Latent Growth Curve Analysis: A Gentle Introduction

Alan C. Acock*

Department of Human Development and Family Sciences

Oregon State University

Fuzhong Li

Oregon Research Institute

*Professor and Chair of the Department of Human Development and Family Sciences, 322 Milam Hall, Oregon State University, Corvallis OR 97331-5102. Telephone is 541.737.4992, fax is 541.737.1076, e-mail is alan.acock@orst.edu

Latent Growth Curve Analysis: A Gentle Introduction

Latent Growth Curve Analysis—A Way to Explain Change

Most of us want to understand the process of change for whatever the topic happens to be. How much do we change? Do we get better? Worse? What explains how much we change? What mitigates adverse changes? What optimizes positive changes? We may use different terms for these issues: covariates, risk factors, protective factors, mediating effects, main effects, interactions, multi-level influences—the list is almost endless. The point is we want to study change!

Traditionally, we have used static data. For example, we try to learn about the effects of divorce by comparing divorced people to married people. Static data will tell us if divorced people are different (happier or less happy) than married people, but it will not tell us if divorced people are happier or less happy than they were before they got divorced. The topic can be divorce effects, likelihood of divorce, delinquency outcomes for adolescents, adjustment to retirement, or whatever you pick. Latent Growth Curve analysis is a method to study change.

Section 1: Conceptual Introduction: Questions We Can Answer

Today we will introduce procedures that are becoming widely used in other fields. These procedures are developed to study change. They need measurements taken at three or more times. They are becoming easier to use, but are still sufficiently complex that a technical expert is needed. Can the well-educated practitioner understand these procedures—we will soon know.

There are two statistical traditions used for studying change. Both refer to this as growth curve analysis. It is important to remember that growth can be positive or negative. One approach to studying growth curves is a structural equation modeling of latent growth curves and the other is called either hierarchical modeling or multi-level modeling. If these names sound reasonable to you, then you are way ahead of me. They confuse me. It is important to recognize them, only so we can use the procedures or evaluate their use by authors. We will focus on latent growth curve modeling using structural equation modeling, but multi-level data analysis is a powerful alternative approach.

How can we approach the study of change using structural equation models of latent growth curves? Whenever we are describing change, we need to identify the form of the change. It is important to remember that growth may be positive or negative. Change may be linear—going up or down in a straight direction—or it can be nonlinear such as going up rapidly and then leveling off. For example, if we were studying the transition from adolescence to adulthood, we might think of delinquency having a nonlinear change. Delinquency increases rapidly from 13 to 18, leveling off from 18 to 20, and then dropping off from 20 to 23. The “good” 13 year old boy becomes the “responsible” 23 year old adult, but everybody seems to suffer in between.

If we measured delinquency at the age of 13, 17, 20, 23, we might think the delinquency growth curve would look like Figure 1. In many cases, it is possible when using a linear model of change that the growth “curve” won’t be a curve at all—it will be a straight line. If you look at Figure 1 for the ages between 13 and 18 it is slightly curved, but nearly straight. We call this **piecewise linear growth** since we can break up the curvilinear growth trajectories in Figure 1 into separate linear components like the age period shown in Figure 2.

We might want to study the increase in delinquency problems among adolescent boys from age 13 to age 17. If so, a linear growth curve might be appropriate.

When using longitudinal data, researchers have plotted change for years. Although not common, charts like those shown in Figures 1 are hardly new. If we stopped with these figures, we would not need to understand latent growth curve analysis. Let’s not stop here. What else can we do with figures like these?

What happens over time—the Shape of the growth curve.

First, we might want to test the shape of the growth curve to see if it is linear or nonlinear. If it is nonlinear we might want to know whether it goes up, then down or goes up and then levels off. Much of the early work on growth curve analysis did this. It tested alternative models of growth and demonstrated which shape was most appropriate for the data at hand.

Suppose we wanted to test a linear model of growth. How many data points do we need? What if we only had data for the adolescent at age 13 and again at age 17. With two data points, a straight line will fit perfectly, every time, because two points determine a line. There is nothing to test. We say there are no degrees of freedom in this case—there is no data that could disprove prove the straight line.

If we have three data points, as we do in Figure 2, we can provide a test of the linear model. We have a degree of freedom, because we have data that could prove the linear model is a poor fit. How? Suppose delinquency remained at a low level from the age of 13 to the age of 15, but then skyrocketed between 15 and 17. Figure 3 shows such an expected result. It is a nonlinear growth curve and an attempt to fit it with a linear growth curve would result in a poor fit. When you read an article that says there is a poor fit, this means the model describes the data poorly. When they say there is a good fit, they mean their model provides a good description of the data. The statistics used to decide whether the fit is good or not are complicated and controversial, but the basic concept is no more complex than what we have described. A straight line fits data like that in Figure 2, but does not fit data like that in Figure 3.

Think of the topics you study and whether describing the growth curve is pivotal to your interests. For example, what happens to marital conflict in the first five years of marriage? What happens to the chances of divorce in the second five years of marriage? What happens to the division of household chores when a women enters the paid labor market on a full time bases? What happens to parents' understanding of sexual orientation in the first 12 months after they learn that their daughter is a lesbian? All of these are appropriate issues for latent growth curves.

Describing the growth curve—the intercept. We can do much more than describe the form of the growth curve. To keep the presentation as accessible as possible, let's focus on a linear growth curve. We know that two points determine a straight line. Similarly, two parameters determine a straight line. One of these is called the intercept or constant. It is the value at the start of the process. We sometimes call it a constant, because it is what we start with and the standard from which change is measured. Looking at Figure 2 we can see that the intercept for the range of data we have is 10. The 13-year-olds have an average delinquency score of 10. (Notice that this intercept is different from what you learned in statistics where we were told the intercept is the value of Y when X is zero.) Now, we are saying the intercept is the value of the outcome, delinquency, when the growth curve begins. We start at 13 years because we first measured their delinquency when they were age 13. I like to call the intercept the “initial level.”

The intercept is not telling us where a particular adolescent (Joe, Samatha, etc.) starts. It is the average or mean delinquency for the 13 year olds. Some children may have zero delinquency at 13. Others may have already engaged in many delinquent activities. We might want to explain the intercept. That is, we might think of the intercept as a dependent variable. Do girls have a lower intercept than boys? Do adolescents raised in “stable” families have a lower score than adolescents raised by continuously single parents? Is mother’s education relevant? It’s up to the researcher’s theoretical ideas to come up with predictors. The point is that we might want to think of the intercept as a variable. Sue has an intercept of zero. John has an intercept of 70. Why some people start (intercept) with a lower delinquency score is a question we should ask. Is this because Sue is a girl? Because she is from a stable family? Because her mother has a college degree? Because she has strong religious beliefs?

Describing the growth curve—the slope. In addition to the intercept, we have a slope. This parameter tells us how much the curve grows each year. Figure 4 shows a steep slope, with delinquency increasing a lot each year. As the adolescent grows from 13 to 17 the mean delinquency score jumps from 10 to 70. This is a 15 point growth each year. The slope, therefore, is 15.

This slope of 15 for Figure 4 is really an average or mean rate of growth. Many adolescents will have much smaller slopes. Sue may start at zero delinquency when she is 13 and never get involved in substantial delinquency rising to only a score of 5 by the time she is 17. Her slope is very small compared to the latent growth curve. This can be seen by comparing the average growth in delinquency shown in Figure 4 to what happened for Sue as appears in Figure 5.

Consider a second adolescent. Juan may start with an intercept of 60, a high level of delinquency, but by the time he is 17 he may have dropped to zero. This is shown in Figure 6. Why does the average adolescent become more delinquent? Why does Sue stay at a low level? Why does Juan become less delinquent? We ask these fundamental questions, regardless of what we are studying. We want to understand the average change, but we also want to understand individual variation in change. By variability in change, I mean why some people change more or less and why some people have get better and others get worse.

Figure 5 shows examples of 5 people who have different growth curves along with the bold red line for the overall growth curve. Sue and Juan are very different than the other three who are all pretty close to the line formed by the intercept and slope. The latent growth curve is a sort of average of all the individual growth curves.

There is another thing we can do using a multi-level approach to latent growth curve modeling. A popular topic among family scholars is how much can be explained by individual effects and how much can be explained by family effects. It is possible to partition the variance in the growth into several components and address this question. In fact, we can show whether there is a family effect that is separate from the individual effects of family members. This is most often done using multi-level analyzes through programs such as HLM (hierarchical linear modeling), but can also be done with some structural equation modeling programs.

We have only scratched the surface of what is possible. It is possible to examine how changes in one or more variables affect the changes in another variable. It is possible to do an exploratory analysis to locate clusters of cases that have very different growth curves. If we had a 1,000 adolescents in our study, there might be a cluster that does not change (adolescents like Sue), a small cluster that changes in the opposite direction of the overall pattern (adolescents like Juan), and a cluster that changes very rapidly.

A former student of mine, Fuzhong Li, is a co-author of a book on Latent Growth Curve Modeling (Duncan, et al., 1999) that illustrates a wide range of applications using family data and you are encouraged to explore these applications. I've included a reference to his work and other work I've done with him at the end. His book has an extensive set of references. We have also prepared a draft of a manual illustrating how to estimate growth curves for using different software. For our purposes, we will limit our goals to what we have already covered:

- What happens over time? Is growth linear? Is there a plateau? Does it go up, then down? Can we develop a growth model that fits the data?
- Where does the process start? What is the initial level (intercept)?
- How rapidly does the process develop? Is there a steep slope? If it is nonlinear, when does the direction change?

- What accounts for the initial level (intercept)?
- What accounts for the rate of growth?
- Does the rate of change in one attribute relate to the rate of change in another?

Section 2: It's About those Ugly Figures

You may have seen papers that use latent growth curve models. These usually present figures that are designed to simplify the presentation, but have the opposite effect for many readers. Many readers jump to the Discussion section immediately—worse, they flip to the next article. If you are going to understand latent growth curve articles, it is helpful to have a basic understanding of the figures. They use a variety of specialized notation that is confusing. You don't really need to understand the notation. Figure 8 shows a typical model of a latent growth curve. If this figure makes even a little bit of sense to you, you will want to come to the workshop I am presenting. Our purpose is to de-mystify the figure.

We will start with the boxes that contain y's (see Figure 9). There are four of these. They represent the variables we actually measure. Figure 9 shows that we have measured it 4 times (y_1, y_2, y_3, and y_4), when the adolescent was 13, then again at 14, 15, and 16. We have a score for each adolescent for each of these four years. This single group of adolescents are a panel. Different authors will use different notation for these measures. Some might be so nice as to use more descriptive names such as delinquency_13, delinquency_14, delinquency_15, and delinquency_16—such kindness is unlikely.

At the bottom of Figure 10 are some e's (e_1, e_2, e_3, and e_4). What are these? A great benefit of a structural approach to latent growth curves is that we can allow each variable to contain some measurement error. We know we cannot measure delinquency or anything else without some error. Studies report the reliability of measures and we have learned that the reliability should be .8. This means that some of the variability in the time-specific measures of delinquency is error. Conventional statistical procedures assume that there is no error! Latent growth curve models can incorporate time-specific measurement error into the model. This is a tremendous advantage compared to traditional procedures.

If everybody used an e for error, this would be a convention worth remembering. Unfortunately, a variety of symbols are used including ε and δ. The important point to know is that these things represent measurement error.

Two ovals are added in Figure 11. One of these is the intercept and the other is the slope. These are what we are most interested in and have been discussing in Section 1. The intercept tells us the initial level at the start of the process and the slope tells us the rate of change. There might be more than one slope for complex nonlinear models.

The lines from the intercept and slope to the y's are the "trick" of the structural equation modeling approach to latent growth curves. The four lines from the intercept to the four y's are all fixed at a constant value of 1. This is why the intercept is sometimes called the constant by statisticians. Since all four values are fixed at the same value, this is the constant level of delinquency, **if there were no growth**. In Figure 2, the intercept will be 10.

Figure 12 shows four lines from the intercept to the observed variables, y_i's. Each of these paths is fixed at the constant value of 1. Other values would work, but this is the convention. Since all of the values are the same, this establishes the initial, or constant, level of delinquency. You can think of it as establishing what level of delinquency adolescents would have, if there were not growth between age 13 and 16.

The four lines from the slope to the four y_i's that appear in Figure 13 and their values may not make any sense initially. The values for the lines are fixed at 0, 1, 2, and 3 respectively. Can you guess why the line from the slope to y_1 is fixed at zero? There has been zero growth at y_1 since this is the initial level. Using values of 1, 2, and 3 results in a linear growth curve, since there is a 1-year difference between each measure. The y_1 is delinquency at the age of 13, the y_2 is one year later at 14 (so its value is 1), the y_3 is one more year latter at 15 (so its value is 1 + 1 = 2), and y_4 is one more year later at 16 (so its value is 1 + 1 + 1 = 3).

Fixing the values from the slope is how you identify model growth. It is fairly complex. Suppose we had no data for when the adolescents were 13. Can you see how we would modify the fixed values? If we had measured delinquency at the ages of 13, 15, and 16 we would use 0, 2, and 3 for the three lines we would have to get a linear growth curve. We would skip the first year data we have fixed at a value of 1 in Figure 13. If the model should have a curve, we would have two slopes. One would be fixed at 0, 1, 2, and 3 as in Figure 13. The Second slope would have the four paths fixed at 0, 1, 4, and 9, the square of the corresponding the linear paths. Some researchers fix only the first two paths of a single slope at 0 and 1 and let the program estimate the values of the remaining paths. This can allow a very complex model of growth.

Figure 14 returns us to our full model. There are two parameters associated with both the intercept and the slope. We use the label Mean μ_i (mean mu sub i) for the mean of the intercept and Var D_i (variance D sub i) for the variance of the intercept. The mean is the intercept. Using Figure 2 this would have a value of 10, the initial value of the growth process. The variance of the intercept reflects the variation of individual intercepts. If everybody had an intercept that was very close to 10, the variance of the intercept would be small. This would be the case for Figure 7 if everybody were like Maria and Cecil. However, some individuals may vary more from the mean intercept as is true for Sue and Robert, and especially true for Juan whose individual intercept is 60. When we try to explain the intercept using other variables (e.g., gender, religiosity), we are trying to explain this variance or why Juan was so delinquent at the age of 13 and why Sue was not delinquent in any way at the same age.

The mean slope is represented by Mean μ_s (mean mu sub s). This is the average slope as is represented by the growth curve in Figures 2 and 7 by the red line. The variance of the slope is represented by Var D_s (variance D sub s). Looking at Figure 4, we can see what this means. The variance of the slope reflects the extent to which individuals have different slopes. Explaining this variance amounts to explaining why some individuals have a more rapid increase in delinquency than others between the age of 13 and 16 and why some individuals may actually have a decrease in their deviance during this period. Different authors use different symbols to represent the mean and variance of the intercept and slope. You can usually recognize what they are by looking at how the figure is drawn.

There is only one more thing to understand about Figure 14. We show a curved line with an arrow on both ends between the variance of the intercept and the variance of the slope. This represents the covariance (correlation) of the two variances. You may draw a blank at what this means, but it can be theoretically important. Suppose we have a positive covariance (correlation) between the two variances. Can you see what this means? It means that individuals who have higher intercepts also have higher slopes and individuals with lower intercepts have lower slopes. In other words, an adolescent who starts out with a high initial level of delinquency will become increasingly more delinquent than others. Alternatively, a person who starts at a low level of delinquency will "grow" more slowly than others. Sue would be an example of this since she starts low and ends low. My hunch is that counselors assume there is a positive covariance. They assume that adolescents who are in trouble at 13 are going to get into serious trouble by the time they are 16—they are "bad kids." and adolescents who are never in trouble at 13 will remain "good kids." Usually we do not have an explicit hypothesis about this covariance. Traditional statistical procedures must assume the intercept and slope are independent. It is curious to hear people say the assumptions of structural equation modeling are too restrictive, when the assumptions of traditional approach are much more restrictive (no measurement error, no correlated variances).

Section 3. An Extension of the Model to Include Predictors

The first extension of the model involves finding predictors. We've already mentioned a few. We've suggested that gender is important and that girls should have both a lower intercept and a flatter slope than boys. Girls should be systematically less delinquent at 13 and increase less in delinquency between 13 and 16 should be gradual. Boys should be systematically more delinquent at 13 and their level of delinquency should increase more rapidly between 13 and 16. Figure 15 illustrates how this would look.

We might believe that church attendance has a positive influence. Perhaps we feel that it will not have much influence on the level of delinquency at 13, but that there should be an increasing effect of church attendance over time. This would mean that church attendance would not influence the intercept, but would influence the slope. Finally, we may believe that mother's education is an important protective factor and will influence the slope. The higher the mothers education is, the lower the slope. In other words, mothers with limited education risk having their children become increasingly delinquent with age. The education disadvantage becomes greater as the adolescent moves from 13 to 17 years of age.

How can we represent such a complex latent growth curve graphically? Figure 16 does this. It adds the three predictors, gender, church attendance, and mothers' education. There is a line with an arrow pointing toward both the intercept and slope that originates at gender. Church attendance and mothers' education have lines pointing to the slope, but not the intercept. These lines appear as bold blue lines to make them easier to see. The rest of the figure is unchanged from Figure 5, except we have labeled the intercept “initial delinquency” and the slope “slope delinquency” to be more descriptive of what we are doing.

Section 4. Interpreting a Complex Example

Daughters’ caregiving for their elderly mothers interested Alexis Walker, Sally Bowman, Fuzhong Li, and me. One question we asked was whether changes for care given influenced changes in the level of care giving satisfaction. That is, does giving more care, reduce the level of satisfaction. This question asks if change in one variable leads to change in another variable—one of the most common questions we ask. Until now, we had very limited ways of answering this question.

We created two latent growth curves, one for changes in caregiver satisfaction and the other for changes for care giving. We have two primary hypotheses about the initial level of care (a) the higher the initial level of care the lower the initial satisfaction and (b) the higher the initial level of care, the greater the decline in satisfaction. We had one hypothesis, our most important one, that asserted that (c) the greater the increase in care given, the greater the decrease in caregiver satisfaction.

Figure 17 shows how we can represent these ideas using a structural equation modeling approach to latent growth curve analysis. The bottom two ovals represent the initial level of satisfaction and the slope, change in satisfaction. Notice how the mean slope is -.230, indicating that there is a decline in satisfaction over time. The growth is negative.

The top two ovals represent the initial level of care giving and the change in care giving. The mean slope for the change in care given is .144. This is not surprising, because care needed often increases with time.

The most important hypothesis is (c); that increases in care giving leads to decreases in care giving satisfaction. The effect shown in Figure 17 is -.638, $p < .01$. This supports our hypothesis. Notice, we are talking about change in one variable leading to change in another variable. Our dependent variable is a slope and so is our independent variable. You can see that our first hypothesis is also supported. The higher the initial level of care provided, the lower the initial level of satisfaction. The coefficient is -.541, $p < .01$. Our other hypothesis is not supported. The initial level of care given does not influence the change in satisfaction (.090, ns). This shows that it is change in the care giving, rather than the level of care giving that explains decreases satisfaction.

This model has two covariates, co-residence and duration of care giving. These are included as controls. The goodness of fit information is provided to the right of the figure. The chi-square is significant. Unlike most uses of chi-square, this is not good news. The significant chi-square means that our model does not perfectly reproduce the data. However, structural equation modeling experts often emphasizes other measures of goodness of fit. We have shown two of these, the normed fit index and the comparative fit index. Since both of these are above .90, the model does a reasonable job of fitting the data, just not a perfect job.

Section 5. What About Qualitative Data

Latent Growth models benefit from qualitative data. Some of the people you select to illustrate individual growth curves (see Figure 7), should be allowed to speak in their own words. The exceptional cases may provide some of the most useful qualitative information. An open-ended interview with Juan, because his delinquency decreased so much, might provide understanding of how delinquency problems may be mitigated among other adolescents.[1]

---

1 In the Walker et al. (1996) paper we discuss several individuals. Here is an actual example for case 9:

Case 9 presents a more positive picture than average. “Ivy,” the daughter, was 72 years old at the onset of

Section 6. What is Next?

The workshop at xxxxx will review these models from a programming point of view. This presentation will be much more technical and deal with the computer software programming. You might want the manual that will be discussed at that session, even if you cannot attend. It may be useful when you get home if you have a technical support person who works with you. The manual and the notes for this presentation may be read and printed from the Web, http://osu.orst.edu/dept/hdfs/papers/paper.html . The manual has more information on interpretation and covers model specification, programming, and interpretation of output for the leading structural equation modeling software.

---

the study, when she had been caring for her mother, "Virginia," for 11 years. She and her mother lived together. Virginia had high blood pressure, heart problems, and arthritis, each of which restricted her activities. At Wave 1, Ivy ran all of her mother's errands, completed less than half of the indoor tasks, prepared most of her meals, and did all of the outdoor maintenance. Ivy also acted as a mediator for Virginia with health care professionals and others. Her caregiving satisfaction was above average. Through both Wave 2 and Wave 3, Ivy's satisfaction declined slightly when Virginia suffered hearing and vision losses. By Wave 2, Ivy had taken on some financial management tasks for the first time, and she began to provide a little ADL help. In Wave 3, Ivy kept those responsibilities and began giving financial aid to her mother as well. By Wave 4, Ivy, who had her own health problems, found a secondary caregiver to take on some of her responsibilities. Virginia's general health improved somewhat so that she no longer needed ADL help and could function more effectively in the kitchen. Furthermore, Ivy no longer did any outdoor work. As revealed by the figure, her caregiving satisfaction in this last wave showed a marked

I think that with this background and The Walker et al. (1996) paper, you should be able to begin to understand and even use latent growth curve models. We have all survived the development of other new analysis strategies. Latent Growth Curves is one of the most powerful procedures I've seen and it is especially salient to the theoretical questions we ask in family studies. I expect you will see a rapid profusion of these models over the next five years. Hopefully, we have done enough today to put you ahead of the game!

improvement relative to the three earlier time periods.

References

Duncan, Terr. E., Duncan, Susan C., Strycker, Lisa A., Li, Fuzhong, and Alpert, Anthony. 1999. An Introduction to Latent Variable Growth Curve Modeling: Concepts, Issues, and Applications. Mahwah, NJ: Lawrence Erlbaum.

McArdle, J. J. (1988). Dynamic but strctural equation modeling of repeated measures data. In R. B. Cattel & J. Nesselroade (Eds.), Handbook of multivariate experimental psychology (2nd ed., pp. 561-614). New York: Plennum.

McArdle, J. J., & Epstein, D. (1987). Latent growth curves within developmental structural equation models. Child Development, 58, 110-133.

Meredith, W., & Tisak, J. (1990). Latent curve analysis. Psychometrika, 55, 107-122.

Muthén, B. (1991). Analysis of longitudinal data using latent variable models with varying parameters. In L. Collins & J. Horn (Eds.), Best methods for the analysis of change: Recent advances, unanswered questions, future directions (pp. 1-17). Washington, DC: American Psychological Association.

Rogosa, D. R., Brandt, D., & Zimowski, M. (1982). A growth curve approach to the measurement of change. Psychological Bulletin, 92, 726-748.

Willett, J. B., & Sayer, A. G. (1994). Using covariance structure analysis to detect correlates and predictors of individual change over time. Psychological Bulletin, 116, 363-381.

Walker, Alexis J., Acock, Alan C., Bowman, Sally R.,, and Li, Fuzhong. 1996. Amount of Care Given and Caregiving Satisfaction: A Latent Growth Curve Analysis. Journal of Gerontology: Psychological Sciences. 51B:3. Pp 130-142.

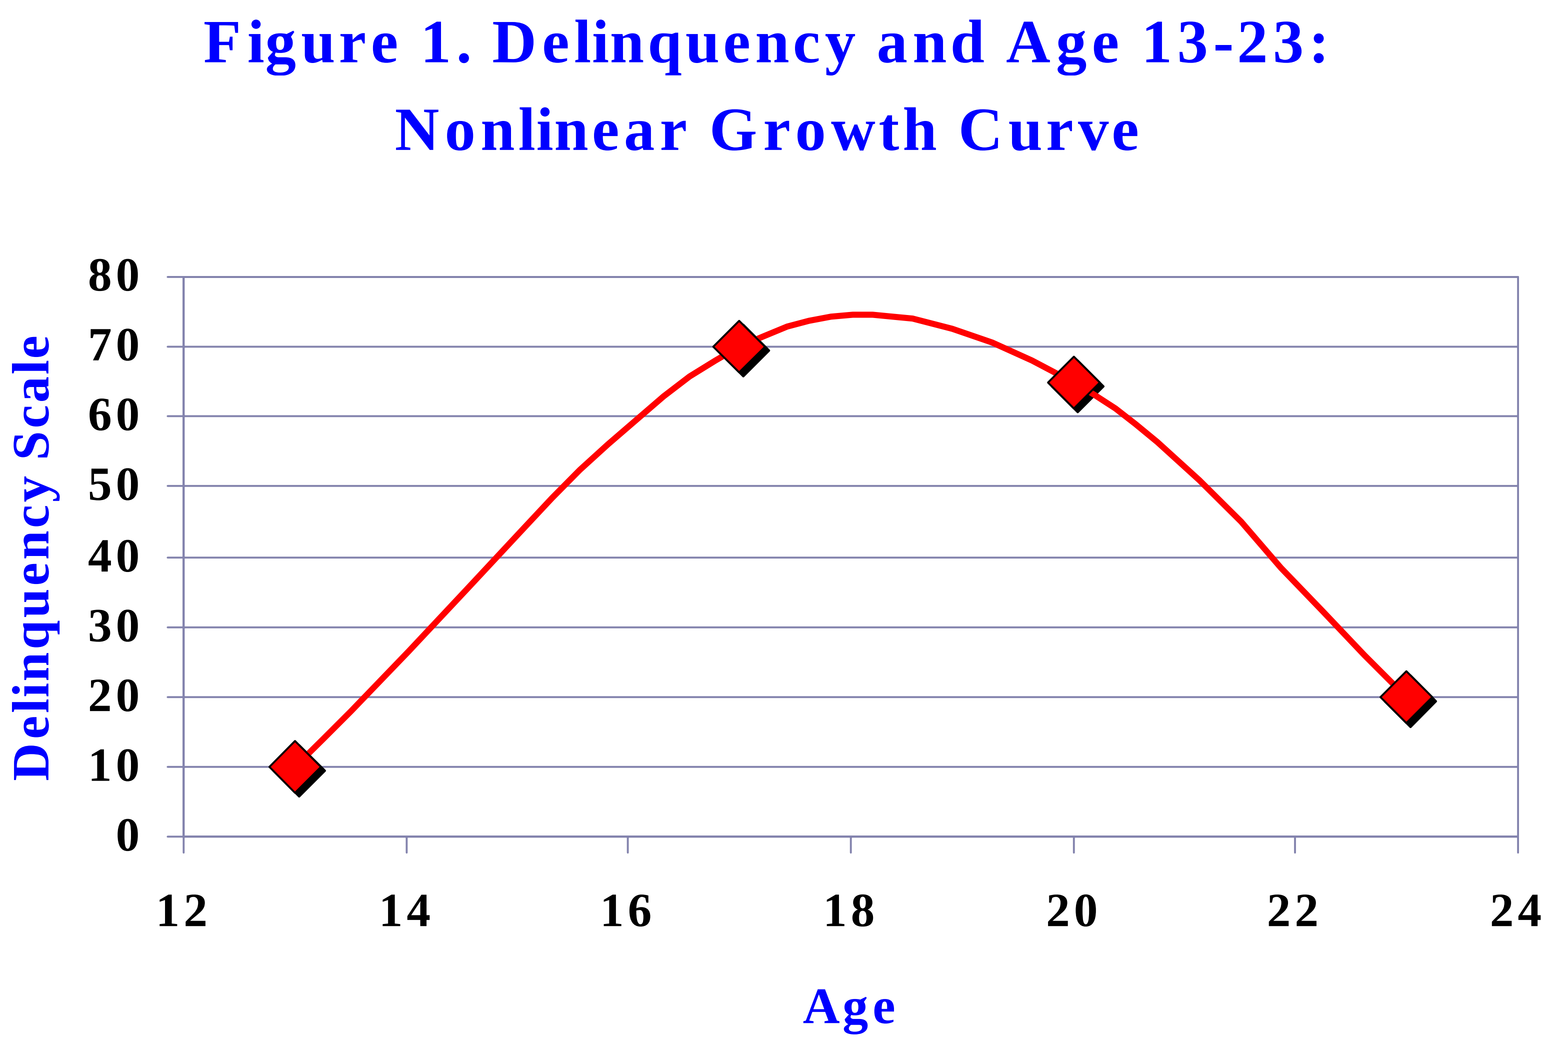
Figure 1. Delinquency and Age 13-23:
Nonlinear Growth Curve
Delinquency Scale
80
70
60
50
40
30
20
10
0
12
14
16
18
20
22
24
Age

Figure 2. Delinquency Growth 13-17: A Linear Growth Curve

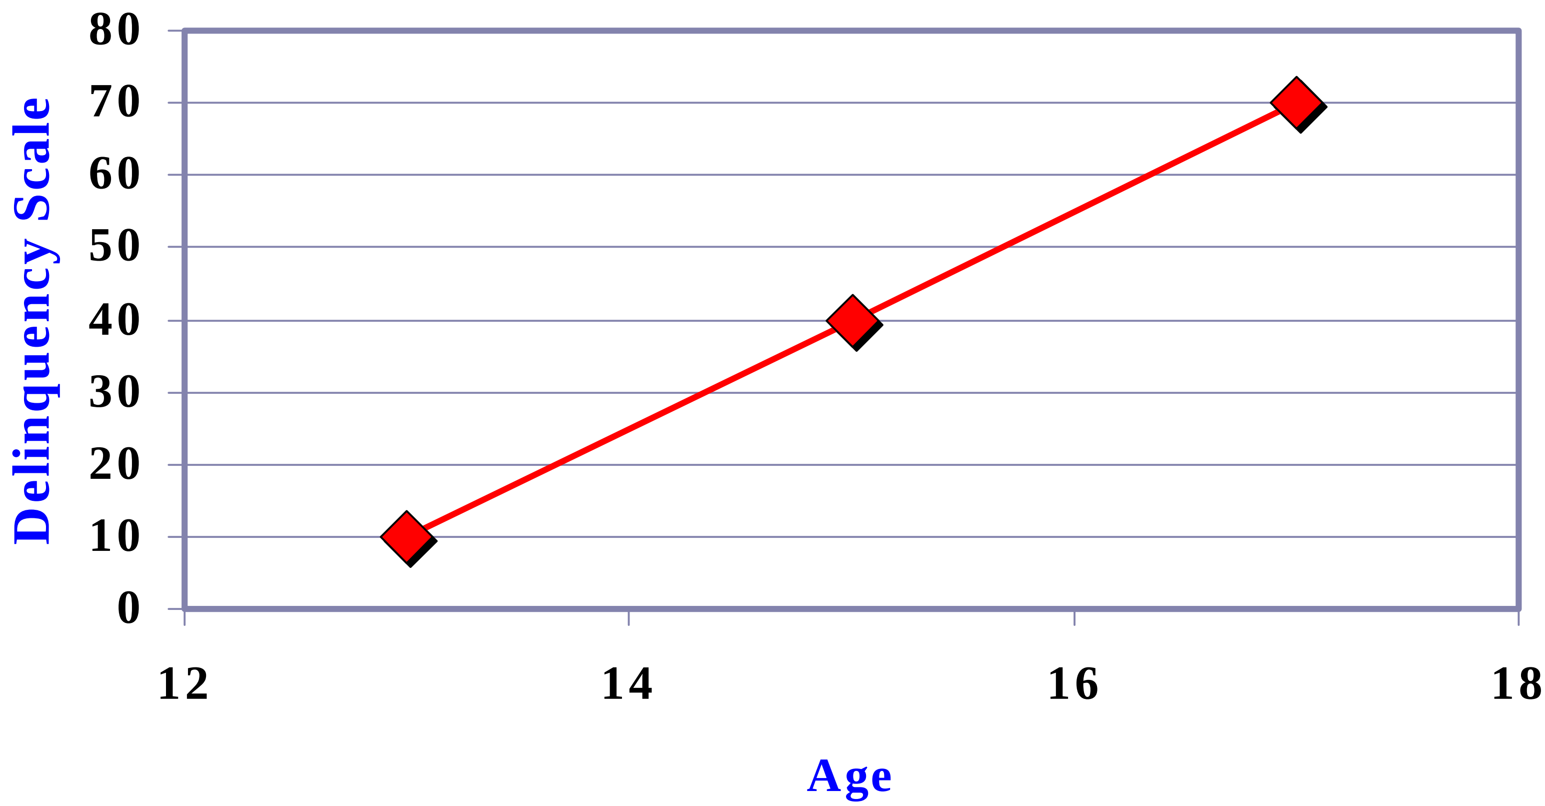

Figure 3. Delinquency Growth 13-17: A Nonlinear Growth Curve

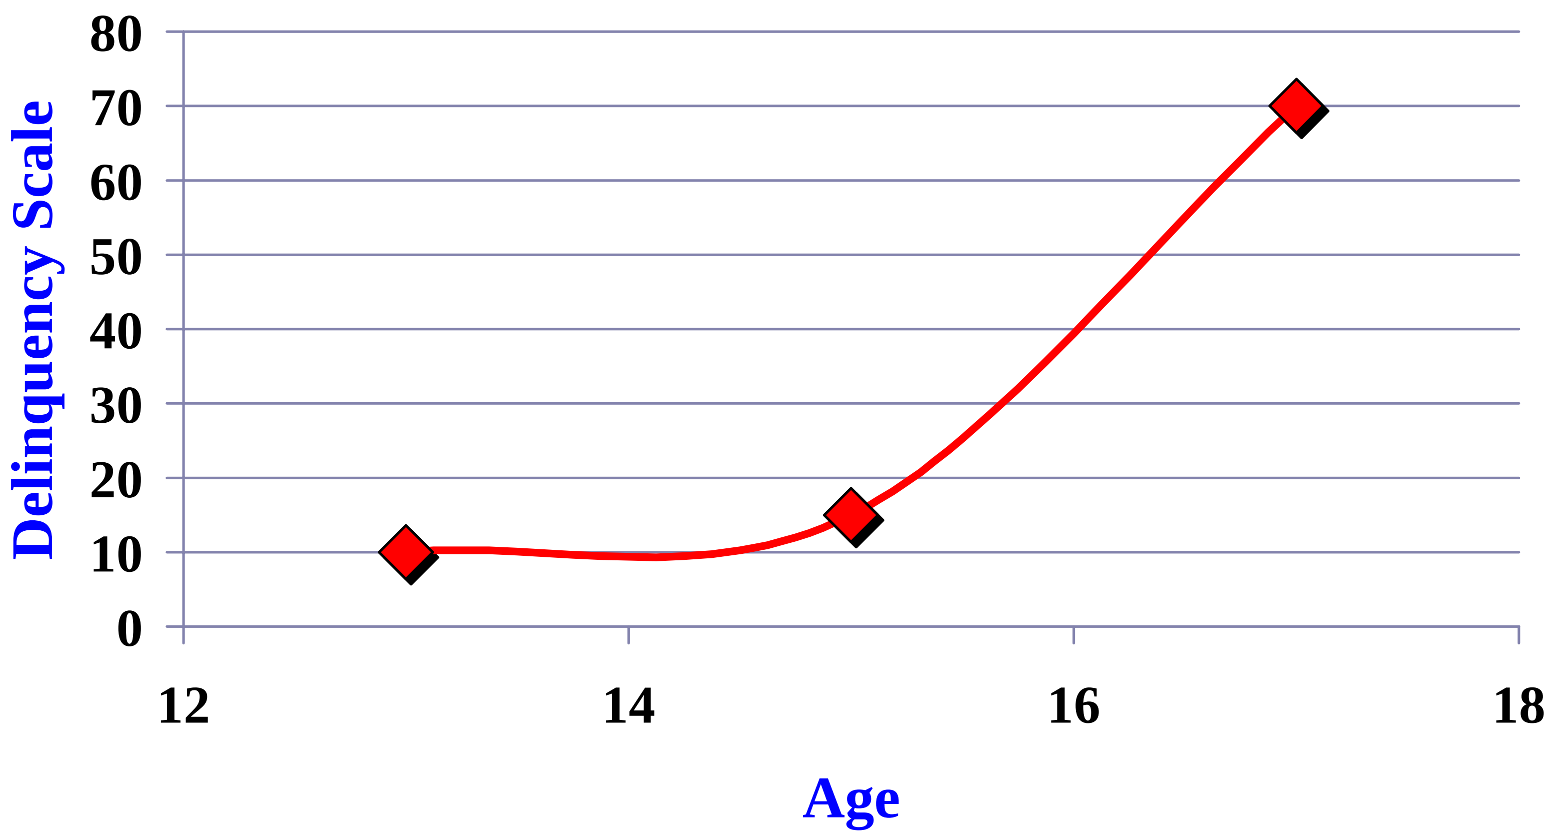

Figure 4. Delinquency Growth 13-17: A Linear Growth Curve

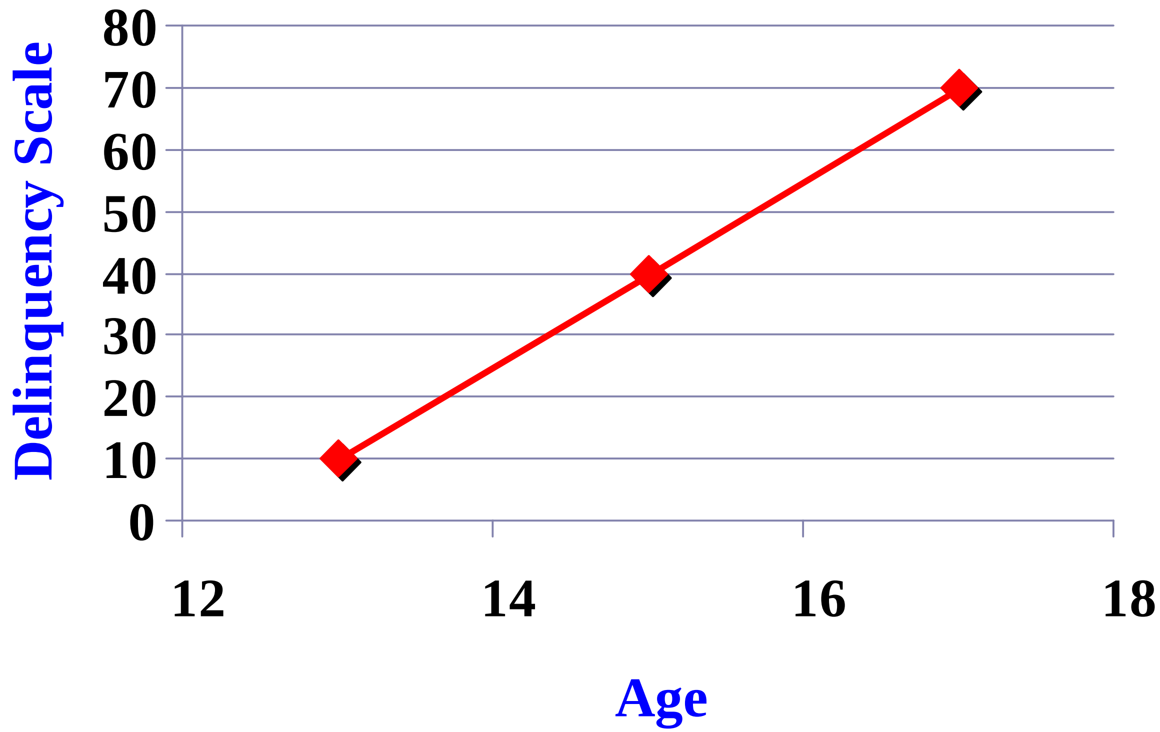

Figure 5. Sue's Growth Curve is Very Different

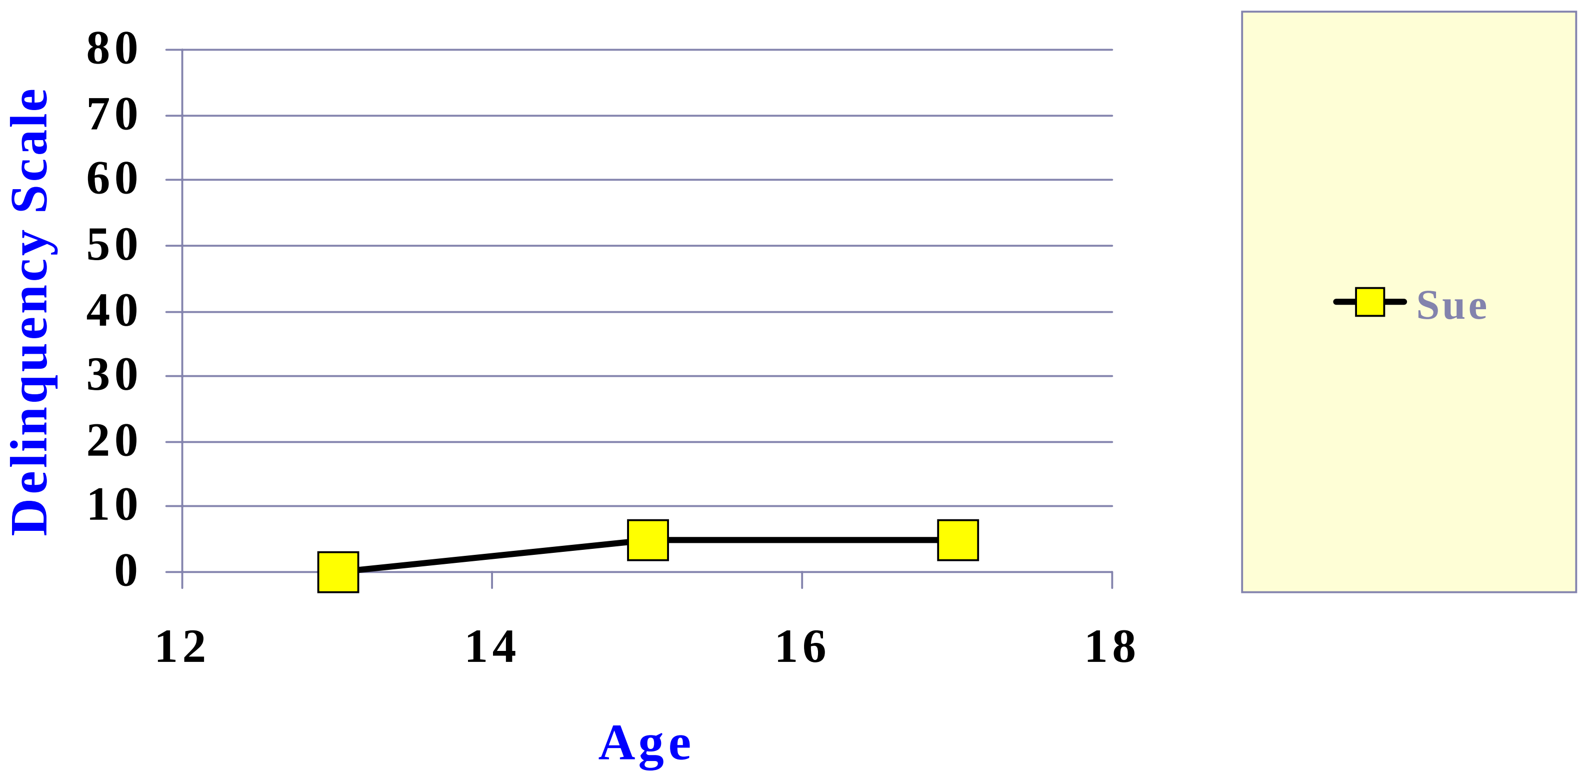

Figure 6. Juan's Growth Curve is in the Opposite Direction of the Overall Curve

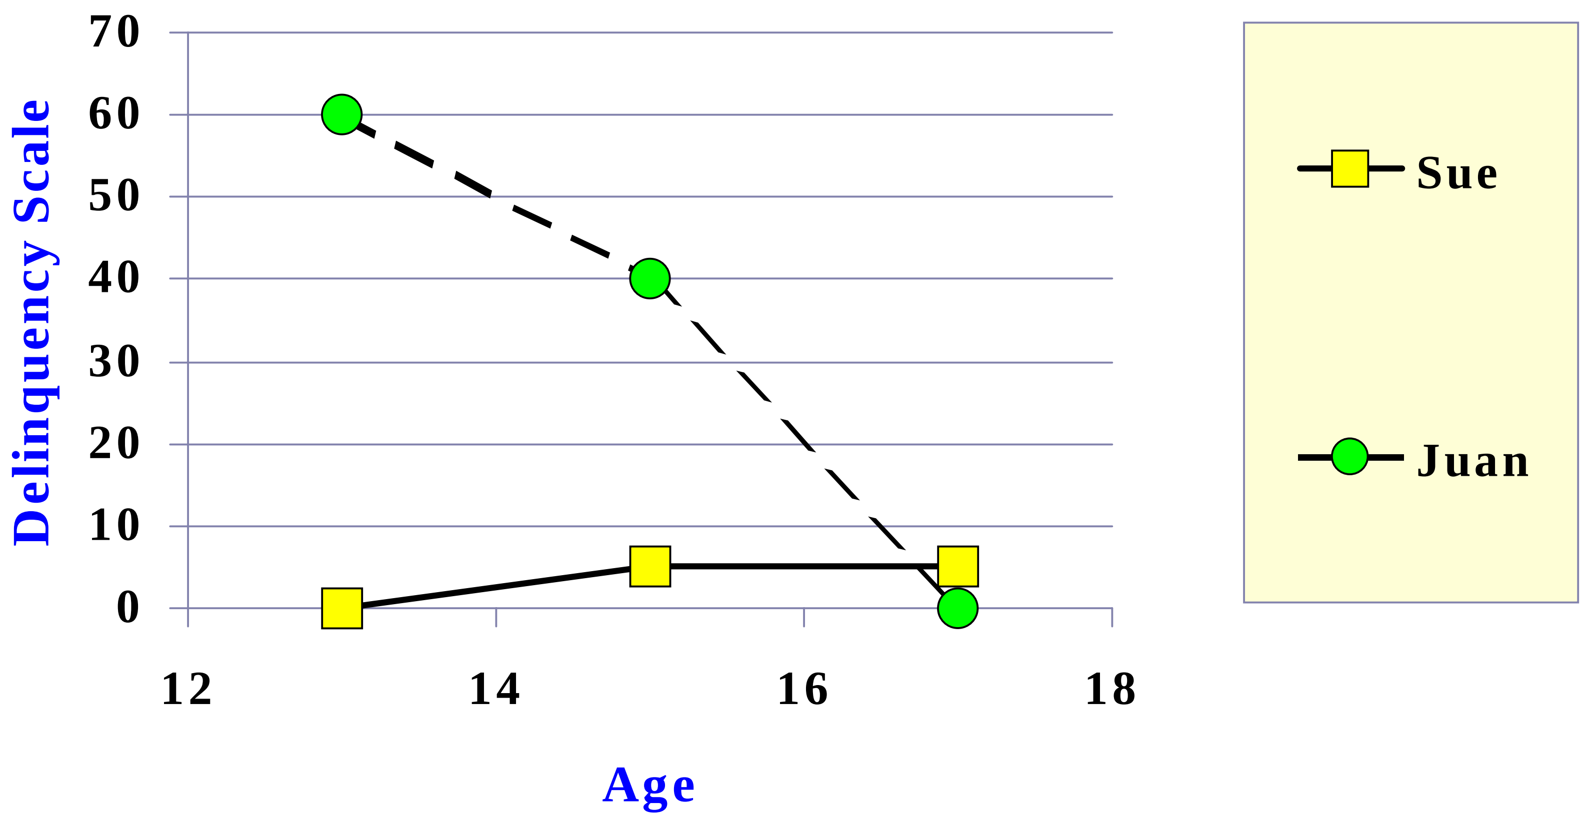

Figure 7. Adding Cases that Follow the Overall Growth Curve

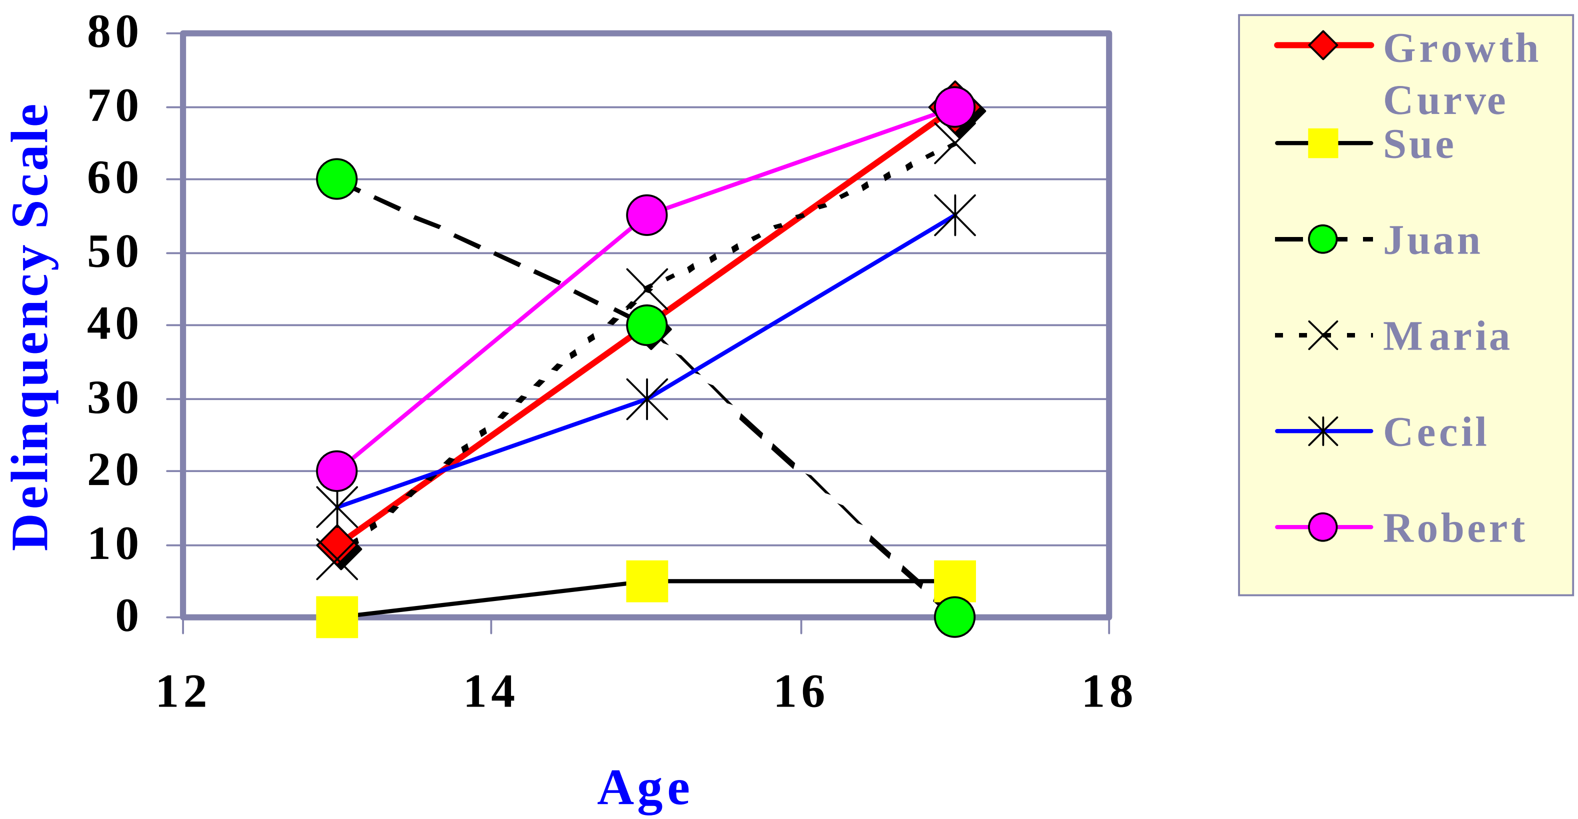

Figure 8. Conceptual Path Diagram

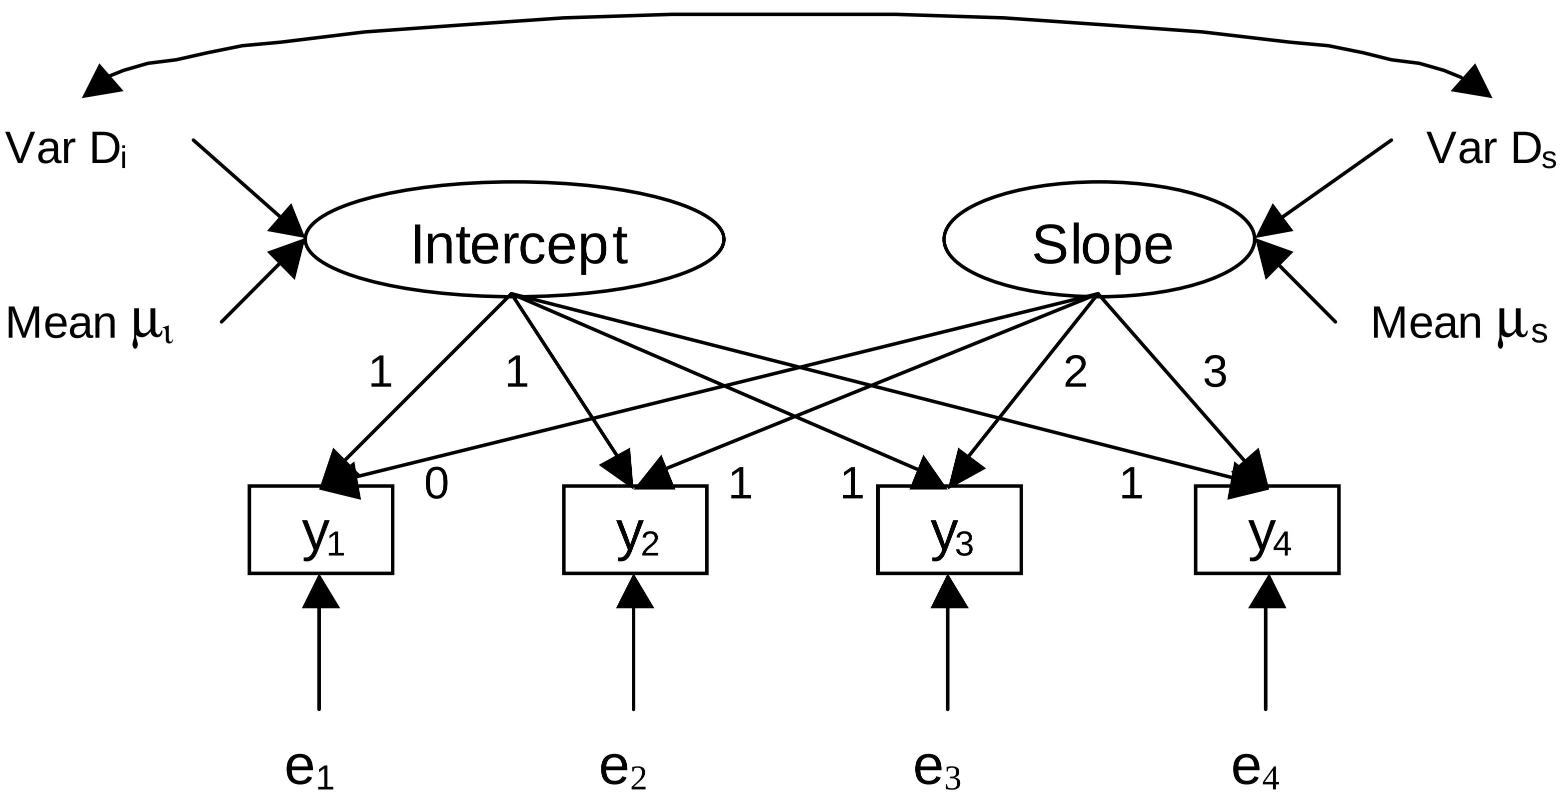

Figure 9. Delinquency Measured at Age 13, 14, 15, and 16

Figure 10. We assume our measures are not perfectly reliable

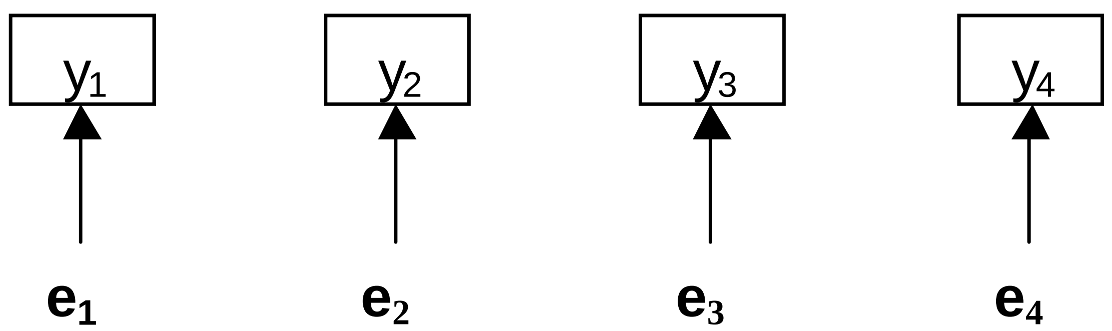

Figure 11. The Intercept and Slope are Latent Variables

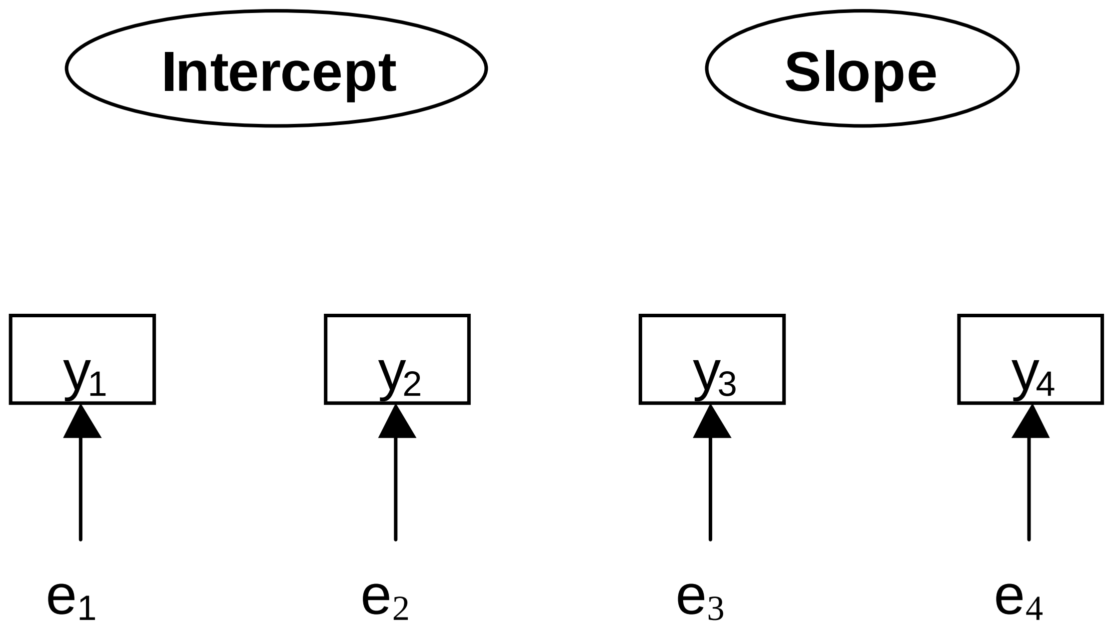

Figure 12. This is How We Identify the Intercept

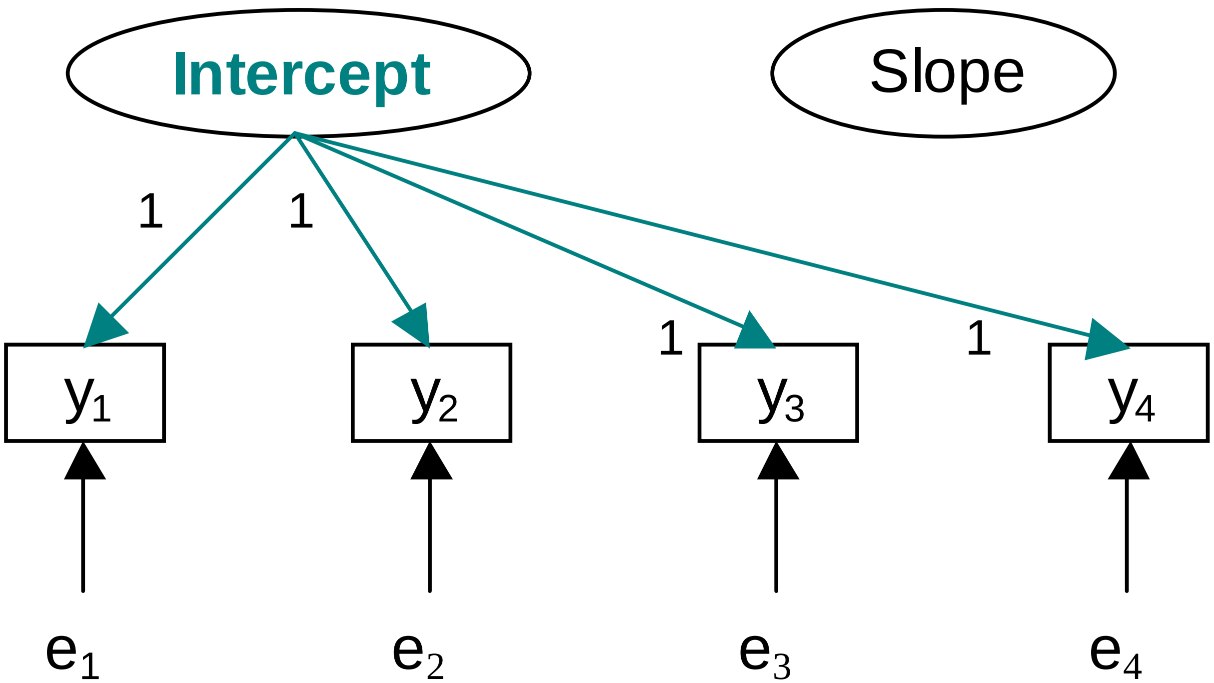

Figure 13. Now We Identify a Linear Slope

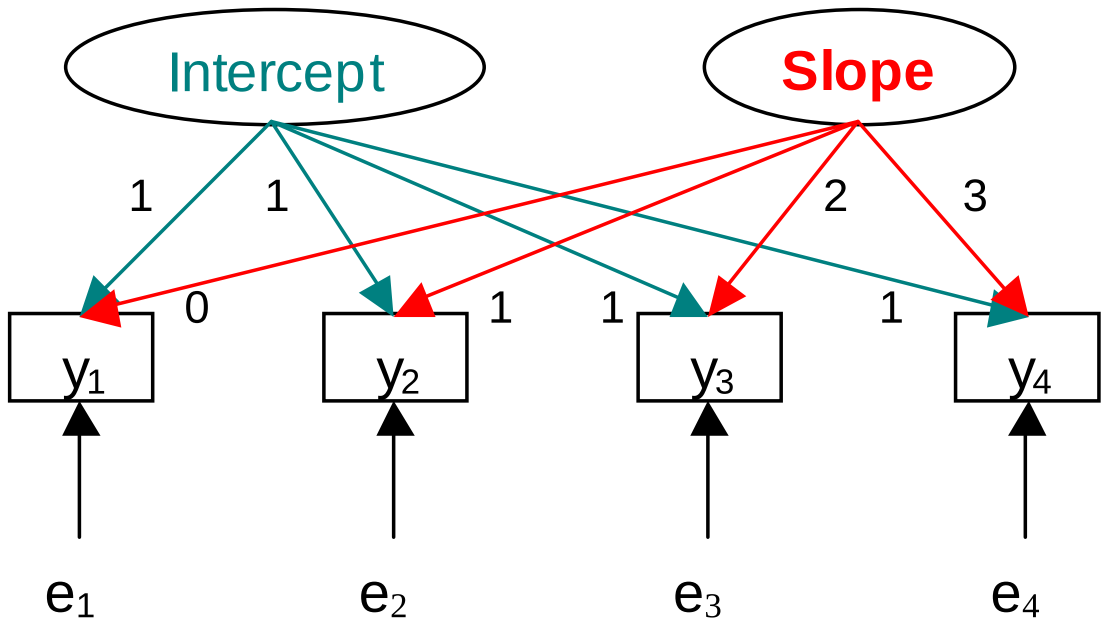

Figure 14. The Intercept and Slope Each have a Mean and Variance

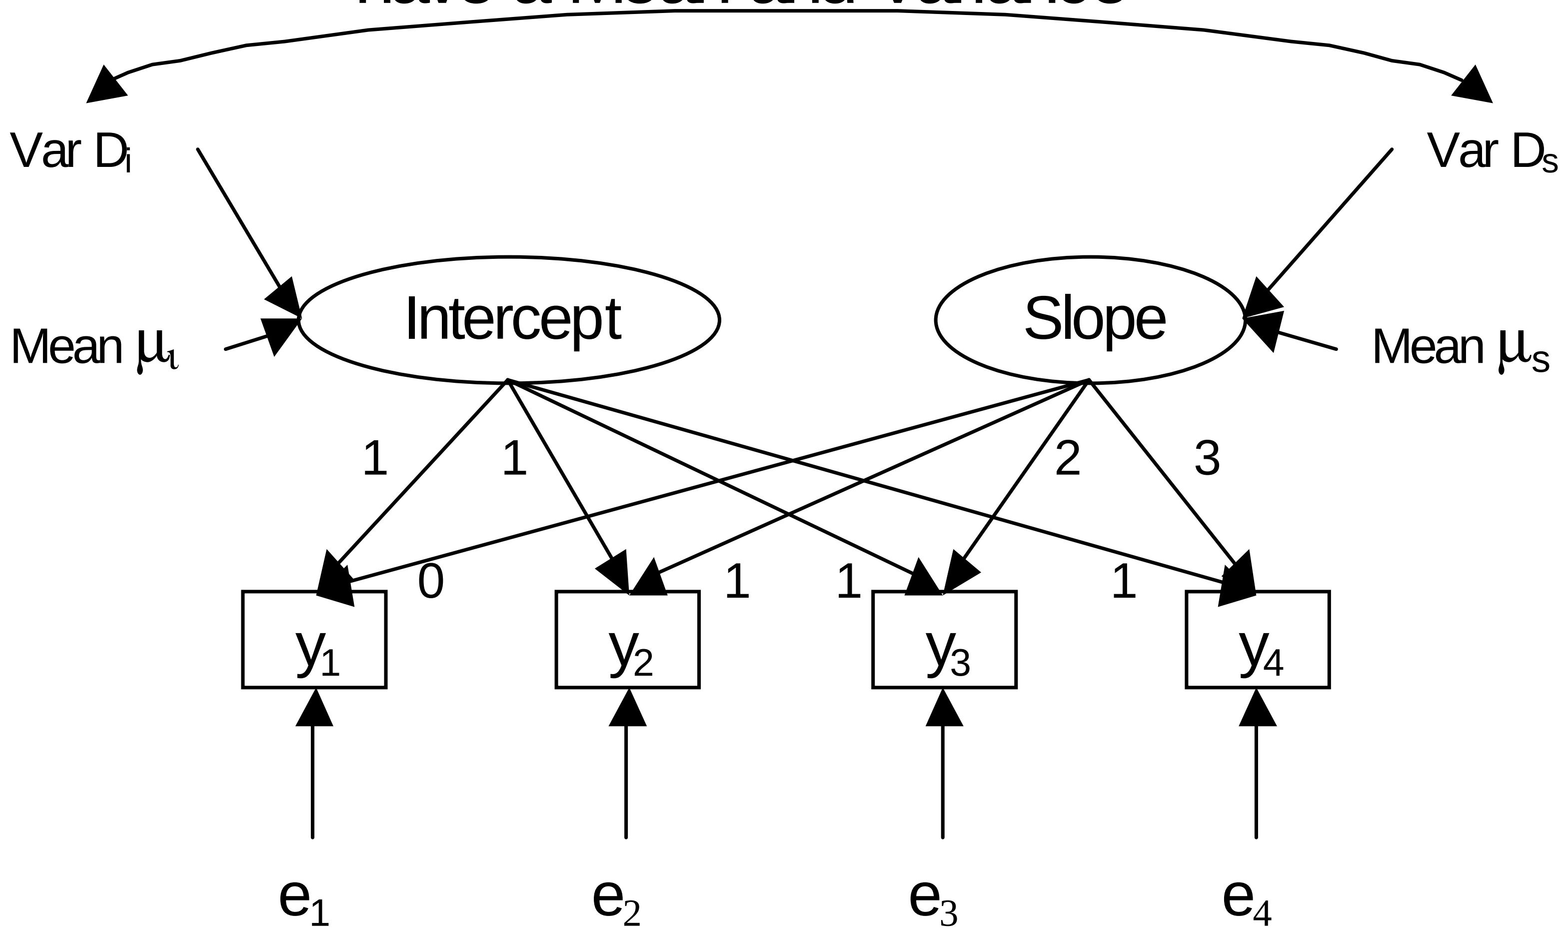

Figure 15. Delinquency Growth 13-17: Comparing Girls and Boys

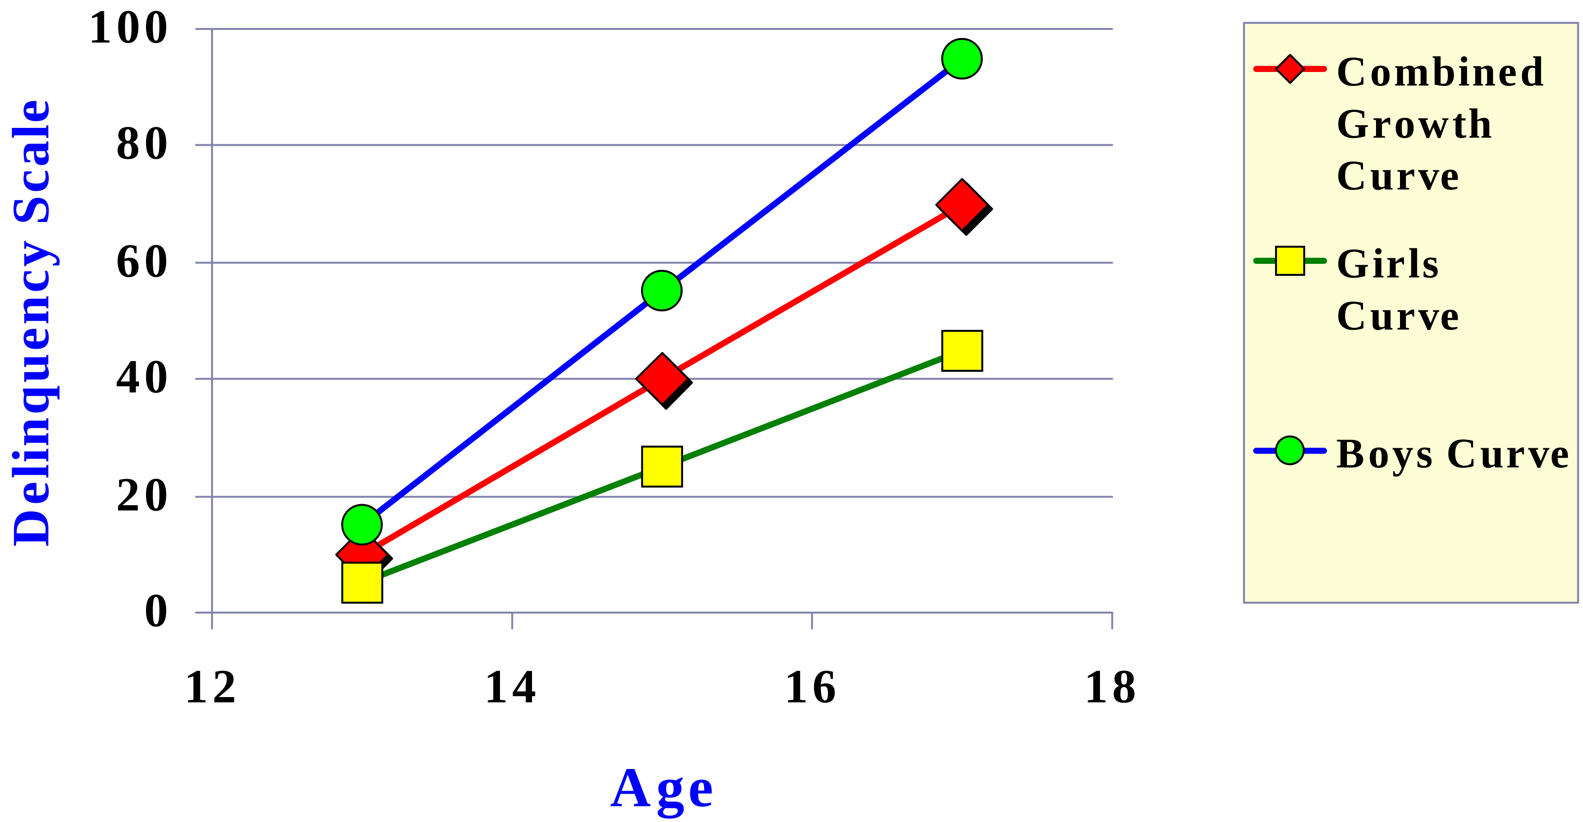

Figure 16. Predicting the Growth of Delinqent Behavior Among Adolescents

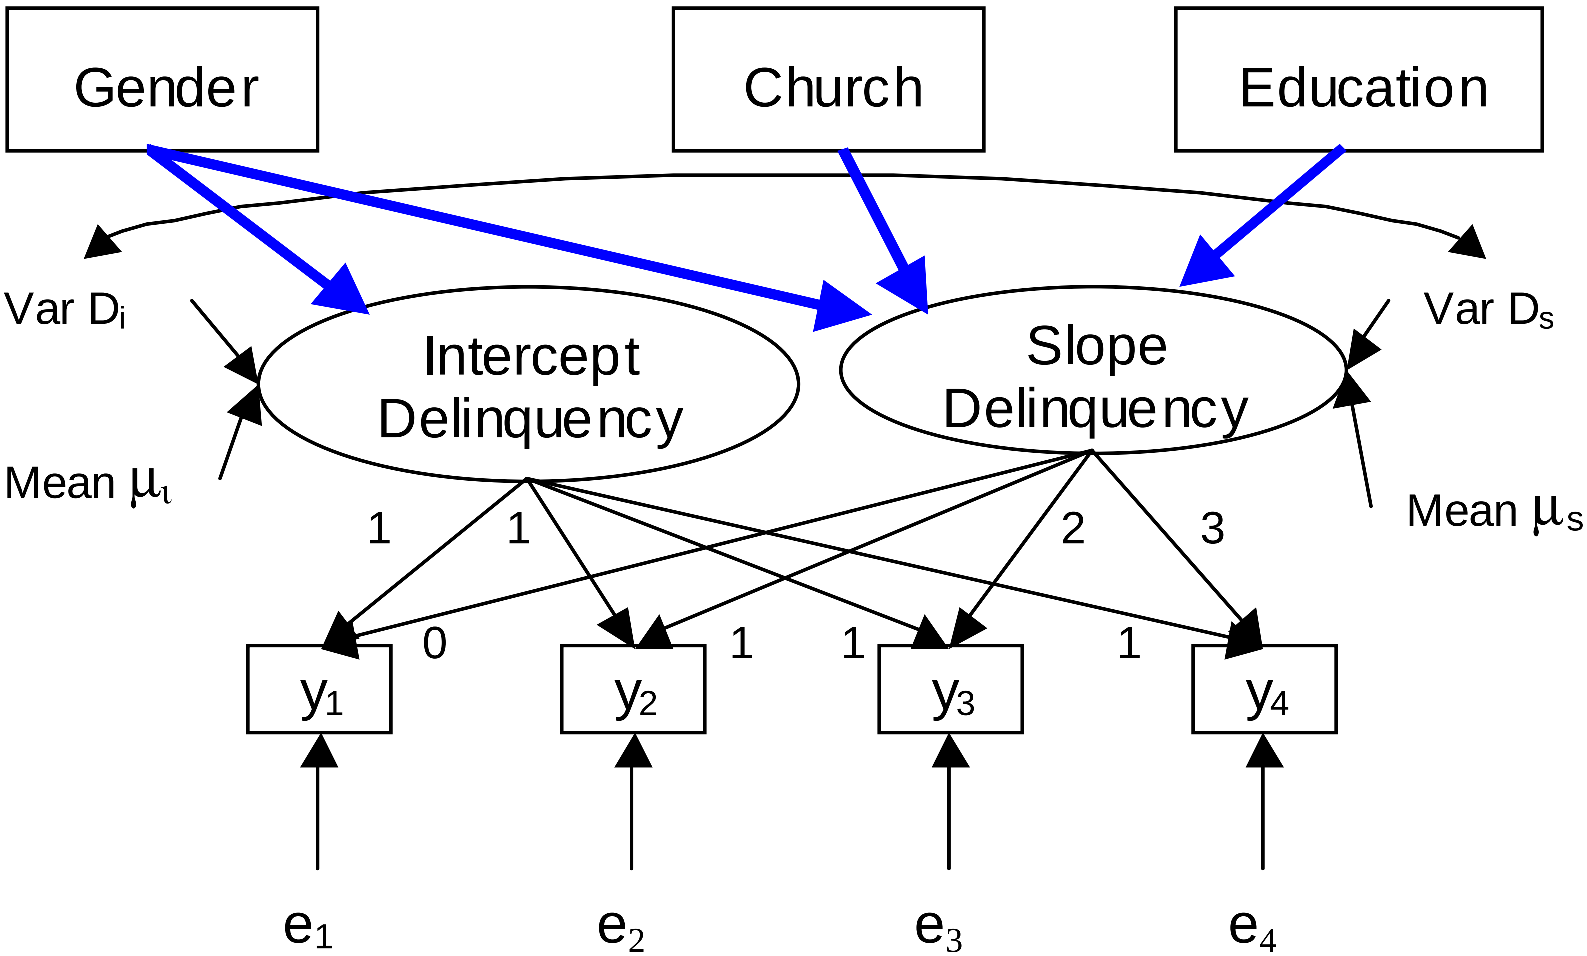

Figure 17: Empirical Example with Two Growth Curves