model, the nonlinear curve fitting approach suggested by Meredith and Tisak would require that the first loading be fixed to zero to estimate the intercept, the second loading would be fixed to one to identify the metric of the slope factor, but the third through fifth loadings would be free. In this case, the time metrics are being empirically determined. When this type of model is estimated, it perhaps makes better sense to refer to the slope factor as a *shape* factor.

An Example of Nonlinear Curve Fitting

In this example, we estimate the science achievement growth model allowing estimation of a general shape factor. As suggested by Meredith and Tisak

(1990), we fix the first and second loadings as in the conventional growth curve modeling case and free the loadings associated with the third, fourth, and fifth waves of the study. The results are displayed in Table 8.4. It is clear from an inspection of Table 8.4 that the nonlinear curve fitting model results in a substantial improvement in model fit. Moreover, we find that there are significant sex differences with respect to the intercept in the nonlinear curve fitted model.

8.4.3 AUTOREGRESSIVE LATENT TRAJECTORY MODELS

Recently, Bollen and Curran (2004) and Curran and Bollen (2001) advocated the blending of an autoregressive structure into conventional growth curve modeling. They refer to this hybrid model as the autoregressive latent trajectory (ALT) model.

Table 8.4 Maximum Likelihood Estimates From Nonlinear Curve Fitting Models

	Model 0 Estimates	*Model 1 Estimates*
Intercept by		
Ach1	1.000	1.000
Ach2	1.000	1.000
Ach3	1.000	1.000
Ach4	1.000	1.000
Ach5	1.000	1.000
Shape by		
Ach1	0.000	0.000
Ach2	1.000	1.000
Ach3	3.351	3.299
Ach4	3.928	3.869
Ach5	5.089	5.004
Ach. intercept	50.360*	49.966*
Ach. shape	1.693*	1.770*
r(shape, intercept)	–0.397*	–0.398*
Intercept on		
Male		0.737*
Shape on		
Male		–0.091
BIC	111240.859	115769.718

*$p < .05$.

It is not difficult to make the case for specifying an ALT model for developmental research studies. Consider the example used throughout this chapter where the focus is on modeling the development of science proficiency through the middle and high school years. We can imagine that interest centers on how change in science proficiency predicts later outcomes of educational relevance—such as majoring in science-related disciplines in college. It is not unreasonable, therefore, to assume that in addition to overall growth in science proficiency prior science scores predict later science scores thus suggesting an autoregressive structure.

In the case of long periods between assessment waves, we might reasonably expect small autoregressive coefficients, as opposed to more closely spaced assessment waves. Nevertheless, if the ALT model represents the true data generating structure, then omission of the autoregressive part may lead to substantial parameter bias. A recent article by Sivo, Fan, and Witta (2005) found extensive bias for all parameters of the growth curve model as well as biases in measures of model fit when a true autoregressive component was omitted from the analysis.

For the purposes of this chapter, we focus on the baseline lag-1 ALT model with a time-invariant predictor. This will be referred to as the ALT(1) model. The ALT(1) specification indicates that the outcome at time t is predicted only by the outcome at time $t - 1$. It should be noted lags greater than one can also be specified. As with conventional growth curve modeling, the ALT model can be extended to include more than one outcome, each having its own autoregressive structure, as well as extensions that include proximal or distal outcomes and time-varying and time-invariant predictors.

To contextualize the study consider the example of an ALT model for the development of reading competencies in young children. The first model is a baseline lag-1 ALT model. This model can be written in structural equation modeling notation as

$$\mathbf{y} = \boldsymbol{\alpha} + \boldsymbol{\Lambda}\boldsymbol{\eta} + \mathbf{B}\mathbf{y} + \boldsymbol{\delta}, \quad [8.8]$$

$$\boldsymbol{\eta} = \boldsymbol{\tau} + \boldsymbol{\Gamma}\boldsymbol{\eta} + \boldsymbol{\zeta}, \quad [8.9]$$

where y is a vector of repeated measures, $\boldsymbol{\Lambda}$ is a matrix of fixed coefficients that specify the growth parameters, $\boldsymbol{\eta}$ is a vector of growth parameters, $\mathbf{B}$ is a matrix of regression coefficients relating the repeated measures to each other, and $\boldsymbol{\delta}$ is a vector of residual variances with covariance matrix $\text{Cov}(\boldsymbol{\delta}\boldsymbol{\delta}') = \boldsymbol{\Theta}$. A path diagram of ALT(1) model is shown in Figure 8.6.

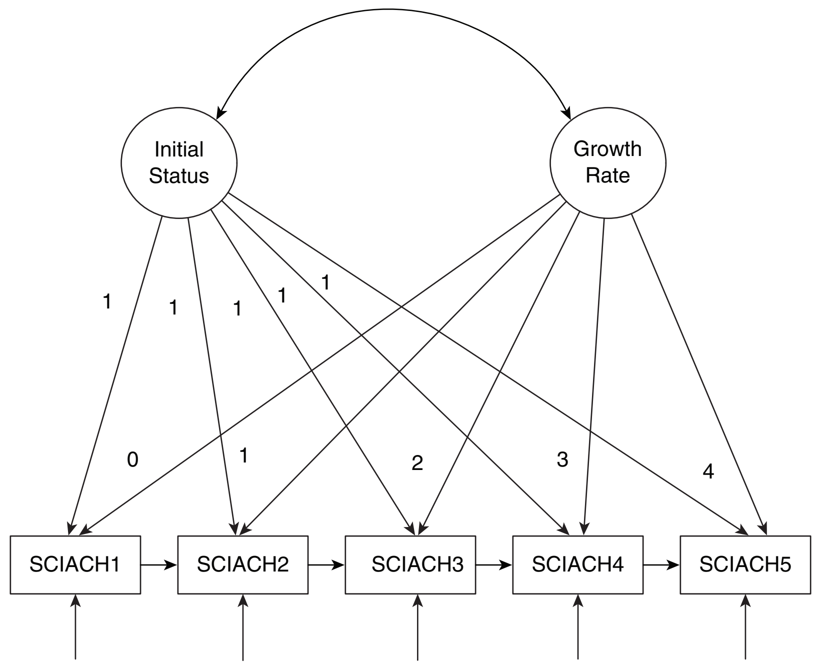

Figure 8.6 Autoregressive Latent Trajectory(1) [ALT(1)] Model of Science Achievement

An Example of an ALT Model

For this example, we estimate an ALT(1) among the science scores from the LSAY example. Model 0 of Table 8.5 displays the results for the ALT(1) model without the addition of gender as a time-invariant predictor. It can be seen that the autocorrelation effects are small but statistically significant. Model 1 under Table 8.5 adds gender to the ALT model. The addition of gender in Model 1 appears to worsen the overall fit of the model as evidenced by the increase in the BIC.

8.4.4 ALTERNATIVE METRICS OF TIME

Up to this point, we have assumed a highly restrictive structure to the data. Specifically, we have assumed that each wave of measurement is equidistant and that we have complete data on all units of analysis at each time point. In many cases, this assumption is too restrictive and we need a way of handling more realistic time structures. For example, in developmental research, the

Table 8.5 Maximum Likelihood Estimates From Autoregressive Latent Trajectory Models

	Model 0 Estimates	*Model 1 Estimates*
Ach5 ON		
Ach4	0.135*	0.135*
Ach4 ON		
Ach3	0.103*	0.103*
Ach3 ON		
Ach2	0.102*	0.102*
Ach2 ON		
Ach1	0.031*	0.031*
Ach. intercept	50.335*	49.911*
Ach. slope	0.246*	0.329*
r(slope, intercept)	–0.575*	–0.573*
Intercept on		
Male		0.814*
Slope on		
Male		–0.156
BIC	111195.653	115723.216

*$p < .05$.

wave of assessment may not be nearly as important as the chronological age of the child. In this case, there might be a quite a bit of variability in chronological ages at each wave of assessment. In other cases, the nature of the assessment design is such that each child has his or her own unique interval between testing. In this section, we introduce two approaches that demonstrate the flexibility in dealing with the time metric in longitudinal studies: The cohort sequential design and the individual varying metrics of time design.

The Cohort Sequential Design. In cohort sequential designs, we consider age cohorts within a particular time period (Bollen & Curran, 2006). Thus, at Wave 1 of the study, we may have children who vary in age from 5 years to 7 years. At Wave 2 of the study, we may have children varying in age from 7 years to 9 years years, and so on. Notice, there is an overlap of ages at each wave.

As Bollen and Curran (2006) point out, there are two ways that this type of data structure can be addressed. First, we can go back to treating wave as the metric of time and use age of respondent as a covariate in the study. The second approach is to exploit the inherent missing data structure. In this case, we could arrange the data as shown in Table 8.6 patterned after Bollen and Curran (2006, p. 77). Notice that there are three cohorts and five time points. Any given child in this example can provide between one and four repeated measures. The pattern of missing data allows estimation using maximum likelihood imputation under the assumption of missing-at-random (Allison, 1987; Arbuckle, 1996; Muthén et al., 1987). Thus, the growth parameters spanning the entire time span can be estimated.

As Bollen and Curran (2006) point out, however, this approach suffers from the potential of cohort effects. That is, children in Cohort 1 may have been 7 years old at the second wave of assessment, but children in Cohort 2 would have been 7 at the first wave of assessment.

Individually Varying Metrics of Time. Perhaps a more realistic situation arises when individuals have their own unique spacing of assessment waves. An example of this would be the situation where a researcher is collecting individual longitudinal assessments in schools. At the beginning of the semester, the researcher and his or her assistants begin data collection. Because it is probably not feasible that every child in every sampled school can be assessed on exactly the same day, the assessment times may spread over, say, a 2-week period. At the second wave of assessment, the first child assessed at Wave 1 is not necessarily the same child assessed at Wave 2. Indeed, in the worst case scenario, if the first child assessed at Wave 1 is the last child assessed at Wave 2, the length of time between assessments will be much greater than if the child is the last one assessed at Wave 1 and the first assessed at Wave 2. Although I have presented the extreme case, the consequences for a study of development, especially in young children, would be profound. A better approach is to mark the date of assessment for each child and use the time between assessments for each child as his or her own unique metric of time. Time can be measured in days, weeks, or months, with the decision based on developmental considerations.

Table 8.6 Cohort Sequential Data Structure

	Age of Assessment			
Cohort	*Time 1*	*Time 2*	*Time 3*	*Time 4*
1	6	7	8	9
2	7	8	9	10
3	8	9	10	11

8.5 Evaluating Growth Curve Models Using Forecasting Statistics

It may be useful to consider if there are aspects of model fit that are pertinent to the questions being addressed via the use latent growth curve models. Clearly, we can apply traditional statistical and nonstatistical measures of fit, such as the likelihood ratio chi-square, RMSEA, NNFI, or the like. In many cases, the Bayesian information criterion is used to compare latent growth curve models as well. However, these measures of fit are capturing whether the restrictions that are placed on the data to provide estimates of the initial status and growth rate are supported by the data. In addition, these measures are assessing whether such assumptions as non-autocorrelated errors are supported by the data.

The application of traditional statistical and nonstatistical measures of fit does provide useful information. However, because growth curve models provide estimates of rates of change, it may be useful to consider whether the model predicted growth rate fits the empirical trajectory over time. So, for example, if we know how science achievement scores have changed over the five waves of LSAY, we may wish to know if our growth curve model accurately predicts the known growth rate. In the context of economic forecasting, this exercise is referred to as ex post simulation. The results of an ex post simulation exercise is particularly useful when the goal of modeling is to make forecasts of future values.

To evaluate the quality and utility of latent growth curve models, Kaplan and George (1998) studied the use of six different ex post (historical) simulation statistics originally proposed by Theil (1966) in the domain of econometric modeling. These statistics evaluate different aspects of the growth curve. The first of these statistics discussed by Kaplan and George was the root mean square simulation error (RMSSE) as

$$\text{RMSSE} = \sqrt{\frac{1}{T}\sum_{t=1}^{T}(y_t^{\text{s}} - y_t^{\text{a}})^2}, \qquad [8.10]$$

where T is the number of time periods, y_t^{s} is the simulated (i.e., predicted) value at time t, and y_t^{a} is the actual value at time t. The RMSSE provides a measure of the deviation of the simulated growth record from the actual growth record and is the measure most often used to evaluate simulation models (Pindyck & Rubinfeld, 1991).

Another measure is the *root mean square percent simulation error* (RMSPE), which scales the RMSSE by the average size of the variable at time t.

The RMSPE is defined as

$$\mathrm{RMSPE} = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(\frac{y_t^{\mathrm{s}} - y_t^{\mathrm{a}}}{y_t^{\mathrm{a}}}\right)}. \qquad [8.11]$$

A problem with the RMSPE is that its scale is arbitrary. Although the lower bound of the measure is zero, the upper bound is not constrained. Thus, it is of interest to scale the RMSSE to lie in the range of 0 to 1. A measure that lies between 0 and 1 is Theil's *inequality coefficient*, defined as

$$U = \frac{\sqrt{\frac{1}{T}\sum_{t=1}^{T}(y_t^{\mathrm{s}} - y_t^{\mathrm{a}})^2}}{\sqrt{\frac{1}{T}\sum_{t=1}^{T}(y_t^{\mathrm{s}})^2} + \sqrt{\frac{1}{T}\sum_{t=1}^{T}(y_t^{\mathrm{a}})^2}}. \qquad [8.12]$$

An inspection of Equation [8.12] shows that perfect fit of the simulated growth record to the actual growth record is indicated by a value $U = 0$. However, if $U = 1$, the simulation adequacy is as poor as possible.

An interesting feature of the inequality coefficient in Equation [8.12] is that can be decomposed into components that provide different perspectives on the quality of simulation performance. The first component of Theil's U is the *bias proportion*, defined as

$$U^{\mathrm{M}} = \frac{(\bar{y}^{\mathrm{s}} - \bar{y}^{\mathrm{a}})^2}{\left(\frac{1}{T}\right)\sum_{t=1}^{T}(y_t^{\mathrm{s}} - y_t^{\mathrm{a}})^2}, \qquad [8.13]$$

where $\bar{y}^{\mathrm{s}}$ and $\bar{y}^{\mathrm{a}}$ are the means of the simulated and actual growth record, respectively, calculated across the T time periods. The bias proportion provides a measure of systematic error because it considers deviations of average actual values from average simulated values (Pindyck & Rubinfeld, 1991).

The ideal would be a value of $U^{\mathrm{M}} = 0$. Values greater than 0.1 or 0.2 are considered problematic.

Another component of Theil's U is the *variance proportion* defined as

$$U^{\mathrm{S}} = \frac{(\sigma_{\mathrm{s}} - \sigma_{\mathrm{a}})^2}{\left(\frac{1}{T}\right)\sum_{t=1}^{T}(y_t^{\mathrm{s}} - y_t^{\mathrm{a}})^2}, \qquad [8.14]$$

where σ_s and σ_a are the standard deviations of the respective growth records calculated across the T time periods. The variance proportion provides a measure of the extent to which the model tracks the variability in the growth record. If U^S is large, it suggests that the actual (or simulated) growth record varied a great deal while the simulated (or actual) growth record did not deviate by a comparable amount.

A final measure based on the decomposition of the inequality coefficient is the *covariance proportion*, defined as

$$U^C = \frac{2(1-\rho)\sigma_s\sigma_a}{\left(\frac{1}{T}\right)\sum_{t=1}^{T}(y_t^s - y_t^a)^2}, \quad [8.15]$$

where ρ is the correlation coefficient between y_t^s and y_t^a. The covariance proportion U^C provides a measure of unsystematic error, that is, error that remains after having removed deviations from average values.

The decomposition of U results in the relation

$$U^M + U^S + U^C = 1, \quad [8.16]$$

and an ideal result for a simulation model would be $U^M = 0$, $U^S = 0$, and $U^C = 1$.

Values greater than zero for U^M and/or U^S are indicative of some problem with the model vis-à-vis tracking the empirical growth record.

8.5.1 COMPARISON OF STANDARD GOODNESS-OF-FIT TESTS AND FORECASTING STATISTICS FOR SCIENCE ACHIEVEMENT GROWTH MODEL

Table 8.7 displays the forecasting statistics described above for the science achievement model. It can be seen that the simple linear trend model (Model 1) demonstrates the best historical forecasting performance as measured by all six forecasting statistics. Model 2 incorporates the time-invariant predictor of gender. Here, it can be seen that historical forecasting performance worsens. When time-varying predictors of teacher push and parent push are added to account for the variability in the growth curve, the historical forecasting performance improves as measured by U^S as expected. Figure 8.7 compares the observed growth in science achievement with the model-predicted growth curves.

Table 8.7 Observed and Predicted Science Achievement Means and Forecasting Statistics for the Science Achievement Model

Grade	*Observed Means*	*Predicted Means*		
		Model 1	*Model 2*	*Model 3*
7	50.345	50.507	50.591	49.953
8	52.037	52.714	52.404	51.748
9	56.194	54.921	54.217	55.722
10	56.840	57.128	56.030	56.511
11	58.970	59.335	57.843	57.021
Forecasting Statistics				
RMSSE		0.681	1.098	0.935
RMSPSE		0.012	0.019	0.016
U		0.006	0.010	0.008
U^{M}		0.004	0.361	0.539
U^{S}		0.014	0.409	0.202
U^{C}		0.982	0.230	0.259

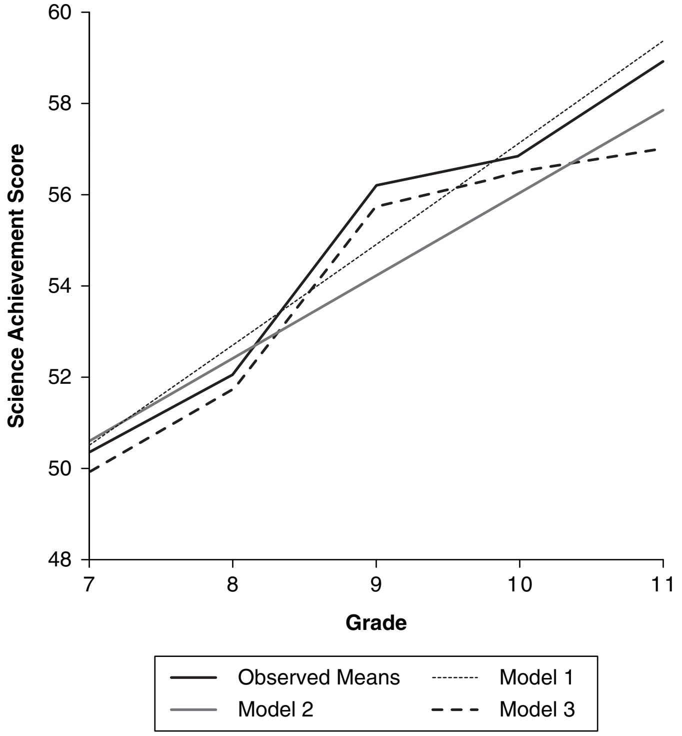

Figure 8.7 Observed Versus Model-Predicted Science Achievement Means

8.6 Conclusion

This chapter focused on the extension of structural equation modeling to the study of growth and change. We outlined how growth can be considered a multilevel modeling problem with intraindividual differences in students modeled at Level 1, individual differences modeled at Level 2, and individuals nested in groups modeled Level 3. We discussed how this specification could be parameterized as a structural equation model and discussed how the general model could be applied to (a) the study of growth in multiple domains, (b) the study of binary outcomes, and (c) intervention studies.

In addition to the basic specification, we also discussed approaches to the evaluation of growth curve models—focusing particularly on the potential of growth curve modeling for prediction and forecasting. We argued that growth curve modeling could be used to develop predictions of outcomes at future time points, and we discussed the use of econometric forecasting evaluation statistics as an alternative to more traditional forms of model evaluation.

Notes

1. Clearly, other choices of centering are possible. Centering will not affect the growth rate parameter but will affect the initial status parameter.

2. LSAY was a National Science Foundation funded national longitudinal study of middle and high school students. The goal of LSAY was to provide a description of students' attitudes toward science and mathematics focusing also on these areas as possible career choices (Miller et al., 1992, p. 1).