



Published in final edited form as:

Dev Psychol. 2009 May ; 45(3): 652–676. doi:10.1037/a0014851.

New Approaches to Studying Problem Behaviors: A Comparison of Methods for Modeling Longitudinal, Categorical Adolescent Drinking Data

Betsy J. Feldman,

Graduate School of Education, University of California, Berkeley

Katherine E. Masyn, and

Department of Human and Community Development, University of California, Davis

Rand D. Conger

Department of Human and Community Development, University of California, Davis

Abstract

Analyzing problem-behavior trajectories can be difficult. The data are generally categorical and often quite skewed, violating distributional assumptions of standard normal-theory statistical models. In this paper, we present several currently-available modeling options, all of which make appropriate distributional assumptions for the observed categorical data. Three are based on the generalized linear model: a hierarchical generalized linear model (HGLM), a growth mixture model (GMM), and a latent class growth analysis (LCGA). We also describe a longitudinal latent class analysis (LLCA), which requires fewer assumptions than the first three. Finally, we illustrate all of the models using actual longitudinal adolescent alcohol-use data. We guide the reader through the model-selection process, comparing the results in terms of convergence properties, fit and residuals, parsimony, and interpretability. Advances in computing and statistical software have made the tools for these types of analyses readily accessible to most researchers. Using appropriate models for categorical data will lead to more accurate and reliable results, and their application in real data settings could contribute to substantive advancements in the field of development and the science of prevention.

Preventing mental health disorders and problem behaviors, such as delinquency, risky sexual behaviors, and substance use, in childhood and adolescence is critically important to the well-being of young people and, ultimately, to our society. Problem behaviors often occur in tandem with one another (Donovan & Jessor, 1985) and are associated with concurrent difficulties, such as family dysfunction, academic failure, and poor peer relationships during childhood and adolescence (e.g., Hawkins, Catalano, & Miller, 1992; Wiesner & Windle, 2004). Later, as these behaviors continue into emerging adulthood, avenues toward a successful life course may be shut off. This can lead to adulthood failures in areas such as work and education, as well as to physical and emotional disorders, all of which are costly to the individual and society as a whole (Hill, White, Chung, Hawkins, & Catalano, 2000; Marmorstein & Iacono, 2005; Wiesner & Silbereisen, 2003). Understanding the etiology of childhood and adolescent problem

Correspondence regarding this manuscript should be addressed to: Betsy Feldman; University of California, Berkeley; Graduate School of Education; 3659 Tolman Hall, #1670; Berkeley, CA 94720-1760. Email: bfeldman@ucdavis.edu.

Publisher's Disclaimer: The following manuscript is the final accepted manuscript. It has not been subjected to the final copyediting, fact-checking, and proofreading required for formal publication. It is not the definitive, publisher-authenticated version. The American Psychological Association and its Council of Editors disclaim any responsibility or liabilities for errors or omissions of this manuscript version, any version derived from this manuscript by NIH, or other third parties. The published version is available at <http://www.apa.org/journals/dev/>.

behaviors, in part through optimal statistical modeling of developmental trajectories, may help point the way toward more successful approaches to, and timing of, interventions (Shaw, Gilliom, Ingoldsby, & Nagin, 2003).

Problem behaviors are problematic to study for a number of reasons. Because psychopathology is a developmental process, proper understanding requires longitudinal data and analyses (Cicchetti & Toth, 1998; Sameroff, 2000). In addition, if researchers want to generalize their results to the population at large, they need to study population-based samples.¹ But problem behaviors are, by definition, rare in the general population, and the data generally, often egregiously, violate assumptions of standard normal-theory linear models. They are frequently measured on a categorical scale, and the categories are generally unevenly spaced (e.g., During the past year, how often did you drink alcohol? 1 = never, 2 = a couple of times, 3 = 1 – 3 times a month, 4 = 1 or more times a week, 5 = daily). The data tend to have large masses in the lowest category, which signifies an absence of the behavior; and, to the extent that people do display the behavior, the distributions tend to be quite skewed. Ordinal data like these are most often treated as though continuous and normally distributed. Although some researchers have suggested that categorical variables with 4 or more categories could reasonably be analyzed using models that assume normal distributions (e.g., Bentler & Chou, 1987), other research suggests that this is not the case. Rather, Dolan (1994) found that this business-as-usual approach can lead to biased estimates, incorrect standard errors, and incorrect fit statistics even when the data are symmetric; and these problems get worse to the extent that the distribution departs from symmetry (also see, e.g., DiStefano, 2002; Feldman & Masyn, 2008; West, Finch, & Curran, 1995). Treating the data as censored normal (censored at zero) is a popular way of accounting for the asymmetry caused by the pileup of zeros, but it still assumes that the data above zero are continuous, so is subject to many of the same problems that occur when the data are treated as normally distributed.

When these ordinal variables are modeled as inherently categorical, a lack of symmetry is not a problem because the proportions in each category are explicitly modeled. Additionally, the mass at the bottom of the scale is only a problem to the extent that it may, with a limited sample size, be responsible for empty cells at the high end of the scale; and too many empty cells can cause estimation difficulties. Although longitudinal or clustered categorical-data models have been available for a number of years and used in fields such as medicine (Harville & Mee, 1984; Hedeker & Gibbons, 1994), economics (Butler & Moffitt, 1982), and education (Bock & Lieberman, 1970); until recently, software that estimated these models was relatively difficult to use and the estimation process was computationally arduous. Because of advances in computers and statistical software, appropriate analysis of longitudinal categorical data has become much more straightforward for applied researchers.

In this paper, we demonstrate some of the methods appropriate for modeling trajectories with longitudinal categorical observed data. Our goal is to place these techniques into a coherent framework. Such a framework will, we believe, help developmental researchers feel more confident in selecting and using the methods. A second important contribution this paper makes is to offer a systematic approach to choosing from among competing models for longitudinal categorical data. Because there is no generally available test of overall fit for most of these models, nor any single widely-accepted criterion for model selection, we present several techniques for assessing the models. We utilize adolescent-alcohol-use data to illustrate the process of model specification, selection, and interpretation; and we guide the reader through the process of choosing an optimal model for the data, from among those investigated, by

¹Results from analyses of clinical samples can only be generalized to that population and may not apply to the general public, and use of other types of non-representative samples may also have problems with respect to inference (Hernán, Hernández-Díaz, & Robins, 2004).

comparing results with respect to fit (using comparative fit statistics and residual analysis), parsimony, and substantive considerations.

Models for Longitudinal Categorical Data

The Generalized Linear Model

In the sections that follow, we describe several models for characterizing individual differences in change or growth over time, given categorical data. Appendix A lists the models we discuss here, along with their distributional assumptions and alternative names and acronyms sometimes seen in the literature.

To introduce models for categorical data, we begin with the most basic and general form of a generalized linear model,² in which the *expected value* (or mean), μ_i , of individual i 's response, y_i (given covariate values $x_{0i}, x_{1i}, \dots, x_{mi}$), is related to the covariates through a *linear predictor*, η_i . η_i has a linear relationship with the covariates which is quantified by regression coefficients $\beta_0, \beta_1, \dots, \beta_m$ (Agresti, Booth, Hobert, & Caffo, 2000; Skrondal & Rabe-Hesketh, 2004):

$$\eta_i = \beta_0 x_{0i} + \beta_1 x_{1i} + \dots + \beta_m x_{mi}; \quad (1)$$

where, usually, $x_{i0} = 1$ (making β_0 the intercept of the equation), and there are m covariates. η_i is related to μ_i (given individual i 's values on covariates) through some *link function*, g :

$$g(\mu_i) = \eta_i. \quad (2)$$

The type of link function used depends on the conditional distribution of the data.³ For instance, when the distribution is normal (Gaussian), g is typically an *identity* function and the equation is a standard linear regression (with an identity link, the estimate, $\hat{\mu}_i$, can be thought of as \hat{y}_i):

$$\eta_i = \mu_i = \beta_0 x_{1i} + \beta_1 x_{2i} + \dots + \beta_m x_{mi}. \quad (3)$$

But if y_i has a different distribution, an identity link may not be appropriate. For example, if y_i is binary, then its distribution is Bernoulli and the expected value, μ_i , is the conditional probability of giving a positive (or correct) response: $pr(y_i = 1 | x_{1i}, \dots, x_{mi}) = \mu_i$. Using an identity link in this case is improper because it would allow the model-predicted response, $\hat{\eta}_i$, to go below zero or above one—outside of the range of probabilities. Because of this, either a logit (or logistic) link or a probit link is usually used with dichotomous data. In the logit-link model for binary data, the natural log of the odds of giving a positive response (often referred to as the log-odds or logit) is modeled as a linear function of the covariates:

$$g(\mu_i) = \ln\left(\frac{\mu_i}{1 - \mu_i}\right) = \eta_i, \text{ and} \quad (4)$$

²For extended discussions of generalized linear models, including methods for various types of non-Gaussian distributions, see, e.g., Agresti (2002), McCullagh and Nelder (1999), and Skrondal and Rabe-Hesketh (2004).

³All of these models assume that the distribution of the observed response variable is a member of the exponential family of distributions; most commonly seen among discrete distributions: multinomial (ordered or unordered categorical data), Poisson (count data), or Bernoulli (dichotomous data); and continuous distributions: gamma and Gaussian (or normal).

$$\eta_i = \beta_0 x_{0i} + \beta_1 x_{1i} + \dots + \beta_m x_{mi}, \text{ as in Equation 3.}$$

Therefore:

$$\ln\left(\frac{\mu_i}{1 - \mu_i}\right) = \beta_0 x_{0i} + \beta_1 x_{1i} + \dots + \beta_m x_{mi}, \quad (5)$$

or, equivalently,

$$\mu_i = \frac{\exp[\beta_0 x_{0i} + \beta_1 x_{1i} + \dots + \beta_m x_{mi}]}{1 + \exp[\beta_0 x_{0i} + \beta_1 x_{1i} + \dots + \beta_m x_{mi}]}. \quad (6)$$

In this paper, the logit-link model is used for ordinal data—an extension of the binary model referred to as the *ordered logit*, or *cumulative logit* model. For an ordinal variable with $j = 1, 2, \dots, J$ categories, the ordered logit model represents the probability of scoring in category j or above, versus any category below j (see, e.g., Agresti, 2002). In discussing ordered logit models, it is helpful to think of the ordinal observed outcome as a coarse categorization of an underlying (latent) continuous variable, y_i^* (see Figure 1), for which the residual errors are assumed to have a logistic distribution (Agresti, 2002; Gurland, Lee, & Dahm, 1960; Hedeker & Gibbons, 1994; Skrondal & Rabe-Hesketh, 2004).⁴ In this characterization, referred to by Skrondal and Rabe-Hesketh (2004) as a *latent response* formulation, the cutpoints that separate the underlying continuous distribution into categories are referred to as *thresholds*, or τ 's, and, the thresholds define the relationship between the categorical observed y_i (in Figure 1, $y_i = 1, 2, 3$, or 4), and the continuous latent y_i^* , such that $y_i \geq j$ if $y_i^* \geq \tau_j$, and $y_i = j$ if both $y_i^* \geq \tau_j$ and $y_i^* < \tau_{j+1}$; $j = 0, 1, 2, \dots, J - 1$, $\tau_0 = -\infty$ and $\tau_J = \infty$.

Imagine, for example, that Figure 1 is showing an ordinal measure used to assess frequency of suicidal ideation in the past year in which 0 = no ideation, 1 = thought of suicide once, 2 = thought of suicide more than once, but not frequently, 3 = contemplated suicide many times. If individual i 's latent suicidal ideation score, y_i^* (latent, because only a category is observed), is greater than or equal to τ_1 (Figure 1), then the observed value for individual i (y_i) is 1, 2, or 3 (this individual has experienced suicidal ideation). However, if y_i^* is greater than or equal to τ_1 and less than τ_2 , then $y_i = 1$ (this individual thought of suicide once). Likewise, if y_i^* is less than τ_1 , then $y_i = 0$ (not suicidal) and if it is equal to or greater than τ_3 , then individual i 's response is 3 (chronically suicidal).

The expected value of y_i^* , μ_i^* , is modeled directly with an identity link, as in Equation 3:

$$\mu_i^* = \beta_0 x_{0i} + \beta_1 x_{1i} + \dots + \beta_m x_{mi} \quad (7)$$

and,

$$y_i^* = \beta_0 x_{0i} + \beta_1 x_{1i} + \dots + \beta_m x_{mi} + \varepsilon_i \quad (8)$$

⁴In the probit model, the errors are assumed to be normally distributed (Agresti, 2002; Skrondal & Rabe-Hesketh, 2004). For an early references to similarly formulated probit models, see, e.g., Ashford (1959), Goldberger (1964).

where $\varepsilon_i \sim \text{Logistic}\left(0, \frac{\pi^2}{3}\right)$ with fixed mean and variance.

It is important to note that the ordered logit model assumes that the multiplicative effect of a covariate (observed or latent) on the odds of being in a category j , is the same for all $j = 0, 1, 2, \dots, J - 1$. This *proportional odds* assumption is necessary to permit the relationship between the ordinal variable and each covariate to be quantified by a single coefficient.⁵

Longitudinal Categorical Data

When researchers gather multiple cross-sectional or longitudinal measures on the same individuals (for example, repeatedly measuring individuals' levels of suicidality over several years), the repeated measures on an individual are correlated with one another. This violates an important assumption of statistical models—that of conditionally independent observations—resulting in standard errors that are too low. Generalized estimating equations (GEE) were developed to solve this problem with categorical data. GEE use a “working” correlation matrix (specified by the data analyst) to represent the correlations between observations, and it uses a special estimator (known as a “sandwich” estimator), to correct the standard errors. However, the GEE approach has a few limitations that hamper its usefulness in modeling psychological data. The estimates that result from GEE are *population averaged* estimates. That is, the estimates are mean effects for the population as a whole, rather than for any given individual (one coefficient applies equally to all members of the population). GEE does not allow the researcher to explicitly model the population heterogeneity, or individual differences. For instance, in the suicidality example, GEE would give only an estimate of the average intercept and slope of suicidality in the population. GEE can estimate mean differences based observed covariates (e.g., the difference between male and female averages), but not, for instance, a specific individual's estimated suicidality level or trajectory (or how it might change based on a covariate value). Nor would it permit individual prediction from the intercept and slopes to a distal outcome.

Hierarchical Generalized Linear Model

A longitudinal model for categorical data that does model the individual differences is the *hierarchical* generalized linear model (HGLM; Raudenbush & Bryk, 2002), also known as a generalized linear mixed model (GLMM; Skrondal & Rabe-Hesketh, 2004).⁶ This approach combines the generalized linear model (GLM) with a hierarchical linear model (HLM; also known as a random-effects, random-coefficients, mixed, or latent growth-curve model), in which repeated measurements on individuals are expressed as a function of time (see Figure 2). Individual differences in the outcome variable when time equals zero and change in the outcome over time are modeled by permitting the intercept and slope coefficients to vary across individuals. The intercept and slope(s) are, therefore, referred to as random coefficients, random effects, or (latent) growth factors. Both time-varying (not discussed in this paper) and time-invariant covariates may be used to explain within- and between-person variability, respectively.

In modeling the categorical drinking data, we use the logit link described in Equation 4 to Equation 6. A categorical random-effects model has been described by Hedeker and colleagues (Hedeker & Gibbons, 1994; Hedeker, Gibbons, & Flay, 1994; Hedeker & Mermelstein, 1998, 2000), and others (Agresti, 2002; Agresti et al., 2000; B. Muthén & Asparouhov,

⁵Other modeling options are possible if the proportional-odds assumption does not hold (see, e.g., Agresti, 2002; Hedeker & Mermelstein, 1998).

⁶For extended discussions of hierarchical generalized linear models, see, e.g., Agresti et al. (2000), Raudenbush, Bryk, & Congdon (2005), Singer & Willett (2003), Skrondal & Rabe-Hesketh, (2004).

2002; Raudenbush & Bryk, 2002; Skrondal & Rabe-Hesketh, 2004; Vermunt, 2006). In the models that follow, μ_{ti}^* (individual i 's expected latent response value at time t ; see Equation 7 & Equation 8), is modeled as a linear function of the random coefficients (growth factors), which comprise the intercept coefficient, β_{0i} , (the expected latent response value for individual i when time = 0) and, usually, one or more slope coefficients, β_{si} ($s = 1, 2, \dots, S$) which are multiplied by powers of time (e.g., $\beta_{1i}a_{ti}$, $\beta_{2i}a_{ti}^2$, $\beta_{3i}a_{ti}^3$, ..., $\beta_{Si}a_{ti}^S$). The slope coefficients describe systematic linear or curvilinear change over time in individual i 's latent responses.⁷ The times of measurement, a_{ti} , are usually centered by subtracting a relevant age or timepoint; often this is the time of, or age at, the first measurement occasion, as is shown in Figure 2. This makes the intercept, β_{0i} , individual i 's expected latent response at the first measurement occasion. The times of measurement may be the same for all individuals, as shown in Figure 2, or may be specified in the model to be unique for each individual (e.g., actual age).

The model that follows has linear slope (higher order polynomials are also common), and is expressed as having two levels. Level 1 characterizes the individual latent responses at each timepoint, and Level 2 characterizes the individual trajectories over time. The generalized linear model (Equation 4 – Equation 8) is at Level 1, embodied in y_{ti}^* :

$$y_{ti}^* = \beta_{0i} + \beta_{1i}a_{ti} + \varepsilon_{ti}; \quad (9)$$

where: $\varepsilon_{ti} \sim \text{Logistic}\left(0, \frac{\pi^2}{3}\right)$ with fixed mean and variance, and are assumed independent between individuals (i.e., $\text{cov}[\varepsilon_{ti}, \varepsilon_{tj}] = 0$), and conditionally independent between times (i.e., $\text{cov}[\varepsilon_{ti}, \varepsilon_{vi}] = 0$; referred to as *local independence*).

At Level 2, individual differences in the random coefficients from Level 1 (β_{0i} , β_{1i}) are represented by variability (u_{0i} , u_{1i}) around the mean intercept (γ_{00}) and mean slope (γ_{10}). The individual differences are modeled as a function of an individual-level, time-invariant covariate, x_i (multiple covariates are possible), quantified by regression coefficients γ_{01} and γ_{11} for intercept and slope, respectively. The conditional joint distribution of the intercept and slope is assumed to be multivariate normal:

$$\begin{aligned} \beta_{0i} &= \gamma_{00} + \gamma_{01}x_{1i} + u_{0i}, \\ \beta_{1i} &= \gamma_{10} + \gamma_{11}x_{1i} + u_{1i}; \end{aligned} \quad (10)$$

where: $\mathbf{u} \sim \mathbf{N}(\mathbf{0}, \mathbf{\Psi})$, $\text{cov}(\varepsilon_t, u_s) = 0$ for all $t = 1, \dots, T$ and $s = 0, 1$, and $\gamma_{00} = 0$ for identification. This identification restriction is necessary because the scale and location of the intercept are arbitrary when this latent response parameterization is used.⁸ The thresholds (τ 's) are assumed to be constant across time.

Utilizing the earlier example of suicidality, this hierarchical generalized linear model (HGLM) characterizes individuals' systematic change over time in the continuous latent response. The τ 's that relate the latent suicidality score to the observed categorical suicidality measures are estimated as part of the model (Equation 5 – Equation 8) and, because they are constant across time, we know that the latent suicidality score, despite changing over time, relates to the observed categories of the ordinal variables in the same way at each measurement occasion. The intercept mean is defined to equal zero ($\gamma_{00} = 0$) on the arbitrary scale of the latent

⁷For models in which change over time is permitted to be non-linear, see, e.g., Blozis (2004, 2007), and Ferrer and McArdle (2003).

⁸For examples of alternative parameterizations, in which different restrictions on the thresholds permit γ_{00} be estimated, see Mehta, Neale, and Flay (2004), Millsap and Yun-Tein (2004), and Skrondal and Rabe-Hesketh (2004).

suicidal-ideation response (necessary in this parameterization); however, the estimated individual intercepts vary around zero. The individual slopes vary about the estimated average slope, as well, and these individual intercepts and slopes are often correlated.

Growth Mixture Model

Several assumptions have been discussed that are very important for the hierarchical generalized linear model (HGLM): (1) the random coefficients (intercept and slopes) are assumed to have a multivariate normal distribution; (2) change is a smooth function of time (most often linear or curvilinear); (3) proportional odds; and (4) measures are independent, conditional on the random effects (local independence). Any of these assumptions might not hold, especially when modeling problem-behavior data. In the models that follow, we discuss methods that relax the first three assumptions.

To begin, when the data are categorical and highly skewed, it is possible that the individual intercepts and slopes are not normally distributed, potentially resulting in biased parameter estimates (B. Muthén & Asparouhov, 2006; Vermunt, 2006). However this assumption is difficult to test (Agresti, 2002, p. 496; Carlin, Wolfe, Brown, & Gelman, 2001; Vermunt, 2006). The Level 2 normality assumption can be relaxed by using a *mixture* of normal distributions, known as mixing components or *classes*, to characterize the joint distribution of the random effects. This is pictured in Figure 3a. This model is a type of finite mixture model (see McLachlan & Peel, 2000), which we refer to here as a (generalized linear) *growth mixture model* (GMM), and it assumes only within-class normality.⁹ The classes have different mean structures at Level 2, and can comprise different proportions of the population. The within-class Level-2 variance-covariance matrix can also differ across classes, however, under some circumstances, this may lead to estimation problems, resulting from an unbounded likelihood (McLachlan & Peel, 2000). The multinomial latent class variable, c_i , which identifies individual i 's class membership, takes on values $k = 1, 2, \dots, K$, where K is the number of classes. Here, the τ 's are assumed to be constant across both time and classes.

Figure 3a shows how the individual differences are, in this model, characterized in part as discrete classes, and in part as continuous distributions within those classes. Because of this within-class variability, covariates can be used to predict individual differences within class, as well as to predict the probability of class membership. The hierarchical generalized linear model (HGLM) can be seen as a special case of the growth mixture model (GMM), with just a single class ($K = 1$). When GMM has two or more classes, Level 1 equations are indexed by class:

$$y_{kti}^* = \beta_{k0i} + \beta_{k1i} a_{ti} + \varepsilon_{kti}; \quad (11)$$

for individual i at time t in class k ; ε_{kti} distributed as before.

The Level 2 equations are also indexed by class, and a multinomial logistic equation is added, which expresses the conditional probability of membership in class k as a function of a class-specific intercept, α_k and the covariate, x_i , multiplied by coefficient, λ_k :

$$\begin{aligned} \beta_{k0i} &= \gamma_{k00} + \gamma_{k01} x_i + u_{k0i}, \\ \beta_{k1i} &= \gamma_{k10} + \gamma_{k11} x_i + u_{k1i}, \end{aligned} \quad (12)$$

⁹Muthén and Asparouhov (2006) also refer to this model as a growth mixture analysis (GMA; for more on GMM with continuous outcomes, see, e.g., Muthén, 2004; Muthén et al., 2002).

$$\text{pr}(c_i=k|x_i) = \frac{\exp[\alpha_k + \lambda_k x_i]}{\sum_{h=1}^K \exp[\alpha_h + \lambda_h x_i]}; \quad (13)$$

where $\gamma_{K00} = 0$ and $\alpha_K = \gamma_K = 0$ for identification; \mathbf{u}_{ki} distributed as before.

This model has a more complicated interpretation than the hierarchical generalized linear model (HGLM). In the suicidality example, if there are, for instance, three classes, it implies three unique sets of intercept and slope means; one for each class. Within each class, individual intercepts and slopes vary about those means and can be predicted by covariates, as can class membership. The grouping of individuals into classes, however, is seldom perfect. Instead, person i has some estimated probability of membership in *each* suicidality class. These are referred to as *posterior probabilities*. For each individual, the posterior probabilities sum to 1 and are based on that individual's observed outcome responses and covariate values. Thus, person i 's estimated trajectory of suicidality is a function of the probability of membership in each of the classes, the class-specific means and covariance structures, and possibly covariates.

Latent Class Growth Analysis

There are several reasons that the within-class normality assumption of the growth mixture model (GMM) may not be sustained by the data. The primary one is that there may be insufficient variability to estimate the within-class variances and covariances in one or more classes. Two or more classes can be collapsed together, but if the classes are well separated (i.e., the mean trajectories are very different), then collapsing them might not be optimal. In this case, it is possible to constrain one or more elements in the within-class covariance matrix (Ψ_k) to zero. When all variances and covariances in a class are fixed at zero, it means, effectively, replacing a continuous, normally-distributed class with a mass point, as shown in Figure 3b. We refer to this model as a (generalized linear) *latent class growth analysis* (LCGA) and, like the hierarchical generalized linear model (HGLM) this model is a special case of GMM.¹⁰ The Level 1 equation is the same as that in GMM, but the Level 2 equations change to reflect the elimination of within-class variances (and with no variation to explain, within-class predictors are also dropped):

Level 2:

$$\begin{aligned} \beta_{k0i} &= \gamma_{k00}, \\ \beta_{ki} &= \gamma_{ki0}. \end{aligned} \quad (14)$$

When all within-class variances and covariances equal zero, individual differences at Level 2 are being modeled solely by class membership, so it is likely that more classes will be needed to adequately characterize the variability in the random effects distribution (B. Muthén & Muthén, 2000b). The probability of membership in each of the K classes may be regressed upon the covariates as in Equation 13, with the same identification constraints. Figure 4a shows a path diagram of a growth mixture model (GMM) and a latent class growth analysis (LCGA). The solid lines belong to both models and dashed lines show the optional parameters that, if any are present, define the model as GMM (for more on the comparisons between these models, see, e.g., Kreuter & Muthén, 2007; Nagin, 2005; Raudenbush, 2005; Reinecke, 2006).

¹⁰This approach is sometimes referred to as “semiparametric group-based modeling” (Nagin, 1999; see Appendix A).

This model has a slightly simpler interpretation than the GMM because all individual differences in estimated suicidality trajectories are characterized by class membership, and those differences can be modeled as a function of predictors of class membership. Within each class, all individuals are assumed to share the same suicidality trajectory. One way to think about the shared within-class trajectory is that it is an approximation of the true within-class individual latent trajectories, similar to the way mass points approximate the true distribution of the random effect shown in Figure 3b. Person i 's predicted trajectory of suicidality is a straightforward function of his or her posterior class membership probabilities and the class intercept and slope means.

Longitudinal Latent Class Analysis

Finally, it is possible that development cannot be characterized as a simple function of time, as is assumed in the models we have been discussing, or that the proportional odds assumption does not hold. All of the models thus far are bound by these two assumptions because they model *scaled* growth (change across time is along a single continuous scale). Longitudinal latent class analysis (LLCA; latent class analysis applied to longitudinal data) is a type of mixture model but it does not belong to the family of growth models we have been discussing. This is because, rather than modeling scaled change, LLCA models patterns of *states* across time.¹¹ If the outcome data are multinomial, LLCA is the only appropriate analytic approach because any changes are necessarily state changes.

Longitudinal latent class analysis (LLCA) models the joint distribution of the repeated outcome measures directly with a latent class variable, which characterizes both the within-person variation that had been handled at Level 1 of the previous growth models, and the between-person differences that had been handled at Level 2 (see, also, Appendix A). Its only assumption is local independence. In LLCA, the time- and class-specific probability of scoring in or above category j is modeled directly, so the thresholds are indexed by time (t) and class (k), in addition to category:

$$\text{pr}(y_{it} \geq j | x, k) = \frac{\exp[-\tau_{tjk}]}{1 + \exp[-\tau_{tjk}]}, \quad (15)$$

and

$$\text{pr}(y_{it} \geq j | x) = \sum_{k=1}^K \left[\left(\frac{\exp[-\tau_{tjk}]}{1 + \exp[-\tau_{tjk}]} \right) \text{pr}(c_i = k | x_i) \right]. \quad (16)$$

Class membership is based on groupings of similar patterns of responses over time, and covariates can be used to predict class membership probabilities, as shown in Equation 13. Note that, because it is the category thresholds that define the differences both across time and across classes in longitudinal latent class analysis (LLCA), they are all estimated. Thus, while LLCA makes no assumptions about the distribution of the observed variables, or the form of change, it trades fewer assumptions for many more estimated parameters (Vermunt, 2006), in a manner similar to spline or piecewise models (e.g., McArdle, 2004). Additionally, the model,

¹¹LLCA differs from latent transition analysis (LTA) because LTA models state changes across consecutive timepoints, rather than patterns that span several timepoints (see, e.g., Collins, Graham, Rousculp, & Hansen, 1997; Graham, Collins, Wugalter, Chung, & Hansen, 1991). As such, it answers different types of questions than LLCA and the other models discussed here, just as autoregressive models answer different types of questions than latent growth curves.

as parameterized here, does not permit unique, individually varying times of measurement. For an application of LLCA to drinking patterns across time, see Lanza and Collins (2006).

In a longitudinal latent class analysis (LLCA) of the suicidality example, patterns of change are defined by the variable categories, rather than estimated intercepts and slopes. Like the latent class growth analysis (LCGA), everybody in a class shares the same parameters, however now it is the category thresholds that are shared within class and different between classes. Because thresholds can also change across time, change can be much more irregular or complex in LLCA. For instance, the course of suicidality for individual i might fluctuate or be intermittent. Differences in suicidality could, perhaps, be associated with influences such as puberty, high-school graduation, or other shared events; and the covariates can be used to predict the individual's probability of membership in each class.

Adding Covariates

The selection and modeling of covariates in mixture models is a complex issue currently under investigation, and an extended discussion of the issues is beyond the scope of this article. However, early work (Nylund & Masyn, 2008) suggests that the number of classes should be determined using an unconditional model, and that substantial changes in parameter values after covariates are included indicate misspecification of their effects (e.g., omitting direct effects of covariates on observed variables). As is the case when estimating any statistical model, it is also important to let substantive theory, along with consideration of parsimony and interpretability, inform the inclusion and specification of covariates. The added complexity of mixture models makes this a particularly critical issue. For example, in a growth mixture model, it is theoretically feasible to allow a covariate to affect class membership as well as one or more growth factor variances in one or more classes. Additionally, all of the standard rules for typical regression and HLM models still pertain (e.g., concern with multicollinearity among covariates, the need to dummy-code nominal predictors).

Applying the Models

Longitudinal mixture models (GMM, LCGA, and LLCA) are often used to find unobserved but distinct groups of individuals (e.g., Lanza & Collins, 2006; B. Muthén & Muthén, 2000a; Shaw, Lacourse, & Nagin, 2005). However, any non-normal distribution can be approximated by a finite mixture of normal distributions (Bauer & Curran, 2003; McLachlan & Peel, 2000), so the extraction of two or more classes does not *necessarily* indicate that distinct groups exist in the population. It is possible that finding multiple classes is an indication of a non-normal distribution of the random effects or, potentially, some other violation or model misspecification (Bauer, 2005). In addition, class membership is seldom unequivocal. Rather, as mentioned earlier, the individual has an estimated probability of belonging to each class.

In this paper, we use actual alcohol-use data, collected from a sample of adolescents when they were in seventh through twelfth grades, to illustrate the similarities and differences between these analytic approaches and to motivate the process of model selection (see Appendix B for annotated Mplus syntax). The value these models is enhanced to the extent that childhood events and circumstances may be used to help to predict later pathological processes, as this can potentially allow targeted interventions to take place before such processes begin. In addition, under some circumstances, covariates can aid in model selection. Thus we include a covariate that has been found to predict adolescent alcohol use—association with alcohol-using peers (e.g., Curran, Stice, & Chassin, 1997; Hawkins et al., 1992)—to demonstrate how a covariate may be specified and interpreted in the models.

Method

Participants and Procedure

Data come from a longitudinal, community epidemiological study of 451 (52% female) target youths and their families from two-parent households in the Midwest. Participants in the study were recruited from eight rural counties and the original sample of families was primarily lower middle- or middle-class. Due to a very small minority population in this area, all participants are of European heritage. Annually from 1989 (seventh grade) through 1992 (tenth grade) and again in 1994 (twelfth grade) the adolescents and participating family members individually completed questionnaires pertaining to subjects such as the demographic characteristics of the family, their own personal characteristics and behavior, and characteristics and behavior of other family members and friends. The retention rate in 1994 (the last wave used in the current study) was 94%. Additional details about the study, which is still ongoing, can be found in Conger and Conger (2002).

The primary outcome for these analyses is the alcohol use measure from the questionnaire data. Target adolescents were asked each year how often they had consumed beer, wine, or hard liquor during the previous year. The drinking items were coded into a single, ordinal alcohol-use variable with four possible responses: 0 = never, 1 = less than weekly, 2 = once or twice a week, and 3 = three or more times a week. Each year, the target adolescents were also asked how many of their close friends used alcohol during the previous year. Answers ranged from 0 = “none of them,” to 4 = “all of them.” The adolescent’s answer to this question in seventh grade was used as a covariate, predicting the intercept and change in the target adolescent’s own drinking.¹²

Statistical Models

We began the unconditional analyses with hierarchical generalized linear model (HGLM; Figure 2), using both linear and quadratic slopes, and utilizing the logit link and model shown in Equation 9 and Equation 10. This model has the simplest interpretation and it gives an overview of the form of change. Following HGLM, we tested the three mixture models, beginning with growth mixture model (GMM; Equation 11 – Equation 13). GMM comprises a very large set of potential models. For example, a quadratic model with three classes has 14 possible ways of specifying the random effects if they are constrained to be equal across classes, and dozens more if they are permitted to vary across classes. To limit the number of models we were testing, we decided to constrain variances equal across classes, as discussed earlier, and to use a systematic approach to relaxing the constraints on the variance/covariance matrix. For each model, we tested the following: (a) intercept variance only, (b) intercept and linear slope variances, and (c) intercept, linear, and quadratic slope variances. We tested each model both with all possible covariances estimated, and with all covariances constrained to zero.

All intercept and slope variances were fixed to zero for latent class growth analysis (LCGA; Equation 14), and, finally, we modeled the joint distribution of the repeated drinking outcomes directly with longitudinal latent class analysis (LLCA; Equation 15 & Equation 16). Intercepts for HGLM, GMM, and LCGA are placed at ninth grade—the first year of high school. Entry to high school begins a time of increased risk for developing problem drinking. Incoming students are exposed to older adolescents, many of whom have already begun to drink, and alcohol is much more available in high school (Johnston, O’Malley, Bachman, & Schulenberg, 2006).

¹²Allowing covariates to affect the slope means that the effect of the covariates on the outcome changes across time; however, the covariate is assumed to affect all individuals in the same fashion and is known as a “fixed effect.” Alternatively, we could constrain the covariate effect to be time-invariant but allow that effect to vary across individuals (as a random effect).

For all analyses, we used a full-information maximum likelihood (FIML) estimator with robust standard errors (MLR), as implemented in Mplus 4.21 (B. Muthén & Muthén, 1998 – 2006), to estimate the parameters (the robust standard errors offer protection against inflated alpha values in case of certain types of model or distributional misspecifications). With any longitudinal study, some individuals will miss some assessments or drop out altogether. These models utilize all of the available data under a missing-at-random (MAR) assumption (as would multiple imputation), which allows that the missingness may be related to variables included in the analysis (Little & Rubin, 2002). Unless missing data are missing completely at random (MCAR; i.e., missingness is unrelated to any data being modeled, present or missing), which is a much stronger assumption, FIML yields less biased estimates than other methods of missing-data methods, such as listwise deletion (which also reduces the power of the analysis by reducing sample size) or mean imputation (Schafer & Graham, 2002). FIML does not, however, guarantee unbiased estimates if missingness is related to the missing values, themselves, and this risk of bias increases with the proportion of missingness (the maximum for a single wave in this study is 10.5%). However, we feel that, given the broad range of topics under study, MAR is a reasonable assumption (for more on mechanisms of missingness and techniques for dealing with missing data, see, e.g., Foster, Fang, & Conduct Problems Prevention Research Group, 2004; Little, 1995; Little & Rubin, 2002; Schafer & Graham, 2002). FIML can be computationally demanding with categorical data, however, it offers advantages over methods such as penalized quasi-likelihood (PQL) and marginal quasi-likelihood (MQL), which, while computationally less demanding, may result in biased estimates.

Model Assessment and Selection

Any given data set can potentially be analyzed using a multitude of statistical models, however, there are generally theoretical and statistical reasons that many possible models can be ruled out. For example, we are not considering any models that require an assumption of continuous observed data (e.g., two-part models; see Olsen & Schafer, 2001; Tooze, Grunwald, & Jones, 2002). Furthermore, because we are interested in modeling change across multiple timepoints, rather than pairs of timepoints, we are not considering any autoregressive-type models, such as latent transition analysis (LTA).

For most longitudinal categorical analyses, there are no readily available statistics that assess overall fit of the model to the data. The two categorical goodness-of-fit statistics, the Pearson chi-square and likelihood chi-square (also called deviance, or G^2) may not perform well when there are many low expected cell frequencies (Agresti, 2002), and this is often the case when longitudinal categorical data are used (e.g., five waves of a 4-category variable results in 4^5 or 1024 cells, most of which are likely to be empty). Nor is there a method for comparing models that is widely accepted as best (B. Muthén & Asparouhov, 2006; Nylund, Asparouhov, & Muthén, 2007). Thus, when assessing models, researchers generally consider both statistical and substantive criteria (e.g., Jackson, Sher, & Schulenberg, 2005; Tucker, Orlando, & Ellickson, 2003). In this section, we will discuss four complementary approaches to model assessment: (1) quality of convergence; (2) comparative fit; (3) residual analysis; and (4) visualization. In addition, issues such as parsimony and, in mixture models, class size should be taken into consideration. When a class comprises a small number of individuals, the parameter estimates may not be very stable or reliable. Often, no single criterion is, alone, sufficient to select a model and two or more of the criteria we discuss must be considered simultaneously.

Convergence—The estimation algorithm searches for the global maximum of the likelihood function, but in complex categorical data models and mixture models, algorithms are more likely to converge on local maxima than with continuous data or in less complicated models,

so the use of multiple starts from random locations in the parameter space is recommended (Hagenaars & McCutcheon, 2002; Hipp & Bauer, 2006; McLachlan & Peel, 2000; B. Muthén & Asparouhov, 2006; Vermunt, 2006; see Appendix B). In some cases, there may not be a global maximum (McLachlan & Peel, 2000). This can result in non-convergence, or failure of the algorithm to replicate a maximum loglikelihood value over many starting values. That is, each start may end up at a different local maximum, suggesting that the model parameter estimates are untrustworthy. A second indication that adequate convergence was not reached is a failure of the algorithm to generate standard errors. This can be brought about by a singular information matrix, which implies that the model is not identified (L. K. Muthén & Muthén, 1998–2006). A model that has not converged, has resulted in a singular information matrix, or has failed to yield a consistent maximum loglikelihood value is considered a failed model, in this paper, and eliminated from further consideration.

Comparative fit—The standard chi-square difference test (likelihood ratio test; LRT) is not helpful in determining which model is best, because the different models are not necessarily nested. It also cannot help choose the optimal number of classes in a mixture model, because regularity conditions of the test are violated when comparing a k -class model to a $(k - 1)$ -class model (McLachlan & Peel, 2000). Often, an information index, such as the Bayesian Information Criterion (BIC; sometimes called Schwartz Information Index) is used in model selection. This index and similar ones (e.g., Akaike Information Criterion) take into account the model loglikelihood (higher is better), and penalize for model complexity (i.e., the number of parameters estimated relative to the sample size). Thus, using them reduces the risk of overfitting the model to a single sample, thereby improving the possibility of replicating the model findings with future samples. In general, a lower value on an information criterion indicates a better model. Mplus reports three information indices, but the most widely used is the BIC (see Hagenaars & McCutcheon, 2002; Nylund et al., 2007; Vermunt, 2006).

Two more statistics, available in Mplus 4.21 (B. Muthén & Muthén, 1998 – 2006), may be useful for determining the optimal number of classes in a mixture model. The Lo-Mendell-Rubin test (LMR; B. Muthén, 2004) analytically approximates the LRT distribution and the bootstrapped LRT (BLRT), suggested by McLachlan and Peel (2000), uses bootstrap samples to empirically derive the sampling distribution of the LRT statistic (Nylund et al., 2007). Both tests compare a k -class model with a $(k - 1)$ -class model and, in both cases, a statistically significant p -value suggests the current model offers improvement over the model with one class fewer.

Also discussed in the literature on mixture models is entropy (B. Muthén, 2004), which, similar to the average posterior probabilities for most likely class membership, serves as a measure of the precision of individual classification. It ranges from 0 (everybody has an equal posterior probability of membership in all classes) to 1 (each individual has posterior probability 1 of membership in a single class, and probability 0 of membership in the remaining classes). Equivalently, high entropy indicates clear class separation. Entropy is not a measure of fit, nor was it originally intended for model selection (Ramaswamy, Desarbo, Reibstein, & Robinson, 1993); however, if it is extremely low, that suggests the model may not be useful for some purposes. For instance, if the model is intended to find homogenous clusters of individuals with distinctive patterns of change (e.g., Nagin, 1999), low entropy indicates that it may be doing a poor job. However, the reverse is not necessarily true; that is, high entropy does not necessarily show the existence of homogeneous clusters of individual trajectories (Feldman & Masyn, 2008).

Residual analysis—The observed and model-predicted cell and marginal proportions can be compared to help assess fit of the model to the data (Carlin et al., 2001; Skrondal & Rabe-Hesketh, 2004). As mentioned earlier, when a model is large and complex, such as a growth

process with multiple timepoints, most cells have expected frequencies that are very small or approach zero. However, often, a few patterns of responses recur relatively frequently. In this case, it may be helpful to inspect standardized Pearson residuals (Haberman, 1973) in the cells with these more common response patterns (B. Muthén & Asparouhov, 2006). In addition, standardized residuals from the univariate margins (i.e., category proportions at each timepoint) and bivariate margins (i.e., proportions in cells of all two-timepoint cross-tabulations) may be inspected (Skrondal & Rabe-Hesketh, 2004). Standardized residuals can be compared with a standard normal distribution (Haberman, 1973); thus, too many standardized residuals greater than 2 suggests poor fit of the model to the data. All of these observed and expected cell proportions, along with standardized residuals, can be generated in Mplus with the TECH10 output command (see Appendix B). Instructions on how to calculate the expected frequencies when other software is used are available from the first author.

Plots—Graphical methods can give a clearer picture than numeric output alone of model-predicted trajectories and patterns of change, but using them effectively for model selection depends upon researcher judgment and substantive knowledge. For instance, plots may make it easier to judge whether results of an analysis make interpretive sense or map reasonably onto theory and previous findings. Mixture models, in particular, must be considered in the context of theory, interpretability, and usefulness, because they tend to be used in an exploratory fashion and are innately data-driven.

Linear (and curvilinear) trajectories based on the (within-class) means of the random effects can be plotted, and these give an easy-to-describe picture of the overall mean change; but because the scale of these parameters is arbitrary, the plots can potentially be difficult to interpret. To facilitate interpretation in our line plots, we include the estimated thresholds (τ_j) that divide the categories of the observed data. In addition, we show plots of the time- and class-specific estimated category probabilities, which give a more nuanced picture of how alcohol use changes over time. The predicted-probabilities plots were created in Excel using the TECH7 output from Mplus, and line plots were generated in Excel by using Mplus estimated intercept and slope means multiplied by time (and the model-estimated thresholds, which are fixed across time).

Results

The sample size for these analyses is 451 (236 females). The observed category proportions at each grade are shown in Figure 5, broken out by gender. Nearly 30% of the target adolescents of both genders reported drinking in seventh grade, and this grew to approximately 65% by twelfth grade (Figure 5). Going along with this increase in the likelihood of drinking is an increase in the proportion of youths drinking heavily. By the time the adolescents graduated from high school, approximately 20% of them (17% of females, 21% of males) reported drinking weekly or more often (Figure 5). In addition, out of 451 seventh graders, approximately 1/3 (27% of females and 34% of males) reported having at least one close friend who drank alcohol and 5% of females and 7% of males reported that half or more of their friends drank. Because drinking did not differ by very much between males and females (Figure 5), we did not use gender as a covariate in the model.

Model Estimates and Fit Statistics

Fit indices are shown in Table 1 for the all models. Underlined values identify the best fitting number of classes within model type, as selected by the statistic in the row.

HGLM—Two unconditional hierarchical generalized linear models were tested. The first characterized in drinking-frequency trajectories over adolescence with an intercept (in ninth

grade) and linear slope, and the second added a quadratic slope. The quadratic slope significantly improved the model fit ($\Delta\chi^2 = 49.83$, $\Delta df = 4$, $p < .001$) and was retained. Based on this result, the mixture models that follow were also specified with quadratic slopes. The intercept mean was fixed to zero for identification (Equation 9 & Equation 10) and the average linear trend was positive and statistically significant ($\beta_{10} = 0.84$, $p < .001$; see Table 2). The mean quadratic slope was not significant, however, there was considerable individual variability around the intercept and both slopes. In addition, the statistically significant covariances suggest that the frequency with which adolescents were drinking in ninth grade was related to changes in their drinking over the 6-year timespan (Table 2). On average, more drinking or more rapidly increasing drinking in ninth grade was associated with trajectories that were flattening out by the end of high school; while less frequent drinking in ninth grade was associated with drinking trajectories that were accelerating in twelfth grade. The BIC for this model was the lowest of all models tested. The other three statistics in Table 1 do not apply unless there are two or more classes.

GMM—In general, when more within-class variances and covariances were estimated in the growth mixture models, fewer classes were successfully extracted. For example, when all three variances and their covariances were estimated even the 2-class model failed to converge, but with the covariances fixed to zero, the same 2-class model did converge (see Table 1) Increasing the number of classes to three resulted in a singular information matrix.

In the next model, the quadratic variance was constrained to zero and only intercept and linear slope variances were estimated. The 2-class models converged, both with and without the covariance, but there was no significant difference in fit between the two ($\Delta\chi^2 = 0.59$, $df = 1$, $p = .44$), so we constrained the covariance to zero. Here, again, the model did not converge with three classes.

In the last set of GMM models only the intercept variance was estimated, and the remaining variances (and, necessarily, all covariances) were constrained to zero. In this case, we were able to estimate up to four classes. The 5-class model did not converge (the information matrix was singular; Table 1) and we ruled out the 4-class model because one class was too small to estimate reliably (4%, representing approximately 18 individuals). This left plausible models with two and three classes. In comparing the 2- and 3-class models, the fit statistics did not agree on which was optimal: the LMR preferred two classes, the BIC three classes, and the BLRT no fewer than four classes (Table 1). It is worth noting that entropy for all of the GMM models was much lower than for it was for any other mixture models (Table 1). We retained all of the models the converged except the 4-class GMM.

LCGA—As mentioned earlier, in latent class growth analysis, because all within-class variances are constrained to zero, more classes may be needed to characterize the same joint distribution of growth factors. Equivalently, as the number of classes increases, the latent class variable is likely to account for more and more between-individual variability. Although a maximum of only four classes could be estimated with GMM, seven classes were successfully extracted with LCGA. The 7-class LCGA had a class that was too small to estimate well (3%, representing approximately 12 individuals). Nonetheless, the BLRT was still significant at seven classes, suggesting that no fewer than seven classes would be needed to model these data (Table 1). The BIC and LMR both selected the 4-class LCGA, but the 3-class model had a BIC value that was nearly equivalent to that of the 4-class model. We retained the 3- and 4-class models.

LLCA—The longitudinal latent class analysis failed to replicate the best loglikelihood when seven classes were specified, and none of the fit statistics selected more than four classes (Table 1). The BLRT suggested that four classes were adequate, and the BIC and LMR both selected

the 2-class model. Because the LLCA models require so many parameters, it is not surprising that no fit index selected a model with more than four classes, nor that its BIC values were consistently higher than for any other models (as was entropy, suggesting cleaner class separation when more parameters were estimated). We retained the models with two, three, and four classes.

Residuals

We inspected the univariate and bivariate residuals for the HGLM and the best 10 out of the 17 mixture models that had converged properly. Table 3 shows the frequencies of univariate and bivariate residuals greater than 2.¹³ All models were able to replicate the univariate observed values adequately, but there was more variability among the bivariate residuals (Table 3). The 3- and 4- class LLCA models fit the bivariate cell counts the most accurately and the worst-fitting were the 2-class models: LLCA and the three GMMs (Table 3).

Table 4 shows comparisons of observed and predicted cell frequencies for the most commonly observed patterns of responses over the five occasions. Note that, although each of the four categories of the outcome were endorsed by at least a few individuals at every timepoint, the most commonly observed patterns comprised only zeros and ones (representing non-drinkers [0] and infrequent drinkers [1]). Here, again, the 3- and 4-class LLCA modeled offer the best fit to the observed data. Only the 2-class LLCA model did a noticeably worse job than the other models under consideration at reproducing the cell counts. It is also worth noting that the frequency of one response pattern (01111) was not reproduced well by any model except the most highly parameterized LLCAs.

In summary, the bivariate residuals suggested that all of the 2-class models fit the data relatively poorly, and the 2-class LLCA also failed to adequately match the common response patterns. We dropped the 2-class models from further consideration. In contrast, 3- and 4-class LLCAs both fit the observed frequencies particularly well, but the 4-class model required 16 more parameters, had a higher BIC (Table 1). Thus, we rejected the 4-class LLCA in favor of the more parsimonious 3-class model. We continued to consider the HGLM, the 3-class GMM with intercept variance, the 3-class LLCA, and 3- and 4-class LCGAs.

Plots

Plots of model-predicted trajectories and patterns of change can also help a researcher choose the most appropriate or reasonable model from several competing ones, and plots are often used this way in model selection (e.g., Greenbaum, Del Boca, Darkes, Wang, & Goldman, 2005; B. Muthén & Muthén, 2000b). Figure 6 shows a comparison of the estimated mean curves, based on the y^* metric, generated by the 3-class growth mixture model (GMM; top) and 3-class latent class growth analysis (LCGA; bottom). Because the log-odds trajectories are on arbitrary scales, the plots show category thresholds as dashed lines to facilitate comparison across analyses. In addition, the GMM plot includes *SD* bars above and below the class means at each timepoint (based on the intercept variance). These *SD* bars illustrate the overlap of classes that accompanies the extremely low entropy. The GMM finds one unexpected class, comprising slightly over 1/3 of the sample, in which average drinking probabilities are constant across adolescence. This class shows approximately 50% of the individuals reporting alcohol use at each time (the first threshold separates non-drinkers from drinkers). Its other two classes are increasing, one early and one later, which is consistent with the covariances found in the hierarchical generalized linear model (HGLM; Table 2). The LCGA plot shows one class with consistent low drinking probabilities, one with rapidly increasing probabilities, and one with

¹³Absolute size of standardized residuals can also be taken into consideration, if they are very large (e.g., > 10), which did not occur in these analyses.

relatively high drinking probabilities across all timepoints. These last two classes are also consistent with the HGLM covariances in Table 2. That is, one class has lower drinking probabilities in ninth grade, but those probabilities are increasing rapidly, and the other, more likely to be drinking (or drinking more) in ninth grade, is increasing more slowly (and shows more quadratic slope; Figure 6).

The second way to plot the trajectories and patterns of change is to show the predicted category probabilities from each time (see Figure 7). This gives a somewhat richer picture of the change because it illustrates the shifts in proportions at each level of a categorical variable. It complements the plots of y^* , shown in Figure 6, and when the categories are not evenly spaced, or have different substantive meanings, this approach can be more informative than a line plot. Figure 7 shows a comparison of the 3-class GMM, LCGA, and longitudinal latent class analysis (LLCA, which *only* yields category probabilities). The LLCA and LCGA classes are very similar to one another, both in class-specific patterns of change and size, suggesting that the extra parameters in LLCA are not yielding much more information. Both are consistent with prior literature (e.g., Hix-Small, Duncan, Duncan, & Okut, 2004; White, Johnson, & Buyske, 2000), in which researchers have found classes with low, high, and increasing alcohol-use probabilities across adolescence.

Only GMM found the class with constant moderate drinking probabilities, and Figure 7 shows that this class groups together individuals who responded in all four drinking categories. Because this class fails to differentiate high versus low drinking probabilities (all categories are represented in the class), it may not be very informative. It is also not consistent with prior literature. In addition to these problems, the overlap of classes in GMM makes it difficult to interpret, and a mixture model with such indistinct classes may offer little if any advantage over the more standard hierarchical generalized linear model (HGLM) without classes. For all of these reasons, we eliminated it from further consideration.

The 4-class LCGA was selected by both LMR and BIC as the best LCGA model, however, the BICs for the 3- and 4-class models were very close (Table 1) and, overall, the two models did approximately equally well with respect to the standardized residuals (Table 3 & Table 4). Figure 8 shows the small added class (12%; labeled “High at A”) which is largely a subset of the low class from the 3-class model, but with higher drinking probabilities in the seventh and eighth grades (such class splitting can be investigated with a cross-tabulation of modal class assignments). By tenth grade, it does not differ substantially from the low class. This type of pattern, in which adolescents drink more in the early and later years, but not in between, is not predicted by theory or prior literature (e.g., Colder, Campbell, Ruel, Richardson, & Flay, 2002; Hill et al., 2000; Orlando, Tucker, Ellickson, & Klein, 2005). It is our feeling that this added class arises as an artifact of a testing effect that is well-established among self-report measures in longitudinal studies; higher levels of negative feelings and behaviors are often reported at the first wave of data collection (see, e.g., Knowles, Coker, Scott, Cook, & Neville, 1996; Twenge & Nolen-Hoeksema, 2002). An inflated first-year report could explain the added class, so we decided to eliminate the 4-class LCGA.

At this point, based on combinations of fit statistics, residuals, considerations of parsimony, and plots, we felt that all 2- and 4-class models, as well as all GMMs could be reasonably ruled out for this data set, leaving the 3-class LCGA and LLCA models and the HGLM still under consideration.

Adding Covariates

After selecting the best unconditional models, we used the adolescent’s seventh-grade report of alcohol consumption by close friend to explain individual differences and to help guide model selection. The covariate was allowed to predict class membership in the mixture models

and the growth-factor variability in the hierarchical generalized linear model (HGLM). If a growth mixture model (GMM) were still under consideration, it would be possible to use the covariate to explain within-class variation, as well as the probability of class membership. In the case of GMM, as in the other cases, covariates are used to explain individual differences. When a covariate is permitted to affect both within-class and between-class variability, it is predicting a complex combination of an individual's probability of class membership and that individual's trajectory, relative to the mean trajectory in each class in which the individual has some non-zero probability of membership. This can make the covariate's effects potentially difficult to interpret.

In the HGLM, targets with more close seventh-grade friends drinking alcohol were, themselves, predicted to be drinking more in ninth grade ($\beta_{01} = 1.36$, $SE = 0.24$, $p < .001$), but it was also predicted that their drinking would be increasing more slowly in ninth grade ($\beta_{11} = -.19$, $SE = 0.07$, $p = .011$). Despite this effect, however, a target who reported that even a few of his or her friends were drinking in seventh grade (i.e., 1 on the 0 to 4 scale) would still end up, on average, drinking more in 12th grade than an adolescent who reported having no drinking friends in seventh grade (see Figure 9).

The regression coefficient for a single class represents the change in the log-odds of membership in that class, versus a reference class, for a one-unit change in the covariate (for members of those two classes). The reference class can be any class, and Mplus 4.21 (B. Muthén & Muthén, 1998–2006) gives results with each class, in turn, treated as the reference. Viewing results this way permits comparison of all pairs of classes with respect to the covariate, which may be more useful than simply using a single reference class. Note, however, that, as the number of classes increases, testing all possible pairs may not be necessary or useful, and could introduce the possibility of an inflated Type-1 error rate. Additionally, global tests of association between the latent class variable (with a given number of classes) and one or more covariates can be conducted using the standard likelihood ratio test for nested models.

Table 5 shows all paired class comparisons in 3-class LCGA and LLCA models. The results across the two models are very similar. Members of the high class reported more alcohol-using friends in seventh grade than either of the remaining two classes, which did not differ significantly from one another. In deciding between LLCA and LCGA 3-class models, we weighed the better fit of the LLCA against its lack of parsimony; as compared with the LCGA (the LLCA requires 34 more parameters; Table 1). If the LLCA patterns of change had shown noticeable signs of non-linearity, then it would have made sense to choose that model over LCGA. However, it did not (Figure 7), and the substantive stories told by the trajectory plots and the relationships between classes and the covariate are the same for the two models. Based on all of these factors, we retained the LCGA and dropped the LLCA. Parameter estimates for the conditional 3-class LCGA are shown in Table 6.

Summary

A single set of categorical adolescent-drinking data was analyzed using a hierarchical generalized linear model (HGLM) with linear and quadratic slopes, and 21 specifications of mixture models. The HGLM showed a statistically significant mean linear increase in drinking (with a non-significant mean quadratic slope), and reflected significant variability around the intercept, in ninth grade, and both linear and quadratic slopes. All of the mixture models also revealed an overall increase in drinking across time. In general, freeing more variances and covariances in the growth mixture models (GMMs) led to fewer classes being estimable, resulting in a range of one class (equivalent to HGLM), when all variances and covariances were estimated, to seven classes, when all variances were constrained to zero (in the latent class growth analysis; LCGA). Based on the convergence properties, fit indices (BIC and LMR

proved most useful), and, in two cases, the extraction of unacceptably small classes, we were able to pare our list down to 10 potential models, including the HGLM.

We next considered the fit in terms of model-predicted versus observed category frequencies by looking at standardized residuals for the univariate and bivariate marginal frequencies (Table 3), and the most commonly occurring patterns of responses (Table 4). Too many standardized residuals larger than 2 is a sign of poor fit of the model to the data. All models fit the univariate marginal frequencies fairly well, and all but one—the 2-class longitudinal latent class analysis (LLCA) model—fit the common-pattern frequencies adequately. We eliminated the 2-class LLCA and also dropped the remaining 2-class models, all of which had 15 or more large bivariate residuals (this cutoff was chosen somewhat arbitrarily). The 3- and 4-class LLCA models both predicted cell frequencies extremely well, so we eliminated the 4-class model, which required 16 more parameters and had a higher BIC.

For the remaining selection decisions, we used plots to help us consider substantive issues: interpretability, usefulness, relationships with covariates, and consistency with theory and previous literature. Based on the plots, we rejected the sole remaining GMM. One of its classes was of little use in terms of explaining outcomes and it lacked a class with consistently low drinking probabilities that is predicted by theory and previous literature (e.g., Colder et al., 2002).¹⁴ Finally, the plots showed how poor the class separation was when the intercept variances were estimated. Given the overlap between classes and the difficulty in estimating and interpreting this model, even with all other things being equal, we would see little reason to select this GMM over the more straightforward and parsimonious HGLM.

We also compared the plots for the 3- and 4-class LCGAs, which had very similar BICs (Table 1). The added class showed a pattern that was neither predicted by theory nor found in previous studies of adolescent alcohol use. Rather, this new class appeared to be an artifact resulting from test-retest effects often found in longitudinal studies comprising self-reported data. We rejected the 4-class model in favor of the 3-class.

The 3-class LCGA and LLCA models were very similar to one another (Figure 7). Both had one large class with low drinking probabilities across adolescence and a small increase toward the end of high school (labeled “low”), a small class with relatively high drinking probabilities (including more frequent drinking) across adolescence (“high”), and a class that started with low drinking probabilities in seventh grade and then increased fairly rapidly, until its drinking probabilities were similar to the “high” class in 12th grade (“increasing”).

Finally, we added a covariate. The three remaining models (LCGA and LLCA 3-class models, and HGLM) all showed that adolescents who reported having close friends who used alcohol in seventh grade were likely to be drinking more frequently themselves across all of the junior high and high school years. Such similar findings across the three models suggest that all offer reasonable representations of drinking across middle and late adolescence, and given the large difference in parsimony between LCGA and LLCA (13 vs. 47 parameters, respectively), it made sense to reject the LLCA in favor of the LCGA. In situations where plots show that change over time is less continuous and linear or curvilinear (e.g., episodic problems like suicidality), or when the proportional odds assumption is not tenable, LLCA may offer considerable advantage over the other models.

The two remaining models, HGLM and the 3-class LCGA, were equivalent in terms of fit to the data and to theory, and they told essentially the same substantive story, so we did not choose

¹⁴It is possible, in GMM and other mixture models, to specify a class with zero probability of drinking at all timepoints if such a class is believed to exist. For details on how this is done, see Kreuter and Muthén (2007).

between them. It is possible that, based on further analyses, one of the two would stand out as a clearly better model. For example, differences might emerge when other covariates are used, or one of the two models could prove a better predictor of individual values on a distal outcome.

Discussion

In any data-analysis situation, there is a multitude of decisions that must be made, from determining a statistical model for the distribution of the data to choosing the structural models that best answer the research questions. The study of problem behaviors can be a particular challenge; the behaviors represent developmental processes, requiring longitudinal analyses, and because the behaviors are rare, their distributions are highly skewed. In addition, most data measuring these behaviors are discrete (binary, ordinal, or count). We have presented a selection of alternative models for analyzing longitudinal categorical data. These approaches assume more appropriate distributions for problem-behavior data than the traditional continuous-data models, permitting researchers to have more faith in the results.

Starting with the concepts involved in the generalized linear model (GLM) for categorical data, we discussed longitudinal extensions, including the hierarchical (or longitudinal) version of the generalized linear model (HGLM), and several specifications of mixture models. While there has been a great deal of excitement recently over using mixture models to find homogeneous unobserved subgroups in the population, we showed how such approaches can also be viewed as extensions of the random-effects model (Bauer & Curran, 2003; B. Muthén & Asparouhov, 2006), which relax some of its strong assumptions. Because of this, they might be the most appropriate choices for certain types of data distributions, irrespective of any question of subpopulations.

Our goal in this paper was to present a group of related models for longitudinal categorical data and to offer substantive researchers a useful guide to testing and selecting between the alternative models. This is a particularly difficult issue with categorical data and complicated longitudinal models, such as mixture models, because, as we find here, there are no completely reliable fit statistics for these types of models, and no readily available measures of absolute fit. We suggested several considerations to use in model selection. The first have to do with convergence properties, measures of comparative fit, and residual analysis. We dismissed models that did not converge well (including both improper solutions and failure to repeat the best loglikelihood value). From the remaining models, we chose the most plausible ones based on the fit statistics, and then ruled out several more because they failed to adequately reproduce the observed cell frequencies or were equivalent to more parsimonious models with respect to fit. Next, we included substantive considerations in model selection. For instance, we rejected a model because its substantive meaning was both obscure, rendering the model less useful, and not a reasonable fit to theory or to prior research. We also selected one of two similar models, despite poorer fit, because it was more parsimonious. Of course, we tested a limited number of models and only a single covariate. Choice and specification of covariates are an important factor in model selection and differences in these could potentially lead to different model choices

Through this process of elimination, we selected the two models that offered the most reasonable combination of parsimony, fit to the data, and fit to theory—the hierarchical generalized linear model (HGLM) and the 3-class latent class growth analysis (LCGA). The two models also use roughly the same number of parameters: 11 and 13, respectively and the substantive findings were similar. Both showed that, while adolescents are increasingly likely to drink as they approach the end of high school, those who, in seventh grade, reported having more friends who were drinking, were more likely to be drinking, themselves, throughout the rest of the secondary school years. The choice between these two models, based on the data

we have presented, is, in our opinion, largely a matter of ease of interpretation, although further investigation could turn up substantive or statistical reasons to choose one over the other, or even select a different model over these two.

This study has the strengths and the weaknesses that inevitably result from using actual data as an example. On the one hand, it shows how the models might look in an actual research situation. On the other hand, with simulated data, we could compare the model results with a known population to discover whether there are any systematic biases in the analyses. It is also possible that different approaches to measuring alcohol you could yield different results (Feldman & Masy, 2008). Finally, there is one more critical piece which warrants further study. Often the purpose of longitudinal studies is to predict what will happen to the young people later in life. The models may not be of much use if they fail in that respect, no matter how well they appear to characterize the data at hand. Future research needs to take distal outcomes into account in assessing the adequacy of the models.

In this paper we present and compare a fairly large set of models, but it is not an exhaustive list of potential models. There is an infinite set of possible models and, with empirical data, we cannot know what the true model is. For instance, there are models outside of the GLM framework that could be considered when growth is not linear or curvilinear, as well as two-part models, in which behavior is modeled as two related processes: (1) exhibiting the behavior versus not doing so, and (2) intensity or extent of the behavior (Olsen & Schafer, 2001; Tooez et al., 2002). However, although there remains some uncertainty with respect to model choice, we feel that all of the methods shown here may offer improvement over approaches that treat all data, irrespective of their actual distributions, as continuous and normally distributed.

In addition, because comparing longitudinal categorical models is not always straightforward, we have offered a systematic approach to assessing and selecting models from among several competing, equally appropriate statistical options. It is our belief that adopting this new business-as-usual approach will improve our ability to understand trajectories of problem behaviors over childhood and adolescence. With improved models, we can better determine risks and protective factors, and this information could potentially be used for early identification of those most at-risk, helping to select the best candidates for targeted preventive interventions.

References

- Agresti, A. Categorical data analysis. Vol. 2nd ed.. Hoboken, NJ: Wiley; 2002.
- Agresti A, Booth JG, Hobert JP, Caffo B. Random-effects modeling of categorical response data. *Sociological Methodology* 2000;30:27–80.
- Ashford JR. An approach to the analysis of data for semi-quantal responses in biological assay. [Article]. *Biometrics* 1959;15(4):573–581.
- Bauer DJ. A semiparametric approach to modeling nonlinear relations among latent variables. *Structural Equation Modeling* 2005;12:513–535.
- Bauer DJ, Curran PJ. Distributional assumptions of growth mixture models: Implications for overextraction of latent trajectory classes. *Psychological Methods* 2003;8:338–363. [PubMed: 14596495]
- Bentler PM, Chou C-P. Practical issues in structural modeling. *Sociological Methods and Research* 1987;16:78–117.
- Bock RD, Lieberman M. Fitting a response model for n dichotomously scored items. [Article]. *Psychometrika* 1970;35(2):179–&.
- Butler JS, Moffitt R. A computationally efficient quadrature procedure for the one-factor multinomial probit model. [Note]. *Econometrica* 1982;50(3):761–764.

- Carlin JB, Wolfe R, Brown CH, Gelman A. A case study on the choice, interpretation and checking of multilevel models for longitudinal binary outcomes. *Biostatistics* 2001;4:397–416. [PubMed: 12933632]
- Cicchetti D, Toth SL. The development of depression in children and adolescents. *American Psychologist* 1998;53(2):221–241. [PubMed: 9491749]
- Colder CR, Campbell RT, Ruel E, Richardson JL, Flay BR. A finite mixture model of growth trajectories of adolescent alcohol use: Predictors and consequences. *Journal of Consulting and Clinical Psychology* 2002;70(4):976–985. [PubMed: 12182281]
- Collins, LM.; Graham, JW.; Rousculp, SS.; Hansen, WB. Heavy caffeine use and the beginning of the substance use onset process: An illustration of latent transition analysis. In: Bryant, KJ.; Windle, M.; West, SG., editors. *The science of prevention: Methodological advances from alcohol and substance abuse research*. Washington, DC: American Psychological Association; 1997. p. 79-99.
- Conger RD, Conger KJ. Resilience in Midwestern families: Selected findings from the first decade of a prospective, longitudinal study. *Journal of Marriage and Family* 2002;64(2):361–373.
- Curran PJ, Stice E, Chassin L. The relation between adolescent alcohol use and peer alcohol use: A longitudinal random coefficients model. *Journal of Consulting and Clinical Psychology* 1997;65(1):130–140. [PubMed: 9103742]
- DiStefano C. The impact of categorization with confirmatory factor analysis. *Structural Equation Modeling* 2002;9(3):327–346.
- Dolan CV. Factor analysis of variables with 2, 3, 5 and 7 response categories: A comparison of categorical variable estimators using simulated data. *British Journal of Mathematical and Statistical Psychology* 1994;47:309–326.
- Donovan JE, Jessor R. Structure of problem behavior in adolescence and young adulthood. *Journal of Consulting and Clinical Psychology* 1985;53:890–904. [PubMed: 4086689]
- Feldman, BJ.; Masyn, KE. The measurement and modeling of alcohol use: A simulation study; Paper presented at the International Meeting of the Psychometric Society; 2008.
- Ferrer E, McArdle JJ. Alternative structural models for multivariate longitudinal data analysis. *Structural Equation Modeling* 2003;10(4):493–524.
- Foster EM, Fang GY. Conduct Problems Prevention Research Group. Alternative methods for handling attrition-An illustration using data from the Fast Track evaluation. [Article]. *Evaluation Review* 2004;28(5):434–464. [PubMed: 15358906]
- Goldberger, AS. *Econometric Theory*. New York: Wiley; 1964.
- Graham JW, Collins LM, Wugalter SE, Chung NK, Hansen WB. Modeling transitions in latent stage-sequential processes: A substance use prevention example. *Journal of Consulting and Clinical Psychology* 1991;59(1):48–57. [PubMed: 2002142]
- Greenbaum PE, Del Boca FK, Darkes J, Wang C-P, Goldman MS. Variation in the drinking trajectories of freshmen college students. *Journal of Consulting and Clinical Psychology* 2005;73(2):229–238. [PubMed: 15796630]
- Gurland J, Lee I, Dahm PA. Polychotomous quantal response in biological assay. [Article]. *Biometrics* 1960;16(3):382–398.
- Haberman SJ. The analysis of residuals in cross-classified tables. *Biometrics* 1973;29(1):205–220.
- Hagenaars, JA.; McCutcheon, AL., editors. *Applied latent class analysis*. New York: Cambridge University Press; 2002.
- Harville DA, Mee RW. A mixed-model procedure for analyzing ordered-categorical data. [Article]. *Biometrics* 1984;40(2):393–408.
- Hawkins JD, Catalano RF, Miller JY. Risk and protective factors for alcohol and other drug problems in adolescence and early adulthood: Implications for substance abuse prevention. *Psychological Bulletin* 1992;112(1):64–105. [PubMed: 1529040]
- Hedeker D, Gibbons RD. A random-effects ordinal regression model for multilevel analysis. *Biometrics* 1994;50(4):933–944. [PubMed: 7787006]
- Hedeker D, Gibbons RD, Flay BR. Random-effects regression models for clustered data with an example from smoking prevention research. *Journal of Consulting and Clinical Psychology* 1994;62(4):757–765. [PubMed: 7962879]

- Hedeker D, Mermelstein RJ. A multilevel thresholds of change model for analysis of stages of change data. *Multivariate Behavioral Research. Special Issue: Innovative methods for prevention research* 1998;33(4):427–455.
- Hedeker D, Mermelstein RJ. Analysis of longitudinal substance use outcomes using ordinal random-effects regression models. *Addiction* 2000;95:S381–S394. [PubMed: 11132364]
- Hernán MA, Hernández-Díaz S, Robins JM. A structural approach to selection bias. *Epidemiology* 2004;15:615–625. [PubMed: 15308962]
- Hill KG, White HR, Chung I-J, Hawkins JD, Catalano RF. Early adult outcomes of adolescent binge drinking: Person-and variable-centered analyses of binge drinking trajectories. *Alcoholism: Clinical and Experimental Research* 2000;24(6):892–901.
- Hipp JR, Bauer DJ. Local solutions in the estimation of growth mixture models. *Psychological Methods* 2006;11(1):36–53. [PubMed: 16594766]
- Hix-Small H, Duncan TE, Duncan SC, Okut H. A multivariate associative finite growth mixture modeling approach examining adolescent alcohol and marijuana use. *Journal of Psychopathology and Behavioral Assessment* 2004;26(4):255–270.
- Jackson KM, Sher KJ, Schulenberg JE. Conjoint developmental trajectories of young adult alcohol and tobacco use. *Journal of Abnormal Psychology. Special Issue: Toward a Dimensionally Based Taxonomy of Psychopathology* 2005;114(4):612–626.
- Johnston, LD.; O'Malley, PM.; Bachman, JG.; Schulenberg, JE. Monitoring the Future national results on adolescent drug use: Overview of key findings, 2005. Bethesda, MD: National Institute on Drug Abuse; 2006. (NIH Publication No. 06–5882)
- Knowles ES, Coker MC, Scott RA, Cook DA, Neville JW. Measurement-induced improvement in anxiety: Mean shifts with repeated assessment. *Journal of Personality and Social Psychology* 1996;71(2):352–363. [PubMed: 8765486]
- Kreuter, F.; Muthén, B. Longitudinal modeling of population heterogeneity; Methodological challenges to the analysis of empirically derived criminal trajectory profiles. In: Hancock, GR.; Samuelsen, KM., editors. *Advances in latent variable mixture models*. Charlotte, NC: Information Age Publishing; 2007.
- Lanza ST, Collins LM. A Mixture Model of Discontinuous Development in Heavy Drinking From Ages 18 to 30: The Role of College Enrollment. *Journal of Studies on Alcohol* 2006;67(4):552–561. [PubMed: 16736075]
- Little RJA. Modeling the drop-out mechanism in repeated-measures studies. *Journal of the American Statistical Association* 1995;90(431):1112–1121.
- Little, RJA.; Rubin, DB. *Statistical analysis with missing data*. Vol. 2nd ed.. Hoboken, NJ: Wiley; 2002.
- Marmorstein N, Iacono WG. Longitudinal follow-up of adolescents with late-onset antisocial behavior: A pathological yet overlooked group. *Journal of the American Academy of Child & Adolescent Psychiatry* 2005;44(12):1284–1291. [PubMed: 16292121]
- McArdle, JJ. Latent growth curve analyses using structural equation modeling techniques. In: Teti, DM., editor. *Handbook of research methods in developmental science*. Malden, MA: Blackwell; 2004. p. 340-366.
- McLachlan, G.; Peel, D. *Finite mixture models*. New York: Wiley; 2000.
- Millsap RE, Yun-Tein J. Assessing Factorial Invariance in Ordered-Categorical Measures. *Multivariate Behavioral Research* 2004;39(3):479–515.
- Muthén, B. Latent variable analysis: Growth mixture modeling and related techniques for longitudinal data. In: Kaplan, D., editor. *Handbook of quantitative methodology for the social sciences*. Newbury Park, CA: Sage; 2004. p. 345-368.
- Muthén B, Asparouhov T. Latent variable analysis with categorical outcomes: Multiple-group and growth modeling in mplus. *Mplus Web Notes: No. 4*. 2002
- Muthén, B.; Asparouhov, T. Growth mixture analysis: Models with non-Gaussian random effects. In: Fitzmaurice, G.; Davidian, M.; Verbeke, G.; Molenberghs, G., editors. *Advances in longitudinal data analysis*. Boca Raton, FL: Chapman & Hall/CRC Press; 2006.
- Muthén B, Brown CH, Masyn K, Jo B, Khoo S-T, Yang C-C, et al. General growth mixture modeling for randomized preventive interventions. *Biostatistics* 2002;3(4):459–475. [PubMed: 12933592]

- Muthén, B.; Muthén, LK. Mplus (Version 5.0) [Computer software]. Los Angeles, CA: Muthén & Muthén; 1998–2006.
- Muthén B, Muthén LK. The development of heavy drinking and alcohol-related problems from ages 18 to 37 in a U. S. national sample. *Journal of Studies on Alcohol* 2000a;61(2):290–300.
- Muthén B, Muthén LK. Integrating person-centered and variable-centered analyses: Growth mixture modeling with latent trajectory classes. *Alcoholism: Clinical and Experimental Research* 2000b;24(6):882–891.
- Muthén, LK.; Muthén, B. Mplus user's guide. Vol. 4th ed.. Los Angeles: Muthén & Muthén; 1998–2006.
- Nagin DS. Analyzing developmental trajectories: A semiparametric, group-based approach. *Psychological Methods* 1999;4:139–157.
- Nagin, DS. Group-based modeling of development. Cambridge, MA: Harvard University Press; 2005.
- Nylund KL, Asparouhov T, Muthén B. Deciding on the number of classes in latent class analysis and growth mixture modeling: A Monte Carlo simulation study. *Structural Equation Modeling* 2007:535–569.
- Nylund, KL.; Masyn, KE. Covariates and latent class analysis: Results of a simulation study; Paper presented at the Society for Prevention Research Annual Meeting; 2008.
- Orlando M, Tucker JS, Ellickson PL, Klein DJ. Concurrent use of alcohol and cigarettes from adolescence to young adulthood: An examination of developmental trajectories and outcomes. *Substance Use & Misuse* 2005;40(8):1051–1069. [PubMed: 16040368]
- Olsen MK, Schafer JL. A two-part random-effects model for semicontinuous longitudinal data. *Journal of the American Statistical Association* 2001;96:730–745.
- Ramaswamy V, Desarbo WS, Reibstein DJ, Robinson WT. An empirical pooling approach for estimating marketing mix elasticities with PIMS data. *Marketing Science* 1993;12(1):103–124.
- Raudenbush SW. How do we study “what happens next”? *Annals of the American Academy of Political and Social Science* 2005;602:131–144.
- Raudenbush, SW.; Bryk, AS. Hierarchical linear models: Applications and data analysis methods. Vol. 2nd ed.. Thousand Oaks, CA: Sage; 2002.
- Raudenbush, SW.; Bryk, AS.; Congdon, R. HLM (Version 6) [Computer software]. Lincolnwood, IL: Scientific Software International, Inc; 2005.
- Reinecke J. Longitudinal Analysis of Adolescents' Deviant and Delinquent Behavior: Applications of Latent Class Growth Curves and Growth Mixture Models. *Methodology: European Journal of Research Methods for the Behavioral and Social Sciences. Special Issue: Mixture Structural Equation Modeling* 2006;2(3):100–112.
- Sameroff AJ. Developmental systems and psychopathology. *Development and Psychopathology. Special Issue: Reflecting on the past and planning for the future of developmental psychopathology* 2000;12(3):297–312.
- Schafer JL, Graham JW. Missing data: Our view of the state of the art. *Psychological Methods* 2002;7(2):147–177. [PubMed: 12090408]
- Shaw DS, Gilliom M, Ingoldsby EM, Nagin DS. Trajectories leading to school-age conduct problems. *Developmental Psychology. Special Issue: Violent children* 2003;39(2):189–200.
- Shaw DS, Lacourse E, Nagin DS. Developmental trajectories of conduct problems and hyperactivity from ages 2 to 10. *Journal of Child Psychology and Psychiatry* 2005;46(9):931–942. [PubMed: 16108996]
- Singer, JD.; Willett, JB. Applied longitudinal data analysis: Modeling change and event occurrence. New York, NY, US: Oxford University Press; 2003. p. 644xx
- Skrondal, A.; Rabe-Hesketh, S. Generalized latent variable modeling: Multilevel, longitudinal, and structural equation models. Boca Raton, FL: Chapman & Hall/CRC; 2004.
- Tooze JA, Grunwald GK, Jones RH. Analysis of repeated measures data with clumping at zero. *Statistical Methods in Medical Research* 2002;11:341–355. [PubMed: 12197301]
- Tucker JS, Orlando M, Ellickson PL. Patterns and correlates of binge drinking trajectories from early adolescence to young adulthood. *Health Psychology* 2003;22(1):79–87. [PubMed: 12558205]

- Twenge JM, Nolen-Hoeksema S. Age, gender, race, socioeconomic status, and birth cohort difference on the children's depression inventory: A meta-analysis. *Journal of Abnormal Psychology* 2002;111(4):578–588. [PubMed: 12428771]
- Vermunt, JK. Growth models for categorical response variables: standard, latent-class, and hybrid approaches. In: Montfort, Kv; Oud, H.; Satorra, A., editors. *Longitudinal Models in the Behavioral and Related Sciences*. Mahwah, NJ: Erlbaum; 2006.
- West, SG.; Finch, JF.; Curran, PJ. Structural equation models with nonnormal variables: Problems and remedies. In: Hoyle, RH., editor. *Structural equation modeling: Concepts, issues, and applications*. Thousand Oaks, CA: Sage; 1995. p. 56-75.
- White HR, Johnson V, Buyske S. Parental modeling and parenting behavior effects on offspring alcohol and cigarette use: A growth curve analysis. *Journal of Substance Abuse* 2000;12(3):287–310. [PubMed: 11367605]
- Wiesner M, Silbereisen RK. Trajectories of delinquent behaviour in adolescence and their covariates: Relations with initial and time-averaged factors. *Journal of Adolescence* 2003;26(6):753–771. [PubMed: 14643745]
- Wiesner M, Windle M. Assessing covariates of adolescent delinquency trajectories: A latent growth mixture modeling approach. *Journal of Youth and Adolescence* 2004;33(5):431–442.

Appendix A

Types of Generalized Linear Growth Models

Name	Acronym	Alternative Name (acronym)	Link	Distributional Assumptions	
				Level 1	Level 2 +
Hierarchical Linear Model	HLM	Random Effects Model Random Coefficients Model Mixed Model	Identity	Normal	Normal
Hierarchical Generalized Linear Model	HGLM	Multilevel Model (MLM) Generalized Linear Mixed Model (GLMM)	May be identity, log, logistic or logit, probit, complementary log-log, etc. ^a	Member of Exponential Family ^b	Normal
Growth Mixture Model	GMM	Growth Mixture Analysis Finite Mixture Model	May be identity, log, logistic or logit, probit, complementary log-log, etc.	Exponential Family	Multinomial ^c & Normal
Latent Class Growth	LCGA	Semi-parametric Group-Based Model (SGBM)	Identity	Exponential Family Normal	Multinomial
Latent Profile Analysis Longitudinal Latent	LLCA	Latent Class Analysis (LCA)	May be identity, logistic or logit, probit, complementary log-log, etc.	Normal Multinomial	Multinomial

^a Depends on the distribution of the observed data.

^b Exponential family of distributions includes normal, exponential, gamma, multinomial, binary or Bernoulli, Poisson, zero-inflated Poisson, negative binomial, etc.

^c Multinomial includes both ordered and unordered categorical

^e Model lacks latent growth factors

Appendix B

Syntax Example 1: Unconditional (i.e., no covariates) Hierarchical Generalized Linear Model (HGLM)

TITLE:
Unconditional hierarchical generalized linear model (HGLM)

DATA:
FILE is drk5yr.dat;

Data file should be in folder with the input file.

VARIABLE:
NAMES are famid sex drk1-drk5
frdrink;

USEVARIABLES are drk1-drk5;
CATEGORICAL are drk1-drk5;

NAMES tells Mplus the names of the variables in the data, USEVARIABLES identifies the variables in the analysis, and CATEGORICAL labels the data as binary or ordinal

MISSING are all .;

This identifies the missing value tag (as .)

ANALYSIS:
TYPE = missing;
STARTS = 100 10;
ALGORITHM=integration;

“TYPE = missing” tells Mplus to use all available data (despite missingness); no longer needed in version 5. “STARTS” gives Mplus instructions about how many randomly selected starting values to use, and ALGORITHM = integration tells Mplus how to estimate the model.

MODEL:
int slope quad | drk1@-2 drk2@-1 drk3@0
drk4@1 drk5@2;

Variances and covariances between growth factors are estimated by default.

Intercept is placed at the third timepoint in this example

OUTPUT: TECH1 TECH4 TECH10;

TECH10 gives residuals for the overall analysis.

PLOT:
TYPE is plot3;
SERIES is drk1(0) drk2(1) drk3(2) drk4(3) drk5(5) ;

Plot commands generate an estimated mean (logit) trajectory.

SAVEDATA:
file = HGLM_isq_isqv_uncond.txt;
save = fscores;

SAVEDATA can be used to save various information from the analysis. This command saves model estimates (individual estimated scores on the growth factors; fscores) in a text file.

Syntax Example 2: Conditional (i.e., with covariates) GMM or LCGA

```

TITLE: Conditional growth mixture model (GMM) or latent class growth analysis
(LCGA)

DATA:
  FILE IS drk5yr.dat;

VARIABLE:
  NAMES ARE famid sex drk1-drk5 frdrink;
  USEVARIABLES ARE drk1-drk5 sex frdrink ;
  CATEGORICAL ARE drk1-drk5;

CLASSES = c(2);

MISSING ARE ALL .;
ID = famid;

ANALYSIS:
  TYPE = mixture missing ;
  STARTS = 500 70;
  ALGORITHM = integration;

MODEL:
  %OVERALL%
  int slope quad | drk1@0 drk2@1 drk3@2 drk4@3 drk5@5;

  int*;
  slope@0;
  quad@0;

  int on sex pdrink;
  c#1 on sex pdrink;

OUTPUT: TECH1 TECH4 TECH7 TECH10;

PLOT:
  TYPE IS plot3;
  SERIES IS drk1(0) drk2(1) drk3(2) drk4(3) drk5(5);

SAVEDATA:
  file = GMM2_isq_iv_pred.txt;
  save = cprob fcores;

```

CLASSES defines the name (c) and number of classes in the mixture model

MISSING identifies the missing value indicator in the data. ID identifies the person identifier and with this command, Mplus will include the variable in output files of individual estimates.

"TYPE = mixture" is required for mixture models.

%OVERALL% is a necessary command for a mixture model. The next line is typical Mplus notation for a random effects model of any type. This shows fixed times of measurement, but they may be unique to each individual. int* frees the intercept variance and the linear and quadratic variances are fixed to zero (@0). If all three variances are fixed at zero, the analysis is LCGA.

Intercept is placed at the first timepoint in this example

"int on..." regresses the intercept on the covariates. "c#1 on..." regresses the first class (as compared with the second class) on the covariates

These last commands save model results in a text file: individual scores on the growth factors (fcores) and individual class posterior probabilities (cprob).

Syntax Example 3: LLCA Without Covariates

```

TITLE:
  Unconditional Longitudinal Latent Class Analysis (LLCA)

DATA: file is drk5yr.dat;

VARIABLE:
  names are famid sex drk1-drk5 frdrink;

  usevariables are drk1-drk5 ;
  categorical are drk1-drk5;

  classes = c(2);
  missing are all .;

ANALYSIS: type = mixture missing ;
  starts = 100 10;
  lrtstarts = 100 10 100 10;

MODEL:
%overall%

OUTPUT: TECH1 TECH4 TECH7 TECH10 TECH11 TECH14 residual;

PLOT:
  type is plot3;
  series is drk1(0) drk2(1) drk3(2) drk4(3) drk5(5) ;

SAVEDATA:
  file = LLCA2_uncond.txt;
  save = cprob;

```

In addition to the random starts (STARTS), this syntax tells Mplus how many random starts to use for the bootstrap LRT (BLRT) test (LRTSTARTS).

The difference from GMM/LCGA is that no linear model is specified after the %OVERALL% statement

TECH11 and TECH14 give the Lo-Mendel-Rubin test (LMR) and Bootstrapped Likelihood Ratio Test (BLRT)

Only the posterior class probabilities are saved (there are no estimated intercepts or slopes in LLCA).

Acknowledgments

Support for this research has come from multiple sources, including the National Institute of Mental Health (MH00567, MH19734, MH43270, MH48165, and MH51361), the National Institute on Drug Abuse (DA05347 and HD047573), the National Institute of Child Health and Human Development (HD027724 and HD047573), the Bureau of Maternal and Child Health (MCJ-109572), the MacArthur Foundation Research Network on Successful Adolescent Development Among Youth in High-Risk Settings, and the Iowa Agriculture and Home Economics Experiment Station (Project No. 3320). We thank Shelley Blozis, Keith Widaman, Sophia Rabe-Hesketh for their helpful comments on this paper. We also thank the editor and anonymous reviewers for many insightful and helpful remarks.

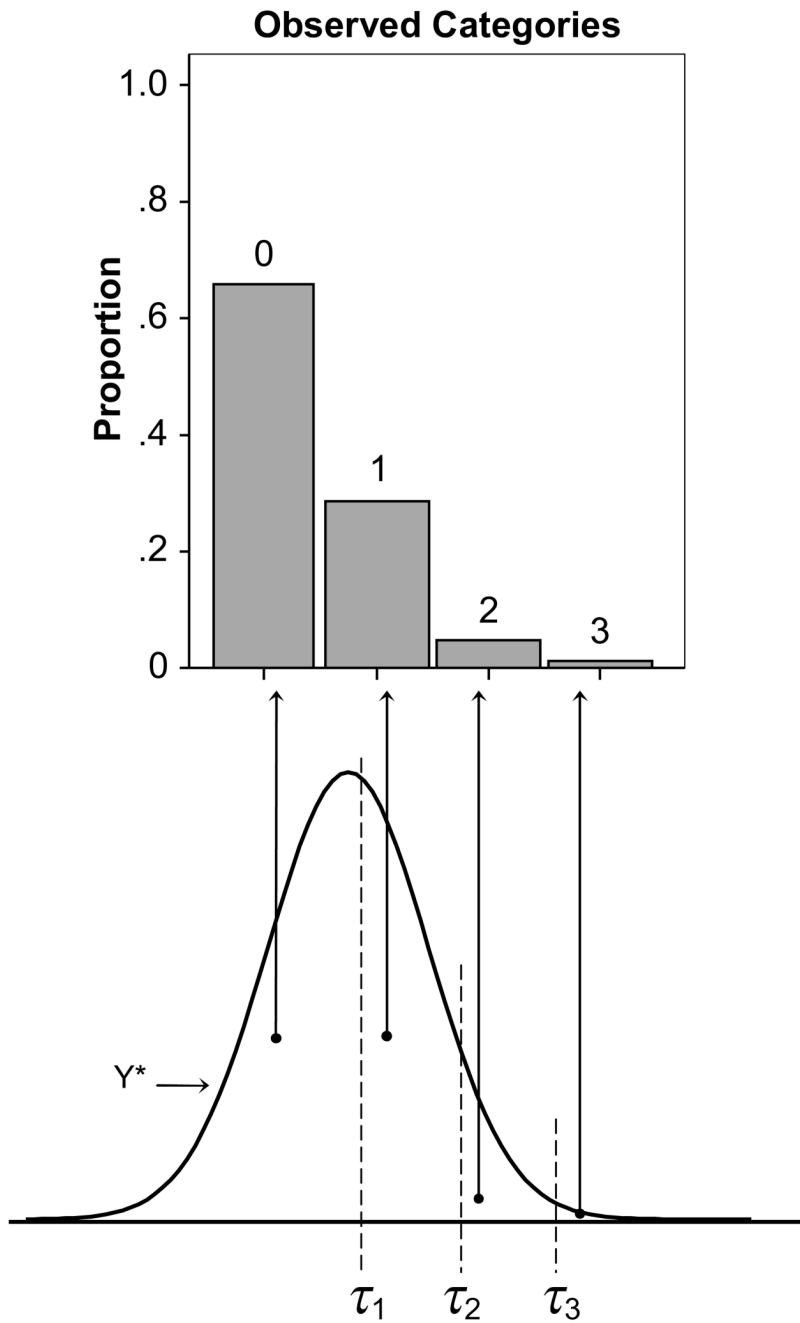


Figure 1. Pictorial representation of how thresholds (τ 's) in an underlying, continuous latent response variable correspond to the categories in an ordinal variable.

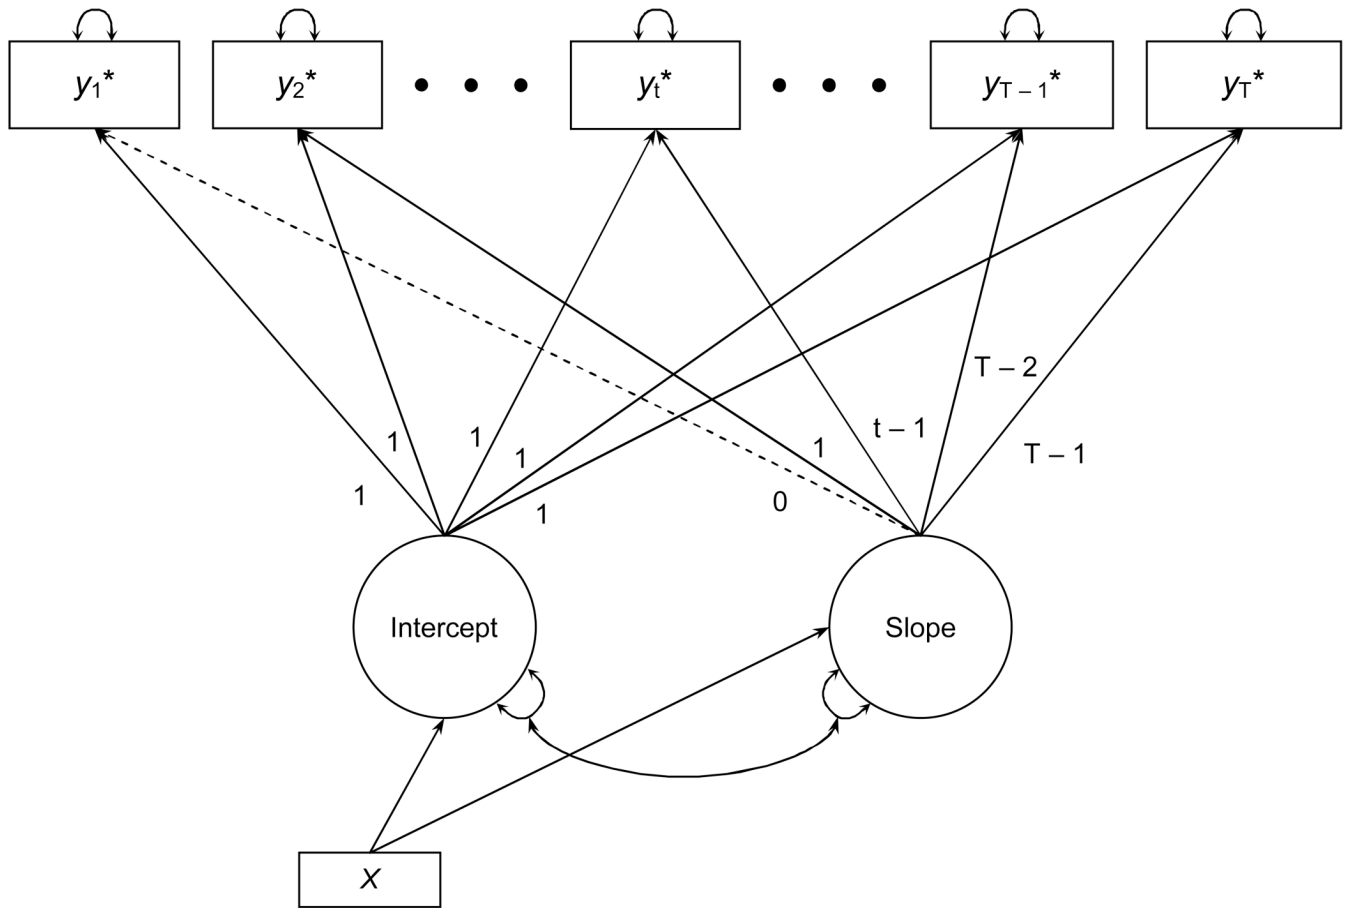



Figure 2. Path diagram of a hierarchical generalized linear model (HGLM) with a covariate, X . Model is shown with fixed times of measurement, but individual times of measurement may be specified in the model.

Figure 3a



Figure 3b

**Figures 3a and 3b.**

A skewed univariate distribution characterized by a mixture of normal distributions on top (generalized linear growth mixture model; GMM), or mass points on the bottom (generalized linear latent class growth analysis; LCGA).

Figure 4a and b

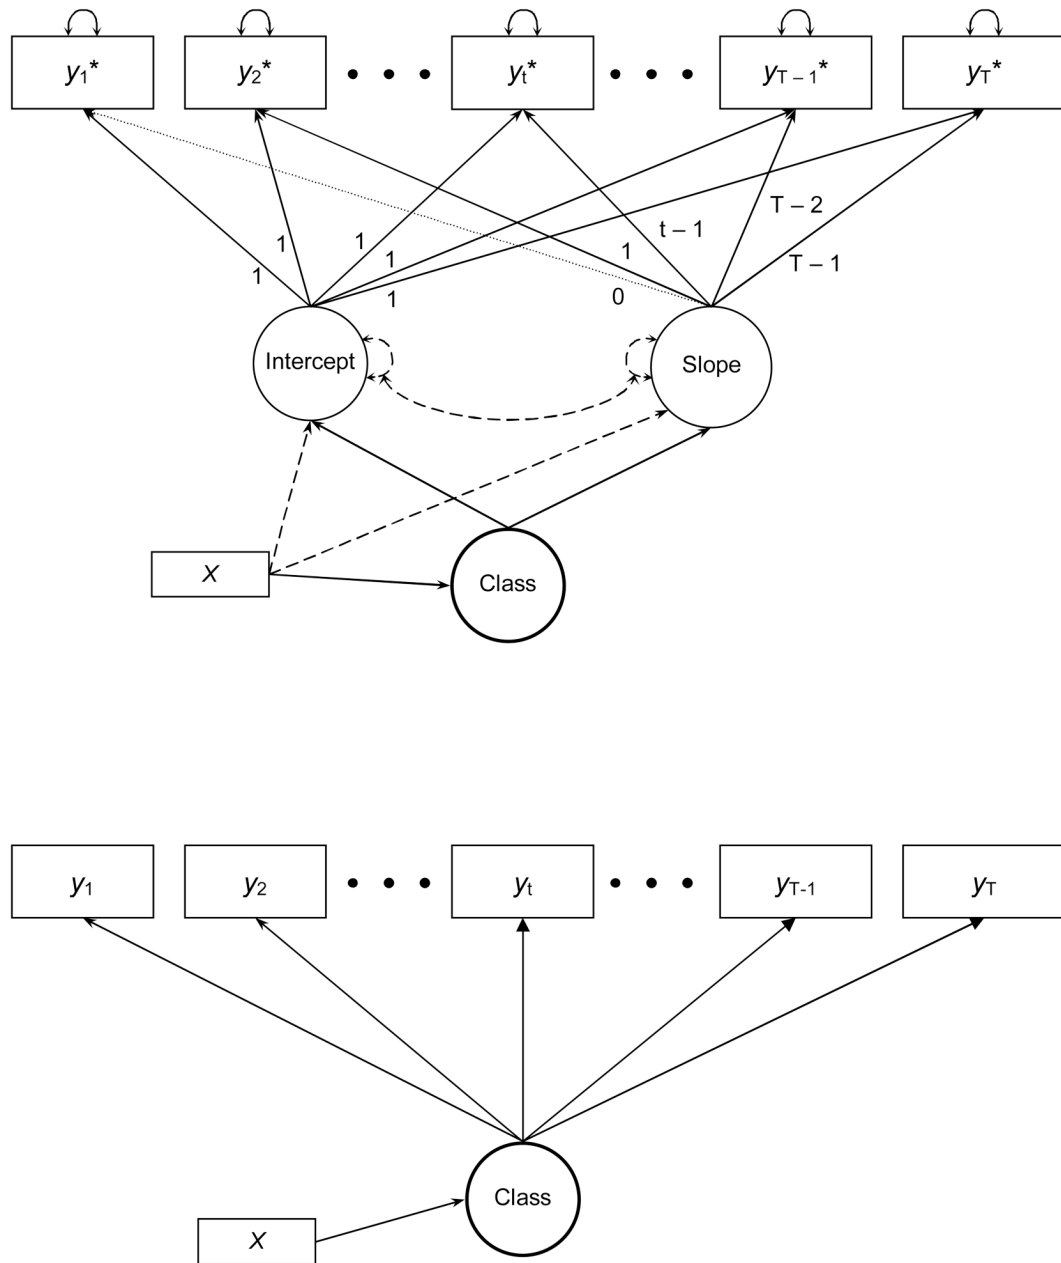


Figure 4a and 4b.

Path diagrams for growth mixture model (GMM) and latent class growth analysis (LCGA) on the top. Solid lines are parameters present in both models and dashed lines represent parameters that, if present, define the model as GMM. Bottom shows the path diagram of longitudinal latent class analysis (LLCA). Both models are shown with covariate, X .

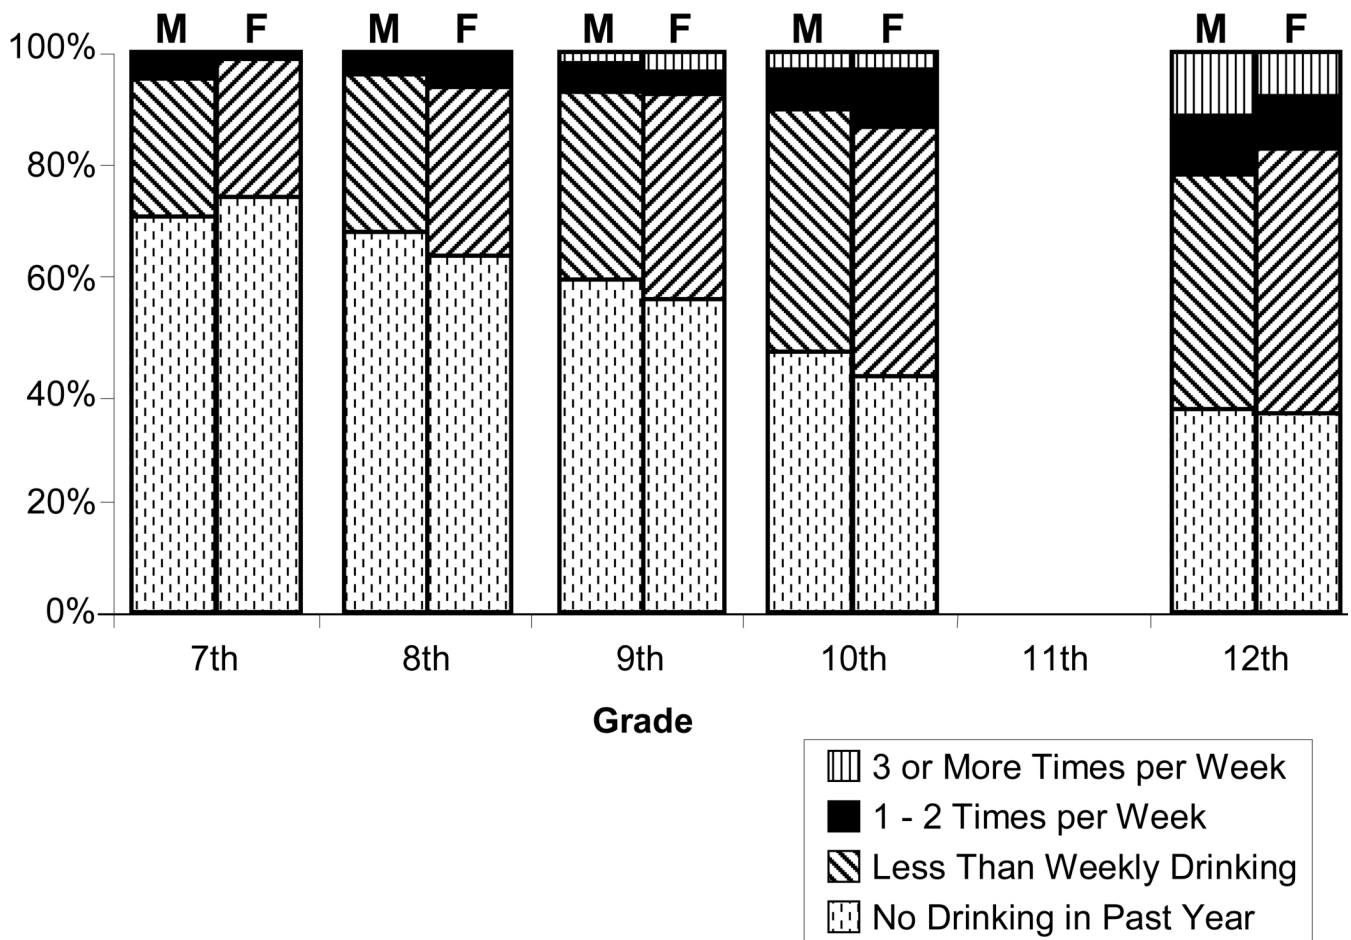


Figure 5.
Observed drinking category probabilities by gender and grade.

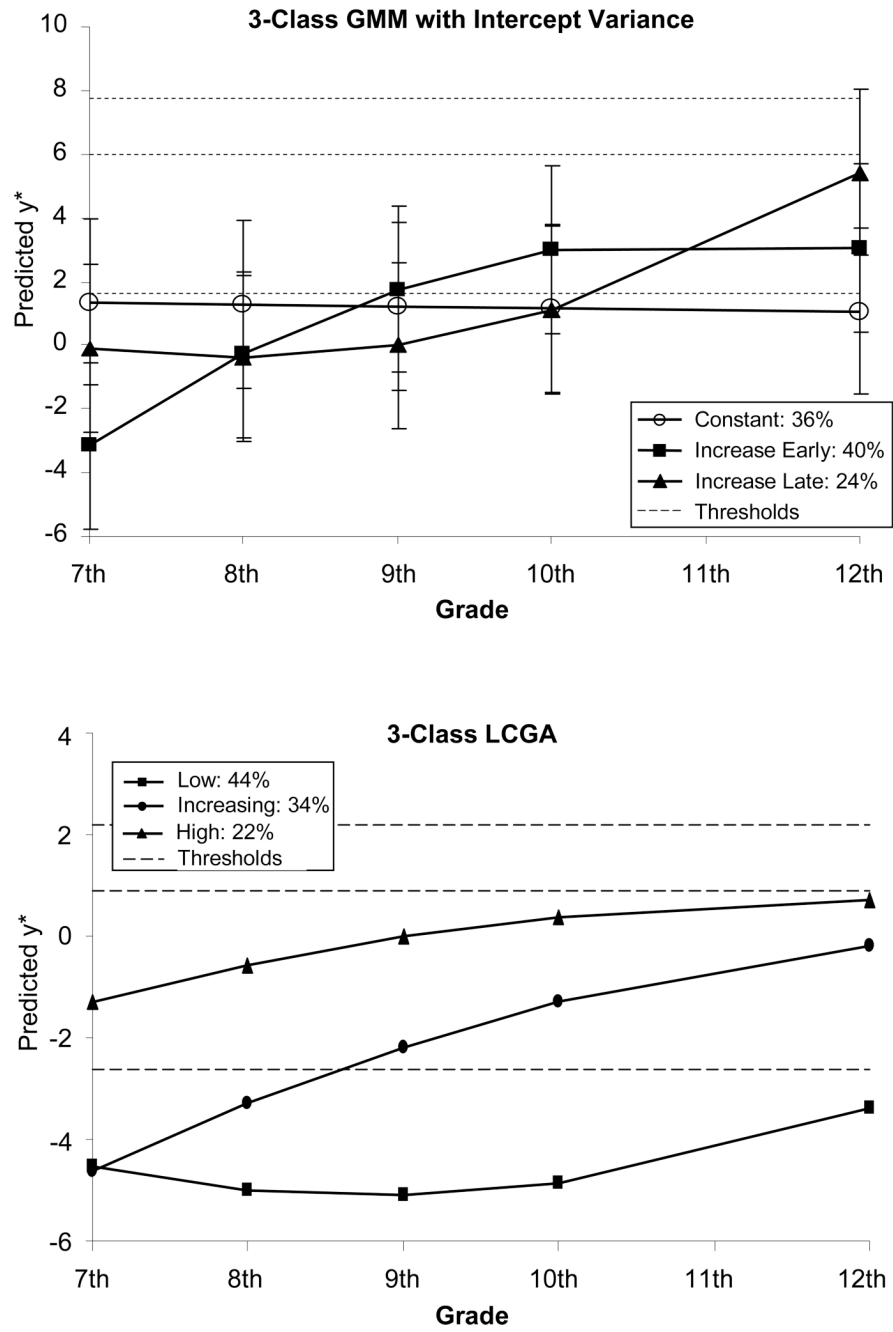


Figure 6. Estimated mean trajectories for 3-class generalized linear growth mixture model (GMM) and generalized linear latent class growth analysis (LCGA). GMM plot includes standard deviation bars at each time and, in both plots, thresholds are shown by dashed lines.




Figure 7. Plots of predicted category proportions for 3-class models: generalized linear growth mixture model (GMM; top), generalized linear latent class growth analysis (LCGA; middle), and longitudinal latent class analysis (LLCA; bottom).




Figure 8. Plots of predicted category proportions for 4-class generalized linear latent class growth analysis (LCGA).

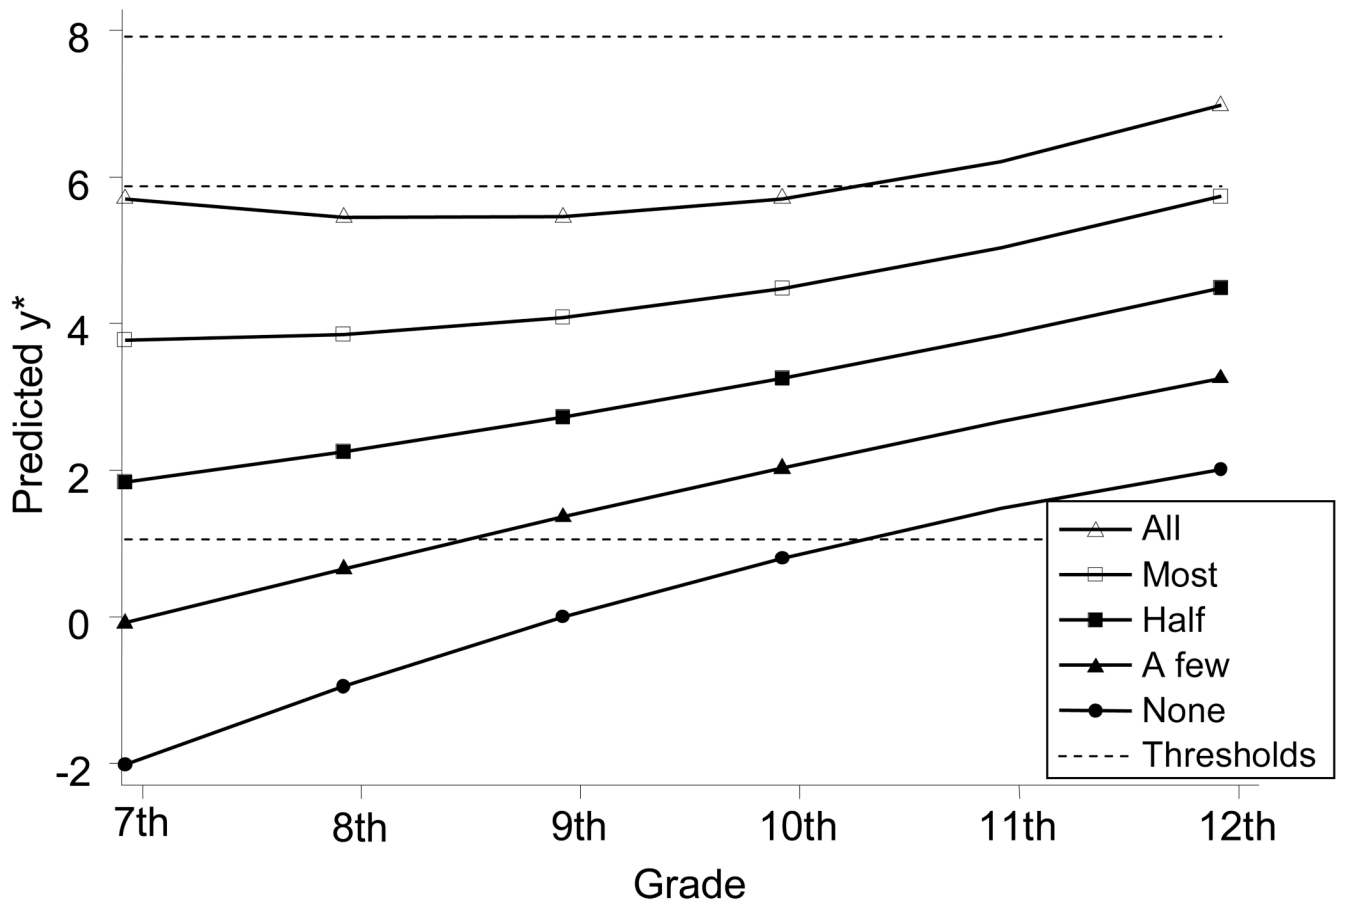


Figure 9. Plot of predicted hierarchical generalized linear model (HGLM) trajectories by levels of the covariate, target's report of the number of alcohol-using friends in 7th grade (range of scale: 0 = none to 4 = all of them).

Table 1

Fit Statistics for All Models

Statistic	1 Class	2 Classes	3 Classes	4 Classes	5 Classes	6 Classes	7 Classes
HGLM							
LL	-1676.73						
# Parameters	11						
BIC	3420.69						
GMM: Int, Lin, Quadratic ^d							
LL		-1678.06	-1676.73 ^b				
# Parameters		12	16				
Entropy		.27					
BIC		3429.46					
BLRT (<i>p</i>)		26.49 (.01)					
LMR (<i>p</i>)		27.58 (<.01)					
GMM: Int, Linear ^d							
LL		-1678.34	-1670.10 ^b				
# Parameters		11	15				
Entropy		.29					
BIC		3423.90					
BLRT (<i>p</i>)		29.09 (<.01)					
LMR (<i>p</i>)		27.95 (.15)					
GMM: Intercept ^d							
LL		-1697.05	-1683.17	-1673.79	-1665.69 ^b		
# Parameters		10	14	18	22		
Entropy		.34	.38	.52			
BIC		3455.21	3451.91	3457.59			
BLRT (<i>p</i>)		44.65 (<.01)	27.45 (<.01)	18.76 (<.01)			
LMR (<i>p</i>)		42.89 (.04)	26.66 (.06)	18.03 (.14)			
LCCA							
LL		-1737.51	-1704.75	-1692.11	-1682.26	-1673.20	-1661.50
# Parameters		9	13	17	21	25	29
Entropy		.75	.69	.67	.68	.69	.72
BIC		3530.02	3488.94	3488.11	3492.85	3499.19	3500.23
BLRT (<i>p</i>)		437.63 (<.01)	65.52 (<.01)	25.28 (<.01)	19.71 (<.01)	18.11 (<.01)	23.41 (<.01)
LMR (<i>p</i>)		420.43 (<.01)	62.95 (<.01)	24.28 (.01)	18.93 (.12)	17.39 (.54)	22.49 (.01)
LLCA							
LL		-1702.59	-1671.25	-1652.50	-1639.53	-1627.09	-1617.49 ^c
# Parameters		31	47	63	79	95	111
Entropy		.81	.73	.79	.86	.79	
BIC		3594.64	3629.74	3690.02	3761.87	3834.76	
BLRT (<i>p</i>)		481.94 (<.01)	62.68 (<.01)	37.50 (<.01)	25.94 (.31)	24.89 (.27)	
LMR (<i>p</i>)		477.06 (<.01)	62.05 (.95)	37.12 (.23)	25.67 (1.00)	24.64 (1.00)	

Note. All models estimated with intercept, and linear and quadratic slopes. GMM = generalized linear growth mixture model, LCCA = generalized linear latent class growth analysis (all variances and covariances estimated), LLCA = longitudinal latent class analysis, BIC = Bayesian information criterion, BLRT = bootstrapped likelihood ratio test, LMR = Lo-Mendel-Rubin test. Underlined values indicate the best of a set of models for that statistic.

^aVariances that were estimated; all covariances constrained to zero

^bSingular information matrix

^cNo repeated loglikelihood

Table 2

Hierarchical Generalized Linear Model Results

	Estimate	SE	<i>p</i>
Means			
β_{00} Intercept	0.00 ^a	—	—
β_{10} Linear Slope	0.84	0.10	< .001
β_{20} Quadratic Slope	-0.06	0.04	.098
Variances and Covariances			
ψ_{11} Intercept	10.28	1.58	< .001
ψ_{22} Linear Slope	0.79	0.17	< .001
ψ_{33} Quadratic Slope	0.14	0.03	< .001
ψ_{12} Intercept with Linear (<i>r</i>)	0.58 (.21)	0.30	.055
ψ_{13} Intercept with Quadratic (<i>r</i>)	-0.52 (-.44)	0.17	.003
ψ_{23} Linear with Quadratic (<i>r</i>)	-0.18 (-.55)	0.05	< .001
Thresholds			
τ_1	0.53	0.19	.006
τ_2	5.34	0.35	< .001
τ_3	7.35	0.46	< .001

^aIntercept mean is fixed to zero for identification

Table 3

Univariate and Bivariate Standardized Residuals

Model	Value of Standardized Residual			Total
	2 – 3	3 – 4	4 – 5	
Univariate (16 Cells)				
HGLM	1	0	0	1
GMM				
Intercept ^a				
2 Classes	0	0	0	0
3 Classes	1	0	0	1
Int, Linear Slope ^a				
2 Classes	2	0	0	2
Int, Lin, Quadratic Slopes ^a				
2 Classes	2	0	0	2
LCGA				
3 Classes	0	0	0	0
4 Classes	0	0	0	0
LLCA				
2 Classes	0	0	0	0
3 Classes	0	0	0	0
4 Classes	0	0	0	0
Bivariate (160 Cells)				
HGLM	6	0	1	7
GMM				
Intercept ^a				
2 Classes	13	3	0	16
3 Classes	6	4	0	10
Intercept, Linear ^a				
2 Classes	14	4	0	18
Intercept, Linear, Quadratic ^a				
2 Classes	12	3	0	16
LCGA				
3 Classes	11	0	1	12
4 Classes	8	0	1	9
LLCA				
2 Classes	12	3	0	15
3 Classes	3	0	0	3
4 Classes	2	0	0	2

Note. All models estimated with intercept, and linear and quadratic slopes. GMM = generalized linear growth mixture model, LCGA = generalized linear latent class growth analysis (all variances and covariances estimated), LLCA = longitudinal latent class analysis,

^aVariances that were estimated; all covariances constrained to zero

Observed and Predicted Frequencies of Most Common Observed Response Patterns

Response Pattern	Observed Frequency	Estimated Frequencies													
		HGLM				GMM: Classes				LLCA: Classes					
		Intercept ^a	Int, Linear ^a	Int, Lin, Quad ^a		Intercept ^a	Int, Linear ^a	Int, Lin, Quad ^a		Intercept ^a	Int, Linear ^a	Int, Lin, Quad ^a			
00000	81	74.40	82.71	82.61	78.99	75.65	78.99	61.51	71.93	73.73					
00001	33	41.60	35.57	31.87	38.23	36.86	38.23	42.45	36.97	35.04					
00011	21	20.22	14.93	17.06	18.18	16.57	18.18	17.34	17.55	22.53					
01111	21	11.94	12.55	11.56	11.64	11.17	11.64	12.96	15.44	16.40					
00111	19	16.79	14.02	16.51	19.25	19.15	19.25	8.00	19.99	18.88					
00010	15	11.22	11.01	12.59	15.48	9.21	15.48	24.36	16.91	14.43					
				Standardized Residuals											
00000	81	.86	-0.21	0.72	.26	.69	.26	2.72	1.19	.95					
00001	33	1.42	-0.45	-1.06	-.89	-.67	-.89	-1.54	-.69	-.36					
00011	21	.18	1.60	0.51	.68	1.11	.68	.90	.85	-.33					
01111	21	2.67	2.43	2.82	2.79	2.99	2.79	2.27	1.45	1.16					
00111	19	.55	1.36	0.63	-.06	-.03	-.06	3.93	-.23	.03					
00010	15	1.15	1.22	0.69	-1.13	1.93	-1.13	-1.96	-.48	.15					

Note. A response *pattern* is the pattern of observed responses across all five waves. In this case, the most common pattern was all zeros (i.e., 81 individuals did not report drinking at any wave). All models estimated with intercept, and linear and quadratic slopes. GMM = generalized linear growth mixture model, LCGA = generalized linear latent class growth analysis (all variances and covariances estimated), LLCA = longitudinal latent class analysis.

^aVariances that were estimated; all covariances constrained to zero

Table 5

Effects of 7th Grade Friends Who Drink on 3-Class LCGA and LLCA Trajectory Class Membership

	LCGA				LLCA			
	Est. (SE)	OR	95% CI	P	Est. (SE)	OR	95% CI	P
High to Low	1.33 (0.48)	3.79	1.48 – 9.67	.005	1.39 (0.35)	4.03	2.03 – 7.99	< .001
Increasing to Low	0.45 (0.41)	1.57	0.71 – 3.48	.266	0.32 (0.40)	1.38	0.63 – 3.01	.423
High to Increasing	0.88 (0.31)	2.41	1.31 – 4.42	.005	1.07 (0.35)	2.93	1.48 – 5.80	.002

Note. LCGA = generalized linear latent class growth analysis, LLCA = longitudinal latent class analysis.

Table 6

Conditional three-class generalized linear latent class growth analysis (LCGA) parameter estimates.

Parameters	Estimate	SE	P
High			
β_{00} Intercept	5.29	0.43	< .001
β_{10} Linear Slope	0.45	0.19	.017
β_{20} Quadratic Slope	-0.08	0.06	.199
Increasing			
β_{00} Intercept	3.07	0.43	< .001
β_{10} Linear Slope	0.86	0.18	< .001
β_{20} Quadratic Slope	-0.08	0.06	.166
Low			
β_{00} Intercept	0.00 ^a	—	—
β_{10} Linear Slope	0.10	0.15	.492
β_{20} Quadratic Slope	0.16	0.06	.001
Covariate Effects (class comparisons)			
High vs. Low	1.33	0.48	.005
Increasing vs. Low	0.45	0.41	.266
High vs. Increasing	0.88	0.31	.005
Thresholds			
τ_1	2.51	0.36	< .001
τ_2	5.99	0.37	< .001
τ_3	7.30	0.38	< .001

^aIntercept mean for one class is fixed to zero for identification