

Statistical power analysis for growth curve models using SAS

ZHIYONG ZHANG AND LIJUAN WANG
University of Notre Dame, Notre Dame, Indiana

Power analysis is critical in research designs. This study discusses a simulation-based approach utilizing the likelihood ratio test to estimate the power of growth curve analysis. The power estimation is implemented through a set of SAS macros. The application of the SAS macros is demonstrated through several examples, including missing data and nonlinear growth trajectory situations. The results of the examples indicate that the power of growth curve analysis increases with the increase of sample sizes, effect sizes, and numbers of measurement occasions. In addition, missing data can reduce power. The SAS macros can be modified to accommodate more complex power analysis for both linear and nonlinear growth curve models.

Power and sample size analysis is an important topic in study designs. The power of a statistical test is usually defined as the probability that the test would reject a false null hypothesis. Power analysis can be prospective or retrospective. The main purpose of a prospective power analysis is to plan the sample size for a future study. Thus, a prospective power analysis is conducted before data collection. A retrospective power analysis is conducted after data collection and analysis, the purpose of which is to determine whether the obtained effect size is conclusive. In this article, we focus only on the prospective power analysis because it can provide us information on planning a study and is more widely used. Power analysis for a variety of statistical tests and models, such as *t* tests, ANOVAs, and multiple regressions, can be conducted easily in statistical software (e.g., SAS PROC POWER).

In the past half-century, growth curve models have become among the most powerful tools in facilitating the analysis of change (e.g., M. W. Browne, 1993; M. W. Browne & Du Toit, 1991; Laird & Ware, 1982; McArdle & Nesselroade, 2003; Meredith & Tisak, 1990; Rao, 1958; Tucker, 1958, 1966). Growth curve models can be fitted in either the latent growth curve modeling framework (a special case of structural equation modeling) or the multilevel/mixed-effects modeling framework. Those two frameworks are essentially equivalent, although the implementations are different. Growth curve models and their estimation methods have been well studied (e.g., Bollen & Curran, 2006; Demidenko, 2004; Hox, 2002; McArdle & Nesselroade, 2003; Zhang, Hamagami, Wang, Grimm, & Nesselroade, 2007), but discussions on the determination of sample sizes or power analysis of growth curve models are still relatively rare, with a few exceptions (see a recent review by Maxwell, Kelley, & Rausch, 2008).

Several studies have investigated the power of latent growth curve models by utilizing the chi-square difference test or the likelihood ratio test from different perspectives. For example, Hertzog, Lindenberger, Ghisletta, and von Oertzen (2006) studied the power of multivariate latent growth curve models in detecting correlated change. Fan (2003) discussed the power of latent growth curve models for detecting group differences in linear growth trajectory parameters. An earlier study by B. O. Muthén and Curran (1997) evaluated the power in detecting treatment effects by latent growth curve models. Several other studies investigated the power of growth curve analysis with missing data and found that the required sample size had to be increased to maintain the same level of power as complete data analysis (Hedeker, Gibbons, & Waternaux, 1999; Jung & Ahn, 2003; Tu et al., 2007).

Meanwhile, there are power-analysis methods available for multilevel or mixed-effects models. These methods can be used to analyze certain types of growth curve models. For example, Raudenbush and Liu (2001) investigated the influence of different sample sizes on the power of group analysis with polynomial changes. Snijders and Bosker (1993) studied power for two-level models and provided software called PINt to conduct the power analysis. Recently, Cools, Van den Noortgate, and Onghena (2008) developed a program called ML-DEs, and W. J. Browne, Golarizadeh Lahi, and Parker (2009) developed a software add-on—MLPowSim—to conduct power analysis for multilevel models through MLwiN (Rasbash et al., 2000). Both ML-DEs and MLPowSim allow simulating unbalanced and incomplete data for power analysis.

Many studies on the power analysis of growth curve models were based on the chi-square difference test that was developed by Saris and Satorra (1993) for covariance

Z. Zhang, zhiyongzhang@nd.edu

structure analysis. Thus, the implementation of a power analysis often requires the use of commercial structural equation modeling or multilevel modeling software, such as Mplus, LISREL, HLM, and MLwiN, to get the chi-square statistics first. Then, the chi-square difference usually has to be processed in another statistical software, such as SAS, to compute power. The two-step, two-software approach makes power analysis cumbersome and hinders, at least partly, the broad adoption of the method. Therefore, we believe that a power analysis procedure that can be implemented in general-purpose software, such as SAS, will benefit researchers more broadly.

Thus, in the present study, we develop a SAS procedure consisting of several SAS macros to conduct power analysis for growth curve models under situations with complete data, missing data, and nonlinear growth trajectories. The procedure can be used to plan the sample size of a longitudinal study. The SAS procedure has advantages that include, but are not limited to, the following: (1) The SAS codes are publicly available and flexible, so users can modify the codes for their own research purposes; (2) it conducts the power analysis without requiring extra software; (3) it can utilize the SAS PROC NL MIXED procedure to conduct power analyses for a wide range of models, including nonlinear growth curve models; and (4) it can deal with missing data seamlessly.

In the following discussion, we will first introduce growth curve models—particularly, the latent linear growth curve model and the exponential growth curve model. We will then discuss the basic ideas of power analysis for growth curve modeling through the simulation-based method. After that, we will present several SAS macros that can implement the power analysis procedure. Finally, we will discuss several examples that utilize the SAS procedure for power analysis.

Growth Curve Models

The latent linear growth curve model (e.g., Laird & Ware, 1982; McArdle & Nesselroade, 2003; Meredith & Tisak, 1990) is probably the simplest, most widely used growth curve model. The path diagram for one form of the latent linear growth curve model is portrayed in Figure 1. The observed variables are drawn as squares. The unobserved or latent variables are drawn as circles, and the constants are represented by triangles. The squares labeled y_1 through y_T are the observed data at Occasions 1 through T , respectively. L_i is the latent initial level of the i th individual, μ_L is the mean of the initial levels across all participants, and σ_L^2 is the variability of the initial levels representing interindividual difference in the latent initial levels. S_i corresponds to the slope, μ_S is the mean of the slopes across all participants, and σ_S^2 represents its variability, or individual differences around the mean slope. S_i can be interpreted as the rate of change and/or growth. The covariance σ_{LS} between the level and the slope is represented by the double-headed arrow between the latent initial level and slope variables. The circles labeled e_1 through e_T are random measurement errors.




Figure 1. Path diagram for a latent linear growth curve model.

In mathematical terms, the latent linear growth curve model in Figure 1 can be expressed as follows:

$$\begin{aligned}
 y_{it} &= L_i + S_i \times (t - 1) + e_{it} \\
 L_i &= \mu_L + v_{0i} \\
 S_i &= \mu_S + v_{1i}
 \end{aligned}
 , i = 1, \dots, n; t = 1, 2, \dots, T, \quad (1)$$

where

$$e_{it} \sim N(0, \sigma_e^2), v_i = (v_{0i}, v_{1i})^t \sim MN \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_L^2 & \sigma_{LS} \\ \sigma_{LS} & \sigma_S^2 \end{pmatrix} \right].$$

It is often convenient to use the correlation between the latent level and slope variables to describe the relationship. The correlation is defined as $\rho = \sigma_{LS}/(\sigma_L\sigma_S)$. For better interpretability, the initial time point was constrained to be 0, so the initial level L_i can be interpreted as the true score at the first occasion of individual i . The residual errors are normally distributed with equal variances and uncorrelated in the path diagram for demonstration. However, the residual variances can vary across time, and the residuals can be correlated across time. The deviances of individual levels and slopes from the mean level and the mean slope follow a bivariate normal distribution. If μ_S is statistically significantly different from 0, there is linear change—either growth or decline—in the observed variable. If σ_S^2 differs significantly from 0, one can conclude that there are significant individual differences in the change.

Nonlinear growth curve models have been used to analyze longitudinal data when the change pattern is not linear. One widely used latent nonlinear growth curve model

is the exponential growth curve model. Using the similar notation in the latent linear growth curve model, one form of the exponential growth curve model (M. W. Browne, 1993) can be expressed as

$$\begin{aligned} y_{it} &= L_i - S_i \times \exp[-(t-1)p] + e_{it} \\ L_i &= \mu_L + v_{0i} \\ S_i &= \mu_S + v_{1i} \end{aligned} \quad , i = 1, \dots, n; t = 1, 2, \dots, T, \quad (2)$$

where L_i estimates the asymptotic level of individual i , S_i estimates the difference between the asymptotic level and the initial level at the first occasion for individual i . p (≥ 0) is the rate of growth or decline. The interindividual variations in L_i and S_i are also estimated in the model. The exponential growth curve model is especially useful for analyzing learning data (e.g., M. W. Browne & Du Toit, 1991). For both the linear and nonlinear latent growth models, although the measurement occasions are often assumed to be equally spaced, the unequally spaced data can be accommodated accordingly.

Power Analysis of Growth Curve Models

Currently, there are two main approaches for evaluating the power of growth curve models, both of which utilize a likelihood ratio test (e.g., Casella & Berger, 2001). The first approach was proposed by Saris and Satorra (1993; see also Satorra & Saris, 1985) for estimating the power of covariance structure analysis and has been widely used in analyzing the power of growth curve models (see, e.g., Fan, 2003; Hertzog et al., 2006; B. O. Muthén & Curran, 1997). The Satorra–Saris method is also known as the *chi-square difference test* and is a large sample theory approach. The result from this approach is accurate when the sample size is large enough. The other approach is the simulation-based method. The result from the simulation-based method is accurate with large enough replications, even when the sample size is relatively small (L. K. Muthén & Muthén, 2002; Satorra & Saris, 1985). In this article, the simulation-based method is used, and the general idea of the simulation-based method is outlined as follows.

Let \mathcal{M}_1 (full model) denote a model with p unknown parameters $\theta : p \times 1$. Assume that \mathcal{M}_0 is a model (null model or reduced model) nested in \mathcal{M}_1 that can be constructed by setting $p - q$ elements in θ to be constants, usually 0. Thus, for \mathcal{M}_0 , we have $\theta : p \times 1 = (\theta^* : q \times 1; c : p - q \times 1)$, with q denoting the number of unknown parameters in \mathcal{M}_0 and c denoting a constant vector. To estimate the power of the test on whether we can put those constraints in \mathcal{M}_0 , we can generate R sets of independent data that are based on \mathcal{M}_1 using predetermined parameter values for θ . For each (the i th) simulated data set, we fit both \mathcal{M}_1 and \mathcal{M}_0 and obtain the log-likelihood value for each model denoting l_{i1} and l_{i0} , respectively. Then the power is estimated as

$$\text{power} = \sum_{i=1}^R H[2(l_{i1} - l_{i0}) - \chi_{\alpha}^2(p - q)] / R, \quad (3)$$

where $\chi_{\alpha}^2(p - q)$ denotes the critical value of the chi-square distribution with degrees of freedom (df) $p - q$ at

the significance level α . The H represents the heavyside function expressed as

$$H[2(l_{i1} - l_{i0}) - \chi_{\alpha}^2(p - q)] = \begin{cases} 1, & \text{if } 2(l_{i1} - l_{i0}) > \chi_{\alpha}^2(p - q) \\ 0, & \text{otherwise} \end{cases}$$

For example, \mathcal{M}_1 is the latent linear growth curve model described in Equation 1 with the parameters $\theta = (\mu_L, \mu_S, \sigma_L, \sigma_S, \rho, \sigma_e)$. To evaluate the power of this model in detecting change, we can fit \mathcal{M}_1 and \mathcal{M}_0 derived from model \mathcal{M}_1 by setting the mean slope $\mu_S = 0$.¹ Then we can compare the 2 times log-likelihood (2log-likelihood) difference to the critical chi-square value from the chi-square distribution with $df = 1$ at the significance level α .

SAS MACROS FOR POWER ANALYSIS

We have developed several SAS macros to conduct the power analysis using the simulation-based likelihood ratio test. The complete macros for obtaining a power curve include three macro functions. The first macro function is used to generate data that are based on a specified full model \mathcal{M}_1 and calculate the 2log-likelihood difference of models \mathcal{M}_0 and \mathcal{M}_1 . The second macro function is used to run the first macro function for a given number of times and estimate the power for a given sample size. The third macro function is used to run the second macro function with different sample sizes to obtain the power curve. The basic macros for a linear growth curve model are provided in Appendix A.

The first macro function is `%MACRO LL(N,T,seed)`. This macro function has three input parameters: N for the sample size, T for the number of measurement occasions, and `seed` for the random number generator seed. This macro function is the most important part of all of the SAS macros. The function can be used to (1) generate the simulated data from a given model \mathcal{M}_1 with given population parameter values and (2) estimate the full model \mathcal{M}_1 and the reduced model \mathcal{M}_0 and calculate the log-likelihood difference between those two models. This macro function can be modified for different power analyses. First, one can change the data generation part according to different forms of full models \mathcal{M}_1 , such as from a linear growth curve model (as in Examples 1 and 2, below) to a nonlinear growth curve model (as in Example 3, below). Second, one can change the model parameter values in \mathcal{M}_1 to generate data for the same model with different effect sizes (as in Examples 1 and 2). Third, one can also change the model estimation part to construct different forms of \mathcal{M}_0 and implement different estimation methods for a desired power analysis (as in Example 3). Fourth, one can change the parameter `miss` to obtain the power of a test with missing data.

The second of the three macro functions is `%MACRO POWER(R,N,T,seed,df)`. This function is designed to estimate power on the basis of a given number of simulation replications (R) for a fixed sample size. The five parameters for this function are the number of replications in the simulation, the sample size, the number of measurement occasions, the random number generator seed, and the difference in the numbers of unknown parameters, in order.

This macro function repeatedly runs the first macro function and obtains the 2log-likelihood difference for each replicate. Power is then obtained using Equation 3. This function should be left untouched. As mentioned earlier, implementing this macro can obtain the power for a specific sample size. For example, `%POWER(1000,100,5,1234,1)` produces the power for a test on a single parameter in the full model with the sample size of 100 and 5 measurement occasions, based on 1,000 replications of simulation.

The third macro function, `%MACRO powercurve(R, seed, st, end, step, T, df)`, estimates the power for a given range of sample sizes—for example, from $N = 50$ to $N = 1,000$ in the examples provided later—and plots the power curve. The minimum sample size is given by *st*, and the maximum sample size is provided by *end*. The interval between two adjacent sample sizes is *step*. The other input parameters are *R*, *seed*, *T*, and *df*, again for the number of replications, the random number generator seed, the number of measurement occasions, and the difference in the numbers of unknown parameters, respectively. This function utilizes the second macro function to estimate power for each sample size and usually does not need to be modified. If the power of only a single sample size is needed, one can use the same sample size value for both *st* and *end* parameters or just simply run the second macro function.

Using these three macro functions, one can obtain the power curve for a designed analysis conveniently. For example, `%powercurve(1000,1234,50,1000,10,4,1)` generates the power curve for a test on a single parameter in a growth curve analysis with sample sizes from 50 to 1,000 and four measurement occasions. Furthermore, we collect all of the parameters that usually should be modified in the beginning of the macros. Thus, users can easily change the parameters without going through the main macros. Even for the power analysis of other different growth curve models such as exponential growth curve models, one usually needs to change only the first macro function.

Availability of the Macros

The SAS macros discussed in this article and the SAS macros for other models, such as quadratic growth curve models and conditional linear growth curve models, are available online at <http://saspower.psychstat.org>. These macros are licensed under the GNU General Public License version 2.0 (www.gnu.org/copyleft/gpl.html).

Examples of How to Use the Macros

In this section, we present several examples that utilize the SAS macros for power analysis. Example 1 illustrates the power analysis for a latent linear growth curve model in detecting change with complete data. Example 2 repeats the power analysis in Example 1, but with missing data. Example 3 provides an example of the power analysis of a nonlinear exponential growth curve model with and without missing data.

Example 1: Power and complete data. In this first example, we investigate the power of a linear growth curve model in detecting change ($H_0: \mu_S = 0$) with complete data. The true model is the one in Figure 1 with the population parameters $\theta = (\mu_L, \mu_S, \sigma_L, \sigma_S, \rho, \sigma_e) = (10, *, 2, 1, 0, 1)$.




Figure 2. Power curves for $T = 3, 4, 5, 6$, with the effect size of $\mu_S = .2$.

Here “*” represents the slope mean value, which is also the effect size in this example. Three factors that can influence power are considered: the sample size (N), the effect size (μ_S), and the number of measurement occasions (T). The sample size ranges from 50 to 1,000, and the number of measurement occasions ranges from three to six. Three effect sizes are evaluated at $\mu_S = .1, .2, \text{ and } .3$.

Figure 2 presents the power curves for different numbers of measurement occasions when $\mu_S = .2$. To obtain the power curve for $T = 3$, one only needs to change $T = 5$ to $T = 3$ in the beginning part of the SAS codes. As in any power analysis, the power is higher with a larger sample size. For the growth curve analysis, the power is also getting higher with the increase in the number of measurement occasions. For example, to obtain a power of .8, one needs a sample size about $N = 300$ for a three-occasion study. However, only about $N = 210$ participants are needed for a study with six measurement occasions.

Figure 3 portrays the power curves for different effect sizes with the same number of measurement occasions




Figure 3. Power curves for the effect sizes of $\mu_S = .1, .2, \text{ and } .3$ when the number of measurement occasions is $T = 5$.

$T = 5$. To obtain the power curves, one needs to change only the parameter value of the slope parameter (μ_S) in the beginning of the SAS codes. Clearly, larger effect size corresponds to larger power. For example, when the effect size is .1, a sample size of $N = 850$ is needed to obtain a power of .8. When the effect size is as large as .3, a sample size of $N = 100$ can achieve the same level of power.

Example 2: Power and incomplete data. Missing data are inevitable, even in very well designed longitudinal studies. Little and Rubin (1987, 2002) have distinguished three kinds of missing data mechanisms: missing completely at random (MCAR), missing at random (MAR), and missing not at random (MNAR) (see also, e.g., Rubin, 1976; Schafer, 1997). The popular statistical modeling techniques and models in general assume that missing data are either MCAR or MAR (e.g., Little & Rubin, 2002; Schafer, 1997).

In this example, we investigate how the MCAR data influence the power of a latent linear growth curve model in detecting change ($H_0 : \mu_S = 0$). The model is the same as the one in Example 1. We further assume that missing data are generated with a constant attrition rate. For example, with a constant missing attrition rate of 10% ($miss = 10\%$), if there are 100 participants at Occasion 1, 90 participants are expected at Occasion 2, 80 participants at Occasion 3, and only 10 participants at Occasion 10.

The following procedure can be used to generate MCAR data. Let $Y^* = (y_{it}^*)$, $i = 1, \dots, N$, $t = 1, \dots, T$ denote the complete data matrix. The MCAR data are then specified according to the missing data indicator matrix $M = (m_{it})$, $i = 1, \dots, N$, $t = 1, \dots, T$. The matrix M takes values of 0 (indicating missing data) and 1 (indicating observed data). In simulating data, the elements of M are given by Equation 4, below. The generated MCAR data $Y = (y_{it})$ are given by

$$y_{it} = \begin{cases} y_{it}^*, & \text{if } m_{it} = 1 \\ \cdot, & \text{otherwise} \end{cases}, \quad i = 1, \dots, N; t = 1, \dots, T,$$

where “.” indicates a missing datum. Because the probability that a datum is missing is independent of Y , the generated data are MCAR.

As in the complete data analysis, we estimate the power with different effect sizes and numbers of measurement occasions. The power can be obtained using the same macros for the complete data example, except for setting the parameter $miss$ to be .1 instead of 0. The power curves are given in Figure 4 and Figure 5. Similar to the complete data results, the power becomes larger as the sample size, number of measurement occasions, and effect size increase.

By comparing Figure 2 with Figure 4 and Figure 3 with Figure 5, we can evaluate the influences of missing data.

$$m_{i1} = 1$$

$$m_{it} = \begin{cases} 1, & \text{if } m_{i,t-1} = 1, \text{ with probability } \frac{1 - miss \times (t-1)}{1 - miss \times (t-2)}, \quad i = 1, \dots, N; t = 2, \dots, T \\ 0, & \text{otherwise} \end{cases} \quad (4)$$




Figure 4. Power curves for $T = 3, 4, 5, 6$, with the effect size of $\mu_S = .2$ in the missing-data condition.




Figure 5. Power curves for the effect sizes of $\mu_S = .1, 0.2, \text{ and } .3$ when the number of measurement occasions is $T = 5$ in the missing-data condition.

Overall, missing data reduce the power. The maximum difference in the power for all conditions considered in this study is about .1446. Thus, to maintain the same power, the required sample size has to be larger than the complete data analysis. For example, with the effect size $\mu_S = .2$ and the number of measurement occasions $T = 5$, a sample size of $N = 210$ at the first occasion is enough to obtain a power of .8. However, a sample size of $N = 270$ at the first occasion has to be used to maintain the same power with the missing data rate of 10%.

Example 3: Power for a nonlinear growth curve model. In this example, we demonstrate how to estimate

the power for a nonlinear growth curve model. Assume that the true model is an exponential growth curve model with the rate of growth $\rho = 1$ and the other model parameters $\theta = (\mu_L, \mu_S, \sigma_L, \sigma_S, \rho, \sigma_e) = (10, 5, 2, .5, .25, 1)$. The power analysis is on testing whether there is interindividual difference in the S_i parameters (the differences between the asymptotic level and the initial level at the first occasion; $H_0: \sigma_S = 0, \rho = 0$). If σ_S is not statistically significantly different from 0, then there is no significant interindividual difference in S_i . To obtain the power, we estimate the full model (freely estimate both σ_S and ρ parameters) and the null model (constrain both parameters to be 0) using PROC NLMIXED to obtain the 2log-likelihood difference. The 2log-likelihood difference can then be compared with the critical value from the chi-square distribution with $df = 2$ at the .05 significance level.

The power for this test is estimated with sample sizes ranging from 100 to 1,000 and six measurement occasions. Power for both the complete and missing data conditions is obtained. The obtained power is plotted in the power curves in Figure 6. Thus, on the basis of the specified model parameters in this example, an approximate sample size of $N = 250$ is required in order to obtain a power of .8 without any missing data. If we assume a missing data rate of 10%, the required sample size increases to about 350 for us to obtain the same power. It seems that the effect of missing data is larger in a nonlinear growth curve model than in a linear growth curve model. We also want to point out that, when one estimates the exponential growth curve models, convergence is a potential problem, as in many other nonlinear multilevel model estimation situations (see, e.g., Wang & McArdle, 2008). In this example, 93% of the simulations with complete data and 75% of the simulations with missing data converged when the sample size was 100. The power curves were generated from data based on the converged simulations. In the SAS macros presented here, convergence is satisfied if the absolute gradient is less than 10^{-5} or the relative gradient is less than 10^{-8} . The convergence is monitored in SAS automatically.

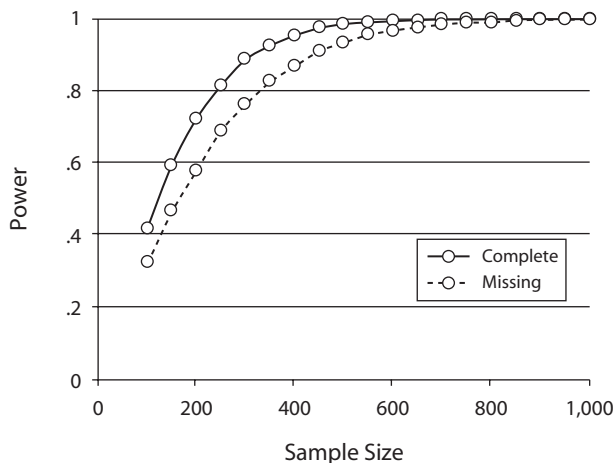


Figure 6. Power curves for the exponential growth curve model example.

DISCUSSION

Here, we developed several SAS macros for estimating power for growth curve analysis. The SAS macros utilized the simulation-based likelihood ratio test. Applications of the SAS macros were demonstrated through several examples. These examples showed that, as in any power analysis, the power became larger with larger sample sizes and effect sizes. Furthermore, for a longitudinal study, increasing the number of measurement occasions also increased the power. Comparing complete data analysis and MCAR data analysis, we found that missing data could reduce power. Given that SAS is widely adopted in academic institutes, the SAS procedure should benefit a considerable audience.

As shown in Figures 2 and 4, sample sizes of $N = 210$ and $N = 270$ are needed to obtain the power of .8 for complete data analysis and MCAR data analysis, respectively. Since missing data may reduce power, the influence of missing data should be considered at the phase of a longitudinal study design. Therefore, when deciding on the sample size, we should consider the effects of attrition rate on the power analysis. Usually, the attrition rate can be determined from a review of previous similar research.

It is also found that, for a fixed sample size, the power increases with larger numbers of measurement occasions, regardless of missing data. Thus, to obtain a desired level of power, we can adopt a design with either a smaller sample size and a larger number of measurement occasions or a larger sample size and a smaller number of measurement occasions. Which design is better depends on many factors, such as costs, time constraints, and the nature of the phenomenon under study. For example, the maximum number of measurement occasions is sometimes limited by the funding period. Therefore, it is sometimes more difficult to collect data with more occasions over a long period of time than to collect data with more individuals. Furthermore, if one wants to study a specific phenomenon, such as child development, one may want to collect data at the critical ages, such as 0, 0.5, 1, 2, 5, 7, and 11–15 years, instead of at equally spaced ages.

Guidelines for the Use of the Macros

Whereas other software and programs for power analysis have certain benefits, our SAS macros have their own merits. First, they are customized for growth curve analysis and require only the minimum input of the model information. Second, the power analysis for both complete data and missing data can be conducted seamlessly. Third, the macros can be slightly modified for a wide variety of linear and nonlinear growth curve models, as demonstrated by the linear and exponential growth curve model examples. The capability of the macros is limited only by the capability of PROC MIXED and PROC NLMIXED.

Experienced SAS users can modify the SAS macros according to their own research purposes. For users who are new to SAS, the following guidelines can be considered. First, if users want to conduct the same power analyses that we have explained in this study, they only need to

modify the parameter values in the beginning of the chosen macro without having to go through all of the macro codes. We recommend leaving the macros unchanged unless users fully understand the consequences of their modifications.

Second, the macros are based on simulations, thus raising an unavoidable question: How many simulation replications are arbitrarily large enough for obtaining valid and reliable power estimates? In our examples with 1,000 replications, we found it possible to obtain reliable estimates of the power for the linear growth curve model with complete data. With missing data, 5,000 replications are needed for obtaining reliable results for the linear model. For the exponential growth curve model, 10,000 replications were implemented. A general suggestion is to conduct the same power analysis multiple times with different random number seeds and then check whether the obtained power values have discrepancies. Moderate discrepancies in the power with different random number seeds signal to the researcher to increase the number of simulation replications.

Third, in the present study, the power is estimated through the simulation-based method. Although it is more accurate than the Satorra–Saris approximation method when the sample size is relatively small (Satorra & Saris, 1985), the simulation-based method is much more time consuming. For example, for each power curve presented in the linear examples, it took a PC with an AMD Opteron 2.6-GHz processor and 1 GB RAM about half an hour to obtain the results. For the nonlinear models, it took significantly more time—several days to obtain a power curve. For the nonlinear models, we suggest obtaining the power curve step by step. For instance, one may start with a single sample size and then gradually increase or decrease the sample size according to the obtained power. However, this amount of computing time may be negligible in comparison with years of time spent on data collection. In addition, as computer hardware improves, the required time may also be reduced.

Future Directions

For the purpose of simplifying the demonstration, the power analysis presented here has focused on the latent linear growth curve model in detecting average intra-individual change. We also demonstrated through an exponential growth curve model example on detecting inter-individual difference in change that the SAS macros can be easily modified for nonlinear growth modeling. However, the use of the power analysis procedure is not limited to the analyses in the examples. The power analysis can be further extended to other longitudinal data analyses. First, researchers are often interested in how well a covariate can predict the individual difference in change. The power for this kind of analysis can also be conducted by specifying the growth curve models with the covariate.² Second, the present study assumes that the missing data are MCAR. The analysis of missing-at-random data will be investigated in the future. Third, for more complex growth curve models, such as growth curve models with nonnormal response variables, one can conduct the power analysis by

using PROC NL MIXED, as in the nonlinear growth curve analysis example. An experienced SAS user can conduct these power analyses by modifying the provided SAS macros. SAS macros for other different kinds of complex longitudinal data analysis will be developed and provided on the Web site at <http://saspower.psychstat.org>.

AUTHOR NOTE

The authors thank the three anonymous reviewers for their helpful suggestions. Correspondence concerning this article should be addressed to Z. Zhang, Department of Psychology, University of Notre Dame, 118 Haggard Hall, Notre Dame, IN 46556 (e-mail: zhiyongzhang@nd.edu).

REFERENCES

- BOLLEN, K. A., & CURRAN, P. J. (2006). *Latent curve models: A structural equation perspective*. Hoboken, NJ: Wiley.
- BROWNE, M. W. (1993). Structured latent curve models. In C. M. Cudeñas & C. R. Rao (Eds.), *Multivariate analysis: Future directions 2* (pp. 171-198). Amsterdam: Elsevier, North-Holland.
- BROWNE, M. W., & DU TOIT, S. H. C. (1991). Models for learning data. In L. M. Collins & J. L. Horn (Eds.), *Best methods for the analysis of change* (pp. 47-68). Washington, DC: American Psychological Association.
- BROWNE, W. J., GOLALIZADEH LAHI, M., & PARKER, R. M. A. (2009). *A guide to sample size calculations for random effect models via simulation and the MLPowSim software package*. University of Bristol.
- CASELLA, G., & BERGER, R. L. (2001). *Statistical inference* (2nd ed.). Duxbury.
- COOLS, W., VAN DEN NOORTGATE, W., & ONGHENA, P. (2008). ML-DES: A program for designing efficient multilevel studies. *Behavior Research Methods*, *40*, 236-249.
- DEMIDENKO, E. (2004). *Mixed models: Theory and applications*. New York: Wiley.
- FAN, X. (2003). Power of latent growth modeling for detecting group differences in latent growth trajectory parameters. *Structural Equation Modeling*, *10*, 380-400.
- HEDEKER, D., GIBBONS, R. D., & WATERNAUX, C. (1999). Sample size estimation for longitudinal designs with attrition: Comparing time-related contrasts between two groups. *Journal of Educational & Behavioral Statistics*, *24*, 70-93.
- HERTZOG, C., LINDENBERGER, U., GHISLETTA, P., & VON OERTZEN, T. (2006). On the power of multivariate latent growth curve models to detect correlated change. *Psychological Methods*, *11*, 244-252.
- HOX, J. (2002). *Multilevel analysis: Techniques and applications*. Mahwah, NJ: Erlbaum.
- JUNG, S.-H., & AHN, C. (2003). Sample size estimation for GEE method for comparing slopes in repeated measurements data. *Statistics in Medicine*, *22*, 1305-1315.
- LAIRD, N. M., & WARE, J. H. (1982). Random-effects models for longitudinal data. *Biometrics*, *38*, 963-974.
- LITTLE, R. J. A., & RUBIN, D. B. (1987). *Statistical analysis with missing data*. New York: Wiley.
- LITTLE, R. J. A., & RUBIN, D. B. (2002). *Statistical analysis with missing data* (2nd ed.). New York: Wiley-Interscience.
- MAXWELL, S. E., KELLEY, K., & RAUSCH, J. R. (2008). Sample size planning for statistical power and accuracy in parameter estimation. *Annual Review of Psychology*, *59*, 537-563.
- MCARDLE, J. J., & NESSELROADE, J. R. (2003). Growth curve analysis in contemporary psychological research. In J. A. Schinka & W. F. Velicer (Eds.), *Handbook of psychology: Vol. 2. Research methods in psychology* (pp. 447-480). New York: Wiley.
- MEREDITH, W., & TISAK, J. (1990). Latent curve analysis. *Psychometrika*, *55*, 107-122.
- MUTHÉN, B. O., & CURRAN, P. J. (1997). General longitudinal modeling of individual differences in experimental designs: A latent variable framework for analysis and power estimation. *Psychological Methods*, *2*, 371-402.
- MUTHÉN, L. K., & MUTHÉN, B. O. (2002). How to use a Monte Carlo study to decide on sample size and determine power. *Structural Equation Modeling*, *9*, 599-620.

- RAO, C. R. (1958). Some statistical methods for comparison of growth curves. *Biometrics*, **14**, 1-17.
- RASBASH, J., BROWNE, W., GOLDSTEIN, H., YANG, M., PLEWIS, I., HEALY, M., ET AL. (2000). *A user's guide to MLwiN*. London: Institute of Education.
- RAUDENBUSH, S. W., & LIU, X.-F. (2001). Effects of study duration, frequency of observation, and sample size on power in studies of group differences in polynomial change. *Psychological Methods*, **6**, 387-401.
- RUBIN, D. B. (1976). Inference and missing data. *Biometrika*, **63**, 581-592.
- SARIS, W. E., & SATORRA, A. (1993). Power evaluations in structural equation models. In K. A. Bollen & J. S. Long (Eds.), *Testing structural equation models* (pp. 181-204). Newbury Park, CA: Sage.
- SATORRA, A., & SARIS, W. E. (1985). Power of the likelihood ratio test in covariance structure analysis. *Psychometrika*, **50**, 83-90.
- SCHAFFER, J. L. (1997). *Analysis of incomplete multivariate data*. Boca Raton, FL: Chapman & Hall/CRC.
- SNIJDERS, T. A., & BOSKER, R. J. (1993). Standard errors and sample sizes for two-level research. *Journal of Educational Statistics*, **18**, 237-259.
- TU, X. M., ZHANG, J., KOWALSKI, J., SHULTS, J., FENG, C., SUN, W., & TANG, W. (2007). Power analyses for longitudinal study designs with missing data. *Statistics in Medicine*, **26**, 2958-2981.
- TUCKER, L. R. (1958). Determination of parameters of a functional relation by factor analysis. *Psychometrika*, **23**, 19-23.
- TUCKER, L. R. (1966). Learning theory and multivariate experiment: Illustration by determination of parameters of generalized learning curves. In R. B. Cattell (Ed.), *Handbook of multivariate experimental psychology* (pp. 476-501). New York: Rand McNally.
- WANG, L., & McARDLE, J. J. (2008). A simulation study comparison of Bayesian estimation with conventional methods for estimating unknown change points. *Structural Equation Modeling*, **15**, 52-74.
- ZHANG, Z., HAMAGAMI, F., WANG, L., GRIMM, K. J., & NESSELROADE, J. R. (2007). Bayesian analysis of longitudinal data using growth curve models. *International Journal of Behavioral Development*, **31**, 374-383.

NOTES

1. The slope variance and the covariance between the initial level and slope are still estimated in this example.
2. An example of the SAS macros for such power analysis can also be found on the SAS Web site at <http://saspower.psychstat.org>.

APPENDIX A

SAS Macros for a Linear Growth Curve Model

```

/*Suppress the output and the log */
options nosource nonotes nosource2 nomprint;

/*Power analysis of growth curve models using SAS*/
/*CHANGE THE PARAMETERS HERE*/
*model parameters;
%LET MuL=10;          *mean level/initial status;
%LET MuS=.2;         *mean slope/rate of change;
%LET Sigma_e=1;      *residual standard deviation;
%LET Sigma_L=2;      *level standard deviation;
%LET Sigma_S=1;      *slope standard deviation;
%LET rho=0;          *correlation between levels and slopes;
%LET miss=0;         *missing data rate, 0: no missing data;
*power parameters;
%LET R=5000;         *number of simulation replications;
%LET T=5;            *number of measurement occasions;
%LET start=50;       *the minimum sample size to consider;
%LET end=1000;       *the maximum sample size to consider;
%LET step=50;        *the step between two sample sizes;
%LET df=1;           *the difference in the numbers of parameters;
%LET seed=123;       *random number generator seed;

/*DO NOT CHANGE CODES BELOW UNLESS YOU KNOW WHAT YOU ARE DOING*/

/*The first Macro: LL*/
/*Calculate the chi-square difference of two nested models*/
%MACRO LL(N,T,seed);
DATA Sim_LinGM;
  N = &N; seed = &seed;
* setup arrays for repeated measures;
  ARRAY y_score{&T} y1-y&T;
  ARRAY M{&T} m1-m&T;
  m1=1;
* generate raw data with considering the missing data rate;
  DO _N_ = 1 TO N;
    e_L=RANNOR(seed);
    e_S=&rho*e_L+SQRT(1-&rho**2)*RANNOR(seed);
    L_score=&MuL+&Sigma_L*e_L;
    S_score=&MuS+&Sigma_S*e_S;
    DO t = 1 TO &T;
      y_score{t} = L_score + (t-1) * S_score + &Sigma_e*RANNOR(seed);
    END;
  END;

```


APPENDIX A (Continued)

```

* include indicator variables to generate missing data;
DO t=2 TO &T;
  m{t}=m{t-1};
  IF m{t-1}=1 AND RANUNI(seed) > (1-&miss * (t-1))/(1-&miss * (t-2))
    THEN m{t} = 0;
  IF m{t}=0 THEN y_score{t}=.;
END;
  KEEP y1-y&T;
  OUTPUT;
END;
RUN;
DATA LingM;
  SET Sim_LinGM;
  %DO t = 1 %TO &T;
    id = _N_; time=&t-1; y=y&t; OUTPUT;
  %END;
  KEEP id time y;
RUN;

/*Fit two nested models to the data*/
ODS OUTPUT FitStatistics(persist=proc)=fit;
*Model 1: the true model;
TITLE1 'Linear Growth Model';
PROC MIXED DATA=LingM NOCLPRINT COVTEST MAXITER=100 METHOD=ML;
  CLASS id;
  MODEL y = time /SOLUTION DDFM=BW CHISQ;
  RANDOM INTERCEPT time / SUBJECT=id TYPE=UN GCORR;
RUN;
*Model 2: mean slope = 0;
PROC MIXED DATA=LingM NOCLPRINT COVTEST MAXITER=100 METHOD=ML;
  CLASS id;
  MODEL y = /SOLUTION DDFM=BW CHISQ;
  RANDOM INTERCEPT time / SUBJECT=id TYPE=UN GCORR;
RUN;

ODS OUTPUT CLOSE;
%MEND LL;

/*The second Macro: POWER*/
/*This Macro calls the first Macro LL for each replication*/
* Calculate power based on R replications;

%MACRO POWER(R,N,T,seed,df);
DATA tempfit;
  DO _N_=1 TO 8;
    tempfit=_N_;
  OUTPUT;
  END;
RUN;

%LL(&N,&T,&seed);
DATA fit;
  MERGE fit tempfit;
RUN;

DATA allfit;
  SET fit;
RUN;

%DO I = 2 %TO &R;
  PROC DATASETS LIBRARY=WORK; DELETE fit; RUN; QUIT;
  %LL(&N,&T,%eval(&seed+&I*1389));
  DATA fit;
    MERGE fit tempfit;

```

APPENDIX A (Continued)

```

RUN;
DATA allfit;
  SET allfit fit;
RUN;
DM 'CLEAR LOG';
%END;

DATA allfit;
  SET allfit;
  IF MOD(_N_,4) ~= 1 THEN DELETE;
  KEEP Value;
RUN;

DATA allfit;
  SET allfit;
  id =INT((_N_-.1)/2)+1;
  modelnum = MOD(_N_+1, 2);
RUN;

PROC TRANSPOSE DATA=allfit OUT=allfit prefix=model;
  BY id;
  ID modelnum;
  VAR Value;
RUN;

DATA allfit;
  SET allfit;
  ss = &N;
  diff = model1 - model0;
  ind = 1;
  IF diff=. THEN DELETE;
  IF diff<0 THEN DELETE;
  IF diff < CINV(.95, &df) THEN ind = 0;
  DROP id _NAME_ model0 model1;
RUN;

PROC MEANS DATA = allfit;
  VAR ss ind;
  OUTPUT OUT=power mean(ss ind)=ss power;
RUN;

%MEND POWER;

/*The third Macro: POWERCURVE*/
/* This Macro calls the second Macro for each sample size*/
%MACRO powercurve(R,seed,st,end,step,T,df);
%POWER(&R, &st, &T, %eval(&seed+&st), &df);
DATA allpower;
  SET power;
  RUN;

%LET st = %eval(&st + &step);
%DO %WHILE (&st <= &end);
  %POWER(&R, &st, &T, %eval(&seed+&st), &df);
  %LET st = %eval(&st + &step);
  DATA allpower;
    SET allpower power;
  RUN;
  DM 'CLEAR LOG';
%END;

* Save the results for possible future use;
DATA allpower;
  SET allpower;

```

APPENDIX A (Continued)

```

FILE "power.txt";
PUT ss power;
RUN;

* Plot the power curve;
ODS PDF FILE='power.pdf' NOTOC;
PROC GPLOT DATA = allpower;
  SYMBOL I=JOIN;
  PLOT power*ss;
RUN;
QUIT;
ODS PDF CLOSE;
%MEND powercurve;

ODS RESULTS OFF;
ODS LISTING CLOSE;
%powercurve(&R, &seed, &start, &end, &step, &T, &df);
ODS RESULTS ON;
ODS LISTING;

```

Appendix B
The SAS LL Macro for the Exponential Growth Curve Model

```

/*CHANGE THE PARAMETERS HERE*/
*model parameters;
%LET MuL=10;          *mean asymptotic level;
%LET MuS=5;          *mean change between the initial status
                    *and the asymptotic level;
%LET Sigma_e=1;      *residual standard deviation;
%LET Sigma_L=2;      *level standard deviation;
%LET Sigma_S=.5;     *change standard deviation;
%LET rho=.25;        *correlation between Li and Si;
%LET miss=0;         *missing data rate, 0: no missing data;
%LET p=1;            *rate of growth/decline;
*power parameters;
%LET R=10000;        *number of simulation replications;
%LET T=6;            *number of measurement occasions;
%LET start=100;      *the minimum sample size to consider;
%LET end=1000;       *the maximum sample size to consider;
%LET step=50;        *the step between two sample sizes;
%LET df=2;           *the difference in the number of parameters;
%LET seed=123;       *random number generator seed;

/*DO NOT CHANGE CODES BELOW UNLESS YOU KNOW WHAT YOU ARE DOING*/

/*Calculate the chi-square difference between two nested models*/
%MACRO LL(N,T,seed);
DATA Sim_ExpGM;
* set statistical parameters;
  N = &N; seed = &seed;
* setup arrays for repeated measures;
  ARRAY y_score{&T} y1-y&T;
  ARRAY M{&T} m1-m&T;
  m1=1;

* generate raw data with considering the missing data rate;
  DO _N_ = 1 TO N;
    e_L=RANNOR(seed);
    e_S=&rho*e_L+SQRT(1-&rho**2)*RANNOR(seed);
    L_score=&MuL+&Sigma_L*e_L;
    S_score=&MuS+&Sigma_S*e_S;
    DO t = 1 TO &T;
      y_score[t] = L_score - S_score*exp(-(t-1)*&p)

```

APPENDIX B (Continued)

```

      + &Sigma_e*RANNOR(seed);
    END;
* include indicator variables to generate missing data ;
DO t=2 TO &T;
  m{t}=m{t-1};
  IF m{t-1}=1 AND RANUNI(seed) > (1-&miss * (t-1))/(1-&miss * (t-2))
    THEN m{t} = m{t-1}*0;
  IF m{t}=0 THEN y_score{t}=.;
END;
  KEEP y1-y&T;
  OUTPUT;
END;
RUN;

DATA ExpGM;
  SET Sim_ExpGM;
  %DO t = 1 %TO &T;
    id = _N_; time=&t-1; y=y&t; OUTPUT;
  %END;
  KEEP id time y;
RUN;

/*Fit two nested models to the data*/
ODS OUTPUT FitStatistics(persist=proc)=fit;
*Model 1: the true model - exponential growth curve model;
PROC NLMIXED DATA = ExpGM;
  traject = level-slope*exp(-p*(time));
  MODEL y ~ NORMAL(traject, v_e);
  RANDOM level slope ~ NORMAL([m_l, m_s], [v_l, c_ls, v_s])
  SUBJECT = id;
  PARS m_l = 10 m_s=5 v_l = 4 c_ls = 0 v_s = .25 v_e = 1 p=1;
RUN;

*Model 2: the null model - no variation in Si;
PROC NLMIXED DATA = ExpGM;
  traject = level-m_s*exp(-p*(time));
  MODEL y ~ NORMAL(traject, v_e);
  RANDOM level ~ NORMAL(m_l, v_l)
  SUBJECT = id;
  PARS m_l = 10 m_s=5 v_l = 4 v_e = 1 p=1;
RUN;
ODS OUTPUT CLOSE;
%MEND LL;

/*The remaining parts are the same as in Appendix A
  and thus are omitted here.*/

```

(Manuscript received November 25, 2008;
 revision accepted for publication May 1, 2009.)