1

Additional Applications of Latent Growth Curve Modeling

In this document, I expand on topics relevant to latent growth curve modeling (LGCM) as applied to RETs. I begin by presenting an example of a traditional two group intervention versus control RET. This contrasts with the example in the main text that compared two intervention conditions but with no control group. I next provide an example of an RET with three treatment conditions, two different intervention types *and* a control condition.

After this, I address SEM-based multi-group approaches to the analysis of LGC RET data followed by non-linear growth functions. I develop the logic of latent basis growth curve modeling, piecewise latent growth curve modeling, and log-based growth curve modeling in this context. I next consider how to deal with binary and ordinal outcomes followed by strategies for dealing with measurement error, including the use of latent variables and sensitivity tests. I then discuss strategies for including covariates in LGC models, both time-varying and time-invariant covariates. Finally, I address alternatives to LGCM including Bollen and Curran's (2022) autoregressive latent trajectory (ALT) modeling. I assume you have read the main text of Chapter 16 and are familiar with the concepts discussed therein. Each section stands independent from other sections, so you can skip text as you deem fit.

TRADITIONAL RCT/RET EXAMPLE

This example uses the same structure as the LGC model in the main text but instead of comparing two treatments, Treatment A and B, it compares a treatment group and a control group. Although the two analytic approaches are quite similar, there are some new lessons to be learned from a more traditional application, hence I consider it here.

In the current example, the intervention group is scored 1 and the control group is scored 0 via a treatment condition dummy variable. There is a single mediator and I use a parallel process LGA model per the main text. Higher scores on the mediator, M, and higher scores on the outcome, Y, are desirable. The intervention sought to increase them both. The mediator and the outcome were each measured at five time points. The baseline measures, M0 and Y0, were obtained prior to the start of treatment with M1 and Y1 representing scores at the immediate posttest. The five assessments of the mediators are referred to as M0 M1, M2, M3, and M4 and the Ys are referred to as Y0, Y1, Y2, Y3, and Y4. After the measurement of M1 and Y1, the three follow-ups were each lagged by 6 months. The metrics of all the variables were such that their standard deviations approximated values of 1.0. The mediator was parameterized as a growth curve model as was the outcome, so the full model contains two growth curves, one for the mediator and one for the outcome. The hypothesis is that the growth curve parameters on Y are linked to the growth curve parameters on M. The relevant influence diagram is in Figure 1. Correlated disturbances between d5 and d6 and between d7 and d8 are omitted from the diagram to reduce clutter but they were included in the Mplus syntax for purposes of model testing. These correlations reflect the fact that the respective latent intercepts and slopes for M and Y respectively may be correlated over and above the determinants of them explicitly shown in the model, i.e., they deal with potential omitted variable confounds.

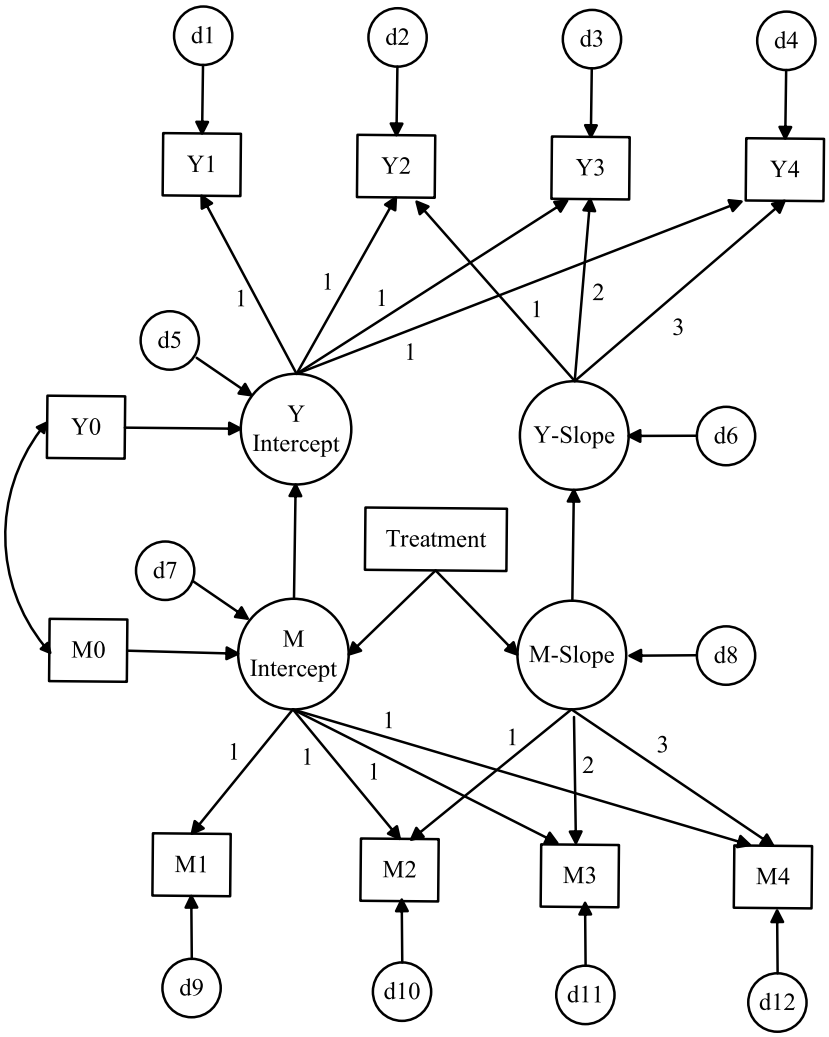

FIGURE 1. RET growth curve model

Table 1 presents the relevant Mplus syntax, all of which should be self-explanatory given you have read Chapter 16. The syntax is identical to the two-treatment LGA example (see Table 16.13 in the main text) with the only difference being the notation I use for the `MODEL CONSTRAINT` commands later. The data I am analyzing here is different from that in the main text and there are some nuanced differences in the parameters I examine relative to the two examples. All of the data are available for download on my website.

Table 1. Mplus Syntax for Traditional RET Example

```
1.  TITLE: Parallel Process Model for Traditional RET  ;
2.  DATA: FILE = LGAdata2M.dat;
3.  VARIABLE:
```

```
4.  DEFINE:
5.   CENTER m0 y0 (GRANDMEAN) ;
6.  NAMES = y1 y2 y3 y4 m1 m2 m3 m4 y0 m0 treat ;
7.  ANALYSIS: ESTIMATOR=MLR ;
8.  MODEL:
9.   !Define latent growth model for y
10.  yi ys | y1@0 y2@1 y3@2 y4@3 ;
11.  !Estimate latent intercepts and assign them labels
12.  [yi] (ya1);      !Intercept for intercept factor
13.  [ys] (ya2);      !Intercept for slope factor
14.  !Estimate variances of intercept and slope factors
15.  yi;       ! Variance of intercept factor
16.  ys;       ! Variance of slope factor
17.  !Estimate covariance between intercept and slope factors
18.  yi WITH ys;
20.  ! Estimate residual variances (error)
21.  y1-y4;
22.  !Define latent growth model for mediator
23.  mi ms | m1@0 m2@1 m3@2 m4@3 ;
24.  !Estimate intercepts and assign them labels
25.  [mi] (ma1) ;      !Intercept for intercept factor
26.  [ms] (ma2) ;      !Intercept for slope factor
27.  !Estimate variances of intercept and slope factors
28.  mi;       ! Variance of intercept factor
29.  ms;       ! Variance of slope factor
30.  !Estimate covariance between intercept and slope factors
31.  mi WITH ms;
32.  !Estimate residual variances (error)
33.  m1-m4;
34.  !Define the regressions and assign labels to coefficients
35.  mi ON treat m0 (p1 p2)  ;
36.  ms ON treat (p3);
37.  yi ON mi y0 (p4 p5);
38.  ys ON ms (p6);
39. MODEL INDIRECT:
40.  yi IND treat ;
41.  mi IND treat ;
42.  ys IND treat ;
43.  ms IND treat ;
44.  OUTPUT: Samp StdYX MOD(4) Residual Tech4  ;
```

When I execute the syntax in Table 1 (removing the line numbers I included for pedagogical purposes), the model fit indices suggested good model fit. The chi square test was statistically non-significant (chi square = 48.22 with df = 44, $p < 0.997$), which is consistent with reasonable data-model correspondence. The RMSEA was 0.010. The upper limit of the 90% confidence interval for it was 0.024. The p value for close fit was $p < 1.00$. The CFI was 1.00 and the standardized RMR was 0.019. For localized fit, there were no

theoretically meaningful modification indices greater than 4 and no meaningful standardized residuals for the predicted and observed covariances on a cell-by-cell basis.

After concluding for reasonable model fit, I re-ran the syntax but added the same `MODEL CONSTRAINT` commands as in the main text but with different notation; variable names that end in `t` refer to the intervention condition and that end in `c` refer to the control condition. I replaced the `OUTPUT` line 44 with the line to remove the request for modification indices, which Mplus does not permit when the `MODEL CONSTRAINT` command is used:

```
OUTPUT: Samp StdYX Residual Tech4  ;
```

Here are the commands I added just before the new Line 44:

```
43a1.  MODEL CONSTRAINT:
43a2.    NEW (mst msc msdiff m1t m2t m3t m4t m1c
43a3.    m2c m3c m4c mdiff1 mdiff2 mdiff3 mdiff4) ;
43a4.    mst = ma2 + p3 ; !Treatment mean of mediator slope factor
43a5.    msc = ma2 ;      !Control mean of mediator slope factor
43a6.    msdiff = mst-msc ;
43a7.    m1t = (ma1+p1) ;        !Treatment mediator value at time 1
43a8.    m2t = m1t + mst*1  ;    !Treatment mediator value at time 2
43a9.    m3t = m1t + mst*2  ;    !Treatment mediator value at time 3
43a10.   m4t = m1t + mst*3  ;    !Treatment mediator value at time 4
43a11.   m1c = ma1 ;             !Control mediator value at time 1
43a12.   m2c = m1c + msc*1 ;     !Control mediator value at time 2
43a13.   m3c = m1c + msc*2 ;     !Control mediator value at time 3
43a14.   m4c = m1c + msc*3 ;     !Control mediator value at time 4
43a15.   mdiff1 = m1t - m1c ;
43a16.   mdiff2 = m2t - m2c ;
43a17.   mdiff3 = m3t - m3c ;
43a18.   mdiff4 = m4t - m4c ;
43b1.    NEW (yst ysc ysdiff y1t y2t y3t y4t y1c
43b2.    y2c y3c y4c ydiff1 ydiff2 ydiff3 ydiff4) ;
43b3.    yst = ya2 + (ma2 + p3)*p6 ;    !Treatment mean of Y slope factor
43b4.    ysc = ya2 + ma2*p6 ;           !Control mean of Y slope factor
43b5.    ysdiff = yst-ysc ;
43b6.    y1t = ya1 + (ma1 + p1)*p4;     !Treatment Y value at time 1
43b7.    y2t = y1t + yst*1;             !Treatment Y value at time 2
43b8.    y3t = y1t + yst*2;             !Treatment Y value at time 3
43b9.    y4t = y1t + yst*3;             !Treatment Y value at time 4
43b10.   y1c = ya1 + ma1*p4;            !Control Y value at time 1
43b11.   y2c = y1c + ysc*1;             !Control Y value at time 2
43b12.   y3c = y1c + ysc*2;             !Control Y value at time 3
43b13.   y4c = y1c + ysc*3;             !Control Y value at time 4
43b14.   ydiff1 = y1t - y1c ;
43b15.   ydiff2 = y2t - y2c ;
43b16.   ydiff3 = y3t - y3c ;
43b17.   ydiff4 = y4t - y4c ;
```

I explained the logic of the lines in the main text for the two treatment comparison RET example. The logic is the same here; only the labels are different. I now consider the core questions of the RET.

Total Effect of the Treatment on the Outcome

To test the effects of the intervention on the outcome, I use the output generated by Lines 40 (`YI IND TREAT`) and 42 (`YS IND TREAT`) from Table 1. The former focuses on the treatment-control mean difference at the immediate posttest and the latter focuses on the treatment-control difference in decay in Y that occurs across the follow-ups. Here is the output for Line 40 that focuses on the difference between the treatment and control groups at the immediate posttest:

```
TOTAL, TOTAL INDIRECT, SPECIFIC INDIRECT, AND DIRECT EFFECTS

                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

Effects from TREAT to YI

  Total               0.532      0.048     11.153      0.000
  Total indirect      0.532      0.048     11.153      0.000

  Specific indirect 1
    YI
    MI
    TREAT             0.532      0.048     11.153      0.000
```

The entry in the first column underneath `Specific indirect 1` identifies the causal chain(s) through which the treatment condition reaches the latent Y intercept. It is `TREAT→MI→YI`. The row of primary interest is the one called `Total` underneath the `Effects from TREAT to YI` label. The estimate in this row is the predicted mean difference between Y at the immediate posttest for the treatment condition minus the control condition. In this case, the adjusted difference is 0.532 ±0.096, which is statistically significant (CR = 11.15, $p < 0.05$). The 95% confidence interval for the difference is 0.43 to 0.63. Suppose that prior to the study, the research team set a standard for a meaningful population mean difference at the immediate posttest as an absolute difference of 0.20 or greater. Because the lower limit of the mean difference confidence interval was greater than this standard, I conclude that the intervention effect for the immediate posttest is meaningful.

If desired, I also can report the predicted means of Y at the immediate posttest for the treatment and control conditions to embellish the above results. These predicted values are labeled `Y1T` and `Y1C` in the `MODEL CONSTRAINT` commands. Here is the edited output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    Y1T                0.557      0.042     13.126      0.000
    Y1C                0.026      0.039      0.651      0.515
```

The estimated Y means for the intervention and control conditions at the immediate posttest were 0.557 ±0.08 and 0.026 ±0.08, respectively.

Next, I examined a second facet of the effect of the intervention on the outcome, namely whether there was meaningful decay in the effect of the intervention over time. As a first step, I isolate the coefficient characterizing the decay curve for the intervention group and also that for the control group. These values are available in the `MODEL CONSTRAINT` commands in the variables labeled `YSTA` and `YSTB`. Here is the edited output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    YST               -0.138      0.020     -6.800      0.000
    YSC                0.048      0.020      2.355      0.019
```

The model implied linear decay in the intervention condition is negative and statistically significant (slope = -0.138 ±0.040, CR = 6.80, $p < 0.05$) indicating that the predicted mean of Y decreases as time from the immediate posttest increases across the 18 month follow-up period. For the control condition, the decay in the predicted mean of Y is relatively flat (because the slope is near a value of zero) and, if anything, the predicted Y mean slightly increases over time compared to its value at the immediate posttest (slope = 0.048 ±0.040, CR = 2.36, $p < 0.05$). The slight positive increase is statistically significant.

I can plot the two estimated decay curves using the information in the output from the `MODEL CONSTRAINT` commands for the predicted means at Y1 though Y4 for the intervention group (labeled `Y1T`, `Y2T`, `Y3T`, and `Y4T`) and for the control group (labeled `Y1C`, `Y2C`, `Y3C`, and `Y4C`). Using the program on my website called *Temporal line plot*, here is the plot:

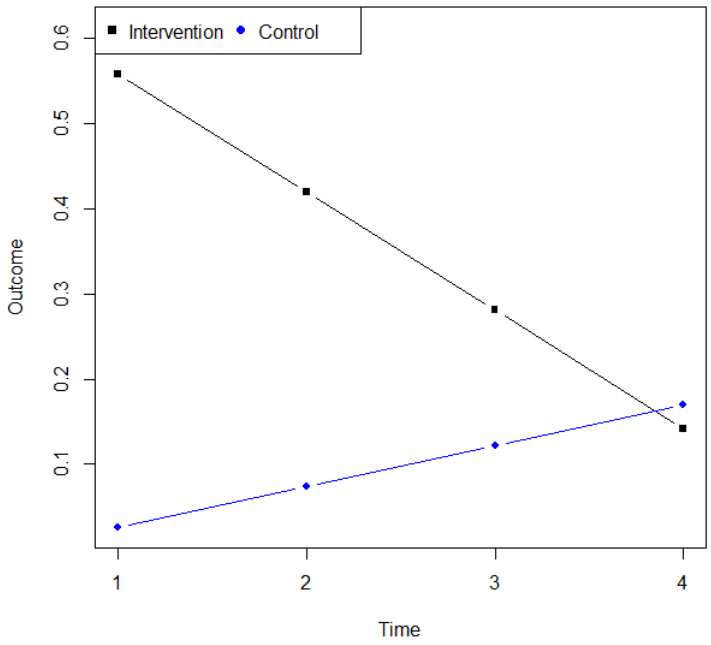

The downward decay in the predicted Y means for the intervention group is evident as is the slight upward trend in the means for the control group. As discussed in the main text, the separation between lines at a given time point reflects the predicted mean difference between the intervention and control groups. It can be seen that the mean difference between the treatment and control conditions decreases as time passes to the point that by month 18 the means are functionally the same.

There are several important lessons in the above results. To make statements about decay of the intervention effect, it is important to define the decay focusing on the *difference* in means between the intervention group and control group. Doing so maintains the spirit of a randomized trial. If you only examine decay in the intervention group *per se*, you are failing to take into account potential across-time artifacts and across-time confounds as discussed in Chapter 4. Stated another way, factors can operate that cause the outcome means in the control group to systematically increase or decrease when people and contexts are left to their own devices. This is the case in the current example. An example with real world data is research with adolescents on risk behaviors such as drinking and smoking. These behaviors often increase as adolescents age, showing the same upward trend evident in the control condition of the current example. This trend needs to be taken into account when evaluating intervention effects over time.

The contrasts that I specified in the `MODEL CONSTRAINT` commands test the outcome predicted mean difference at each of the four time points using the model-informed predicted means. The contrasts are labeled `YDIFF1`, `YDIFF2`, `YDIFF3`, `YDIFF4` for times 1, 2,

3 and 4, respectively (keep in mind that time 1 is the immediate posttest). Here is the relevant output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    YDIFF1             0.532      0.048     11.153      0.000
    YDIFF2             0.345      0.049      7.060      0.000
    YDIFF3             0.159      0.062      2.586      0.010
    YDIFF4            -0.027      0.080     -0.342      0.733
```

The intervention group predicted mean is statistically significantly larger than the control group predicted mean at each of the time points except the last one. Suppose a meaningfulness standard is set by the research team to be a population absolute difference of 0.20 or greater. The 95% confidence interval for `YDIFF1` is 0.44 to 0.638 and for `YDIFF2` it is 0.257 to 0.44 (these intervals can be approximated by adding and subtracting double the estimated standard error for the estimated mean difference; they are formally reported on the Mplus output in a special section for confidence intervals). In both cases, the lower limit of the confidence interval exceeds the meaningfulness standard; I can conclude the differences at the immediate posttest and the 6 month follow-up are meaningful.

For the third time point, the 95% confidence interval is 0.035 to 0.283. The confidence interval overlaps the meaningfulness standard. Given this, I can't confidently conclude the predicted mean difference is meaningful given sampling error. To be sure, the difference is non-zero by virtue of the statistical significance ($p < 0.05$) of the contrast. However, I can't say with confidence the predicted mean difference at time 3 is meaningful.

Finally, the 95% confidence interval for the predicted mean difference at the fourth time point is -0.187 to 0.133. This interval is outside the meaningfulness standard and is statistically non-significant. I can't conclude the effect is non-zero nor that it is meaningful.

Note that these mean difference tests are not tests of the observed mean differences but rather focus on the predicted means by the model. Such is the nature of full information SEM (FISEM); the SEM model is assumed to capture the underlying data generating process, which is why we test for and insist on good fit between the model and the data. The derived coefficients and parameters are model-driven and our conclusions assume the model captures the underlying dynamics. As long as there is close fit between the data and the model, the tests of mean differences from a model-informed perspective is reasonable.

In sum, the intervention had a meaningful overall effect on the outcome at the immediate posttest. The predicted mean for Y for the intervention group decreased over time while for the control group, it showed a modest increase. In the main text, when we compared two treatments to one another, we preferred the treatment with the flatter decay

curve. In a traditional RET with an intervention and control group, it is not uncommon for the decay curve in the control group to be flat because there is no change at the immediate posttest in the control group to "decay." It is not enough to simply examine slope differences between groups to make substantive conclusions. Rather, one must take into account the experimental context to know how to interpret differences. In the current study, the intervention group maintained a meaningful predicted mean difference relative to the control group through follow-up month 6, but for follow-up month 12, although the difference was non-zero, it could not unambiguously be said to be meaningful when sampling error is taken into account. By month 18, the predicted mean Y difference between the intervention and control groups was clearly trivial.

Effect of the Treatment on the Mediator

The analysis of the effect of the intervention on the mediator follows the same structure as that for the total effect on Y. I focus on the output generated by Lines 41 (`MI IND TREAT`) and 43 (`MS IND TREAT`) from Table 1. Here is the output for Line 41 of Table 1 for the mediator difference between the treatment and control groups at the immediate posttest:

```
TOTAL, TOTAL INDIRECT, SPECIFIC INDIRECT, AND DIRECT EFFECTS

                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

Effects from TREAT to MI

  Total                0.728      0.054     13.399      0.000
```

The adjusted mean difference is 0.728 ±0.108, which is statistically significant (CR = 13.40, $p < 0.05$). Suppose that prior to the study, the research team set a standard for a meaningful population mean difference in M at the immediate posttest as an absolute difference of 0.20 or greater. The 95% CI for the treatment difference was 0.620 to 0.836. The lower limit of this interval exceeds the meaningfulness standard. I conclude that the predicted mean difference between the intervention and control groups is meaningful.

If desired, I also can report the predicted means of the mediator at the immediate posttest for the treatment and control conditions to embellish the above results. These values are labeled `M1T` and `M1C` in the output section of the `MODEL CONSTRAINT` commands. The estimated mediator means for the treatment and control conditions at the immediate posttest were 0.759 ±0.082 and 0.031 ±0.076, respectively.

Next, I examined a second facet of the effect of the intervention on the mediator, namely whether there was meaningful decay in the effect over time. As a first step, I isolate the mediator decay curve for the intervention group and also for the control group so that I

can compare them in ways that are of substantive interest. These values are available in the MODEL CONSTRAINT commands in the variables labeled MST and MSC. Here is the output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    MST               -0.222      0.020    -11.003      0.000
    MSC                0.019      0.019      1.012      0.311
```

The model implied linear decay in the intervention condition is negative and statistically significant (slope = -0.222 ± 0.040, CR = 11.00, $p < 0.05$) indicating that the predicted mean of the mediator decreases as time from the immediate posttest increases across the 18 months of the study. For the control condition, the decay in the predicted mean of the mediator is relatively flat because the slope is near a value of zero.

I can plot the two curves using the information in the output section from the MODEL CONSTRAINT commands using the predicted means for the intervention group (labeled M1T, M2T, M3T, and M4T) and for the control group (labeled M1C, M2C, M3C, and M4C). Using the program on my website called *Temporal line plot*, here is the plot:

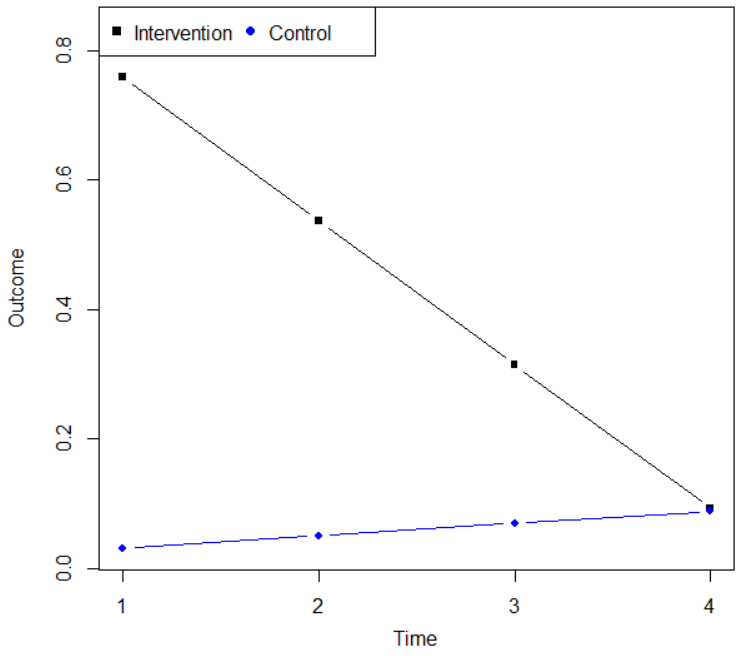

The downward decay in the predicted M means for the intervention group is evident as is the relatively flat slope for the predicted means for the control group. It can be seen

that the mean difference between the treatment and control conditions decreases as time passes to the point that by month 18, the means are functionally the same.

The contrasts I specified in the `MODEL CONSTRAINT` commands also test the mediator mean difference at each of the four time points using the model-informed predicted means. The contrasts are labeled `MDIFF1`, `MDIFF2`, `MDIFF3`, `MDIFF4` for times 1, 2, 3 and 4, respectively (keep in mind that time 1 is the immediate posttest). Here is the relevant output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    MDIFF1             0.728      0.054     13.399      0.000
    MDIFF2             0.487      0.053      9.231      0.000
    MDIFF3             0.245      0.064      3.854      0.000
    MDIFF4             0.004      0.082      0.048      0.962
```

The intervention group predicted mean is statistically significantly larger than the control group predicted mean at each of the time points except the last one. Suppose a meaningfulness standard is set by the research team to be an absolute difference of 0.20 or greater. The 95% confidence interval for `MDIFF1` is 0.62 to 0.83 and for `MDIFF2` it is 0.38 to 0.59. In both cases, the lower limit of the confidence interval exceeds the meaningfulness standard. Given this, I conclude that the differences at the immediate posttest and the 6 month follow-up are meaningful. For the third time point, the 95% confidence interval is 0.12 to 0.37. The confidence interval overlaps the meaningfulness standard. Given this, I can't confidently conclude the predicted mean difference is meaningful when sampling error is taken into account. To be sure, the difference is non-zero by virtue of its statistical significance ($p < 0.05$). However, I can't say with confidence that the predicted mediator mean difference at time 3 is meaningful. Finally, the 95% confidence interval for the predicted mean difference at the fourth time point is -0.16 to 0.17. This confidence interval is outside the meaningfulness standard and it also is statistically non-significant. I can't conclude the effect is non-zero nor that it is meaningful.

In sum, the intervention had a meaningful overall effect on the mediator at the immediate posttest. The predicted mean for M for the intervention group decreased over time while for the control group it was relatively constant. The intervention group maintained a meaningful predicted difference relative to the control group through follow-up month 6, but for follow-up month 12, although the difference was non-zero, it could not unambiguously be said to meaningful when sampling error is taken into account. By month 18, the predicted mean M difference between the intervention and control groups was trivial.

Effect of the Mediator on the Outcome

The analysis of the estimated effect of the mediator on the outcome does not distinguish the intervention and control groups because the guiding model in Figure 1 assumes that the effect is the same in both treatment conditions. The relevant coefficients for evaluating the mediator effect on the outcome are the coefficient from MI to YI and the coefficient from MS to YS. Here are the results from the output:

```
MODEL RESULTS

                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

YI       ON
    MI             0.731      0.041     17.627      0.000
    Y0             0.118      0.028      4.161      0.000

 YS       ON
   MS              0.772      0.057     13.478      0.000
```

The coefficient for MI→YI is the path coefficient for the outcome at the immediate posttest regressed onto the mediator at the immediate posttest controlling for the baseline outcome. The coefficient was 0.73 ±0.08 (CR =17.63, $p < 0.05$). For every one unit that the mediator increased across individuals, the outcome was predicted to increase by 0.73 units. Suppose the research team decided prior to the study that a population coefficient of 0.20 or greater would be deemed meaningful. The 95% confidence interval for the coefficient was 0.66 to 0.89. Because the lower limit of this interval is greater than meaningfulness standard, we can declare the result meaningful.

The coefficient for MS→YS is the path coefficient for the decay slope of the outcome regressed onto the decay slope for the mediator. A positive coefficient implies that people with increasingly more positive slopes on the mediator tend to have increasingly more positive slopes on the outcome, i.e., that there is an association between the two. The path coefficient was 0.77 ± 0.11 (CR = 13.48, $p < 0.05$). For every one unit that the mediator slope increased across individuals, the outcome slope was predicted to increase by 0.77 units. Suppose the research team decided prior to the study that a population coefficient of 0.20 or greater would be deemed meaningful. The 95% confidence interval for the coefficient was 0.66 to 0.88. Because the lower limit of this interval is greater than the meaningfulness standard, we can declare the result meaningful.

In the analyses, I did not include a direct effect of the treatment condition on YI or YS over and above the effect of the treatment condition on MI and MS. This presumes that such

a direct effect is functionally zero. If desired, a researcher can include such effects and many methodologists routinely do so. In the current case, the research team saw no conceptual reason to include them. When I did so, both the estimated direct effect of the treatment on YI (coefficient = 0.03 ± 0.15, CR = 0.41, *ns*) and on YS (coefficient = -0.06 ± 0.06, CR = 0.14, *ns*) were trivial. The chi square test of fit for the model that included them was 45.10, df = 42, p < 0.35). A nested chi square difference test with the model that excluded the two effects yielded a statistically nonsignificant chi square difference of 3.12, df = 2, *ns*).

Omnibus Mediation Analysis

Omnibus mediation tests are of lower priority for me. I instead prefer to focus on the individual links in mediational chains, evaluating their effects sizes. If I want an overall omnibus test of statistical significance, I rely on the joint significance test. In LGC models, evaluations of mediation take two forms. The first focuses on treatment differences in the latent intercept for the mediator and whether these differences, in turn, translate into treatment differences on the outcome latent intercept. The second mediation form focuses on the effect of the latent slopes of the mediator on the latent slopes of the outcome. The RET found meaningful treatment versus control differences for both facets of mediation analysis. The pattern of results is consistent with omnibus mediation.

THREE TREATMENT GROUP RET EXAMPLE

This example is the same as the RET latent growth curve example described in the main text but it includes a control group. The treatment condition variable has three levels, Treatment A, Treatment B, and a control group. The data set I use for this example is different from that in the main text, but the essential structure of the RET is the same, just with three treatment groups, not two. I repeat the influence diagram from the main text in Figure 2 as well as the description of the RET.

The treatment variable is represented by two dummy variables. The first dummy variable assigns a 1 to participants exposed to Treatment A and zeros to everyone else. The second dummy variable assigns a 1 to participants exposed to Treatment B and zeros to everyone else. The use of these two dummy variables treats the control group as the reference group. Coefficients associated with the first dummy variable contrast those in Treatment A with those in the control group. Coefficients associated with the second dummy variable contrast those in Treatment B with those in the control group. Ultimately, I conduct all pairwise comparisons between the three groups.

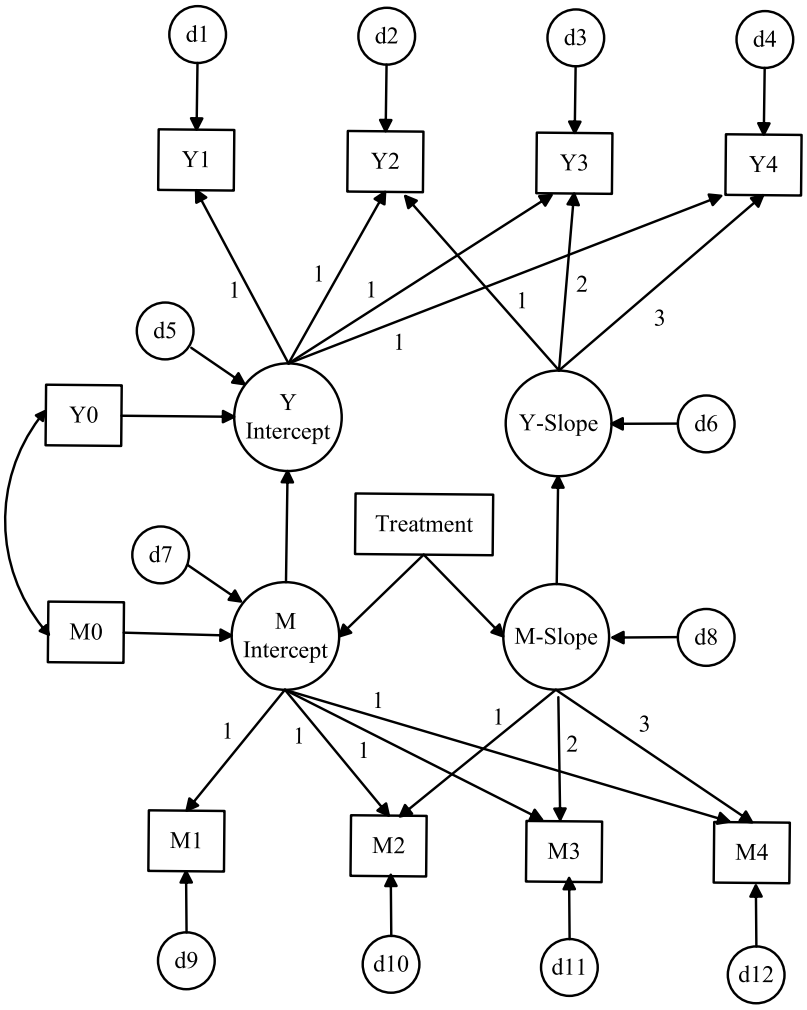

FIGURE 2. RET growth curve model with three treatment conditions

There is a single mediator and I use a parallel process LGC model per the main text. Higher scores on the mediator, M, and higher scores on the outcome, Y, are desirable. The intervention sought to increase them both. The mediator and the outcome were each measured at five time points. The baseline measures, M0 and Y0, were obtained prior to the start of treatment with M1 and Y1 representing scores at the immediate posttest. The five assessments of the mediators are referred to as M0 M1, M2, M3, and M4 and the Ys are referred to as Y0, Y1, Y2, Y3, and Y4. After the measurement of M1 and Y1, the three follow-ups were each lagged by 6 months. The metrics of all the variables were such that their standard deviations approximated values of 1.0. The mediator was parameterized as a growth curve model as was the outcome, so the full model contains two growth curves, one for the mediator and one for the outcome. The hypothesis is that the growth curve

parameters on Y are linked to the growth curve parameters on M. For Figure 2, correlated disturbances between d5 and d6 and between d7 and d8 are omitted from the diagram to reduce clutter but they were included in the Mplus syntax for purposes of testing the model. These correlations reflect the fact that the respective latent intercepts and slopes for M and Y respectively may be correlated over and above the determinants of them explicitly shown in the model.

Table 2 presents the relevant Mplus syntax, all of which should be self-explanatory. I highlight in red the main differences between this syntax and the syntax from the main text LGA (see Table 16.13 in the main text). You will see more substantial changes via the `MODEL CONSTRAINT` commands later.

Table 2. Mplus Syntax for RET Example with Three Treatment Conditions

```
1.  TITLE: Parallel Process Model for Traditional RET  ;
2.  DATA: FILE = LGAdata2M.dat;
3.  VARIABLE:
4.  DEFINE:
5.   CENTER m0 y0 (GRANDMEAN) ;
6.  NAMES = y1 y2 y3 y4 m1 m2 m3 m4 y0 m0 treat dtreata dtreatb dcontrol ;
7.  ANALYSIS: ESTIMATOR=MLR ;
8.  MODEL:
9.   !Define latent growth model for y
10.  yi ys | y1@0 y2@1 y3@2 y4@3 ;
11.  !Estimate latent intercepts and assign them labels
12.  [yi] (ya1);      !Intercept for intercept factor
13.  [ys] (ya2);      !Intercept for slope factor
14.  !Estimate variances of intercept and slope factors
15.  yi;       ! Variance of intercept factor
16.  ys;       ! Variance of slope factor
17.  !Estimate covariance between intercept and slope factors
18.  yi WITH ys;
20.  ! Estimate residual variances (error)
21.  y1-y4;
22.  !Define latent growth model for mediator
23.  mi ms | m1@0 m2@1 m3@2 m4@3 ;
24.  !Estimate intercepts and assign them labels
25.  [mi] (ma1) ;      !Intercept for intercept factor
26.  [ms] (ma2) ;      !Intercept for slope factor
27.  !Estimate variances of intercept and slope factors
28.  mi;       ! Variance of intercept factor
29.  ms;       ! Variance of slope factor
30.  !Estimate covariance between intercept and slope factors
31.  mi WITH ms;
32.  !Estimate residual variances (error)
33.  m1-m4;
34.  !Define the regressions and assign labels to coefficients
```

```
35.  mi ON dtreata dtreatb m0 (p1 p2 p3)  ;
36.  ms ON dtreata dtreatb (p4 p5);
37.  yi ON mi y0 (p6 p7);
38.  ys ON ms (p8);
39. MODEL INDIRECT:
40.  yi IND dtreata ;
41.  yi IND dtreatb ;
42.  ys IND dtreata ;
43.  ys IND dtreatb ;
44.  mi IND dtreata ;
45.  mi IND dtreatb ;
46.  ms IND dtreata ;
47.  ms IND dtreatb ;
48.  OUTPUT: Samp StdYX MOD(4) Residual Tech4  ;
```

When I executed the syntax in Table 2, the model fit indices suggested good model fit. The chi square test of perfect model fit in the population was statistically non-significant (chi square = 66.13 with 50 degrees of freedom, $p < 0.07$), which is consistent with the data being in accord with the model. The RMSEA was 0.015. The upper limit of the 90% confidence interval for it is 0.063. The p value for close fit was not statistically significant ($p < 1.00$). The CFI is 0.99 and the standardized RMR was 0.017. For localized fit, there were no theoretically meaningful modification indices greater than 4 and no meaningful residuals between the predicted and observed covariances on a cell-by-cell basis.

After concluding in favor of a reasonable model fit, I re-ran the syntax but added the same `MODEL CONSTRAINT` commands from the example in the main text but with different notation; variable names that end in `t` refer to the intervention condition and those that end in `c` refer to the control condition. I replaced the `OUTPUT` line 48 with the line:

```
OUTPUT: Samp StdYX Residual Tech4  ;
```

Here are the commands I added just before the new Line 48 (remember. Exclude the line numbers):

```
47a1.  MODEL CONSTRAINT:
47a2.    NEW (msa msb msc msdiffac msdiffbc msdiffab m1a m2a m3a m4a
47a3.    m1b m2b m3b m4b m1c m2c m3c m4c mdiff1ac mdiff1bc mdiff1ab
47a4.    mdiff2ac mdiff2bc mdiff2ab mdiff3ac mdiff3bc mdiff3ab
47a5.    mdiff4ac mdiff4bc mdiff4ab) ;
47a6.    msa = ma2 + p4 ;      !Treatment A mean of mediator slope factor
47a7.    msb = ma2 + p5 ;      !Treatment B mean of mediator slope factor
47a8.    msc = ma2 ;           !Control A mean of mediator slope factor
47a9.    msdiffac = msa-msc ;
47a10.   msdiffbc = msb-msc ;
47a11.   msdiffab = msa-msb ;
```

```
47a12.   m1a = (ma1+p1) ;      !Treatment A mediator value at time 1
47a13.   m2a = m1a+msa*1 ;     !Treatment A mediator value at time 2
47a14.   m3a = m1a+msa*2 ;     !Treatment A mediator value at time 3
47a15.   m4a = m1a+msa*3 ;     !Treatment A mediator value at time 4
47a16.   m1b = (ma1+p2) ;      !Treatment B mediator value at time 1
47a17.   m2b = m1b+msb*1 ;     !Treatment B mediator value at time 2
47a18.   m3b = m1b+msb*2 ;     !Treatment B mediator value at time 3
47a19.   m4b = m1b+msb*3 ;     !Treatment B mediator value at time 4
47a20.   m1c = ma1 ;           !Control mediator value at time 1
47a21.   m2c = m1c+msc*1 ;      !Control mediator value at time 2
47a22.   m3c = m1c+msc*2 ;      !Control mediator value at time 3
47a23.   m4c = m1c+msc*3 ;      !Control mediator value at time 4
47a24.   mdiff1ac = m1a - m1c ;
47a25.   mdiff1bc = m1b - m1c ;
47a26.   mdiff1ab = m1a - m1b ;
47a27.   mdiff2ac = m2a - m2c ;
47a28.   mdiff2bc = m2b - m2c ;
47a29.   mdiff2ab = m2a - m2b ;
47a30.   mdiff3ac = m3a - m3c ;
47a31.   mdiff3bc = m3b - m3c ;
47a32.   mdiff3ab = m3a - m3b ;
47a33.   mdiff4ac = m4a - m4c ;
47a34.   mdiff4bc = m4b - m4c ;
47a35.   mdiff4ab = m4a - m4b ;
47b1.    NEW (ysa ysb ysc ysdiffac ysdiffbc ysdiffab y1a y2a y3a y4a
47b2.    y1b y2b y3b y4b y1c y2c y3c y4c ydiff1ac ydiff1bc ydiff1ab
47b3.    ydiff2ac ydiff2bc ydiff2ab ydiff3ac ydiff3bc ydiff3ab
47b4.    ydiff4ac ydiff4bc ydiff4ab) ;
47b5.    ysa = ya2 + (ma2 + p4)*p8 ;     !Treatment A mean of y slope factor
47b6.    ysb = ya2 + (ma2 + p5)*p8 ;     !Treatment B mean of y slope factor
47b7.    ysc = ya2 + ma2*p8 ;            !Control mean of y slope factor
47b8.    ysdiffac = ysa-ysc ;
47b0.    ysdiffbc = ysb-ysc ;
47b10    ysdiffab = ysa-ysb ;
47b11.   y1a = ya1 + (ma1 + p1)*p6;      !Treatment A y value at time 1
47b12.   y2a = y1a + ysa*1;              !Treatment A y value at time 2
47b13.   y3a = y1a + ysa*2;              !Treatment A y value at time 3
47b14.   y4a = y1a + ysa*3;              !Treatment A y value at time 4
47b15.   y1b = ya1 + (ma1 + p2)*p6;      !Treatment B y value at time 1
47b16.   y2b = y1b + ysb*1;              !Treatment B y value at time 2
47b17.   y3b = y1b + ysb*2;              !Treatment B y value at time 3
47b18.   y4b = y1b + ysb*3;              !Treatment B y value at time 4
47b19.   y1c = ya1 + ma1*p6;             !Control y value at time 1
47b20.   y2c = y1c + ysc*1;              !Control y value at time 2
47b21.   y3c = y1c + ysc*2;              !Control y value at time 3
47b22.   y4c = y1c + ysc*3;              !Control y value at time 4
47b23.   ydiff1ac = y1a - y1c ;
47b24.   ydiff1bc = y1b - y1c ;
47b25.   ydiff1ab = y1a - y1b ;
```

```
47b26.   ydiff2ac = y2a - y2c ;
47b27.   ydiff2bc = y2b - y2c ;
47b28.   ydiff2ab = y2a - y2b ;
47b29.   ydiff3ac = y3a - y3c ;
47b30.   ydiff3bc = y3b - y3c ;
47b31.   ydiff3ab = y3a - y3b ;
47b32.   ydiff4ac = y4a - y4c ;
47b33.   ydiff4bc = y4b - y4c ;
47b34.   ydiff4ab = y4a - y4b ;
```

All of the syntax should be self-explanatory given your exposure to material in the main and the prior example. I now consider the core questions of the RET but adapted to the comparison of the three treatment condition instead of the more traditional treatment versus control conditions.

Effects of the Treatments on the Outcome

To evaluate the effect of Treatment A on the outcome, I used the output generated by Line 40 (`YI IND TREATA`) from Table 2. The output focuses on the mean difference at the immediate posttest for Treatment A minus the control group mean. Here are the results:

```
TOTAL, TOTAL INDIRECT, SPECIFIC INDIRECT, AND DIRECT EFFECTS

                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

Effects from DTREATA to YI

  Total                0.561      0.046     12.132      0.000
```

The row of interest is the one called `Total` underneath the `Effects from DTREATA to YI` label. The estimate in this row is the predicted difference between the mean Y at the immediate posttest for Treatment A minus the control group. In this case, the adjusted difference is 0.561 ±0.092, which is statistically significant (CR = 12.13, $p < 0.05$). Suppose that prior to the study, the research team set a standard for a meaningful population mean difference at the immediate posttest as an absolute population difference of 0.20 or greater. The 95% confidence interval for the treatment difference was 0.47 to 0.65. Because the lower limit for this interval exceeds the meaningfulness standard, I can conclude that Treatment A had a meaningful effect on the immediate posttest relative to the control group.

I also can obtain the predicted means of Y at the immediate posttest for Treatment A and for the control group. These values are labeled `Y1A` and `Y1C` within the `MODEL CONSTRAINT` command. Here is the edited output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.    Est./S.E.    P-Value

Y1A                    0.604      0.038     15.760      0.000
Y1C                    0.043      0.037      1.158      0.247
```

Next, I evaluated the effect of Treatment B on the outcome. I used the output generated by Line 41 (`YI IND DTREATB`) from Table 2. The output focuses on the mean difference at the immediate posttest between Treatment B minus the control group. Here are the results:

```
TOTAL, TOTAL INDIRECT, SPECIFIC INDIRECT, AND DIRECT EFFECTS

                                                  Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

Effects from DTREATB to YI

  Total                0.503      0.048     10.530      0.000
```

The adjusted difference is 0.503 ±0.096, which is statistically significant (CR = 10.53, $p < 0.05$). The 95% confidence interval for the treatment difference was 0.41 to 0.60. Because the lower limit for this interval exceeds the meaningfulness standard of 0.20, I can confidently conclude that Treatment B had a meaningful effect on the immediate posttest relative to the control group. The predicted mean of Y at the immediate posttest for Treatment B (taken from `Y1B` in the `MODEL CONSTRAINT` command) was 0.545 ±0.078. The predicted mean for the control group as reported earlier was 0.043 ±0.037.

Finally, I examined the relative difference between Treatment A and Treatment B on the adjusted outcome as taken from the `MODEL CONSTRAINT` commands. Here is the output, where `Y1A` is the adjusted outcome mean for Treatment A and `Y1B` is the adjusted outcome mean for Treatment B:

```
                                                  Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    Y1A                0.604      0.038     15.760      0.000
    Y1B                0.545      0.039     13.818      0.000
    YDIFF1AB           0.058      0.045      1.302      0.193
```

The adjusted mean difference between the two treatments was 0.058 ±0.09, which was statistically non-significant (CR = 1.30, *ns*). To assert treatment equivalence between Treatment A and Treatment B, I need to take into account the latitude of no effect per my

discussion in Chapter 10. Suppose the research team deemed that any population absolute difference less than 0.15 would be deemed to be *functionally* zero. Stated another way, if the population mean difference falls anywhere between -0.15 and +0.15, the effect is deemed trivial. The 95% confidence interval for the mean difference was -.04 to 0.154. Because the interval is not fully contained within the latitude of no effect, I cannot confidently conclude that the two treatments are functionally equivalent taking into account sampling error (because the upper limit 0.154 of the confidence interval exceeds the no effect standard of 0.15). This test is a form of **equivalence testing** or **non-inferiority testing** in randomized trials. Conclusion-wise, I am in a state of suspended judgment. I can't confidently conclude that the treatments are meaningfully different in terms of influencing the outcome but I also can't confidently conclude their difference is functionally zero. The confidence interval is a bit too wide to make a firm conclusion. Methodologically I need to reduce its width, perhaps by increasing the sample size.

Turning to the decay curves for the outcome, I gain a sense of them for each of the three groups by plotting the estimated outcome means as taken from the `MODEL CONSTRAINT` commands (variables `Y1A, Y2A, Y3A, Y4A, Y1B, Y2B, Y3B, Y4B, Y1C, Y2C, Y3C, Y4C`). Here is the plot that I constructed using the *Temporal line plot* program from my website:

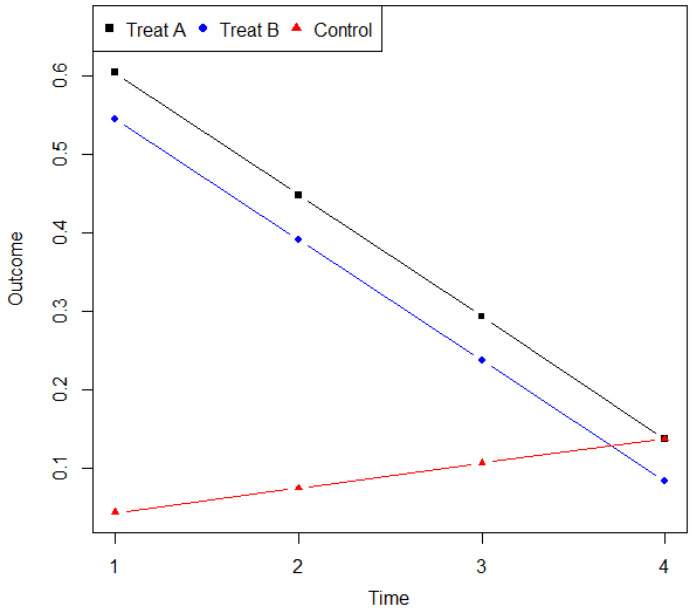

There is a steady, parallel decline in the mean outcome for Treatment A and for Treatment B over time. The control group shows a slight increase in the mean Y across

time, although it is not statistically significant (see below). Here are the estimated slopes for the three groups taken from the MODEL CONSTRAINT commands:

```
MODEL RESULTS

                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

New/Additional Parameters

    YSA               -0.156      0.018     -8.419      0.000
    YSB               -0.154      0.019     -8.273      0.000
    YSC                0.031      0.020      1.589      0.112
```

,

The estimated slopes for Treatment A and Treatment B were negative and statistically significant (Treatment A = -0.156 ±0.036, CR = 8.41, $p < 0.05$; Treatment B = -0.154± 0.038, CR = 8.27, $p < 0.05$). The estimated slope for the control group was 0.031 ±0.040, (CR = 1.59, *ns*). There is little change in Y across time for the control group. By contrast, the two treatment conditions showed nontrivial decreases in Y across time.

The MODEL CONSTRAINT commands also provide pairwise contrasts between the three slopes; at the end of DIFF, the letters A = Treatment A, B = Treatment B and C = control:

```
                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

    YSDIFFAC          -0.187      0.024     -7.715      0.000
    YSDIFFBC          -0.186      0.025     -7.516      0.000
    YSDIFFAB          -0.001      0.021     -0.066      0.947
```

The estimated slope difference between Treatment A and the control group was -0.187 ±0.048 (CR = 7.72, $p < 0.05$). The estimated slope difference between Treatment B and the control group was -0.186 ±0.050 (CR = 7.52, $p < 0.05$). The estimated slope difference between the two treatments was -0.001 ±0.042 (CR = 0.07, *ns*). This difference was functionally zero.

A final set of contrasts that is informative are again from the MODEL CONSTRAINT commands. They examine pairwise mean differences between the three groups at each time point, i.e., the immediate posttest, the 6 month follow-up, the 12 month follow-up and the 18 month follow-up. In the output below the notation after the root term DIFF lists the timepoint (1, 2, 3 or 4) followed by letters representing the two groups that are contrasted where A = Treatment A, B = Treatment B, and C = control.

	Estimate	S.E.	Est./S.E.	Two-Tailed P-Value
YDIFF1AC	0.561	0.046	12.132	0.000
YDIFF1BC	0.503	0.048	10.530	0.000
YDIFF1AB	0.058	0.045	1.302	0.193
YDIFF2AC	0.374	0.047	7.994	0.000
YDIFF2BC	0.317	0.047	6.738	0.000
YDIFF2AB	0.057	0.043	1.334	0.182
YDIFF3AC	0.187	0.058	3.197	0.001
YDIFF3BC	0.131	0.058	2.264	0.024
YDIFF3AB	0.055	0.050	1.098	0.272
YDIFF4AC	0.000	0.076	-0.004	0.997
YDIFF4BC	-0.054	0.076	-0.716	0.474
YDIFF4AB	0.054	0.065	0.833	0.405

None of the contrasts comparing Treatment A to Treatment B at any of the time points were statistically significant. There were no meaningful mean differences between Treatment A and Treatment B. At the immediate posttest, both Treatment A and Treatment B are statistically significantly and meaningfully different from the control group. However, by the 18 month follow-up, these differences have dissipated to the point that all three group contrasts are statistically non-significant. At the 6 month follow-up, both Treatment A and Treatment B differences from the control group are diminished somewhat but still statistically significant and meaningful. At the 12 month follow-up, Treatment A and Treatment B remain statistically significantly different from the control group, but their 95% confidence intervals for the differences overlap with the meaningfulness standard indicating I cannot conclude the separation of the two treatments from the control group is meaningful.

In sum, when evaluating the effect of the intervention on the outcome, there are multiple facets to examine. The overall pattern of results suggest that the two treatment conditions produce comparable and meaningful effects on the outcome at the immediate posttest but that this effect diminishes across time to the point that by the 18 month follow-up, the treatments have lost their effects. On the positive side of things, the treatments still have meaningful (but diminished effects) on the outcome after 6 months.

Effects of the Treatments on the Mediator

The analysis of the effects of the treatments on the mediator follows the same structure as that for the intervention effect on Y. To test the effect of Treatment A on the outcome, I used the output generated by Line 44 (`MI IND DTREATA`) from Table 2. The output focuses

on the treatment mean difference at the immediate posttest between Treatment A and the control group. Here are the results:

```
TOTAL, TOTAL INDIRECT, SPECIFIC INDIRECT, AND DIRECT EFFECTS

                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

Effects from DTREATA to MI

  Total                0.722      0.054     13.355      0.000
```

The estimate is the predicted difference between the mediator mean at the immediate posttest for Treatment A minus the control group. In this case, the adjusted difference is 0.722 ±0.108, which is statistically significant (CR = 13.36, $p < 0.05$). Suppose that prior to the study, the research team set a standard for a meaningful population mean difference at the immediate posttest for the mediator as an absolute population difference of 0.20 or greater. The 95% confidence interval for the treatment difference was 0.61 to 0.83. Because the lower limit for the interval exceeds the meaningfulness standard, I conclude that Treatment A had a meaningful effect on the immediate posttest relative to the control group.

I also want to know the values of the predicted mediator means at the immediate posttest for Treatment A and for the control group. These values are labeled `M1A` and `M1C` within the `MODEL CONSTRAINT` command. Here is the edited output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    M1A                0.740      0.040     18.379      0.000
    M1C                0.019      0.037      0.502      0.616
```

Next, I evaluated the estimated effect of Treatment B on the outcome, I used the output generated by Line 45 (`MI IND DTREATB`) from Table 2. The output focuses on the treatment mean difference at the immediate posttest between Treatment B and the control group:

```
TOTAL, TOTAL INDIRECT, SPECIFIC INDIRECT, AND DIRECT EFFECTS

                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

Effects from DTREATB to MI

  Total                0.647      0.055     11.693      0.000
```

The adjusted difference is 0.647 ±0.110, which is statistically significant (CR = 11.69, $p < 0.05$). The 95% confidence interval for the treatment difference was 0.407 to 0.599. Because the lower limit of this interval exceeds the meaningfulness standard of 0.20, I can confidently conclude that Treatment B had a meaningful effect on the immediate posttest relative to the control group. The predicted mean of Y at the immediate posttest for Treatment B (taken from `Y1B` in the `MODEL CONSTRAINT` command) was 0.545± 0.078. The predicted mean for the control group was already reported above when I evaluated the effects of Treatment A on M.

Finally, I examined the relative difference between Treatment A and Treatment B on the adjusted outcome as taken from the relevant `MODEL CONSTRAINT` commands. Here is the output, where `M1A` is the adjusted outcome mean for Treatment A and `M1B` is the adjusted outcome mean for Treatment B:

	Estimate	S.E.	Est./S.E.	Two-Tailed P-Value
M1A	0.740	0.040	18.379	0.000
M1B	0.666	0.042	15.661	0.000
MDIFF1AB	0.075	0.058	1.298	0.194

The adjusted mediator mean difference between the two treatments was 0.075 ±0.12, which was statistically non-significant (CR = 1.30, *ns*). To assert treatment equivalence, I need to take into account the latitude of no effect per Chapter 10. Suppose the research team deemed that any population absolute difference on the mediator less than 0.15 would be deemed to be *functionally* zero. Stated another way, if the population mean difference falls anywhere between -0.15 and +0.15, the effect would be deemed trivial. The 95% confidence interval for the mean difference was -.04 to 0.19. Because the interval is not fully contained within the latitude of no effect, I cannot confidently conclude that the two treatments are functionally equivalent taking into account sampling error (because the upper limit 0.19 of the confidence interval exceeds the no effect standard of 0.15). Conclusion-wise, I am in a state of suspended judgment. I can't confidently conclude that the treatments are meaningfully different in terms of influencing the mediator but I also can't confidently conclude their difference is functionally zero. The confidence interval is a bit too wide to make a firm conclusion.

To gain a sense of the mediator decay curves for the two treatments and the control groups, I plot the adjusted estimated mediator means for the three groups as taken from the `MODEL CONSTRAINT` commands (variables `M1A, M2A, M3A, M4A, M1B, M2B, M3B, M4B, M1C, M2C, M3C, M4C`). Here is the plot using the *Temporal line plot* program from my website:

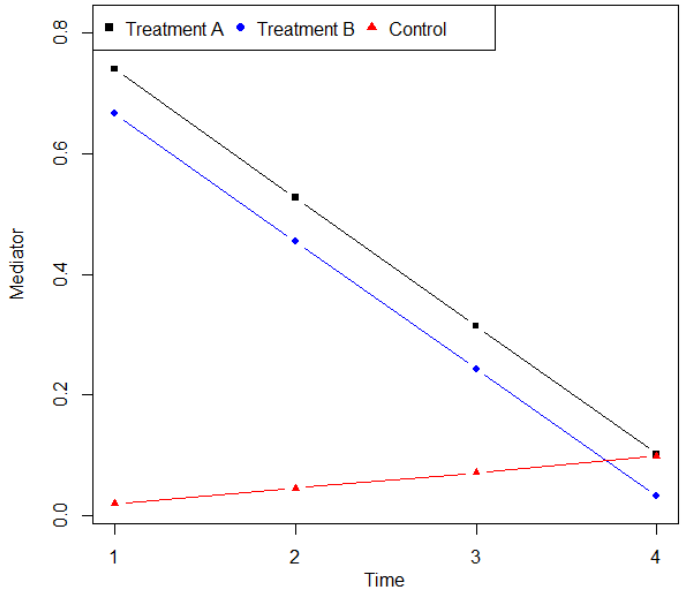

There is a steady, parallel decline in the mediator mean outcome for Treatment A and Treatment B over time. The control group shows a slight increase in the mean Y over time (although it is not statistically significant – see below). Here are the estimated slopes for the three groups taken from the `MODEL CONSTRAINT` commands:

```
MODEL RESULTS

                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

New/Additional Parameters

    MSA             -0.213      0.020    -10.687      0.000
    MSB             -0.211      0.020    -10.653      0.000
    MSC              0.026      0.019      1.396      0.163
```

The estimated slopes for Treatment A and Treatment B were negative and statistically significant (Treatment A = -0.213 ±0.040, CR = 10.69, $p < 0.05$; Treatment B = -0.211 ±0.040, CR = 10.65, $p < 0.05$). The estimated slope for the control group was 0.026 ±0.039, (CR = 1.40, *ns*). As in the analyses of the outcome, Y, there is little change in the mediator means across time for the control group. By contrast, the two treatment conditions showed nontrivial decreases in M across time.

The `MODEL CONSTRAINT` commands also provide the pairwise mean contrasts between the three slopes:

	Estimate	S.E.	Est./S.E.	Two-Tailed P-Value
MSDIFFAC	-0.240	0.027	-8.916	0.000
MSDIFFBC	-0.238	0.027	-8.894	0.000
MSDIFFAB	-0.002	0.027	-0.066	0.947

The estimated slope difference between Treatment A and the control group was -0.240 ±0.054 (CR = 8.92, $p < 0.05$). The estimated slope difference between Treatment B and the control group was -0.238 ±0.054 (CR = 8.89, $p < 0.05$). The estimated slope difference between the two treatments was -0.002 ±0.054 (CR = 0.07, *ns*). This difference was functionally zero.

A final set of contrasts that are informative are again from the `MODEL CONSTRAINT` commands. They examine pairwise estimated mean differences between the three groups at each time point, i.e., the immediate posttest, the 6 month follow-up, the 12 month follow-up and the 18 month follow-up. [remember that the output notation after the root term `DIFF` is followed by the timepoint (1, 2, 3 or 4) and then letters representing the two groups that are compared where A = Treatment A, B = Treatment B, and C = control):

	Estimate	S.E.	Est./S.E.	Two-Tailed P-Value
MDIFF1AC	0.722	0.054	13.355	0.000
MDIFF1BC	0.647	0.055	11.693	0.000
MDIFF1AB	0.075	0.058	1.298	0.194
MDIFF2AC	0.482	0.053	9.181	0.000
MDIFF2BC	0.409	0.053	7.773	0.000
MDIFF2AB	0.073	0.055	1.333	0.183
MDIFF3AC	0.243	0.064	3.820	0.000
MDIFF3BC	0.172	0.063	2.743	0.006
MDIFF3AB	0.071	0.065	1.100	0.271
MDIFF4AC	0.003	0.082	0.041	0.968
MDIFF4BC	-0.066	0.080	-0.822	0.411
MDIFF4AB	0.069	0.083	0.837	0.403

None of the contrasts comparing Treatment A to Treatment B at any of the time points are statistically significant. At the immediate posttest, both Treatment A and Treatment B are statistically significantly and meaningfully different from the control group. However, by the 18 month follow-up, these differences have dissipated to the point that all three group contrasts are statistically non-significant. At the 6 month follow-up, the treatment group differences from the control group are diminished somewhat but still statistically significant

and meaningful. At the 12 month follow-up, Treatment A and Treatment B remain statistically significantly different from the control group, but their 95% confidence intervals overlap with the meaningfulness standard indicating I cannot conclude the separation of them from the control group is meaningful.

In sum, when evaluating the effect of the intervention on the mediator, there are multiple facets to examine. The overall pattern of results suggest that the two treatment conditions produce comparable and meaningful effects on the mediator at the immediate posttest but this effect diminished across time to the point that by the 18 month follow-up, the treatments have lost their effects. On the positive side of things, the treatments still have meaningful effects on the mediator after 6 months.

Effect of the Mediator on the Outcome

The analysis of the estimated effect of the mediator on the outcome does not distinguish the three treatment groups because the guiding model assumes the effect is the same in all conditions. The coefficients for evaluating the mediator effect on the outcome are the coefficient from `MI` to `YI` and the coefficient from `MS` to `YS`. Here are the results:

```
MODEL RESULTS
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

YI       ON
    MI                 0.777      0.035     22.246      0.000
    Y0                 0.113      0.023      4.913      0.000

YS       ON
    MS                 0.781      0.048     16.371      0.000
```

The coefficient for MI→YI is the path coefficient for the outcome at the immediate posttest regressed onto the mediator at the immediate posttest controlling for the baseline outcome. The coefficient was 0.78 ±0.07 (CR = 22.25, $p < 0.05$). For every one unit that the mediator increased across individuals, the outcome was predicted to increase by 0.78 units. The 95% confidence interval for the coefficient was 0.71 to 0.85. Because the lower limit of this interval is greater than meaningfulness standard of 0.20, I can declare the result meaningful.

The coefficient for MS→YS is the decay path coefficient for the outcome regressed onto the decay coefficient for the mediator. A positive coefficient implies that people with increasingly more positive slopes on the mediator tend to have increasingly more positive slopes on the outcome, i.e., that there is an association between the two. The path coefficient was 0.78 ±0.10 (CR = 16.37, $p < 0.05$). For every one unit that the mediator slope increased

across individuals, the outcome slope was predicted to increase by 0.78 units. Suppose the research team decided prior to the study that a population coefficient of 0.20 or greater would be deemed meaningful. The 95% confidence interval for the coefficient was 0.68 to 0.88. Because the lower limit of this interval is greater than the meaningfulness standard, we can declare the result meaningful.

Omnibus Mediation Analysis

Omnibus mediation tests are of lower priority for me. I instead prefer to focus on the individual links in mediational chains, evaluating their effects sizes. If I want an overall omnibus test of statistical significance, I rely on the joint significance test. In LGC models, evaluations of mediation take two forms. The first focuses on treatment differences in the latent intercept for the mediator and whether these differences, in turn, translate into treatment differences on the outcome latent intercept. The second mediation form focuses on the effect of the latent slopes of the mediator on the latent slopes of the outcome. The RET found meaningful treatment versus control differences for both facets of mediation analysis. The pattern of results is consistent with omnibus mediation.

Concluding Comments for the Example

Although I do not pursue them here, it is good practice to evaluate the model for sensitivity to measurement error and omitted variable bias as well as the presence of outliers. Methods for doing so are described below. As I mentioned in the main text, correctly specifying the growth function is particularly important. In the current example, I assumed a linear function. If the function is misspecified, then the fit of the SEM model should be poor and this should be reflected in the fit diagnostics. This is an advantage of using SEM. Despite this, I like to augment the diagnostics with other checks when possible. For example if I plot the observed means (not the predicted means) for the outcome across the different time points for the three groups separately, do I observe reasonable approximations to linearity? Using the *Temporal line plot* program on my website, here is the relevant plot for the outcome Y from the immediate posttest through the 18 month follow-up:

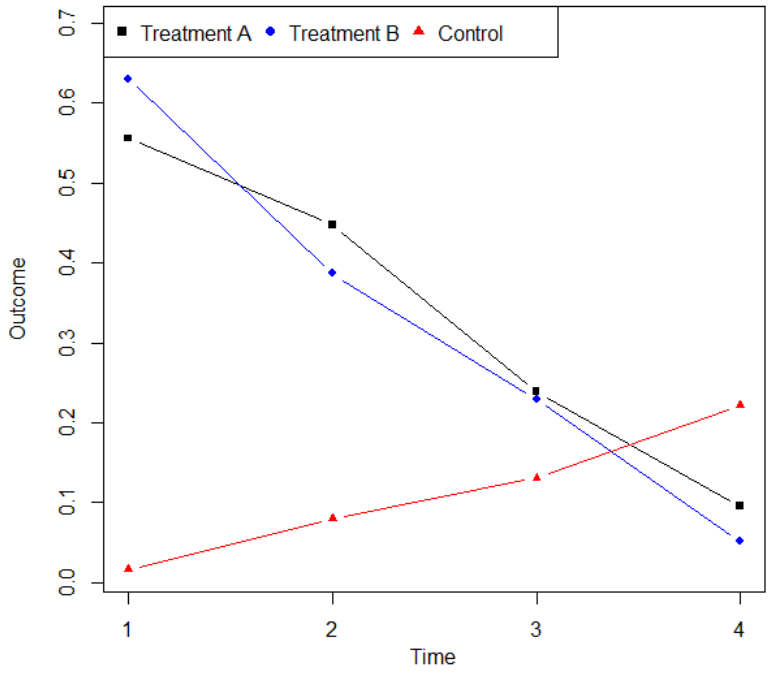

The trends seem close enough to linearity to justify a linear growth curve model. This also was true for the mediator observed means. Granted, there are some deviations but these deviations likely are due to random noise. The plot differs from those I presented earlier because the earlier plots are model-based, i.e., they assume the model in Figure 2 faithfully represents the true data generating process.

MULTIPLE GROUP ANALYSIS OF LGCM IN RETs

It is possible to perform LGCM in RETs using SEM-based multigroup analyses instead of the strategies described above and in the main text. The advantage of doing so is that it offers more flexibility and is simpler for introducing the `MODEL CONSTRAINT` commands. The primary disadvantage is that the multi-group approach is somewhat more sample size demanding. Muthén and Curran (1997) advocate for the multi-group approach and recommend a sequence of steps to use when pursuing it, such as conducting preliminary analyses on the intervention and control groups separately and testing for treatment-mediator interactions via across-group contrasts. My focus here is on explicating the basic set-up of the multigroup approach. See Muthén and Curran (1997) for elaborations.

I discuss SEM-based multiple group modeling in general and for moderator analysis in Chapter XX. To illustrate the approach for LGCM, I use the first example I presented in this document for the treatment-control RET that relies on Figure 1. I repeat the figure here for convenience; see Figure 3.

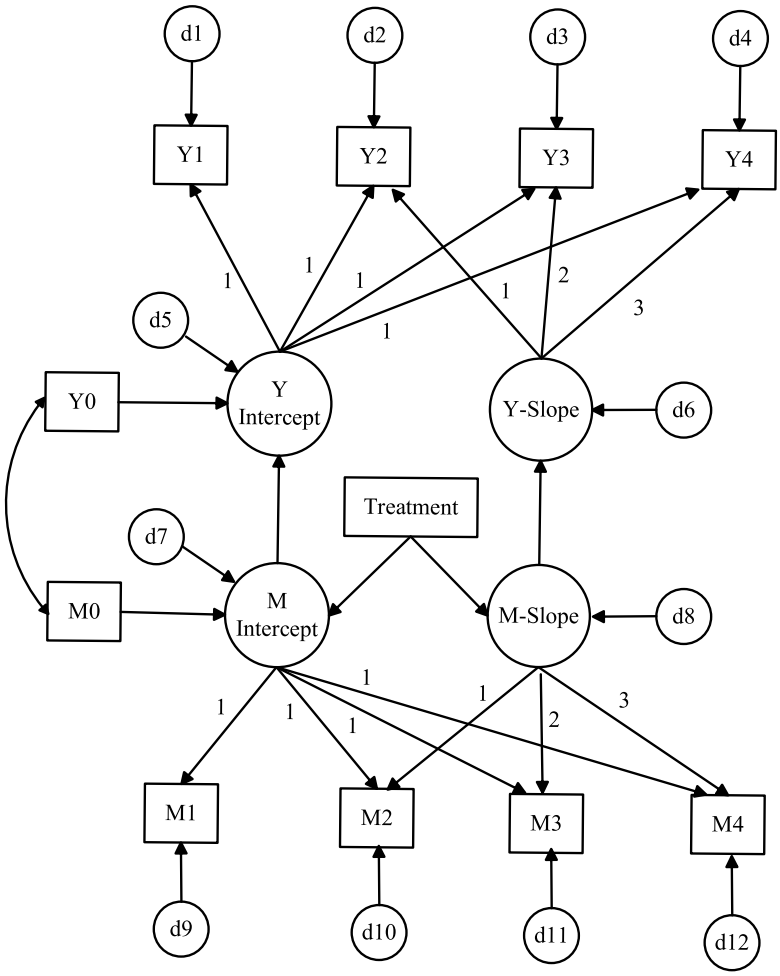

FIGURE 3. RET growth curve model

In a multiple group analysis, I split the sample in two, half consisting of everyone in the intervention condition and the other half consisting of everyone in the control condition. I then apply the model in Figure 3 to each half but leaving out the treatment dummy variable because dividing the sample in half renders it irrelevant for within-group analyses. The treatment condition is essentially held constant in each of the groups. To be sure, I need to conduct comparative analyses across the intervention and control groups to answer the three core questions of RETs. I show how to do this shortly. You can think of Figure 3 as representing the theory that guides the analysis rather than the analytic model per se.

Instead of physically conducting separate analyses on the two halves of the sample, SEM offers an option to combine the groups into a single analysis while still respecting the group differences, hence the name **multi-group analysis**. Mplus reports the results for the parameter estimates and a subset of fit indices for each group separately, but it also presents

fit statistics for the two groups considered simultaneously. For example the overall chi square statistic for the simultaneous analysis equals the value of the chi square statistic for the first group plus the chi square statistic for the second group. The degrees of freedom for the overall chi square statistic equals the degrees of freedom associated with the chi square statistic for the first group plus the degrees of freedom for the chi square statistic for the second group. Simultaneous fit indices are also calculated for most of the other commonly reported fit indices, as I show below.

Table 3 has the first portion of the Mplus syntax I used to execute the analysis. Most of this syntax is the same as that of Table 1 for the analysis of a single group that included the `TREAT` variable and that merged the two groups into one large sample. I will add more commands to the Table 3 syntax shortly. I highlight the lines that are different from the Table 1 syntax in red:

Table 3. Mplus Syntax for Multi-Group Analysis

```
1.  TITLE: LGCM multiple group analysis of RET  ;
2:  FILE = LGAdata2M.dat;
3.  DEFINE:
4.    CENTER m0 y0 (GRANDMEAN) ;
5.  VARIABLE:
6.    NAMES = Y1 Y2 Y3 Y4 M1 M2 M3 M4 Y0 M0 DTREAT ;
7.    GROUPING IS DTREAT (0=control, 1=treat) ;
8.  ANALYSIS: ESTIMATOR=MLR ;
9.  MODEL:
10.  yi ys | y1@0 y2@1 y3@2 y4@3 ; Define latent growth model for y
11.  !Estimate latent intercepts and assign them labels
12.  [yi] (ya1);      !Intercept for intercept factor
13.  [ys] (ya2);      !Intercept for slope factor
14.  !Estimate variances of intercept and slope factors
15.  yi;       ! Variance of intercept factor
16.  ys;       ! Variance of slope factor
17.  !Estimate covariance between intercept and slope factors
18.  yi WITH ys;
20.  ! Estimate residual variances (error)
21.  y1-y4;
22.  !Define latent growth model for mediator
23.  mi ms | m1@0 m2@1 m3@2 m4@3 ;
24.  !Estimate intercepts and assign them labels
25.  [mi] (ma1) ;      !Intercept for intercept factor
26.  [ms] (ma2) ;      !Intercept for slope factor
27.  !Estimate variances of intercept and slope factors
28.  mi;       ! Variance of intercept factor
29.  ms;       ! Variance of slope factor
30.  !Estimate covariance between intercept and slope factors
31.  mi WITH ms;
```

```
32.  !Estimate residual variances (error)
33.  m1-m4;
34.  !Define the regressions and assign labels to coefficients
35.  mi ON m0   ;
37.  yi ON mi y0 ;
38.  ys ON ms ;
```

Most of the syntax should be familiar. On Line 7, I added new syntax to identify for Mplus the variable used to define the groups and in parentheses, what the numbers are on that variable that define the groups (in this case, 0 and 1) and then the labels to assign to the numbers (in this case "control" and "treat"). On the NAMES line I renamed the input variable TREAT to DTREAT so that Mplus does not confuse it with the label on Line 7. I delete all references to the TREAT or DTREAT variable throughout the syntax that I had programmed in Table 1. This includes all of the MODEL INDIRECT commands from Table 1 because they all included reference to the variable TREAT.

In multiple group SEM, it is understood that all of the lines under the word MODEL (Lines 9 to 38) are applicable to all of the groups being analyzed unless you specifically tell Mplus to override some of them. For example, if you want to assign unique labels to a subset of the parameters in each group, you need to tell Mplus that such is the case. Or, you might want to fix a parameter to a specific value in one group but not the other. The way we communicate this to Mplus is to add to the above a separate MODEL command for each group containing the commands we want to take precedence over the more general MODEL commands, like this:

```
39. MODEL control:
40.  [yi] (ya1gp0) ;
41.  [ys] (ya2gp0);
42.  [mi] (ma1gp0) ;
43.  [ms] (ma2gp0) ;
44.  yi ON mi y0 (p1gp0 p2gp0) ;
45.  ys ON ms (p3gp0);
46. MODEL treat:
47.  [yi] (ya1gp1) ;
48.  [ys] (ya2gp1);
49.  [mi] (ma1gp1) ;
50.  [ms] (ma2gp1) ;
51.  yi ON mi y0 (p1gp1 p2gp1) ;
52.  ys ON ms (p3gp1);
53. OUTPUT: Samp StdYX MOD(4) Residual Tech4  ;
```

Lines 39 and 52 specify the group specific command overrides using the specialized labels I assigned to the group names in Line 7. In this case, the only overrides I invoked was

to add group specific labels to a subset of the commands. For example, I overrode Line 15 in the general model command which was `[yi];` to now have a group specific label `[yi] (ya1gp0);`. I make use of the group specific label later when I add `MODEL CONSTRAINT` commands.

I execute the above syntax so that I can evaluate the overall fit of the model. Once I am convinced I have a reasonable fitting model, I add `MODEL CONSTRAINT` commands to explore specific parameter estimates and to make substantively interesting contrasts. I then execute the augmented syntax again but with changes to the `OUTPUT` line to eliminate the request for modification indices. This is because if I add `MODEL CONSTRAINT` commands, Mplus forces me to remove the request for modification and I do not want to sacrifice them in my initial analysis when evaluating model fit. Let's examine the simultaneous fit indices. Here is the relevant output:

```
MODEL FIT INFORMATION

Chi-Square Test of Model Fit

          Value                             74.441*
          Degrees of Freedom                    76
          P-Value                           0.5292
          Scaling Correction Factor         0.9931
            for MLR

Chi-Square Contribution From Each Group

          CONTROL                           32.532
          TREAT                             41.909

RMSEA (Root Mean Square Error Of Approximation)

          Estimate                           0.000
          90 Percent C.I.                    0.000  0.024
          Probability RMSEA <= .05           1.000

CFI/TLI

          CFI                                1.000

SRMR (Standardized Root Mean Square Residual)

          Value                              0.024
```

The overall chi square index is for the two groups combined. The p value for it was statistically non-significant, suggesting a reasonable fitting model. As noted, the total chi

square statistic is an additive function of the separate chi square indices for each group. In the section `Chi-Square Contribution From Each Group`, Mplus reports the separate chi square values for each group. One typically does not want to see a sizeable disparity in these values because it would then suggest the model is fitting much better in one of the groups compared to the other. If the disparity is large, the question becomes why is the model fitting worse in one group than the other and you will want to resolve this.

The remaining indices should all be familiar to you but keep in mind that these statistics are for the simultaneous analysis, not the per group analyses. All of them appear to be in order for the current example. For the localized fit indices, Mplus reports modification indices, cell-by-cell z tests of the difference between the predicted and observed covariances, as well as estimates of the differences between the predicted and observed correlations *per se*. I do not show these here in the interest of space, but they follow the same format and interpretation as prior chapters. In the current case, there were no theoretically meaningful problematic localized fit indices for either of the two groups.[1]

In general, if you make a modification in one of the groups to free-up estimation of a given parameter that was previously constrained to zero in the model (e.g., introduce a correlated disturbance that had before been set to zero), you will need to also free up that particular parameter in the other group even if the local fit indices for the other group suggest the parameter can be left alone. In most multiple group analyses, we want the full model to be parallel in all the groups for the analytics to work properly. This is not an absolute requirement but it is common in practice. To allow different model configurations in the multiple groups requires specialized programming and non-trivial statistical issues that are beyond the scope of my book. If a parameter is freed up in a group where it did not need to be but is freed for purposes of creating parallel configurations, then all that is lost is a degree of freedom in that group that otherwise could have been spared. The value of the parameter estimate itself when the modified program is estimated likely will just be a value near zero, which is what you assumed anyway when you initially omitted the parameter.

Given a reasonable fitting multi-group model, I then remove the request for modification indices from the `OUTPUT` line and just before that line, add my `MODEL CONSTRAINT` commands. Here are the commands I add just before Line 53 in Table 3 (note: some of the lines make use of the group specific labels I introduced earlier):

```
52a1.  MODEL CONSTRAINT:
52a2.   NEW (mst msc msdiff m1t m2t m3t m4t m1c
52a3.   m2c m3c m4c mdiff1 mdiff2 mdiff3 mdiff4) ;
```

[1] You can obtain separate indices for each group of the RMSEA, CFI, and SRMR indices but to do so you must execute separate analyses on each group via a traditional Mplus program. These indices do not have the summative feature that the chi square statistic has.

```
52a4.   mst = ma2gp1 ;      !Treatment mean of mediator slope factor
52a5.   msc = ma2gp0 ;      !Control mean of mediator slope factor
52a6.   msdiff = mst-msc ;
52a7.   m1t = ma1gp1 ;             !Treatment mediator value at time 1
52a8.   m2t = m1t + mst*1  ;       !Treatment mediator value at time 2
52a9.   m3t = m1t + mst*2  ;       !Treatment mediator value at time 3
52a10.  m4t = m1t + mst*3  ;        !Treatment mediator value at time 4
52a11.  m1c = ma1gp0 ;              !Control mediator value at time 1
52a12.  m2c = m1c + msc*1 ;         !Control mediator value at time 2
52a13.  m3c = m1c + msc*2 ;         !Control mediator value at time 3
52a14.  m4c = m1c + msc*3 ;         !Control mediator value at time 4
52a15.  mdiff1 = m1t - m1c ;
52a16.  mdiff2 = m2t - m2c ;
52a17.  mdiff3 = m3t - m3c ;
52a18.  mdiff4 = m4t - m4c ;
52b1.   NEW (yst ysc ysdiff y1t y2t y3t y4t y1c
52b2.   y2c y3c y4c ydiff1 ydiff2 ydiff3 ydiff4) ;
52b3.   yst = ya2gp1 + (ma2gp1)*p3gp1 ; !Treatment mean of Y slope factor
52b4.   ysc = ya2gp0 + ma2gp0*p3gp0 ;   !Control mean of Y slope factor
52b5.   ysdiff = yst-ysc ;
52b6.   y1t = ya1gp1 + ma1gp1*p1gp1;     !Treatment Y value at time 1
52b7.   y2t = y1t + yst*1;               !Treatment Y value at time 2
52b8.   y3t = y1t + yst*2;               !Treatment Y value at time 3
52b9.   y4t = y1t + yst*3;               !Treatment Y value at time 4
52b10.  y1c = ya1gp0 + ma1gp0*p1gp0;     !Control Y value at time 1
52b11.  y2c = y1c + ysc*1;               !Control Y value at time 2
52b12.  y3c = y1c + ysc*2;               !Control Y value at time 3
52b13.  y4c = y1c + ysc*3;               !Control Y value at time 4
52b14.  ydiff1 = y1t - y1c ;
52b15.  ydiff2 = y2t - y2c ;
52b16.  ydiff3 = y3t - y3c ;
52b17.  ydiff4 = y4t - y4c ;
```

I described how to derive the above expressions in the main text example in Chapter 16 and in the Appendix of Chapter 16. These terms are key to answering the first two questions of the three core RET questions.

Before I review the results of the analysis, there is an important feature of multi-group SEM that I should mention. In traditional single group RET models, we obtain a single estimate of each model parameter. For example, the estimated effect of the mediator slope factor on the outcome slope factor in the single group analysis of the data was 0.77 ± 0.11. In the multigroup-analysis we obtain two parallel parameter estimates, one for the intervention group and the other for the control group. The values for the aforementioned path, as you will see below, were 0.77 ±0.19 in the intervention group and 0.70 ±0.15 in the control group. When you conduct the single-group analysis, you implicitly make the assumption that the coefficients in the two groups are equal and that any difference between

them when analyzed separately reflects random noise. The multi-group solution estimates them separately such that you can empirically evaluate this homogeneity assumption. In the current example, the coefficients were close (0.77 versus 0.70). If I add a contrast comparing the two coefficients using the `MODEL CONSTRAINT` command, the test of the difference in this case was statistically non-significant (CR = 0.49, *ns*).

Although the multi-group analysis has the ability to test the underlying homogeneity assumption, it comes at some cost because the separate parameter estimates for each group are based on a smaller N than the single group analysis, namely they each are based on the separate group sizes. You can see the sample size differential at work in the estimates reported above because the multi-group estimates have somewhat larger margins of error than the single group estimate.

A related difference between the single and multigroup approaches is that the latter allows all of the coefficients in the model for the two groups to vary between the groups. Not only did the MS→YS coefficient differ between the groups pe the above, albeit only slightly and probably because of random noise, the same is true for the other coefficients in the model. Some would argue that this property is a plus as the multi-group analysis is more nuanced given this property. Others argue the property can lead to sample-specific overfitting. Of course, you can always impose across-group equality constraints across on whatever coefficients you wish but this should be done cautiously. Strict coefficient equality is a very stringent condition that rarely holds in the real-world. Perhaps it is better, the argument goes, to allow for minor between-group deviations to more accurately reflect a world in which minor differences likely exist. If I have strong theory that suggests across group equivalence for a given coefficient, then I let that theory guide my equality constraint decisions. Absent that, my preference often leans toward the single group analysis as outlined at the beginning of this document because it tends to have more statistical power and it can better accommodate smaller N. If my sample size is large, the power advantage of the single group approach dissipates. Sometimes I pursue sensitivity analyses where I compare a multi-group model with equality constraints to a multi-group model without them as well as a single group model. If my substantive conclusions are the same in all models, then the matter is moot. If the conclusions differ, then I need to track down the source of the differences.

Total Effect of the Treatment on the Outcome in the Multi-Group Analysis

To test the effects of the intervention on the outcome, I first examine treatment condition differences for Y at the immediate posttest. I obtain this contrast and the separate means for the intervention and control groups from the `MODEL CONSTRAINT` commands. Here is the relevant (edited) output:

```
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

New/Additional Parameters
    YDIFF1             0.560      0.066      8.429      0.000
    Y1T                0.573      0.048     11.827      0.000
    Y1C                0.013      0.045      0.276      0.783
```

The adjusted mean difference is 0.56 ±0.13, which is statistically significant (CR = 8.43, $p < 0.05$). The 95% confidence interval is 0.43 to 0.69. Suppose that the standard for a meaningful population mean difference at the immediate posttest is a population absolute difference of 0.20 or greater. Because the lower limit of the mean difference confidence interval (which roughly equals the `Estimate` minus double the `S.E.`) was greater than this standard, I conclude that the intervention effect for the immediate posttest is meaningful. The estimated Y means for the intervention and control conditions at the immediate posttest were 0.57 ± 0.10 and 0.013 ± 0.09, respectively.

Next, I examined a second facet of the effect of the intervention on the outcome, namely whether there was meaningful decay in the intervention effect over time. As a first step, I isolate the coefficient characterizing the decay curve for the intervention group and that for the control group. I then examine a contrast comparing the two coefficients. The relevant values are obtained in the `MODEL CONSTRAINT` command. Here is the edited output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    YST               -0.159      0.023     -6.988      0.000
    YSC                0.066      0.022      2.940      0.003
    YSDIFF            -0.225      0.032     -7.037      0.000
```

The model implied linear decay in the intervention condition is negative and statistically significant (slope = -0.159 ± 0.046, CR = 6.99, $p < 0.05$) indicating that the predicted mean of Y decreases as time from the immediate posttest increases across the 18 months of the study. For the control condition, the change in the mean of Y across time is relatively flat although the slight positive increase is statistically significant (slope = 0.066 ±0.044, CR = 2.94, $p < 0.05$). The difference between the two coefficients (-0.225 ±0.06) was statistically significant (CR = 7.04, $p < 0.05$).

I can plot the two curves using the information in the output section from the `MODEL CONSTRAINT` commands for the predicted means of Y1 through Y4 for the intervention

group (labeled Y1T, Y2T, Y3T, and Y4T) and for the control group (labeled Y1C, Y2C, Y3C, and Y4C). Using the program on my website called *Temporal line plot*, here is the plot:

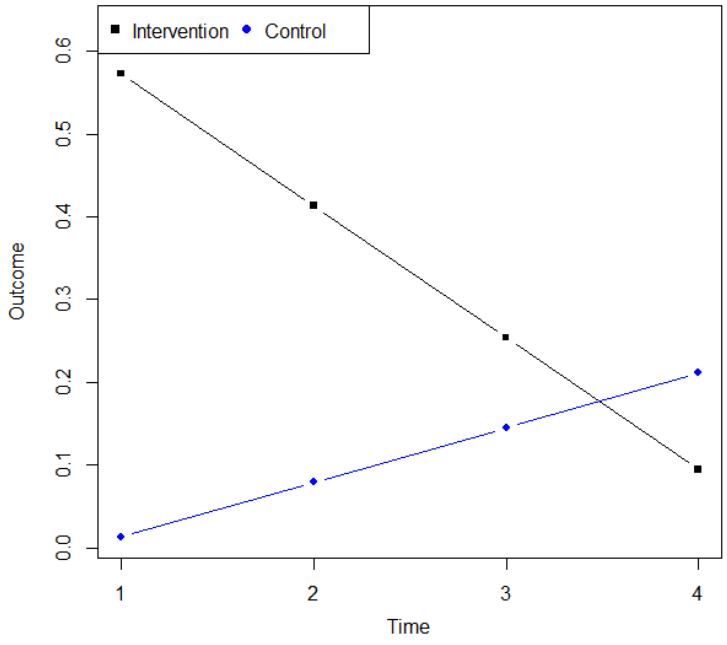

The downward decay in the predicted Y means for the intervention group is evident as is the slight upward trend in the means for the control group. The separation between lines at a given time reflects the predicted mean difference between the intervention and control groups. It can be seen in the plot that the mean difference between the treatment and control conditions decreases with time to the point that by month 18 the means are close.

The contrasts that I specified in the MODEL CONSTRAINT commands include tests of the outcome mean difference at each of the four time points using the model-informed predicted means. The contrasts are labeled YDIFF1, YDIFF2, YDIFF3, YDIFF4 for times 1, 2, 3 and 4, respectively (keep in mind that time 1 is the immediate posttest). Here is the output:

```
New/Additional Parameters
                                                    Two-Tailed
                   Estimate       S.E.  Est./S.E.    P-Value

    YDIFF1            0.560      0.066      8.429      0.000
    YDIFF2            0.335      0.064      5.265      0.000
    YDIFF3            0.110      0.076      1.449      0.147
    YDIFF4           -0.116      0.097     -1.190      0.234
```

The intervention group predicted mean is statistically significantly larger than the control group mean at the first two time points but not the last two. Suppose a

meaningfulness standard is set by the research team to be a population absolute difference of 0.20 or greater. The 95% confidence interval for YDIFF1 is 0.43 to 0.69 and for YDIFF2 it is 0.21 to 0.46. In both cases, the lower limit exceeds the meaningfulness standard. Given this, I conclude these Y predicted mean differences are meaningful.

Effect of the Treatment on the Mediator

To test the effects of the intervention on the mediator, I follow the same structure as that for the total effect on Y. I first examine treatment condition differences for the mediator at the immediate posttest. I obtain this contrast and the separate means for the intervention and control groups from the MODEL CONSTRAINT commands. Here is the relevant output:

```
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

New/Additional Parameters
    MDIFF1             0.707      0.056     12.577      0.000
    M1T                0.748      0.041     18.150      0.000
    M1C                0.041      0.038      1.068      0.286
```

The adjusted mean difference is 0.71 ±0.11, which is statistically significant (CR = 12.58, $p < 0.05$). The 95% confidence interval is 0.60 to 0.82. Suppose that the standard for a meaningful mean difference for the mediator at the immediate posttest is a population absolute difference of 0.20 or greater. Because the lower limit of the mean difference confidence interval was greater than this standard, I conclude that the intervention effect for the immediate posttest is meaningful. The mediator means for the intervention and control conditions at the immediate posttest were 0.75 ±0.08 and 0.04 ±0.08, respectively.

Next, I examined a second facet of the intervention effect on the mediator, namely whether there was meaningful decay in the effect of the intervention on the mediator over time. As a first step, I isolate the coefficient for mediator decay for the intervention group and that for the control group. I then examine a contrast comparing the two coefficients. The relevant values are from the MODEL CONSTRAINT command. Here is the edited output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    MST               -0.214      0.021    -10.375      0.000
    MSC                0.013      0.020      0.642      0.521
    MSDIFF            -0.227      0.029     -7.929      0.000
```

The model implied linear decay for the mediator in the intervention condition is negative and statistically significant (slope = -0.214 ±0.042, CR = 10.38, $p < 0.05$) indicating that the predicted mean for the mediator decreases as time from the immediate posttest increases across the 18 follow-up months of the study. For the control condition, the decay in the mean of M is relatively flat (because the slope is near a value of zero) and not statistically significant (slope = 0.013 ±0.04, CR = 0.64, *ns*). The difference between the two coefficients (-0.227 ±0.06) was statistically significant (CR = 7.92, $p < 0.05$).

I can plot the two curves using the information in the output section from the `MODEL CONSTRAINT` commands for the predicted means at M1 though M4 for the intervention group (labeled `M1T`, `M2T`, `M3T`, and `M4T`) and for the control group (labeled `M1C`, `M2C`, `M3C`, and `M4C`). Using the program on my website called *Temporal line plot*, here is the plot:

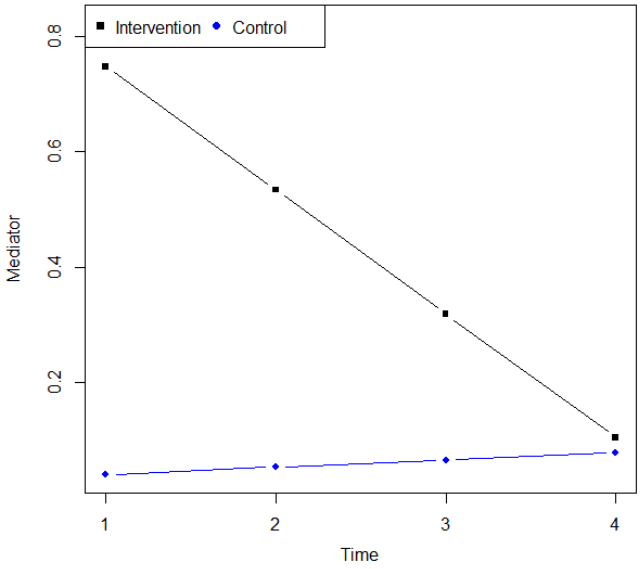

The downward decay in the predicted mediator means for the intervention group is evident as is the slight upward trend in the means for the control group. The separation between lines at a given time point reflects the predicted mean difference between the intervention and control groups. The mean difference between the treatment and control conditions decreases as time passes to the point that by month 18 the means are close.

The contrasts that I specified in the `MODEL CONSTRAINT` commands also test the mediator mean difference at each of the four time points using the model-informed predicted means. The contrasts are labeled `MDIFF1`, `MDIFF2`, `MDIFF3`, `MDIFF4` for times 1, 2, 3 and 4, respectively (keep in mind that time 1 is the immediate posttest). Here is the relevant output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    MDIFF1             0.707      0.056     12.577      0.000
    MDIFF2             0.480      0.053      9.043      0.000
    MDIFF3             0.253      0.064      3.943      0.000
    MDIFF4             0.026      0.084      0.307      0.759
```

The intervention group predicted mediator mean is statistically significantly larger than the control group predicted mean at the first three time points but not the last one. Suppose a meaningfulness standard is set by the research team to be a population absolute difference of 0.20 or greater. The 95% confidence interval for `MDIFF1` is 0.64 to 0.76, for `MDIFF2` it is 0.37 to 0.59, and for `MDIFF3` it is 0.12 to 0.36. For the first two time periods, the lower limit of the confidence interval exceeds the meaningfulness standard. Given this, I conclude that the differences at the immediate posttest and the 6 month follow-up are meaningful. For the 12 month follow-up, although the result is statistically significant, I cannot confidently conclude the effect is meaningful because the confidence interval overlaps the meaningfulness standard.

Effect of the Mediator on the Outcome

For the multi-group model, the analysis estimates the effect of the mediator on the outcome separately for the intervention and control groups. The relevant coefficients are taken from the `MODEL RESULLTS` section and appear as follows (edited):

```
MODEL RESULTS

                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

Group CONTROL

  YI ON MI             0.620      0.073      8.523      0.000
     ON M0             0.124      0.040      3.100      0.002

  YS ON MS             0.772      0.057     13.478      0.000

Group TREAT

  YI ON MI             0.864      0.070     12.434      0.000
     ON M0             0.116      0.039      2.942      0.003

  YS ON MS             0.760      0.096      7.879      0.000
```

The coefficient for MI→YI is the path coefficient for the outcome at the immediate posttest regressed onto the mediator at the immediate posttest controlling for the baseline outcome. It was statistically significant and meaningful in both the intervention and control conditions. The coefficient for MS→YS is the path coefficient for the decay slope of the outcome regressed onto the decay slope for the mediator. A positive coefficient implies that people with more positive slopes on the mediator tend to have more positive slopes on the outcome, i.e., that there is an association between the two. The path coefficient was statistically significant and meaningful in both the intervention and control conditions.

Omnibus Mediation Analysis

Omnibus mediation tests are of lower priority for me. I instead prefer to focus on the individual links in mediational chains, evaluating their effects sizes. If I want an overall omnibus test of statistical significance, I rely on the joint significance test. In LGC models, evaluations of mediation take two forms. The first focuses on treatment differences in the latent intercept for the mediator and whether these differences, in turn, translate into treatment differences on the outcome latent intercept. The second mediation form focuses on the effect of the latent slopes of the mediator on the latent slopes of the outcome. The RET found meaningful treatment versus control differences for both facets of mediation analysis. The pattern of results is consistent with omnibus mediation.

NON-LINEAR GROWTH FUNCTIONS

The analysis of linear trajectories in LGCM is straightforward given the widespread use of linear models in the social sciences. With nonlinear growth functions LGCM becomes more complex. Some researchers use polynomial modeling to capture non-linearity dynamics but this approach is limited in the types of non-linearity it can accommodate (see Chapter 15) and requires some often overlooked assumptions that I consider later. I consider three non-linear approaches that are more flexible than polynomial modeling and that are somewhat easier to work with. The first approach uses basis growth curve models. The second approach uses piecewise models. The third approach uses log transformed variables. In the remainder of this section, I assume you have read and mastered the material in the prior section on SEM based multi-group LGCM.

Basis LGCM

A core feature for the implementation of linear growth curve models is the fixing of path coefficients to equal 1.0 from the latent intercept factor to the growth-based observed variables and the latent slope factor to values that map onto linear functions of time. In the

examples I have considered thus far that use five time points for Y, the baseline measure has been treated as a covariate and the latent slope factor coefficients have been fixed at values of 0, 1, 2, and 3 where 0 reflects the immediate posttest. In **latent basis growth modeling** (also known as **shape latent growth modeling**), nonlinear growth functions are accommodated by empirically estimating a subset of the slope factor coefficients rather than fixing them at *a priori* values. For such a non-linear model to be identified, we need to fix one of the slope coefficients to 0 and a second coefficient to 1 but we estimate the remaining coefficients empirically. Researchers differ on which slope factor coefficients they choose to fix to 0 and 1. A common practice is to fix the first two time points in the growth curve to 0 and 1, respectively. I orient my discussion towards this strategy for reasons that become apparent shortly. A popular alternative is to fix the first time period at 0 and the last time interval at 1.0. For details see Preacher et al. (2008) and Bollen & Curran (2006).

Another variation I implement relative to prior examples is to include the baseline measure as part of the growth curve rather than as a covariate. The example I analyze, as in previous examples, has five time points in which I measure the outcomes Y0, Y1, Y2, Y3, Y4 and the mediators M0, M1, M2, M3, and M4. I model these variables using a parallel process growth model. I model the RET using SEM-based multiple group analyses as described earlier. This is because the nature of the non-linearity in the two groups could, in principle, be quite disparate in form for one group compared to the other. Such a state of affairs can create convergence issues. Multiple group modeling is more forgiving.

The influence diagram for the model is in Figure 4. I again omit selected correlated disturbances from the diagram to reduce clutter but they are modeled in the Mplus syntax. As with the prior multi-group analysis, I do not formally include the treatment variable in the within-group analyses. You can think of Figure 4 as the theory that guides the analysis rather than the analytic model per se. The variable metrics are such that their standard deviations are near 1.0. There is a six month lag between the immediate posttest and the first follow-up as well as six month lags between subsequent follow-up assessments. Instead of referring to the latent slope factor with the term "slope," I refer to it as a latent basis factor because in a basis growth model, it does not necessarily represent a slope. The path coefficient from the latent Y-basis to Y0 is fixed at zero which is equivalent to omitting the path. I omit the path from the figure to avoid clutter but the implication is that the path coefficient is fixed at zero. The same is true for the path coefficient from M-Basis to M0. I place an asterisk next to the paths from the latent Y-basis variable to Y2, Y3 and Y4 and from the latent M-Basis variable to M2, M3, and M4 to reflect they are estimated rather than fixed at a priori values. This means that instead of equal change being estimated between two consecutive time points with the same lag time per a linear model, growth can accelerate or decelerate across consecutive time periods.

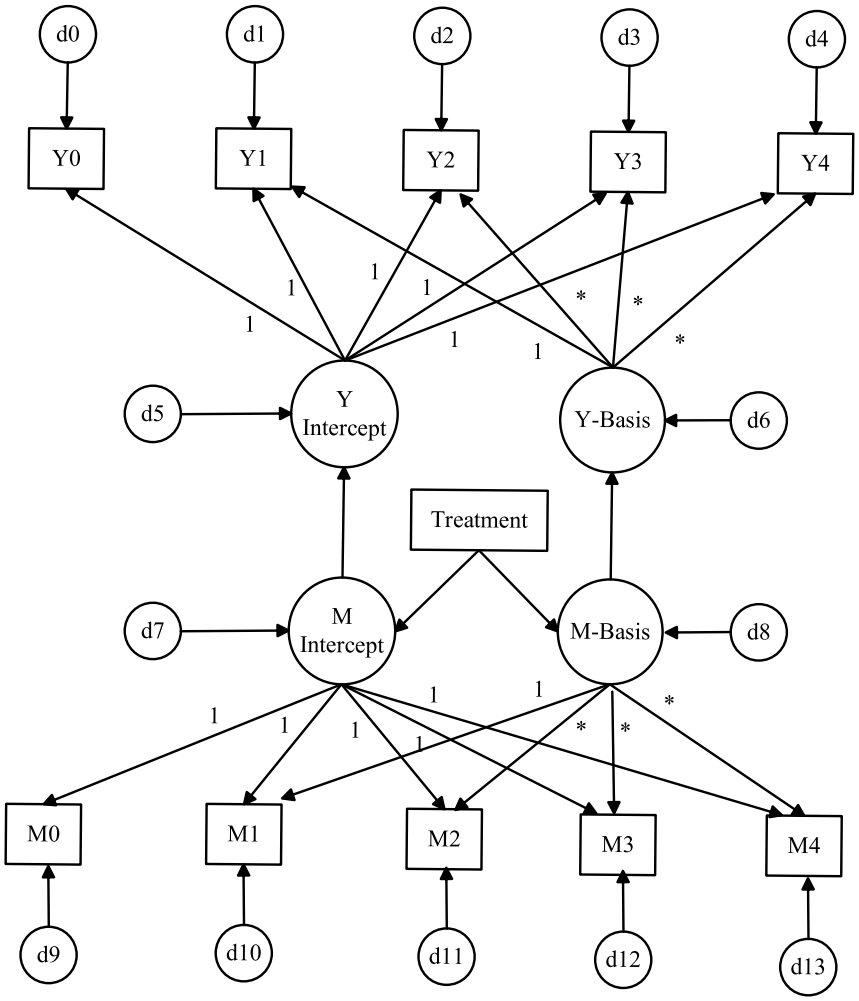

FIGURE 4. RET basis growth curve model

The basis growth curve model as implemented has interesting properties. The latent intercept factor for Y and M reflects the estimated baseline values of Y and M, respectively. These values generally are not of substantive interest. The latent basis factors have different interpretations depending on which paths you fix at 0 and 1. For the case where the first two time points are fixed, the mean of the mediator basis factor equals the estimated mean difference between the mediator mean at Time 1 minus the mediator mean at Time 0.

Table 4 presents the initial Mplus syntax. It is similar to the syntax in Table 3 for the prior multi-group example. I have highlighted in red in Table 4 where I made changes to the Table 3 syntax. Normally, instead of using YS to refer to the latent basis factor I would use YBASIS because I am not estimating a classic slope for growth. However, to maintain continuity with the syntax in Table 3, I will still use YS to refer to the latent basis factor in the Mplus syntax.

Table 4. Mplus Syntax for Latent Basis Growth Model

```
1.  TITLE: LGCM basis growth curve  ;
2:  FILE = LGAdata4M.dat;
3.  VARIABLE:
4.    NAMES = id Y0 Y1 Y2 Y3 Y4 M0 M1 M2 M3 M4 DTREAT ;
5.    GROUPING IS DTREAT (0=control, 1=treat) ;
6.  ANALYSIS: ESTIMATOR=MLR ;
7.  MODEL:
8.   yi BY y0-y4@1;
9.   ys BY y0@0 y1@1 y2* y3* y4*;
10.  [y0-y4@0 yi ys];
11.  !Estimate latent intercepts and assign them labels
12.  [yi] (ya1);      !Intercept for intercept factor
13.  [ys] (ya2);      !Intercept for slope factor
14.  !Estimate variances of intercept and slope factors
15.  yi;       ! Variance of intercept factor
16.  ys;       ! Variance of slope factor
17.  !Estimate covariance between intercept and slope factors
18.  yi WITH ys;
20.  ! Estimate residual variances (error)
21.  y1-y4;
22.  !Define latent growth model for mediator
23.  mi BY m0-m4@1;
24.  ms BY m0@0 m1@1 m2 m3 m4;
25.  [m0-m4@0 mi ms];
26.  !Estimate intercepts and assign them labels
27.  [mi] (ma1) ;      !Intercept for intercept factor
28.  [ms] (ma2) ;      !Intercept for slope factor
29.  !Estimate variances of intercept and slope factors
30.  mi;       ! Variance of intercept factor
31.  ms;       ! Variance of slope factor
32.  !Estimate covariance between intercept and slope factors
33.  mi WITH ms;
34.  !Estimate residual variances (error)
35.  m1-m4;
36.  !Define the regressions and assign labels to coefficients
37.  mi ON m0   ;
38.  yi ON mi y0 ;
39.  ys ON ms ;
40. MODEL control:
41.  ys BY y0@0 y1@1 y2 y3 y4 (yslc0-yslc4) ;
42.  ms BY m0@0 m1@1 m2 m3 m4 (mslc0-mslc4) ;
43.  [yi] (ya1gp0) ;
44.  [ys] (ya2gp0);
45.  [mi] (ma1gp0) ;
46.  [ms] (ma2gp0) ;
47.  yi ON mi (p1gp0) ;   !remove y0 from prior program and label
48.  ys ON ms (p2gp0);    !change label from prior program
```

```
49. MODEL treat:
50.  ys BY y0@0 y1@1 y2 y3 y4 (yslt0-yslt4) ;
51.  ms BY m0@0 m1@1 m2 m3 m4 (mslt0-mslt4) ;
52.  [yi] (ya1gp1) ;
53.  [ys] (ya2gp1);
54.  [mi] (ma1gp1) ;
55.  [ms] (ma2gp1) ;
56.  yi ON mi (p1gp1) ;     !remove y0 from prior program and label
57.  ys ON ms (p2gp1);      !change label from prior program
58. OUTPUT: Samp StdYX MOD(4) Residual Tech4  ;
```

In prior examples, I used an Mplus syntax shortcut for specifying the growth model that had the following format:

```
yi ys | y1@0 y2@1 y3@2 y4@3 ;
```

When I free instead of fix the coefficients in this statement, I need to assign labels to the them for use in later `MODEL CONSTRAINT` commands. However, Mplus does not allow the use of labels with this shortcut. I therefore must write separate syntax that underlies the above shortcut and that allows me to provide the labels I need when I specify the separate group commands. Lines 8 to 10 do so. Line 8 uses the `BY` command to specify the path coefficients from the latent intercept variable, `yi`, to the indicators y0, y1, y2, y3 and y4 fixing each of these coefficients to 1.0. Line 9 uses the `BY` command in the same way to fix the coefficients for y) and y1 to 0 and 1, respectively, and to freely estimate the remaining three coefficients by specifying * in place of fixed values. Line 10 fixes the measurement intercepts of y0 to y4 to zero. Lines 23 to 25 repeat these lines but for the mediator. Lines 41-42 and 50-51 assign labels to the latent basis factor coefficients.

I executed the above syntax to evaluate model fit. Given good model fit, I then add the `MODEL CONSTRAINT` commands and re-execute the program. Here is the fit output:

```
MODEL FIT INFORMATION

Chi-Square Test of Model Fit

          Value                             61.074*
          Degrees of Freedom                    74
          P-Value                           0.8590


  Chi-Square Contribution From Each Group

          CONTROL                           34.141
          TREAT                             26.933
```

```
RMSEA (Root Mean Square Error Of Approximation)

          Estimate                           0.000
          90 Percent C.I.                    0.000  0.015
          Probability RMSEA <= .05           1.000

CFI/TLI

          CFI                                1.000

SRMR (Standardized Root Mean Square Residual)

          Value                              0.020
```

The overall chi square index was 61.07 with 74 degrees of freedom, $p < 0.86$, which is consistent with satisfactory fit. The separate group chi square indices also were reasonable. The RMSEA was <0.001. The upper limit of the 90% confidence interval for it was 0.015. The p value for close fit was $p < 1.00$. The CFI was 1.00 and the standardized RMR was 0.020. For localized fit, there were no theoretically meaningful modification indices greater than 4 and no meaningful standardized residuals.

Here are the `MODEL CONSTRAINT` commands I then added, which follow directly from the commands I explained in the prior multi-group example:

```
57a1.  NEW (mst msc msdiff m0t m1t m2t m3t m4t m0c m1c
57a2.    m2c m3c m4c mdiff0 mdiff1 mdiff2 mdiff3 mdiff4) ;
57a3.    mst = ma2gp1 ;
57a4.    msc = ma2gp0 ;
57a5.    msdiff = mst-msc ;
57a6.    m0t = ma1gp1 ;
57a7.    m1t = m0t + mst*1  ;
57a8.    m2t = m0t + mst*mslt2  ;
57a9.    m3t = m0t + mst*mslt3  ;
57a10.   m4t = m0t + mst*mslt4  ;
57a11.   m0c = ma1gp0 ;
57a12.   m1c = m0c + msc*1 ;
57a13.   m2c = m0c + msc*mslc2 ;
57a14.   m3c = m0c + msc*mslc3 ;
57a15.   m4c = m0c + msc*mslc4 ;
57a16.   mdiff0 = m0t - m0c ;
57a17.   mdiff1 = m1t - m1c ;
57a18.   mdiff2 = m2t - m2c ;
57a19.   mdiff3 = m3t - m3c ;
57a20.   mdiff4 = m4t - m4c ;
57b1.  NEW (yst ysc ysdiff y0t y1t y2t y3t y4t y0c y1c
57b2.    y2c y3c y4c ydiff0 ydiff1 ydiff2 ydiff3 ydiff4) ;
57b3.    yst = ya2gp1 + (ma2gp1)*p2gp1 ;
57b4.    ysc = ya2gp0 + ma2gp0*p2gp0 ;
```

```
57b5.     ysdiff = yst-ysc ;
57b6.     y0t = ya1gp1 + ma1gp1*p1gp1;
57b7.     y1t = y0t + yst*1;
57b8.     y2t = y0t + yst*yslt2;
57b9.     y3t = y0t + yst*yslt3;
57b10.    y4t = y0t + yst*yslt4;
57b11.    y0c = ya1gp0 + ma1gp0*p1gp0;
57b12.    y1c = y0c + ysc*1;
57b13.    y2c = y0c + ysc*yslc2;
57b14.    y3c = y0c + ysc*yslc3;
57b15.    y4c = y0c + ysc*yslc4;
57b16.    ydiff0 = y0t - y0c ;
57b17.    ydiff1 = y1t - y1c ;
57b18.    ydiff2 = y2t - y2c ;
57b19.    ydiff3 = y3t - y3c ;
57b20.    ydiff4 = y4t - y4c ;
```

It helps my discussion to plot the predicted Y means across time for the intervention and control conditions. The predicted means are taken from the output for the `MODEL CONSTRAINT` commands. For the treatment group they are notated on the output as Y0T, Y1T, Y2T, Y3T and Y4T. For the control group they are Y0C, Y1C, Y2C, Y3C and Y4C. Here is the plot using the *Temporal line plot* program on my webpage:

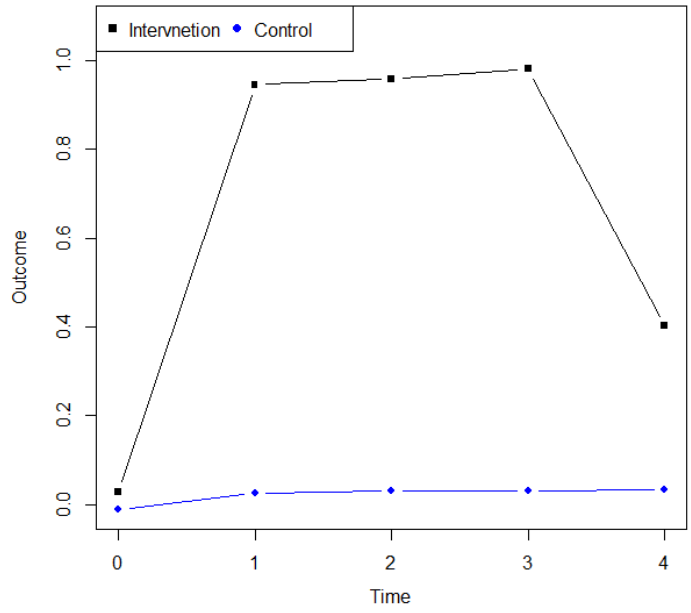

Here is the plot of the actual means rather than the predicted means for the two groups:

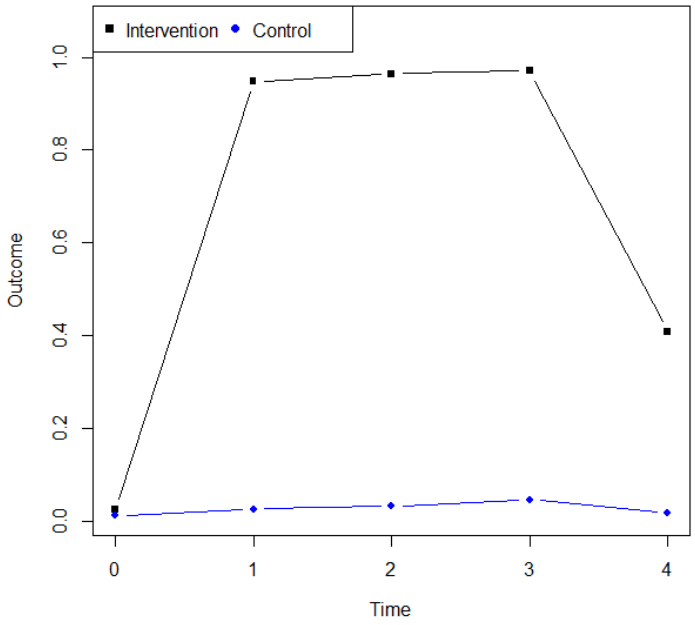

The model captured the pattern of the observed sample means well, which is not surprising given the favorable fit indices that emerged in the model tests.

The temporal patterning of the means is quite distinct for the two groups. In the control condition, the pattern is basically flat and linear with trivial change from baseline through the 18 month follow-up. This is not unexpected given that this group did not receive an intervention. The curve for the intervention group is decidedly non-linear. There is a large positive change in Y from the baseline to the immediate posttest and this likely is due to the intervention. The induced change in Y persists for two follow-up periods and then between the 12 month follow-up and 18 month follow-up, there is a non-trivial decay in Y. Clearly, a linear growth curve for the intervention group is misspecified and this would be true even if I omitted the baseline from the analysis and focused only on the decay curve. The basis growth curve modeling can accommodate this non-linearity. If asked to characterize the "typical" trajectory for the control group using the predicted means for that group, I would characterize it as linear and flat. The "typical" trajectory for the intervention group using the predicted means has an inverted U shape with the induced change by the intervention persisting for a period of about a year and then decaying by about half over the next 6 months. Note that these statements for the intervention group are constrained by having only 4 follow-ups over an 18 month period. More fine-grained statements could be made, for example, if I had 18 monthly follow-up assessments across the 18 month period following the immediate posttest. If describing the more fine-grained curvature is my goal, then my research design should include more frequent assessments. If I am primarily interested in the three six-month follow-up periods in their own right, the current design is fine.

The latent basis growth curve model as implemented has additional noteworthy properties. For the intervention group, the predicted Y baseline mean (for Y0) was 0.029 and the immediate posttest mean (for Y1) was 0.946, a difference of 0.946 – 0.029 = 0.917. Recall that the coefficients for these two time points were fixed at 0 and 1, respectively, on the latent Y basis. I will refer to the difference between the two fixed coefficient time points as the **referent change**. It equaled 0.917. At the first follow-up following the immediate posttest, the predicted Y mean (for Y2) was 0.958. The change from baseline for this follow-up is 0.958 – 0.029 = 0.929, which is close in value to the referent change, namely the change from the baseline to the immediate posttest. Recall that I estimated the coefficient from the latent basis factor to Y2 in the model. I will refer to this coefficient as the basis coefficient for Y2. As you will see below, it equaled 1.014. It turns out that the value of this coefficient reflects a comparison of the change in Y from the baseline to Y2 relative to the referent change, namely the effect of the intervention at the immediate posttest is as follows:

Basis coefficient for Y2 = (Y2 – Y0) / (Y1 – Y0) = 0.929 / 0.917 = 1.014

If there is no decay in change from baseline to the 6 month follow-up relative to the intervention effect at the immediate posttest (as indexed by the referent change), then the basis coefficient for Y2 will equal 1.0. If the basis coefficient for Y2 equals 0.50, this means that the change from baseline to Y2 is half the size of the referent change. If the basis coefficient for Y2 equals 0.33, this means that the change from baseline to Y2 is a third the size of the referent change. If the basis coefficient for Y2 equals 2.0, this means that the change from baseline to Y2 is twice as large as the size of the referent change, which would indicate accelerating change rather decelerating change. And so on.

When I learn that the basis coefficient for Y2 is 1.014, I can deduce from this value that there has been little decay in the initial change created by the intervention. The basis coefficient for Y3 for the intervention group was 1.039. This also indicates that by the 12 month follow-up, there is still relatively little decay in change relative to the initial change created by the intervention. The basis coefficient for Y4 was 0.407. This value suggest that by 18 months, the amount of change in Y was less than half of the initial change created by the intervention, or, more technically it is about four tenths of the initial change. If you examine the curvature in the temporal plots presented earlier, you will see that these characterizations comport well with what you see graphically.

For the control group, the three estimated basis coefficients were Y2 = 1.184, Y3 = 1.146, and Y4 = 1.214. These values suggest that the change from baseline to Y1 (which was near 0) remained that way across the different follow-up assessments.

If I had not fixed the first two time points at 0 and 1 but instead used a different coding scheme, the above properties would be lost. This is why I included the baseline assessment in the growth curve analysis and why I used the coding scheme I did. For randomized trials with baseline assessments, the coding strategy I used is often substantively meaningful.

With this as background, let's turn to answering our fundamental questions of an RET.

Total Effect of the Treatment on the Outcome in the Basis Multi-Group Analysis

To test the effects of the intervention on the outcome, I first examine treatment condition differences for Y at the immediate posttest. I obtain this contrast and the separate means for the intervention and control groups from the `MODEL CONSTRAINT` commands. Here is the relevant (edited) output:

```
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

New/Additional Parameters
    YDIFF1             0.921      0.071     12.886      0.000
    Y1T                0.946      0.054     17.441      0.000
    Y1C                0.025      0.047      0.535      0.592
```

The adjusted mean difference is 0.92 ±0.14, which is statistically significant (CR = 12.89, $p < 0.05$). The 95% confidence interval is 0.78 to 1.06. Suppose the research team decided *a priori* that a meaningful population mean difference at the immediate posttest is a population absolute difference of 0.20 or greater. Because the lower limit of the mean difference confidence interval was greater than this standard, I conclude that the intervention effect for the immediate posttest is meaningful. The estimated Y means for the intervention and control conditions at the immediate posttest were 0.95 ± 0.11 and 0.03 ± 0.09, respectively.

Next, I examined a second facet of the effect of the intervention on the outcome, namely whether there was decay in the intervention effect at the follow-ups. For linear growth curves, this analysis was straightforward because the slopes characterizing decay have a single numerical index reflected by the latent slope factor. For complex non-linear functions, a single index is not available. I plotted the curves for the intervention and control groups earlier in which the decay trends were evident for the two group. The separation between lines at a given time reflects the predicted mean difference between the intervention and control groups. The contrasts I specified in the `MODEL CONSTRAINT` commands include tests of the outcome mean difference between the intervention and control groups at each of the follow-up time points using the model-informed predicted means. The contrasts are labeled `YDIFF2`, `YDIFF3`, and `YDIFF4` for times 2, 3 and 4, respectively. Here is the output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    YDIFF2             0.927      0.073     12.629      0.000
    YDIFF3             0.951      0.073     12.997      0.000
    YDIFF4             0.369      0.070      5.280      0.000
```

The intervention group predicted mean is statistically significantly larger than the control group mean at all three follow-ups. The 95% confidence interval for `YDIFF2` is 0.78 to 1.07, for `YDIFF3` it is 0.81 to 1.10 and for `YDIFF4` it is 0.23 to 0.51. Using a meaningfulness standard of 0.20, in all cases the lower limit of the interval exceeds the meaningfulness standard. Given this, I conclude the Y predicted mean differences between the intervention and control groups are meaningful at each follow-up.

Additional insights can be gained by examining the estimated loadings/coefficients linking the latent basis factor to the indicators Y2, Y3, and Y4 for the intervention group. Here is the relevant (edited) output:

```
MODEL RESULTS

                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value
 Group TREAT

 YS       BY
    Y0                0.000      0.000    999.000    999.000
    Y1                1.000      0.000    999.000    999.000
    Y2                1.014      0.048     21.080      0.000
    Y3                1.039      0.049     21.346      0.000
    Y4                0.407      0.041      9.874      0.000
```

The basis coefficient for Y4 was 0.41 ±0.08. This indicates that the Y change from baseline to the 18th month follow up was equal to about 40% of the initial change that the intervention created between Y0 and Y1.

Effect of the Treatment on the Mediator in the Basis Multi-Group Analysis

To examine the effect of the intervention on the mediator, I use the same approach as that for evaluating the total effect of the intervention on Y. It helps to orient my discussion if I plot the predicted mediator means as a function of time. The predicted means are taken from the output for the `MODEL CONSTRAINT` commands. For the treatment group they are notated on the output as M0T, M1T, M2T, M3T and M4T and for the control group they are M0C, M1C, M2C, M3C and M4C. Here is the plot using the line plot program on my webpage:

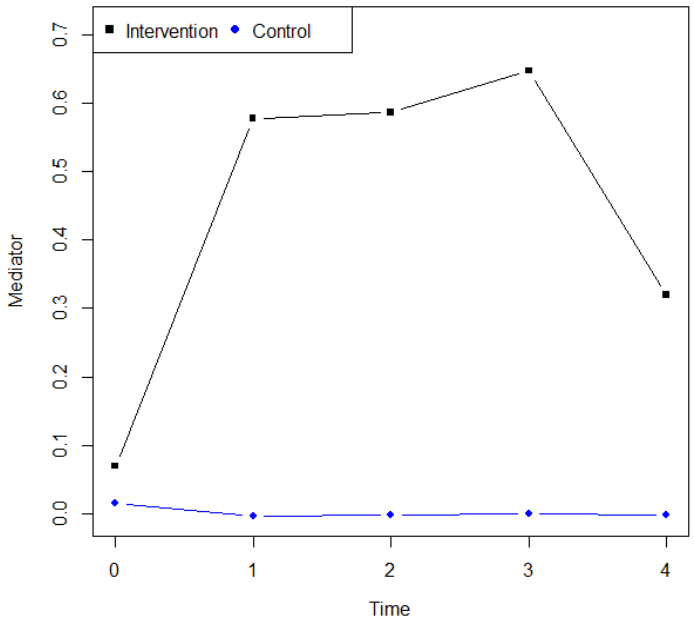

The temporal patterning of the predicted mediator means is similar to that for the corresponding predicted outcome means. In the control condition, the pattern is basically flat and linear with trivial change from baseline through the 18 month follow-up. The curve for the intervention group is again decidedly non-linear. There is a large positive change in M from the baseline to the immediate posttest. The induced change in M persists for two follow-up periods and then between the 12 month follow-up and 18 month follow-up, there is a non-trivial decay in M.

I first examine treatment condition differences for M at the immediate posttest. I obtain this contrast and the separate means for the intervention and control groups from the `MODEL CONSTRAINT` commands. Here is the relevant (edited) output:

```
                                                    Two-Tailed
                   Estimate       S.E.  Est./S.E.    P-Value

New/Additional Parameters
   MDIFF1             0.582      0.064      9.106      0.000
   M1T                0.578      0.047     12.410      0.000
   M1C               -0.004      0.044     -0.087      0.930
```

The adjusted mean difference is 0.58 ±0.13, which is statistically significant (CR = 9.11, $p < 0.05$). The 95% confidence interval is 0.45 to 0.71. Suppose the research team decided *a priori* that a meaningful population mean difference at the immediate posttest is a population absolute difference of 0.20 or greater. The lower limit of the mean difference confidence interval was greater than this standard, so I conclude the intervention effect at

the immediate posttest is meaningful. The estimated mediator means for the intervention and control conditions at the immediate posttest were 0.58 ± 0.09 and -0.004 ± 0.09.

Next, I examined a second facet of the effect of the intervention on the mediator, namely whether there was meaningful decay in the intervention effect over time. In addition to the prior plot, I examined the contrasts on the `MODEL CONSTRAINT` commands for the tests of the mediator mean difference between the intervention and control conditions at each of the follow-up time points using the model-informed predicted means. The contrasts are labeled `MDIFF2`, `MDIFF3`, and `MIFF4` for times 2, 3 and 4, respectively. Here is the output:

```
                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

    MDIFF2           0.589      0.060      9.731      0.000
    MDIFF3           0.648      0.060     10.720      0.000
    MDIFF4           0.322      0.060      5.390      0.000
```

The intervention group mediator predicted mean is statistically significantly larger than the control group mean at all three of the follow-ups. The 95% confidence interval for `YDIFF2` is 0.47 to 0.71, for `YDIFF3` it is 0.53 to 0.77 and for `YDIFF4` it is 0.20 to 0.44. Using a meaningfulness standard of 0.20, in all cases the lower limit of the interval exceeds the meaningfulness standard (although they are quite close for the 18 month follow-up). Given this, I conclude the mediator predicted mean differences between the intervention and control groups are meaningful at each follow-up.

Additional insights can be gained by examining the estimated loadings/coefficients linking the latent basis factor to the indicators Y2, Y3, and Y4 for the intervention group. Here is the relevant (edited) output:

```
MODEL RESULTS

                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value
 Group TREAT

 MS       BY
    M0             0.000      0.000    999.000    999.000
    M1             1.000      0.000    999.000    999.000
    M2             1.018      0.080     12.730      0.000
    M3             1.137      0.085     13.339      0.000
    M4             0.493      0.070      7.032      0.000
```

The basis coefficient for M4 was 0.49 ±0.14. This indicates that the mediator change from baseline to the 18th month follow up was equal to about 49% of the initial change that the intervention created between M0 and M1.

Effect of the Mediator on the Outcome in the Basis Multi-Group Analysis

As noted in the prior multi-group example, multi-group latent growth analysis estimates the effect of the mediator on the outcome separately for the intervention and control groups. For basis modeling, there are two facets of the analysis that provide perspectives on the M→Y link, (1) the path coefficient that links the latent mediator intercept factor to the latent outcome intercept factor, and (2) the path coefficient that links the latent mediator basis factor to the latent outcome basis factor. The former is a straightforward analysis regressing the baseline outcome onto the baseline mediator per traditional growth curve analysis. In this case, it estimates the M→Y coefficient before the intervention has been introduced. For the latter and for the particular coding scheme I used (fixing the basis loadings for the first two time points to 0 and 1, respectively), the M→Y coefficient is analogous to the coefficient one would obtain when regressing the Y1-Y0 difference scores onto the M1-M0 difference scores (Wu & Lang, 2016). In the current example, we expect this coefficient to be positive. Here are the relevant coefficients taken from the Mplus output:

```
MODEL RESULTS

                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

Group CONTROL

 YI ON  MI             0.635      0.075      8.439      0.000
 YS ON  MS             0.626      0.179      3.488      0.000

Group TREAT

 YI ON   MI            0.709      0.062     11.431      0.000
 YS ON   MS            0.621      0.156      3.982      0.000
```

The coefficient for MI→YI was statistically significant and meaningful in both the intervention and control conditions. This also was true for the coefficient for MS→YS. These results are supportive of a link between the mediator and the outcome.

Omnibus Mediation Analysis

Omnibus mediation tests are of lower priority for me. I instead prefer to focus on the individual links in mediational chains, evaluating their effects sizes. If I want an overall omnibus test of statistical significance, I tend to rely on the joint significance test. In basis LGC models one isolates results that evaluate the T→M link and results that evaluate the M→Y. If both these links are statistically significant and meaningful, then one concludes

for omnibus mediation. In the current example, the results were consistent with both links being viable, which provides a sense of what is operating at an omnibus level.

Shortcomings of Latent Basis Growth Curve Modeling

Growth modeling has many strengths but I personally am reluctant to use it. The main limitations of the approach include a frequent problems with convergence and the need for at least four time points to estimate the growth parameters. Also problematic is the issue of overfitting, which can reduce generalizability of the results. Consider the plot of the mediator predicted means, which I reproduce here by I add a dashed red line to capture the linear trend of the means at Times 1, 2 and 3.

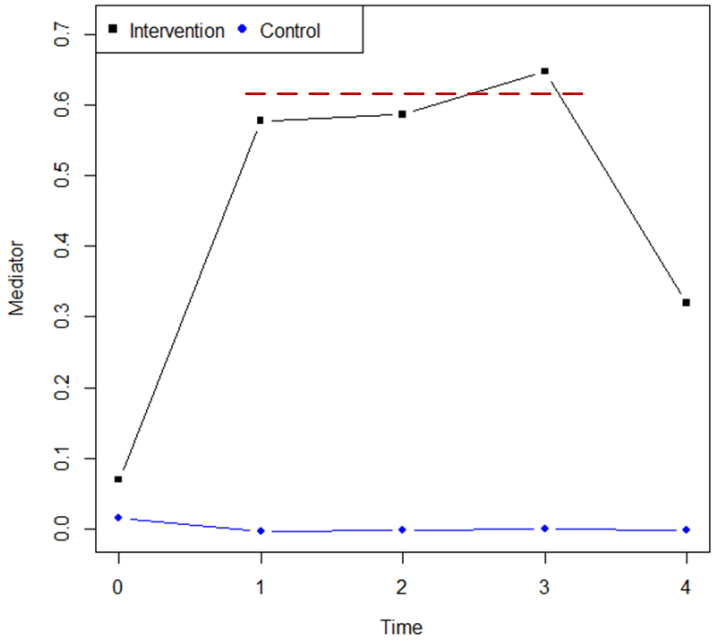

The three means show a distinct linear trend with deviations from the line likely reflecting noise or random error. The basis model actively models these deviations to better reproduce the data. My own preference is to describe this trend as being flat and linear. Indeed, as I pointed out in the main chapter, some argue that the primary goal of growth modeling is to summarize a set of repeated measures with a smoothed underlying function (such as a linear function) that parsimoniously represents how a variable changes over time. Basis modeling, in some respects, abandons this goal. If your interest is not in describing the mathematical functions of trajectories but instead describing changes that occur at specific time points, then there are likely better ways of doing so than basis growth modeling (see, for example, the SEM-based fixed effects analyses described in the main text).

Compared to standard linear latent growth curve models or the non-linear models I discuss below, the basis growth model often has lower statistical power for the types of contrasts that frequently are of interest in RETs.

Another shortcoming of latent basis growth modeling is its reliance on a property known as proportionality in the form of what is known as the **proportionality assumption** (Wu & Lang, 2016). The assumption is that the proportion of change between any given time interval relative to the referent change is constant across individuals. The assumption derives from the fact that although the estimated basis coefficients are allowed to vary across time, the loadings cannot vary across individuals. If the change from time 0 to time 1 is the referent change and the basis coefficient at, say, time 3 is 0.75, then for all individuals, the change from time 0 to time 3 must be 75% of the referent change even though there may be variability in this value (0.75) across individuals. As long as there is not much between-individual variation in this proportion, the latent basis growth model will yield reasonable estimates. However, if there is substantial variability in the proportion, then this introduces bias into the estimates, namely estimates of mean changes, variances of those changes, and covariate effects on those changes (see Wu & Lang, 2016).

McNeish (2020) developed a multi-level strategy using Bayes estimation (see Chapters 8 and 25) to evaluate the degree of variability in each of the estimated basis coefficients. His method allows one to relax the proportionality assumption such that one can then gain perspectives on whether the above latent basis growth curve model is likely to produce biased estimates by assuming proportionality. Table 5 presents the relevant Mplus syntax as applied to the data on which I conducted the more traditional latent basis growth curve model but just for the intervention group. I then perform a second analysis for just the control group. I use this separate group strategy because Mplus does not directly support multiple-group two-level modeling for Bayes estimation. Despite this, the separate group analysis provides me with sufficient information to gain a sense of violations of the proportionality assumption.

Table 5. Mplus Syntax for Bayesian Basis Growth to Adress Proportionality

```
1.  TITLE: LGM Test run for RET example ;
2.  DATA: FILE = LGAdata4M.dat;
3.  VARIABLE:
4.  NAMES = id Y0 Y1 Y2 Y3 Y4 M0 M1 M2 M3 M4 DTREAT ;
5.  USEVARIABLES ARE Y0 Y1 Y2 Y3 Y4 M0 M1 M2 M3 M4 ;
6.  USEOBSERVATIONS DTREAT EQ 1 ;   !select only the intervention group
7.  CLUSTER = id;
8.  WITHIN=y0-y4 m0-m4;
9.  ANALYSIS:
10. ESTIMATOR = BAYES;
```

```
11. !ESTIMATOR = BAYES; BITERATIONS=100000 (50000); BCONVERGENCE =.01;
12. TYPE= TWOLEVEL RANDOM;
13. MODEL:
14. %WITHIN%
15.   yint BY y0@1 y1@1 y2@1 y3@1 y4@1;
16.   ybasis BY y0@0 y1@1;
17.   yf1| ybasis BY y2;
18.   yf2| ybasis BY y3;
19.   yf3| ybasis BY y4;
20.   yint; ybasis;
21.   yint WITH ybasis;
22.   y0-y4;
23.   [yint];[ybasis];
24.   [y0-y4@0];
25.   mint BY m0@1 m1@1 m2@1 m3@1 m4@1;
26.   mbasis BY m0@0 m1@1;
27.   mf1| mbasis BY m2;
28.   mf2| mbasis BY m3;
29.   mf3| mbasis BY m4;
30.   mint; mbasis;
31.   mint WITH mbasis;
32.   m0-m4;
33.   [mint];[mbasis];
34.   [m0-m4@0];
35.   ybasis ON mbasis ;
36.   yint ON mint ;
37. %BETWEEN%
38.   yf1;yf2;yf3;
39.   [yf1];[yf2];[yf3];
40.   mf1;mf2;mf3;
41.   [mf1];[mf2];[mf3];
42. OUTPUT: RESIDUAL CINTERVAL(HPD) TECH8 TECH4 ;
43. PLOT: TYPE = PLOT2;
```

To fully understand this syntax, you need to master MSEM programming in Mplus as applied to clustered RETs. I cover this in Chapter 25. You also need to understand Bayes programming which I cover in Chapter 8. I assume that you are familiar with the relevant material from those chapters. I make use of the Mplus defaults for Bayes modeling. You can use Line 11 instead of Line 10 by commenting out Line 10 and removing the comment demarcation from Line 11 if you want to override some of the Mplus defaults to conduct a more accurate but computer-intense variant. However, if you do, prepare yourself for slow computer execution. I find for current purposes, Line 10 usually suffices but in this case, Mplus recommended on the initial run that I increase the number of `BITERATIONS` to at least 27000 suggesting I should use Line 11. Lines 17-19 define the three basis loadings for the outcome as random rather than fixed, thereby permitting across-individual variability in

them. The same is true for Lines 27-29 for the mediator. I remove `SAMP` and `MOD(4)` from the output line because these options are not allowed with Bayesian models. I add the option `TECH8` to the output line, which then produces the PSR and convergence statistics, which are important for Bayesian modeling. I also change the `CINTERVAL` option on the `OUTPUT` line to read `CINTERVAL(HPD)` to obtain asymmetric credible intervals. Finally, on Line 43. I add the `PLOT` line to generate plots, as discussed in Chapter 8. The remaining syntax should be self-explanatory assuming you have familiarized yourself with Chapters 8 and 25.

In the interest of space, I do not review all of the Bayesian output and fit statistics. See Chapter 8 for the statistics to examine. Model convergence and model fit both were satisfactory. Here is the output that reports the basis coefficients and their variances:

```
Between Level
```

		Posterior	One-Tailed	95% C.I.		
	Estimate	S.D.	P-Value	Lower 2.5%	Upper 2.5%	Significance
Means						
YF1	1.033	0.048	0.000	0.947	1.132	*
YF2	1.057	0.055	0.000	0.959	1.160	*
YF3	0.411	0.042	0.000	0.340	0.500	*
MF1	1.002	0.072	0.000	0.883	1.142	*
MF2	1.120	0.078	0.000	0.999	1.274	*
MF3	0.470	0.072	0.000	0.331	0.609	*
Variances						
YF1	0.106	0.066	0.000	0.018	0.242	*
YF2	0.028	0.031	0.000	0.004	0.104	*
YF3	0.054	0.055	0.000	0.006	0.189	*
MF1	0.047	0.065	0.000	0.002	0.205	*
MF2	0.040	0.080	0.000	0.002	0.244	*
MF3	0.355	0.171	0.000	0.093	0.696	*

The mean basis coefficient estimates for the outcome growth curve model were 1.03, 1.06 and 0.41 as compared to the original analysis estimates of 1.01, 1.04, and 0.41. The two sets are quite close. The variances are fairly small. The square root of the variance entries are estimates of the across-individual standard deviations of the basis coefficients. The same is true for the mean basis coefficient estimates for the mediator growth curve model. The three coefficients were 1.01, 1.14 and 0.49 as compared to the original analysis estimates of 1.01, 1.04, and 0.41. Although the variance estimates all were notated by Mplus as being "statistically significant, this simply means that the population variances of a given basis coefficient was not exactly zero. The small values of the variances and the yet smaller values of the 95% credible intervals associated with them, suggest the across-individual variability in the basis coefficients is not large. Taken together, these results give me some confidence that violations of the proportionality assumption are not problematic for my original analysis. I found comparable correspondence for the analysis of the control group.

Concluding Comments on Latent Basis Growth Curve Modeling

Latent basis growth modeling is a viable strategy for dealing with non-linear growth functions but I tend to shy away from it. It is sample size demanding, can have convergence problems, often has lower statistical power compared to other non-linear alternatives, does not yield the kind of smooth curves that parsimoniously characterize data trends, and it makes strong assumptions about proportionality. Nevertheless, there might be situations where it will be of use. It is a reasonable tool to have in your analytic tool box.

Piecewise LGCM

Another popular method for modeling non-linear growth functions is **piecewise latent growth curve analysis**. This strategy adapts spline regression to LGCM, which I introduced in Chapters 6 and 15. Consider a two group (intervention versus control) RET in which I obtain a baseline measure of a mediator, M, and an outcome, Y, as well as at an immediate posttest, and 5 follow-up assessments at 3 month intervals after the immediate posttest, i.e., at 3, 6, 9, 12, and 15 months afterwards. The outcome measures are signified as Y0, Y1, Y2, Y3, Y4, Y5 and Y6 for the baseline and 6 post-intervention assessments, with comparable notation for the mediator but substituting M for Y. The metrics range roughly from -2 to +2 with standard deviations of approximately 1.0. Higher scores are desired on both M and Y. Here is a line plot of the *observed* outcome means at the different time points for the intervention and control groups:

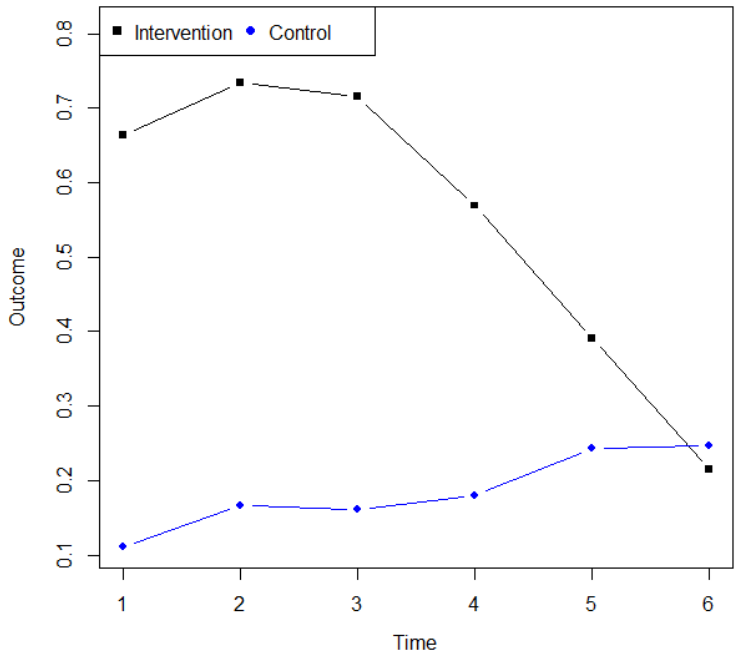

The (decay) curve for the control group is fundamentally linear and relatively flat but with a slight upward trend. The curve for the intervention group is decidedly non-linear and shows intervention persistence through the third time point after which the intervention effect appears to decay linearly.

Piecewise LGCM divides a non-linear curve into segments in such a way that the within-segment data can be reasonably characterized via linear relationships. One then calculates the slopes or path coefficients within each segment. This strategy has the advantage of allowing one stay in the familiar territory of linear models when characterizing non-linear relationships. Here is the prior plot but where I break the curve into two segments using a change knot at Time 3, a point where meaningful change in the curve is clearly defined for the intervention group but only vaguely so for the control group:

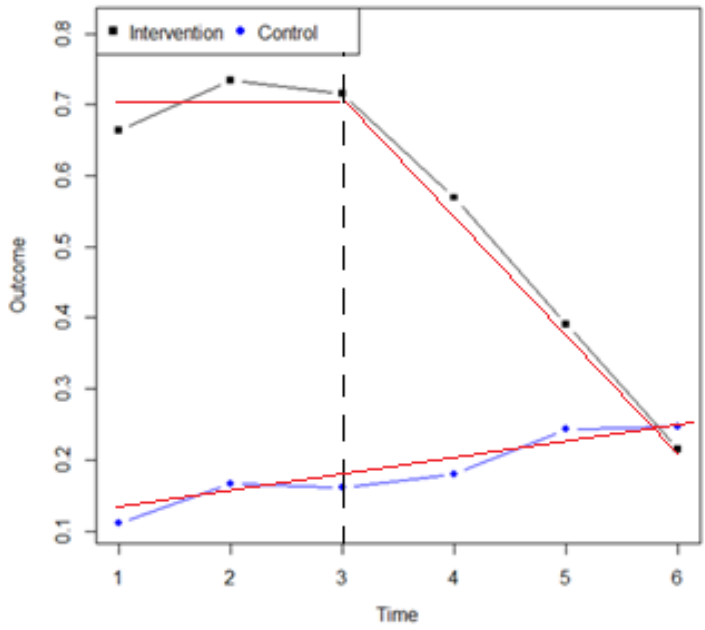

Piecewise growth curve modeling estimates the relevant intercepts and slopes in the segments for the intervention group and for the control group. I then use these estimates and other facets of the analysis to address our three standard RET questions.

Figure 5 presents the influence diagram for data I analyze. I omit disturbances and correlations to reduce clutter. The analysis uses a parallel process framework. I include the coefficient values from the latent intercept and latent slopes to the observed variables.

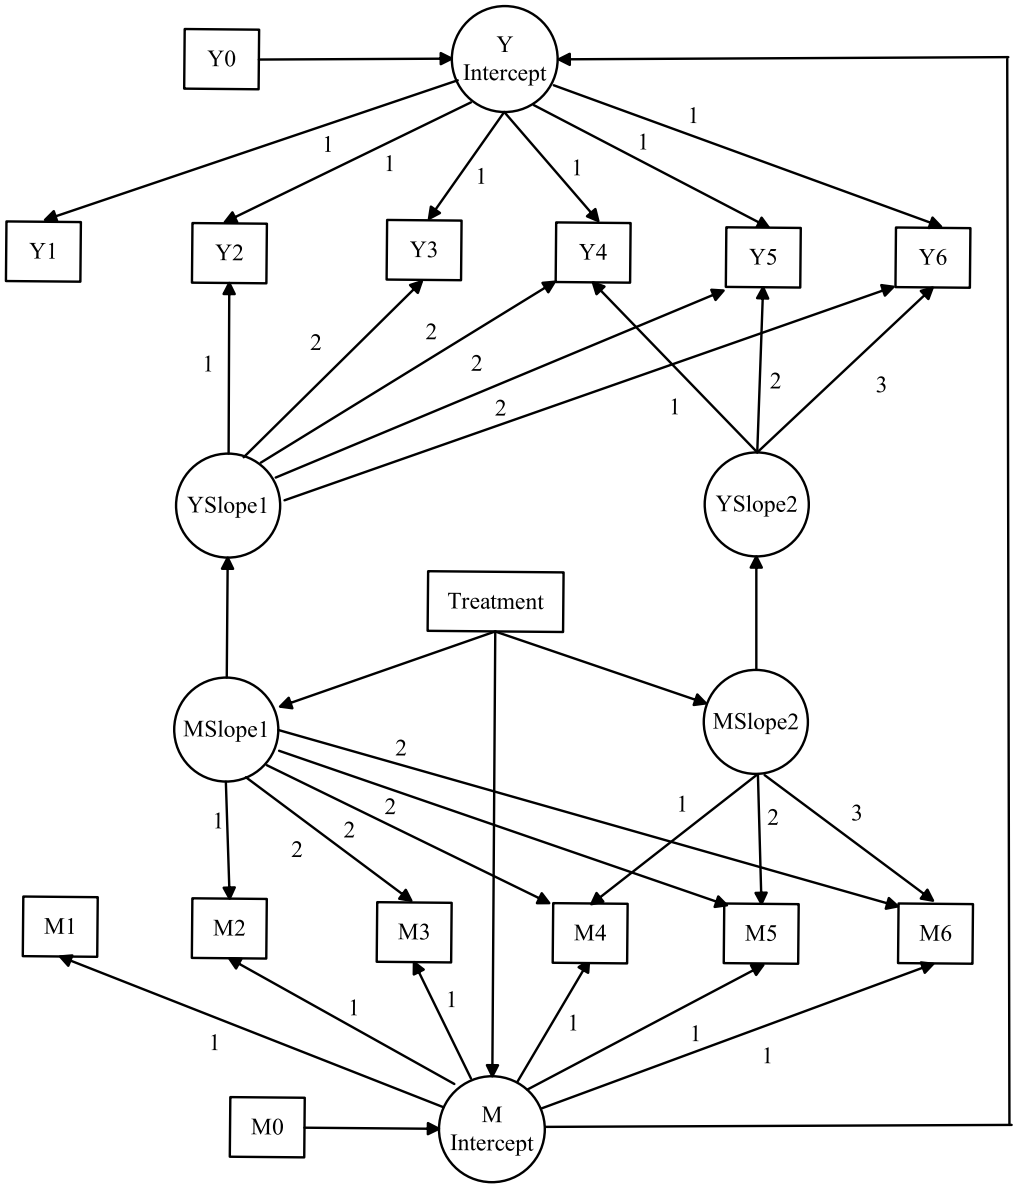

FIGURE 5. RET piecewise growth curve model

There are several noteworthy properties in the Figure. First, despite the fact that I have two latent slopes (one for each segment), there is only one latent intercept. This is because the two segment lines characterized by the slopes are joined together at the knot at Time 3. In this sense, there is no unique intercept for the line/slope in segment 2. If I thought there was a "break" at the start of the second segment such that the line characterizing segment 2 was abruptly higher or lower than the endpoint of the first segment line, then I would model this "break" by introducing a separate intercept for segment 2. This is not the case here.

For the slope factor for the first segment, Figure 5 specifies the standard linear

coefficients of 0, 1, and 2 up until I reach the knot at time 3. After that point, I assign the value of 2, namely the last value I assigned at the end of segment 1, to all subsequent coefficients. For the second slope factor, I assign all zero values to `y1`, `y2` and `y3` in the first segment, and then I assign the linear coefficients of 1, 2, and 3 to `y4`, `y5`, and `y6` thereafter. This approach ensures that `y1`, `y2`, and `y3` contribute in a linear fashion to the first segment and `y4`, `y5` , and `y6` contribute in a linear fashion to the second segment. There is an interesting property to the coefficients for the two segments. The first segment used the coefficients 0, 1, 2, 2, 2, 2 and the second segment used the coefficients 0, 0, 0, 1, 2, 3. If I add the separate cells of these two vectors together, I obtain the usual 0, 1, 2, 3, 4, 5 coefficient values for a linear growth curve. In this sense, the two slope factors can be thought of as a form of decomposition of the overall growth trajectory (Newsom, 2023).

Table 6 presents the initial Mplus syntax I execute for overall model evaluation without using the `MODEL CONSTRAINT` commands.

Table 6. Mplus Syntax for Piecewise Latent Growth Curve Model

```
1.  TITLE: LGM piecewise analysis ;
2.  DATA: FILE = LGAdata5M.dat;
3.  DEFINE:
4.    CENTER m0 y0 (GRANDMEAN) ;
5.  VARIABLE:
6.    NAMES = Y1 Y2 Y3 Y4 Y5 Y6 M1 M2 M3 M4 M5 M6
7.            Y0 M0 TREAT ;
8.  ANALYSIS: ESTIMATOR=MLR ;
9.  MODEL:
10.   yi ys1 | y1@0 y2@1 y3@2 y4@2 y5@2 y6@2 ;  !y segment 1
11.   yi ys2 | y1@0 y2@0 y3@0 y4@1 y5@2 y6@3 ;  !y segment 2
12.   [yi]  (ayi) ;      !intercept for yi
13.   [ys1] (ays1) ;     !slope 1
14.   [ys2] (ays2) ;     !slope 2
15.   yi;                !variance of y intercept
16.   ys1;               !variance of y slope
17.   ys2;               !variance of y slope
18.   yi WITH ys1;  !correlations for intercepts and slopes
19.   yi WITH ys2;
20.   ys1 WITH ys2 ;
21.   y1-y6;             !disturbance variances
22.   mi ms1 | m1@0 m2@1 m3@2 m4@2 m5@2 m6@2 ;  !do mediator
23.   mi ms2 | m1@0 m2@0 m3@0 m4@1 m5@2 m6@3 ;
24.   [mi]   (ami);
25.   [ms1] (ams1);
26.   [ms2] (ams2);
27.   mi;
28.   ms1;
```

```
29.   ms2;
30.   mi WITH ms1;
31.   mi WITH ms2;
32.   ms1 WITH ms2 ;
33.   m1-m6;
34.   mi ON treat m0 (p1 b0) ;  !do the regressions
35.   ms1 ON treat (p2) ;
36.   ms2 ON treat (p3) ;
37.   yi ON mi y0 (p4 b1) ;
38.   ys1 ON ms1 (p5) ;
39.   ys2 ON ms2 (p6) ;
40. MODEL INDIRECT: !examine treatment differences
41.   yi IND treat ;
42.   mi IND treat ;
43.   ys1 IND treat ;
44.   ys2 IND treat ;
45.   ms1 IND treat ;
46.   ms2 IND treat ;
47.   y1 IND treat ;
48.   y2 IND treat ;
49.   y3 IND treat ;
50.   y4 IND treat ;
51.   y5 IND treat ;
52.   y6 IND treat ;
53.   m1 IND treat ;
54.   m2 IND treat ;
55.   m3 IND treat ;
56.   m4 IND treat ;
57.   m5 IND treat ;
58.   m6 IND treat ;
59. OUTPUT: Samp StdYX MOD(4) Residual Tech4  ;
60. !OUTPUT: Samp StdYX Residual Tech4  ;
```

Most of the syntax should be familiar based on the previous sections of this document. Lines 10 and 11 specify the elements of the two separate growth curves for each segment of the outcome variable. Line 10 is for segment 1 and line 11 is for segment 2. Note that the intercept has the same label in each command, namely `yi`. This is because there is one intercept, as discussed earlier. If I wanted to introduce a second intercept, for the two segments, I would use different intercept labels on the two lines, such as `yi1` on Line 10 and `yi2` on Line 11. Here is an example plot of a traditional piecewise growth curve model for an outcome for one group:

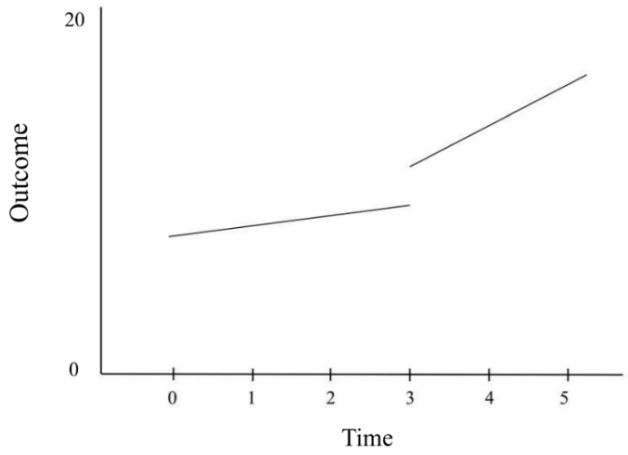

When I executed the syntax in Table 6, the model fit indices suggested good model fit. The chi square test of perfect model fit in the population was statistically non-significant (chi square = 87.13 with 88 degrees of freedom, $p < 0.51$),suggesting the data are in accord with the model. The RMSEA was <0.001. The upper limit of the 90% confidence interval for it is 0.017. The p value for close fit was not statistically significant ($p < 1.00$). The CFI is 1.00 and the standardized RMR was 0.020. For localized fit, there were no theoretically meaningful modification indices greater than 4 and no meaningful residuals between the predicted and observed covariances on a cell-by-cell basis.

After concluding in favor of a reasonable model fit, I re-ran the syntax but added the following `MODEL CONSTRAINT` commands prior to Line 59 to isolate core parameters of interest (I also comment out Line 59 and uncomment Line 60):

```
MODEL CONSTRAINT:
NEW (mic mit ms1c ms2c ms1t ms2t mdiff1 mdiff2 mdiff3 ms2s1c ms2s1t );
mic = ami ;          !M mean of latent intercept for control grp
mit = ami + p1 ;     !M mean of latent intercept for treat grp
ms1c = ams1 ;        !M mean of latent slope 1 for control grp
ms1t = ams1+ p2;     !M mean of latent slope 1 for treat grp
ms2c = ams2 ;        !M mean of latent slope 2 for control grp
ms2t = ams2 + p3;    !M mean of latent slope 2 for treat grp
mdiff1 = mit-mic ;   !M test of differences in above
mdiff2 = ms1t - ms1c ;
mdiff3 = ms2t - ms2c ;
ms2s1c = ms2c - ms1c ;  !M slope 2 - M slope 1 for controls
ms2s1t = ms2t - ms1t ;  !M slope 2 - M slope 1 for treat
NEW (yic yit ys1c ys1t ys2c ys2t ydiff1 ydiff2 ydiff3 ys2s1c ys2s1t) ;
yic = ayi + mic*p4 ;     !Y mean of latent intercept for control grp
yit = ayi + mit*p4 ;     !Y mean of latent intercept for treat grp
ys1c = ays1 + ms1c*p5 ;  !Y mean of latent slope 1 for control grp
ys1t = ays1 + ms1t*p5 ;  !Y mean of latent slope 1 for treat grp
ys2c = ays2 + ms2c*p6 ;  !Y mean of latent slope 2 for control grp
```

```
ys2t = ays2 + ms2t*p6 ;  !Y mean of latent slope 2 for treat grp
ydiff1 = yit - yic ;     !Y test of differences in above
ydiff2 = ys1t - ys1c ;
ydiff3 = ys2t - ys2c ;
ys2s1c = ys2c - ys1c ;   !Y slope 2 – Y slope 1 for controls
ys2s1t = ys2t - ys1t ;   !Y slope 2 – Y slope 1 for treat
NEW (m1c m2c m3c m4c m5c m6c m1t m2t m3t m4t m5t m6t
m1tcdiff m2tcdiff m3tcdiff m4tcdiff m5tcdiff m6tcdiff ) ;
m1c = mic + ms1c*0 ;             !control M est mean at t1
m2c = mic + ms1c*1 ;             !control M est mean at t2
m3c = mic + ms1c*2 ;             !control M est mean at t3
m4c = mic + ms1c*2 + ms2c*1 ;    !control M est mean at t4
m5c = mic + ms1c*2 + ms2c*2 ;    !control M est mean at t5
m6c = mic + ms1c*2 + ms2c*3;     !control M est mean at t6
m1t = mit + ms1t*0 ;             !treat M est mean at t1
m2t = mit + ms1t*1 ;             !treat M est mean at t1
m3t = mit + ms1t*2 ;             !treat M est mean at t1
m4t = mit + ms1t*2 + ms2t*1 ;    !treat M est mean at t1
m5t = mit + ms1t*2 + ms2t*2 ;    !treat M est mean at t1
m6t = mit + ms1t*2 + ms2t*3;     !treat M est mean at t1
m1tcdiff = m1t-m1c ;             !M test of difference in above
m2tcdiff = m2t-m2c ;
m3tcdiff = m3t-m3c ;
m4tcdiff = m4t-m4c ;
m5tcdiff = m5t-m5c ;
m6tcdiff = m6t-m6c ;
NEW (y1c y2c y3c y4c y5c y6c y1t y2t y3t y4t y5t y6t
y1tcdiff y2tcdiff y3tcdiff y4tcdiff y5tcdiff y6tcdiff ) ;
y1c = yic + ys1c*0 ;             !control Y est mean at t1
y2c = yic + ys1c*1 ;             !control Y est mean at t2
y3c = yic + ys1c*2 ;             !control Y est mean at t3
y4c = yic + ys1c*2 + ys2c*1 ;    !control Y est mean at t4
y5c = yic + ys1c*2 + ys2c*2 ;    !control Y est mean at t5
y6c = yic + ys1c*2 + ys2c*3;     !control Y est mean at t6
y1t = yit + ys1t*0 ;             !treat Y est mean at t1
y2t = yit + ys1t*1 ;             !treat Y est mean at t2
y3t = yit + ys1t*2 ;             !treat Y est mean at t3
y4t = yit + ys1t*2 + ys2t*1 ;    !treat Y est mean at t4
y5t = yit + ys1t*2 + ys2t*2 ;    !treat Y est mean at t5
y6t = yit + ys1t*2 + ys2t*3;     !treat Y est mean at t6
y1tcdiff = y1t-y1c ;             !Y test of difference in above
y2tcdiff = y2t-y2c ;
y3tcdiff = y3t-y3c ;
y4tcdiff = y4t-y4c ;
y5tcdiff = y5t-y5c ;
y6tcdiff = y6t-y6c ;
```

The contrast specifications follow the logic outlined in the main text Appendix of Chapter 16 and the prior examples in this document. I now address the three core questions for RET analyses.

Total Effect of the Treatment on the Outcome

To test the effects of the intervention on the outcome, I initially use the output generated by Line 41 (`YI IND TREAT`) from Table 6. The output from this line focuses on the treatment-control outcome predicted mean difference at the immediate posttest:

```
TOTAL, TOTAL INDIRECT, SPECIFIC INDIRECT, AND DIRECT EFFECTS

                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

Effects from TREAT to YI

  Total               0.585      0.043     13.520      0.000
  Total indirect      0.585      0.043     13.520      0.000
```

The row of interest is the one called `Total` underneath the `Effects from TREAT to YI` label. The estimate in this row is the predicted difference between Y at the immediate posttest for the treatment condition minus the control condition. In this case, the adjusted outcome difference is 0.58 ±0.09, which is statistically significant (CR = 13.52, $p < 0.05$). The 95% confidence interval for the difference is 0.49 to 0.67. Suppose that prior to the study, the research team set a standard for a meaningful population outcome mean difference at the immediate posttest as an absolute difference of 0.20 or greater. Because the lower limit of the mean difference confidence interval was greater than this standard, I conclude that the intervention effect for the immediate posttest is meaningful.

If desired, I also can report the predicted means of Y at the immediate posttest for the treatment and control conditions to embellish the above results. These predicted values are labeled `Y1T` and `Y1C` in the `MODEL CONSTRAINT` commands. Here is the edited output:

```
New/Additional Parameters
                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

    YIC               0.107      0.036      2.936      0.003
    YIT               0.692      0.035     19.519      0.000
```

The estimated Y means for the intervention and control conditions at the immediate posttest were 0.69 ±0.07 and 0.11 ±0.07, respectively.

Next, I examine a second facet of the effect of the intervention on the outcome, namely whether there was meaningful decay in the effect of the intervention over time. It helps to put the relevant statistical tests in context if I plot the predicted Y means for the intervention and control groups across the six time points:

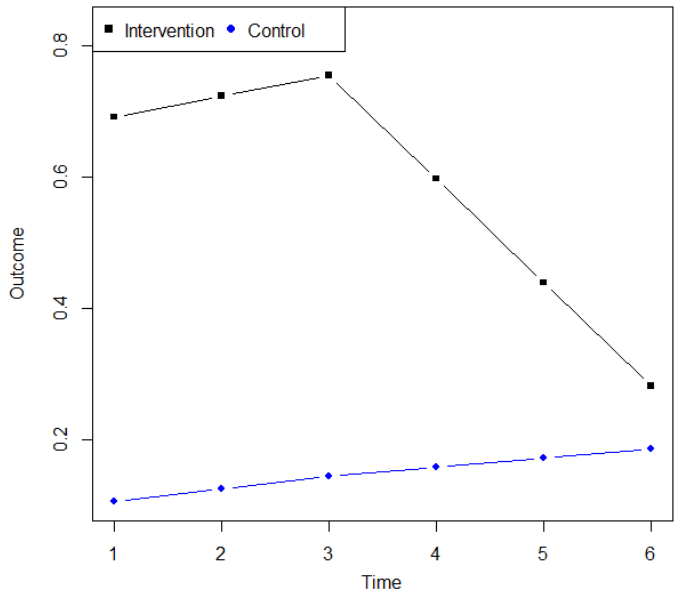

The control group slope is relatively flat across the six time points and this also is true for the first three time points (segment 1) for the intervention condition. For the second segment of the intervention condition, there is a decided linear decay in the predicted Y means.

Line 43 of Table 6 tests the segment 1 slope difference between the intervention and control conditions. Here are the results:

	Estimate	S.E.	Est./S.E.	Two-Tailed P-Value
Effects from TREAT to YS1				
Total	0.013	0.026	0.477	0.634

The difference between the slopes is trivial and statistically nonsignificant (difference = 0.013 ± 0.05, CR = 0.48, *ns*). This is not surprising based on the visual examination of the Y mean line plot. To be sure, the two slopes are widely separated on the Y axis reflecting the immediate effect of the intervention, but the slopes themselves are similarly flat. Here are the values of segment 1 slopes for each group:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    YS1T               0.031      0.026      1.219      0.223
    YS1C               0.019      0.025      0.740      0.459
```

The slopes are both positive indicating slight increases in Y means from Time 1 to Time 3, but the neither slope is statistically significant.

Here are the comparable slope analyses for the second time segment:

```
                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

Effects from TREAT to YS2

  Total             -0.172      0.026     -6.666      0.000

New/Additional Parameters

    YS2T            -0.157      0.023     -6.819      0.000
    YS2C             0.014      0.023      0.626      0.531
```

As evident in the plot, the slope for the second segment for the intervention group is decidedly negative (-0.16 ±0.05) and statistically significant (CR = 6.82, $p < 0.05$). It remains flat and not significant for the control group.

A final detail with respect to the slope segments is to formally test if the slope in segment 2 is significantly different than that in segment 1. Such a difference justifies the piecewise analysis. If the slopes are equivalent, then there is no need to introduce a knot into the analysis. The relevant contrasts occur in the `MODEL CONSTRAINT` section:

```
                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

    YS2S1T          -0.189      0.037     -5.082      0.000
    YS2S1C          -0.004      0.037     -0.118      0.906
```

The segment slope difference for the intervention group was -0.19 ±0.07 (CR = 5.08, $p < 0.05$), which is statistically significant and meaningful. For the control group, the slope segment difference was trivial and statistically non-significant (-0.004 ±0.07 (CR = 0.91, *ns*). Technically, I can omit the knot for the controls but doing so introduces analytic complications for piecewise growth modeling at little gain, so I refrain from doing so.

The contrasts that I specified in the MODEL CONSTRAINT command also test the outcome mean difference at each of the six time points using the model-informed predicted Y means. The contrasts are labeled Y1TCDIFF, Y2TCDIFF, Y3TCDIFF, Y4TCDIFF, Y5TCDIFF, and Y6TCDIFF6 for times 1, 2, 3, 4, 5 and 6, respectively (keep in mind that time 1 is the immediate posttest). Here is the relevant output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    Y1TCDIFF           0.585      0.043     13.520      0.000
    Y2TCDIFF           0.598      0.047     12.796      0.000
    Y3TCDIFF           0.610      0.062      9.811      0.000
    Y4TCDIFF           0.439      0.064      6.808      0.000
    Y5TCDIFF           0.267      0.076      3.519      0.000
    Y6TCDIFF           0.095      0.093      1.024      0.306
```

The intervention group predicted mean is statistically significantly larger than the control group predicted mean at each of the time points except the last one. Suppose a meaningfulness standard is set by the research team to be a population absolute difference of 0.20 or greater. The 95% confidence interval for YDIFF1 is 0.49 to 0.68. Because the lower limit of the confidence interval exceeds the meaningfulness standard, I conclude that the differences at the immediate posttest are meaningful. This same pattern holds for Times 2 through 4. For the fifth time point, the 95% confidence interval is 0.11 to 0.41. The confidence interval overlaps the meaningfulness standard, so I can't confidently conclude the predicted mean difference is meaningful when sampling error is taken into account. To be sure, the difference is non-zero by virtue of the statistical significance ($p < 0.05$) of the contrast. However, I can't say with confidence that the predicted mean difference at time 3 is meaningful.

Finally, the 95% confidence interval at time 6 (15 months after the immediate posttest) was -0.09 to 0.28. This interval overlaps the meaningfulness standard and it also is statistically non-significant. I can't conclude the effect is non-zero nor that it is meaningful.

In sum, the intervention had a meaningful overall effect on the outcome at the immediate posttest. The predicted outcome mean difference for the intervention group minus the control group mean remained relatively stable across the next 9 months after which it showed a steady decline. By follow-up month 15, the predicted difference was not statistically significant. Noteworthy in these analyses is the negative trajectory of -0.16 ±0.05 for segment 2 in the intervention group.

Effect of the Treatment on the Mediator

The test of the effect of the intervention on the mediator follows the same format as the test of the intervention on the outcome, I initially use the output generated by Line 42 (`MI IND TREAT`) from Table 6. The output from this line focuses on the treatment-control mediator predicted mean difference at the immediate posttest:

```
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

Effects from TREAT to MI

  Total                0.839      0.049     16.984      0.000
```

The adjusted predicted mean difference is 0.84 ±0.10, which is statistically significant (CR = 16.98, $p < 0.05$). The 95% confidence interval for the difference is 0.74 to 0.04. Suppose that prior to the study, the research team set a standard for a meaningful population mean difference at the immediate posttest for the mediator as an absolute difference of 0.20 or greater. Because the lower limit of the mean difference confidence interval was greater than this standard, I conclude that the intervention effect for the immediate posttest is meaningful.

If desired, I also can report the predicted means of the mediator at the immediate posttest for the treatment and control conditions to embellish the above results. These predicted values are labeled `M1T` and `M1C` in the `MODEL CONSTRAINT` commands. Here is the edited output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    MIT                0.873      0.036     24.510      0.000
    MIC                0.034      0.035      0.986      0.324
```

The estimated mediator means for the intervention and control conditions at the immediate posttest were 0.87 ±0.07 and 0.03 ±0.07, respectively.

Next, I examine a second facet of the effect of the intervention on the mediator, namely whether there was meaningful decay in the effect of the intervention over time. It helps to put the relevant statistical tests in context if I plot the predicted mediator means for the intervention and control groups across the six time points:

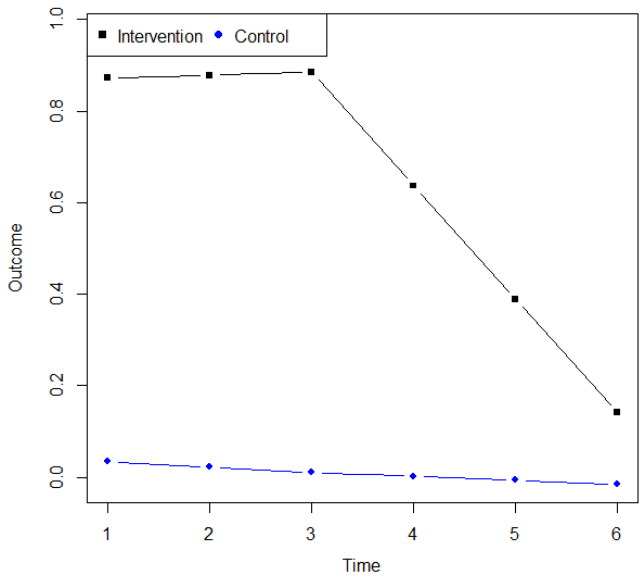

The control group slope is relatively flat across the six time points and this also is true for the first three time points (segment 1) for the intervention condition. For the second segment of the intervention group, there is a decided linear decay in the predicted means.

Line 45 of Table 6 tests the segment 1 slope difference between the intervention and control conditions. Here are the results:

	Estimate	S.E.	Est./S.E.	Two-Tailed P-Value
Effects from TREAT to MS1				
Total	0.017	0.037	0.475	0.635

The difference between the slopes is trivial and statistically nonsignificant (difference = 0.017 ± 0.07, CR = 0.48, *ns*). This is not surprising based on the visual examination of the mediator mean line plot. To be sure, the two slopes are widely separated on the Y axis reflecting the immediate effect of the intervention, but the slopes themselves are similarly flat. Here are the values of segment 1 slopes for each group:

New/Additional Parameters

	Estimate	S.E.	Est./S.E.	Two-Tailed P-Value
MS1T	0.005	0.027	0.200	0.841
MS1C	-0.012	0.027	-0.454	0.650

The slopes are both negligible and not statistically significant. Here are the comparable slope analyses for the second time segment:

```
                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

Effects from TREAT to MS2

 Total              -0.239      0.033     -7.317      0.000

New/Additional Parameters

   MS2T             -0.247      0.024    -10.522      0.000
   MS2C             -0.008      0.023     -0.357      0.721
```

As evident in the plot, the slope for the second segment for the intervention group is decidedly negative (-0.25 ±0.05) and statistically significant (CR = 10.52, $p < 0.05$). It remains flat and not significant for the control group.

A final detail is to formally test if the slope in segment 2 is significantly different than that in segment 1. The relevant contrasts occur in the `MODEL CONSTRAINT` section:

```
                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

   MS2S1T           -0.253      0.039     -6.493      0.000
   MS2S1C            0.004      0.039      0.098      0.922
```

The segment slope difference for the intervention group was -0.25 ±0.08 (CR = 6.49, $p < 0.05$), which is statistically significant and meaningful. For the control group, the slope segment difference was trivial and statistically non-significant (0.004 ±0.08 (CR = 0.98, *ns*). Technically, I can omit the knot for the controls but, as mentioned earlier, doing so introduces analytic complications at little gain.

The contrasts that I specified in the `MODEL CONSTRAINT` command also test the mediator mean difference at each of the six time points using the model-informed predicted mediator means. The contrasts are labeled `M1TCDIFF`, `M2TCDIFF`, `M3TCDIFF`, `M4TCDIFF`, `M5TCDIFF`, and `M6TCDIFF6` for times 1, 2, 3, 4, 5 and 6, respectively. Here is the output:

```
New/Additional Parameters
                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

    M1TCDIFF         0.839      0.049     16.984      0.000
    M2TCDIFF         0.856      0.050     17.251      0.000
```

```
M3TCDIFF           0.873      0.072     12.165      0.000
M4TCDIFF           0.634      0.073      8.678      0.000
M5TCDIFF           0.395      0.088      4.512      0.000
M6TCDIFF           0.156      0.110      1.415      0.157
```

The intervention group predicted mean is statistically significantly larger than the control group predicted mean at each of the time points except the last one. Suppose a meaningfulness standard is set by the research team to be a population absolute difference of 0.20 or greater. The 95% confidence intervals for all of the contrasts with the exception of time 6 yield a lower limit of the 95% confidence interval that exceeds the meaningfulness standard. I conclude that all the differences at these time points are meaningful. The 95% confidence interval at time 6 (15 months after the immediate posttest) was -0.06 to 0.37. This interval overlaps the meaningfulness standard and it also is statistically non-significant. I can't conclude the effect is non-zero nor that it is meaningful.

In sum, the intervention had a meaningful overall effect on the mediator at the immediate posttest. The predicted mediator mean difference for the intervention group minus the control group mean remained relatively stable across the next 9 months after which it showed a steady decline. By follow-up month 15, the predicted difference was not statistically significant. Noteworthy in these analyses is the negative trajectory of -0.25 ±0.08 for segment 2 in the intervention group.

Effect of the Mediator on the Outcome

The analysis of the estimated effect of the mediator on the outcome does not distinguish the intervention and control groups because the guiding model assumes that the effect is the same in both treatment conditions. The relevant coefficients for evaluating the mediator effect on the outcome are the coefficient from `MI` to `YI`, the coefficient from `MS1` to `YS1` for the first slope segment and the coefficient from `MS2` to `YS2`for the second slope segment. Here is the output:

```
MODEL RESULTS

                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

YI ON  MI            0.698      0.043     16.209      0.000
YS1 ON MS1           0.721      0.050     14.568      0.000
YS2 ON MS2           0.717      0.041     17.372      0.000
```

The coefficient for MI→YI is the path coefficient for the outcome at the immediate posttest regressed onto the mediator at the immediate posttest controlling for the baseline outcome. The coefficient was 0.70 ±0.09 (CR =16.21, $p < 0.05$). For every one unit that the mediator

increased across individuals, the outcome was predicted to increase by 0.03 units. Suppose the research team decided prior to the study that a population coefficient of 0.20 or greater would be deemed meaningful. The 95% confidence interval for the coefficient was 0.52 to 0.88. Because the lower limit of this interval is greater than meaningfulness standard, we can declare the result meaningful.

The coefficient for MS1→YS1 is the path coefficient for the segment 1 slope of the outcome regressed onto the segment 1 slope for the mediator. A positive coefficient implies that people with increasingly more positive slopes on the mediator tend to have increasingly more positive slopes on the outcome, i.e., that there is an association between the two. The path coefficient was 0.72 ± 0.10 (CR = 14.57, $p < 0.05$). For every one unit that the mediator segment 1 slope increased across individuals, the outcome segment 1 slope was predicted to increase by 0.72 units. Suppose the research team decided prior to the study that a population coefficient of 0.20 or greater would be deemed meaningful. The 95% confidence interval for the coefficient was 0.66 to 0.88. Because the lower limit of this interval is greater than the meaningfulness standard, we can declare the result meaningful. The corresponding coefficient for the segment 2 slope was 0.72 ± 0.08 (CR = 17.37, $p < 0.05$), which also supports a link between the mediator and the outcome.

In the analyses, I did not include a direct effect of the treatment condition on YI, YS1, or YS2 over and above the effect of the treatment condition on MI, MS1, and MS2. This presumes that such direct effects are functionally zero, which the investigators felt was *a priori* reasonable. When I re-ran the models but included the different direct effects, none of them yielded statistically significant or meaningful path coefficients. This supports the decision to exclude them as did chi square nested model tests.

Omnibus Mediation Analysis

Omnibus mediation tests are of lower priority for me. I instead prefer to focus on the individual links in mediational chains, evaluating their effects sizes. If I want an overall omnibus test of statistical significance, I rely on the joint significance test. In LGC models, evaluations of mediation take two forms. The first focuses on treatment differences in the latent intercept for the mediator and whether these differences, in turn, translate into treatment differences on the outcome latent intercept. The second mediation form focuses on the effect of the latent slopes of the mediator on the latent slopes of the outcome. The RET found meaningful treatment versus control differences for both facets of mediation analysis. The pattern of results is consistent with omnibus mediation.

Concluding Comments on Piecewise LGCMs

When pursuing piecewise growth curve modeling for an RET, you may encounter

convergence issues or cases where you want flexibility to assign different numbers of knots in the treatment and control conditions or different intercept structures. Such matters often can be addressed using the multiple group SEM strategy for growth curve modeling that separates the treatment and control groups, as illustrated earlier in this document. A strength of piecewise LGCM is that it can be applied to a wide variety of non-linear growth curves. Figure 6 presents four examples that I located on the internet. Each panel shows either theoretical curves or data points in a scatter plot that are distinctly non-linear. The lines that define the segments are superimposed, using dashed lines in the top row of the figure and red lines in the bottom row. The segments are segregated with the symbol X. In the top left panel, the researcher was liberal about defining segments (there are 8 segments), whereas in the top right panel, the researcher was less so (there are only 3 segments).

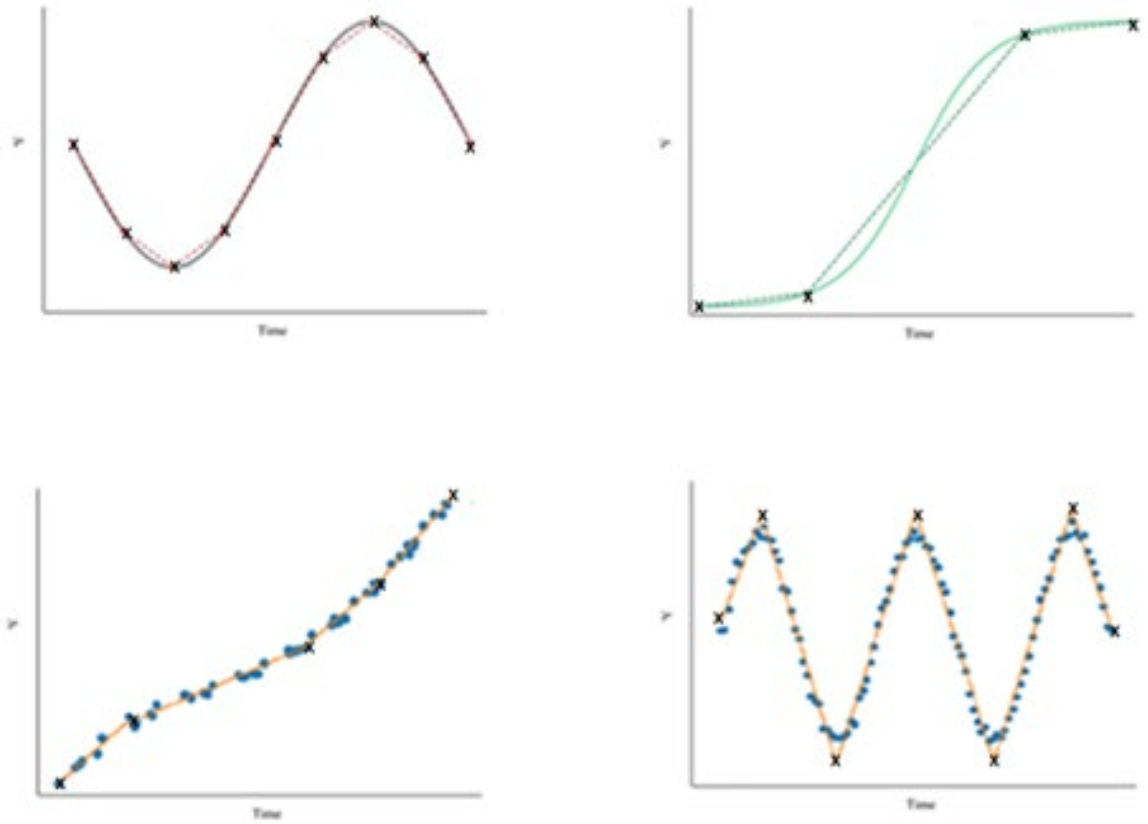

FIGURE 6. Example curves for reciprocal transformation of outcomes

In traditional spline regression, the knots that define segments are either chosen a priori based on theory or they are treated as parameters to be estimated empirically via

various statistical criteria (see Chapter 15). Harring, Strazzeri & Blozis (2021) describe statistical approaches for the empirical estimation of knot parameters for latent growth curve models, including a strategy that allows the knot parameters to vary across individuals. With empirical estimation of the knots, because the results of the piecewise analyses are dependent on knot choice, the p values and confidence intervals for coefficients often are treated as approximate.

Piecewise latent growth curve analysis is straightforward when there are at least three time points within a segment. When a given segment has only two time points, you may encounter identification issues. In such cases, researchers often fix the latent variance for that specific slope to zero and the covariance between the latent intercept and that slope to zero. This removes the under-identification but at the cost of assuming that everyone changes at the same rate during that period. If one of the points in the two-point segment is part of a separate segment with 3 or more points (because it occurs at the beginning or end of the larger segment), then under-identification usually is moot.

Transformation-Based LGCM

The final approach to addressing non-linear growth functions that I consider relies on variable transformations. The idea is to transform either the target variable or the time loadings so that the relationship between the transformed variable and time becomes linear thereby permitting one to apply linear growth curve models to the transformed data. Usually one seeks to transform Y rather than the loadings, although both strategies are viable (Bollen & Curran, 2006). Y transformations that might be used include a log transform (log(Y), an inverse or reciprocal transform (1/Y), a squared transform (Y^2), and a square root transform. A popular transformations is the log transform, which I discussed in some depth in Chapter 15 and that deals with exponential trajectories. I recommend you consult Chapter 15 for the various ways and interpretations of using log transforms on either the outcome, the predictor (which is time in this case)) or both. Examples of the kinds of curve shapes that can be linearized via log transforms of the Y variable are shown in Figure 7.

A strategy for exploring possible Y power transforms of monotonic (as opposed to non-monotonic) changes in Y across time is the Box-Cox transformation algorithm (Box & Cox, 1964; Atkinson, Riani & Corbellini, 2021). This strategy only applies to Y variables that have positive values. If your outcome has zero values or negative values then consider adding a constant, *c*, to the Y scores at all time points to redefine the metric to have positive numbers. The strategy was designed with a broad range of statistical goals in mind but it often helps identify transformations that linearize relationships.

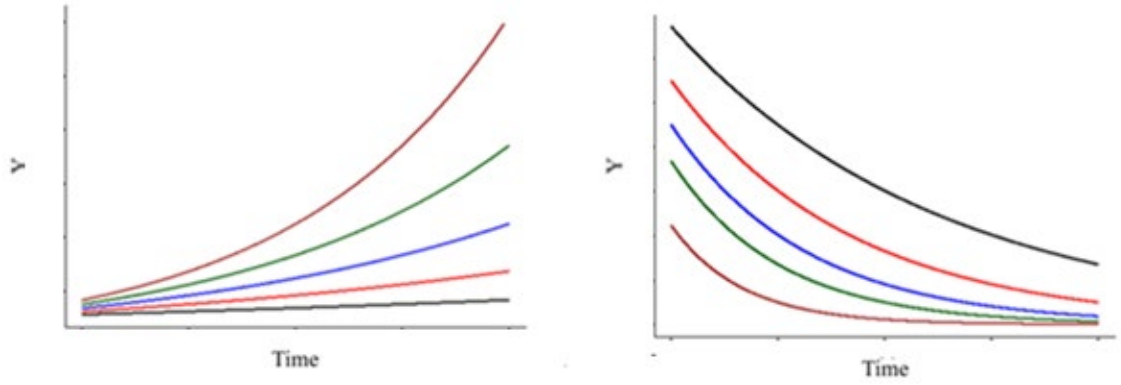

FIGURE 7. Example curves for logged outcomes of target variable

The Box-Cox transformation strategy relies on a power parameter known as lambda. The Y variable is transformed by raising it to the power of lambda which, in turn, stretches or compresses the Y scores. The transformation can linearize the relationship between Y and time, facilitate homoscedasticity, and make Y more normally distributed, but it is not always successful at any of these. The transformation preserves the order of data points, which is why it is not appropriate for non-monotonic functions. The technique assumes the disturbances in the regression of Y on time are uncorrelated or, if they are correlated, inconsequentially so. Latent growth curve modeling typically assumes uncorrelated disturbances for the Y across-time. The most common justification for this is that the time lags between the Y assessments is long enough that correlated disturbances are not expected. Of course, in SEM formulations of growth curve modeling, it is possible to add such correlations to the model if desired.

The Box-Cox formulation suggests a value for lambda to use in one's transformation but it is common for researchers to round lambda to reflect a well-known, established transformation. For example, a lambda value near 2 might suggest a square transformation (Y^2), a lambda value near -1.00 might suggest a reciprocal transformation (Y^{-1}), a lambda value near 0.5 might suggest a square root transformation ($Y^{.5}$), a lambda value near 3 might suggest a cubed transformation (Y^3), a lambda value near zero might suggest a natural log transformation, and a value near 1.00 might suggest not transforming Y. I provide on my website a program called *Box-Cox for LGCM* that facilitates the application of the Box-Cox transformation strategy to latent growth curve modeling. It computes values of lambda to consider.

An important concept when choosing transformation for growth curve modeling is that of **dynamic consistency** (Singer & Willit, 2003). One can summarize "average" growth curves across individuals in one of two ways. First, you can estimate the average

outcome on each measurement occasion and then plot a curve through the averages. Second, you can estimate the growth parameters for each individual and then average these values and derive a plot accordingly. When working with linear growth curves, the two approaches yield the same result. In addition, the average trajectory will have the same functional form as the constituent individual trajectories; the average of a heterogeneous group of straight lines will be a straight line. These characteristics constitute the property of dynamic consistency. It is why when choosing a transformation, we lean towards invoking transformations that linearize the growth curve. The log and reciprocal transforms of the outcome variable have such properties when properly applied. Trajectories that are not dynamically consistent tend to be less tractable.

Parenthetically, the property of dynamic consistency is related to but not the same as that of proportionality in basis latent growth models described earlier. Dynamic consistency is a broader property held by other growth curve models that do not enforce strict proportionality. Proportionality is a specific, restrictive way of achieving dynamic consistency.

Another point to keep in mind when using transformations is that when outcomes are transformed, the coefficients in the transformed equation will not necessarily equal the coefficients derived from pure non-linear strategies that do not rely on transformations. The transformation versus purely non-linear approaches minimize different residuals which can result in different coefficient estimates, although often the differences are not of consequence. Also, log transformations can only be applied to variables that are non-zero and positive. If your scores contain zeros or negative numbers, then it often is recommended to consider adding a small constant to the variable to make the metric such that it is non-zero and positive. Such transformations usually do not affect substantive interpretations.

As discussed in Chapter 15 and as you will see in the RET example, we often calculate the exponent of a mean log(Y) or of a mean log(M) to convert the mean to a value on the variable's original metric. When we do so, the result is indeed a mean on the original metric, but it is a geometric mean not an arithmetic mean. In general, the arithmetic mean adds together all the scores of individuals and then divides the sum by the number of scores, N. A geometric mean multiplies all the scores by one another and then calculates the Nth root of the product. Usually the arithmetic and geometric means will be similar, but not always. Each index has its strengths and weaknesses. Statisticians often prefer geometric means when dealing with rates, ratios, or percentages, as I show below. Note that this is another reason why the variables we work with when using logs cannot be zero in value. In the RET example below, I will use exponents to describe means but will not qualify each mean as reflecting a geometric mean. Keep in mind that this is implied by virtue of my using exponents.

Also note that when I calculate the exponent of the difference between two means of logged variables, such as the mean log(Y) for the intervention group minus the mean log(Y) for the control group, that the exponent of the difference indexes the disparity between the two geometric means on the original metrics of the variable. However, rather than the disparity representing the difference between the geometric means, it represents the ratio of the two means. For example, if the geometric mean of the intervention group is 10 and the geometric mean of the control group is 5, then the exponentiated index when the mean log difference is taken is 10/5 = 2.0; that is the geometric mean for the intervention group is twice as large as the geometric mean for the control group. If the null hypothesis is true and the two geometric means are equal, the ratio will equal 1.0, not 0.

In traditional regression modeling, we assume that for every one unit increase in the predictor, the outcome changes by a constant amount. For example, if I regress annual income onto the number of years of education people have in a given population, the regression coefficient might yield a value of $3,000. This suggests that an additional year of education is worth $3,000 no matter where on the education dimension that increase occurs. A change from 7 years to 8 years of education is predicted to yield an annual income increase of $3,000 as is a change from 13 to 14 years of education. For log based models that reflect exponential relationships, one instead believes that it is the percent change that remains constant across levels of education, not the absolute amount of income change. For example, suppose the percent change in income for a year increase in education is said to be 4%. If the typical income for those with 7 years of education is $15,000, then for those with 8 years of education, the model predicts that the typical income is 15,000*1.04 = $15,600, an increase of $600. If the typical income for those with 13 years of education is $18,980, the model predicts that an additional year of education should increase the typical income to 18,980*1.04 = $19,739, an increase of $759. Note that in contrast to the additive model, the absolute increase differs in the two groups ($600 versus $759) even though the percent increase from one year to the next is the same. The RET example is of this nature.

Another concept we will encounter in the RET example is that of elasticity, which I introduced in Chapter 15. Elasticity as applied to the case of the M→Y link refers to the percent by which the geometric mean of Y changes given a 1% change in the geometric mean value of M. Let p_1 be the path coefficient that results when regressing the log of Y onto the log of M. To calculate the elasticity using a 1% referent change in M, I invoke the rate multiplier value of 1.01 and raise it to the power of the value of p_1. If $p_1$equals 3.3, then this yields $1.01^{3.3} = 1.03$. The result means that for a 1% increase in M, the outcome is predicted to increase by a factor of 1.03. Stated more intuitively, a 1% increase in the value of M leads to a 3% increase in the value of Y. For a 10% increase in M, I use the rate multiplier of 1.10 and raise it to the power of p_1, in this case $1.10^{3.3} = 1.37$. A 10% increase

in the value of M leads to a 37% increase in the value of Y. When a target multiplicative rate of 1.01 or a factor of 1% is used, the result is referred to as an **elasticity**. For more details, see Chapter 15. I use this concept in the RET example.

An RET Example

The RET hypothetical example is an intervention designed to increase trajectories for reading abilities for elementary school children over a two year period. Assessments were taken at 6 month intervals, at baseline and at four post-intervention time points. The baseline measure was signified as y0, with the four follow-ups being y1, y2, y3 and y4. The baseline measure was used as a covariate and the time metric starting with the immediate posttest was quantified in the metric of six month intervals. Given the time it takes for the intervention to take hold, the assessment y1 was obtained 6 months after the intervention was complete and was assigned a slope loading of 0. The y2 assessment was obtained half a school year later and was assigned a slope loading of 0.5. The y3 assessment was obtained half a school year after that and was assigned a slope loading of 1.0. The y4 assessment was obtained half a school year after that and was assigned a slope loading of 1.5. This metric maps onto the strategy used by Willett and Bub (2005) in their growth curve modeling of child reading trajectories, after which the current exampled was modeled. The reading scores ranged from 0 to 100 with a standard deviation of about 10.

The intervention sought to improve the mediator of phonological awareness and phonics, namely the ability to manipulate sounds and connect them to letters. Scores tended to vary between 0 and 100 with a standard deviation near 10. The measures were obtained concomitant with reading ability scores. The analysis documented the trajectory of phonological awareness and its presumed impact on the reading ability trajectory. I used a parallel process growth analysis of mediation as presented earlier in this document.

Based on past research. I expected the reading trajectories to be non-linear. At young ages, increments often are exponential, but not dramatically so. I applied the Box-Cox transformation program to my website to the y1 to y4 reading scores as a function of time of assessment. The resulting value of lambda was 0.15, a result close enough to zero to justify using a log transformation of the Y. Here are plots generated by the program of the mean untransformed Y as a function of time and the transformed Y using lambda.

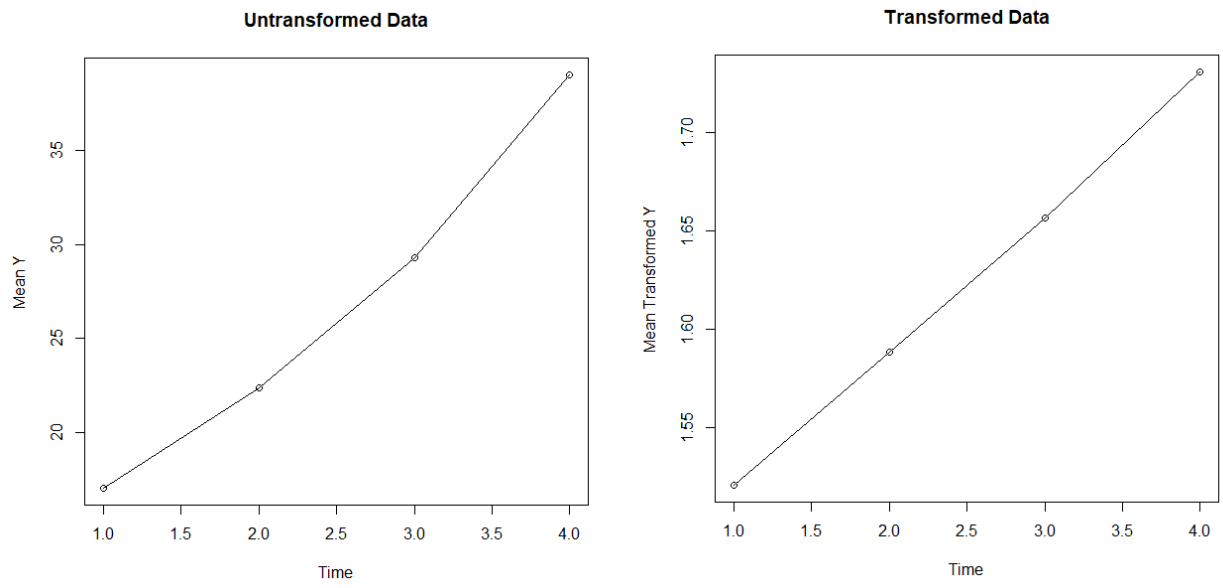

The exponential trend is evident in the untransformed data. The transformed data is linear. The data for the mediator showed the same pattern, so I move forward with log transforms on the mediator and the outcome given this also makes theoretical sense.

I now address the three core questions of an RET: (1) does the intervention meaningfully impact the outcome, (2) does the intervention meaningfully impact the mediator, and (3) does the mediator meaningfully impact the outcome. Table 7 presents the relevant Mplus syntax. I omit the `MODEL CONSTRAINT` commands to obtain a more nuanced evaluation of model fit but add them to a second syntax execution. I log transformed the data before entering it into Mplus. If I had not, I could apply the transformation in the `DEFINE` command using `y1=log(y1)` and repeating this for each variable to log transform.

Table 7. Mplus Syntax for Log-Based Non-Linear Growth Curve Model

```
1.  TITLE: Log-Based Model   ;
2.  DATA: FILE = LGAlogM.dat;
3.  VARIABLE:
4.  DEFINE:
5.   CENTER m0 y0 (GRANDMEAN) ;
6.  NAMES = y1 y2 y3 y4 m1 m2 m3 m4 y0 m0 treat ;
7.  ANALYSIS: ESTIMATOR=MLR ;
8.  MODEL:
9.   !Define latent growth model for y
10.  yi ys | y1@0 y2@1 y3@2 y4@3 ;
11.  !Estimate latent intercepts and assign them labels
12.  [yi] (ya1);      !Intercept for intercept factor
13.  [ys] (ya2);      !Intercept for slope factor
14.  !Estimate variances of intercept and slope factors
```

```
15.  yi;       ! Variance of intercept factor
16.  ys;       ! Variance of slope factor
17.  !Estimate covariance between intercept and slope factors
18.  yi WITH ys;
20.  ! Estimate residual variances (error)
21.  y1-y4;
22.  !Define latent growth model for mediator
23.  mi ms | m1@0 m2@1 m3@2 m4@3 ;
24.  !Estimate intercepts and assign them labels
25.  [mi] (ma1) ;      !Intercept for intercept factor
26.  [ms] (ma2) ;      !Intercept for slope factor
27.  !Estimate variances of intercept and slope factors
28.  mi;       ! Variance of intercept factor
29.  ms;       ! Variance of slope factor
30.  !Estimate covariance between intercept and slope factors
31.  mi WITH ms;
32.  !Estimate residual variances (error)
33.  m1-m4;
34.  !Define the regressions and assign labels to coefficients
35.  mi ON treat m0 (p1 p2)  ;
36.  ms ON treat (p3);
37.  yi ON mi y0 (p4 p5);
38.  ys ON ms (p6);
39. MODEL INDIRECT:
40.  yi IND treat ;
41.  mi IND treat ;
42.  ys IND treat ;
43.  ms IND treat ;
44.  OUTPUT: Samp StdYX MOD(4) Residual Tech4  ;
```

When I execute the syntax in Table 1 (removing the line numbers I included for pedagogical purposes), the model fit indices suggested good model fit. The chi square test was statistically non-significant (chi square = 50.53 with df = 44, $p < 0.24$), which is consistent with reasonable data-model correspondence. The RMSEA was 0.012. The upper limit of the 90% confidence interval for it was 0.025. The p value for close fit was $p < 1.00$. The CFI was 1.00 and the standardized RMR was 0.034. For localized fit, there were no theoretically meaningful modification indices greater than 4 and no meaningful standardized residuals for the predicted and observed covariances on a cell-by-cell basis.

After concluding for reasonable model fit, I re-ran the syntax but added the same `MODEL CONSTRAINT` commands as in the traditional two group example with a few additions. I also removed the request for modification indices on the e `OUTPUT` line. Here are the commands I added just before the new Line 44:

```
43a1.  MODEL CONSTRAINT:
43a2.    NEW (mst msc msdiff m1t m2t m3t m4t m1c
43a3.    m2c m3c m4c mdiff1 mdiff2 mdiff3 mdiff4) ;
```

```
43a4.    mst = ma2 + p3 ; !Treatment mean of mediator slope factor
43a5.    msc = ma2 ;      !Control mean of mediator slope factor
43a6.    msdiff = mst-msc ;
43a7.    m1t = (ma1+p1) ;        !Treatment mediator value at time 1
43a8.    m2t = m1t + mst*1  ;    !Treatment mediator value at time 2
43a9.    m3t = m1t + mst*2  ;    !Treatment mediator value at time 3
43a10.   m4t = m1t + mst*3  ;    !Treatment mediator value at time 4
43a11.   m1c = ma1 ;             !Control mediator value at time 1
43a12.   m2c = m1c + msc*1 ;     !Control mediator value at time 2
43a13.   m3c = m1c + msc*2 ;     !Control mediator value at time 3
43a14.   m4c = m1c + msc*3 ;     !Control mediator value at time 4
43a15.   mdiff1 = m1t - m1c ;
43a16.   mdiff2 = m2t - m2c ;
43a17.   mdiff3 = m3t - m3c ;
43a18.   mdiff4 = m4t - m4c ;
43b1.    NEW (yst ysc ysdiff y1t y2t y3t y4t y1c
43b2.    y2c y3c y4c ydiff1 ydiff2 ydiff3 ydiff4) ;
43b3.    yst = ya2 + (ma2 + p3)*p6 ;    !Treatment mean of Y slope factor
43b4.    ysc = ya2 + ma2*p6 ;           !Control mean of Y slope factor
43b5.    ysdiff = yst-ysc ;
43b6.    y1t = ya1 + (ma1 + p1)*p4;     !Treatment Y value at time 1
43b7.    y2t = y1t + yst*1;             !Treatment Y value at time 2
43b8.    y3t = y1t + yst*2;             !Treatment Y value at time 3
43b9.    y4t = y1t + yst*3;             !Treatment Y value at time 4
43b10.   y1c = ya1 + ma1*p4;            !Control Y value at time 1
43b11.   y2c = y1c + ysc*1;             !Control Y value at time 2
43b12.   y3c = y1c + ysc*2;             !Control Y value at time 3
43b13.   y4c = y1c + ysc*3;             !Control Y value at time 4
43b14.   ydiff1 = y1t - y1c ;
43b15.   ydiff2 = y2t - y2c ;
43b16.   ydiff3 = y3t - y3c ;
43b17.   ydiff4 = y4t - y4c ;
43c1.    NEW(oy1t oy2t oy3t oy4t oy1c oy2c oy3c oy4c om1t om2t om3t om4t
43c2.    om1c om2c om3c om4c expyst expysc expysdif expyidif expmst
43c3.    expmsc expmsdif expmidif elastici elastics) ;
43c4.    oy1t = exp(y1t) ;
43c5.    oy2t = exp(y2t) ;
43c6.    oy3t = exp(y3t) ;
43c7.    oy4t = exp(y4t) ;
43c8.    oy1c = exp(y1c) ;
43c9.    oy2c = exp(y2c) ;
43c10.    oy3c = exp(y3c) ;
43c11.   oy4c = exp(y4c) ;
43c12.   om1t = exp(m1t) ;
43c13.   om2t = exp(m2t) ;
43c14.   om3t = exp(m3t) ;
43c15.   om4t = exp(m4t) ;
43c16.   om1c = exp(m1c) ;
43c17.   om2c = exp(m2c) ;
```

```
43c18.   om3c = exp(m3c) ;
43c19.   om4c = exp(m4c) ;
43c20.   expyst = exp(yst) ;
43c21.   expysc = exp(ysc) ;
43c22.   expysdif = exp(ysdiff);
43c23.   expyidif = exp(ydiff1) ;
43c24.   expmst = exp(mst) ;
43c25.   expmsc = exp(msc) ;
43c26.   expmsdif = exp(msdiff);
43c27.   expmidif = exp(mdiff1) ;
43c28.   elastici = 1.10^p4 ;
43c29.   elastics = 1.10^p6 ;
```

I explained the logic of most of these lines in prior examples. The new lines are in statements 43c1 to 43c29 which I explain below.

Total Effect of the Treatment on the Outcome

To test the effects of the intervention on the outcome, I use the output generated by Lines 40 (`YI IND TREAT`) and 42 (`YS IND TREAT`) from Table 7. The former focuses on the treatment-control Y mean difference at the immediate posttest and the latter focuses on the treatment-control differences in the trajectories of Y that occur across the follow-ups. Here is the output for Line 40 that focuses on the mean log(Y) difference between the treatment and control groups at the immediate posttest:

```
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

Effects from TREAT to YI

  Total                0.059      0.015      3.896      0.000
```

The estimated mean difference between log(Y) at the immediate posttest for the treatment condition minus the control condition is 0.059 ±0.03, which is statistically significant (CR = 3.90, $p < 0.05$). The 95% confidence interval for the difference is 0.03 to 0.09. The (geometric) mean YI in its original metric for the treatment group is in the `New/Additional Parameters` section under the entry `OY1T`. It equaled 16.07 ±0.49. For the control group, the corresponding original metric mean for the entry `OY1C` is 15.14 ±0.42. If I calculate the exponent of 0.059, which I did in the `New/Additional Parameters` section for the entry `EXPIDOFF`, I obtain 16.07/15.14 = 1.06 ±0.03. The intervention group geometric mean at the immediate posttest is about 6% larger than that for the control group.

Suppose that prior to the study, the research team set a standard for a meaningful population outcome mean difference at the immediate posttest on the original metric as

being 5 units or more. The mean difference was just under one unit, 16.07-15.14 = 0.93. Even though the difference is non-zero (by virtue of the significance test for the logged mean difference), it does not meet the meaningfulness standard. It turns out the researchers did not expect much difference in the mean outcome at the immediate posttest between the intervention and control groups; they felt instead that the intervention would more likely lead to group differences in reading *trajectories* across time.

The effect of the intervention on the Y trajectories, is reflected in the test of the treatment minus control difference in the latent Y slopes from Line 42 of the output. Here are the relevant results from the Mplus output:

```
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

Effects from TREAT to YS

  Total               -0.006      0.008     -0.705      0.481
```

The (log-based) slope difference was small and statistically non-significant (difference = 0.006 ±0.016, CR = 0.71, *ns*). The hypothesis that the reading trajectories would meaningfully differ between the treatment and control groups was not supported. To better appreciate this dynamic, it is helpful to plot the original metric geometric mean values for the Y across the four time periods for the treatment and control groups. I used the exponentiated means from the `New/Additional Parameters` section to create the plot:

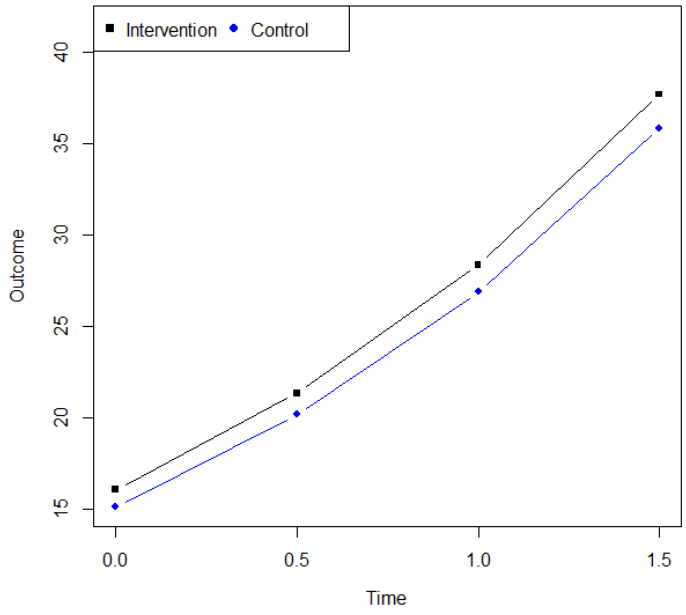

The trajectories appear quite similar. The exponent of the trajectory disparity of -0.006 taken from the term `EXPSDIFF` from the output section `New/Additional Parameters` was 0.994. This indicates the mean trajectory for the intervention group was 1 – 0.994 = 0.006 or about 1% *lower* than that for the control group (0.006 multiplied by 100 rounds to 1%).

Parenthetically, the exponent of the latent slope for the intervention and control groups from the `New/Additional Parameters` section (`EXPYST` and `EXPYSC`) equaled 1.765 and 1.775, respectively. Note that 1.765/1.775 = 0.994, the result reported above. The value 1.765 for the intervention group indicates that the value of the geometric mean increases by a multiplicative factor of 1.765 (or by an increase of 76.5%) for every one full unit increase in time. The value 1.775 for the control group indicates that the value of the outcome geometric mean increases by a multiplicative factor of 1.775 (or by an increase of 77.5%) for every one full unit increase in time.

In the final analysis, the intervention was not successful in meaningfully improving reading trajectories across time. The question then becomes why? Was it because the intervention failed to affect the presumed mediator of phonological awareness and phonics? Was it because the presumed mediator was not meaningfully impactful on the outcome? Or was it both of these dynamics? Analysis of the second two questions for RETs provide informative perspectives on these questions.

Effect of the Treatment on the Mediator

To test the effects of the intervention on the mediator, I use the same strategy as that in the prior section for the effect of the intervention on the outcome. I use the output generated by Lines 41 (`MI IND TREAT`) and 43 (`MS IND TREAT`) from Table 7. The former focuses on the treatment-control mediator mean difference at the immediate posttest and the latter focuses on the treatment-control difference in mediator trajectories. Here is the output for Line 41 that focuses on the mediator difference between the treatment and control groups at the immediate posttest:

```
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

Effects from TREAT to MI

  Total                0.077      0.020      3.945      0.000
```

The estimated mean difference between log(M) at the immediate posttest for the treatment condition minus the control condition is 0.077 ±0.04, which is statistically significant (CR = 3.95, $p < 0.05$). The 95% confidence interval for the difference is 0.04 to 0.12. The (geometric) mean MI in its original metric for the treatment group is in the

New/Additional Parameters section under the entry OM1T. It equaled 27.29 ±0.80. For the control group, the corresponding original metric mean for the entry OM1C is 25.27 ±0.66. If I calculate the exponent of 0.077, which I did in the New/Additional Parameters section for the entry EXPMIDIF, I obtain 27.29/25.27 = 1.08 ±0.04. The intervention mediator mean at the immediate posttest is about 8% larger than that for the control group.

Suppose that prior to the study, the research team set a standard for a meaningful population mediator mean difference at the immediate posttest on the original metric as being 5 units or more. The mean difference was just over two units, 27.29 - 25.27 = 2.02. Even though the difference is non-zero (by virtue of the significance test for the logged mean difference), it does not meet the meaningfulness standard. As with the Y outcome, the researchers did not expect much difference in the mediator mean at the immediate posttest between the intervention and control groups; they felt instead that the intervention would more likely lead to group differences in phonological awareness *trajectories* across time.

The effect of the intervention on these trajectories, is reflected in the test of the treatment minus control difference in the latent M slopes from Line 43 of the output. Here are the relevant results from the Mplus output:

	Estimate	S.E.	Est./S.E.	Two-Tailed P-Value
Effects from TREAT to MS				
Total	-0.007	0.010	-0.716	0.474

The results mirrored those for Y. The (log-based) slope difference was small and statistically non-significant (difference = -0.007 ±0.02, CR = 0.72, *ns*). The hypothesis that the phonological awareness trajectories would meaningfully differ between the treatment and control groups was not supported. To better appreciate this dynamic, it again is helpful to plot the original metric geometric mean values for the mediator across the four time periods for the treatment and control groups. I used the exponentiated means from the New/Additional Parameters section to create the plot:

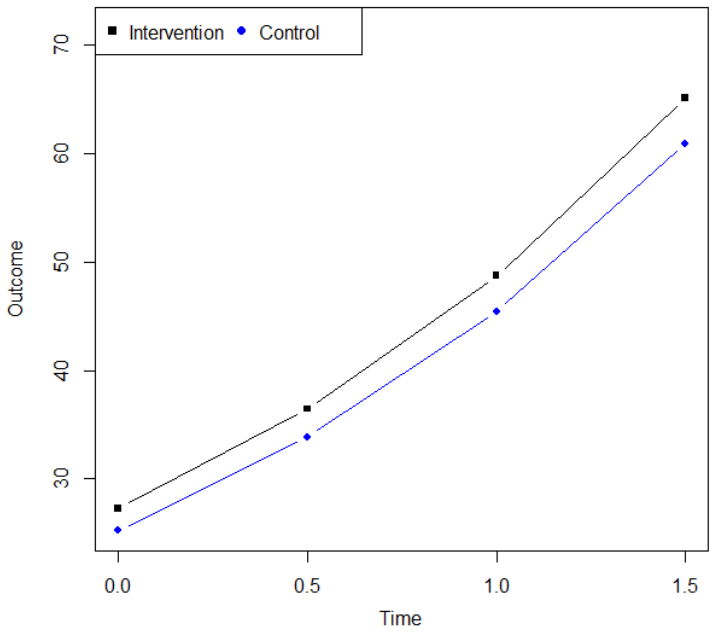

Like the reading outcome, the trajectories appear quite similar. The exponent of the trajectory disparity of -0.007 taken from the term `EXPSDIFF` from the output section `New/Additional Parameters` was 0.993. This indicates the mean trajectory for the intervention group was 1 – 0.993 = 0.007 or about 1% *lower* than that for the control group (0.007 multiplied by 100 rounds to 1%).

The exponent of the mediator latent slope for the intervention and control groups from the `New/Additional Parameters` section (`EXPMST` and `EXPMSC`) equaled 1.785 and 1.798, respectively. Note that 1.785/1.798 = 0.993, the result reported above. The value 1.785 for the intervention group indicates that the value of the geometric mean increases by a multiplicative factor of 1.785 (or by an increase of 78.5%) for every one full unit increase in time. The value 1.798 for the control group indicates that the value of the outcome geometric mean increases by a multiplicative factor of 1.798 (or by an increase of 79.8%) for every one full unit increase in time.

In sum, the intervention was not successful in meaningfully improving phonological awareness trajectories across time. The researchers need to re-examine the intervention activities they used that they thought would impact phonological awareness trajectories because the current activities appear not to work.

Effect of the Mediator on the Outcome

The analysis of the estimated effect of the mediator on the outcome does not distinguish the intervention and control groups because the guiding model assumes that the effect is the same in both treatment conditions. The relevant coefficients for evaluating the mediator

effect on the outcome are the coefficient from MI to YI and the coefficient from MS to YS. Here are the results from the output:

```
MODEL RESULTS

                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

YI       ON
    MI                 0.771      0.035     22.040      0.000
    Y0                 0.038      0.026      1.432      0.152

 YS       ON
    MS                 0.792      0.080      9.892      0.000
```

The coefficient for MI→YI is the path coefficient for the outcome at the immediate posttest regressed onto the mediator at the immediate posttest controlling for the baseline outcome. The coefficient was 0.771 ±0.07 (CR =22.04, $p < 0.05$), which is statistically significant. The coefficient for MS→YS is the path coefficient for the latent Y trajectory regressed onto the latent M trajectory. The path coefficient was 0.79 ± 0.16 (CR = 9.89, $p < 0.05$), which also is statistically significant. Both the outcome and the mediator were logged, so these coefficients reflect a form of log-log regression. As discussed earlier and in Chapter 15, I can express these coefficients as elasticities. I find it easiest to interpret them if I use an elasticity factor grounded in a 10% change of the mediator. In the MODEL CONSTRAINT command I calculate this type of elasticity for MI→YI as 1.10^p4 (see elastici) and for MS→YS as 1.10^p6 (see elastics). Here are the results:

```
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

New/Additional Parameters

    ELASTICI           1.076      0.004    299.746      0.000
    ELASTICS           1.078      0.008    131.015      0.000
```

For the intercept, a 10% change in the value of the mediator is thought to produce a 7.6% ±0.8 change in the value of the outcome. For the slope, a 10% change in the value of the mediator is expected to produce a 7.8% ±1.8 change in the value of the outcome. Suppose the research team *a priori* decided that a meaningful elasticity expressed as percents and using a 1.10 multiplier was 5%. In this case, both of the estimated effects are judged to be meaningful.

Omnibus Mediation Analysis and Conclusions

Because the results suggest that the T→M link for intervention effects on the trajectory for phonetic awareness was "broken," I would not conclude that there is mediation of the T→Y effect by it. However, phonetic awareness seems to be a viable target for improving reading trajectories moving forward. It is just that the current approach used by the investigators to impact the trajectory for phonetic awareness seems deficient.

Concluding Comments on Non-Linear LGCM

The above strategies for addressing non-linear functions in LGCM are but a subset of available strategies. Conspicuous is my omission of quadratic and cubic growth modeling using polynomials. I do not find these latter methods as useful as the ones I have discussed. The nature and type of non-linearity that quadratic and cubic growth modeling is somewhat limiting. Both functions must eventually turn upward or downward which sometimes yields non-sensical curvature shifts. Convergence problems are not uncommon. However, if you want to master quadratic and cubic LGCM, it should be relatively easy to do with the background provided here.

With all of the methods discussed, you should be cautious about generalizing a function form beyond the range of scores in the data. I also highly recommend you evaluate confidence intervals and significance tests using multiple robust methods, such as bootstrapping and robust maximum likelihood.

BINARY AND ORDINAL OUTCOMES IN LGCM

Program evaluations sometimes involve scenarios where the outcome in the longitudinal RET is a binary or ordinal outcome. In this section, I describe the specialized methods needed to address such scenarios. My focus is on binary outcomes. I use the general logic of multi-group SEM, which I introduced in a prior example in this document. I recommend you review that section before starting this one.

The Latent Propensity Framework

As necessary background, I briefly review from Chapter 5 the latent propensity model for binary outcomes for logistic and probit regression. Logistic and probit regression have been conceptualized in different ways, one of which is known as a **latent propensity approach**. Consider the case where the outcome variable is dichotomous and focuses on a behavior we seek to have people perform (e.g., obtain a vaccination). The post-treatment outcome is scored 1 = performed the behavior versus 0 = did not perform the behavior. In the latent

response approach, there is said to exist an underlying, unmeasured latent propensity, that I will signify as y*, that is continuous and reflects the propensity to engage or not engage in the dichotomous behavior.[2] When an individual crosses a threshold value on y*, represented by the parameter τ, his or her value on the observed dichotomous variable, y, changes from 0 (non-performance) to 1 (performance), i.e., s/he performs the behavior. For example, y* might range from -2 to +2 and the threshold value might be 0; if a person's location on y* is below 0, then the person does not perform the behavior. If the person's location on y* is zero or above, the person performs the behavior. Of theoretical interest is what the value of τ is for y* because it differentiates behavioral performance from behavioral non-performance.

In theory, we can construct causal models of the determinants of y* and, with concomitant knowledge of τ, we can then build an understanding of the dichotomous outcome. Like multiple regression, given a predictor set X, y* is assumed to be a linear function of the X:

$$y^* = \gamma_0 + \gamma_1 X_1 + \gamma_2 X_2 + \ldots + \gamma_k X_k + \varepsilon \quad [1]$$

In this equation, I use γ to represent a regression coefficient in place of the usual symbol β to signify that I am working with the latent response formulation. γ_0 represents an intercept. The last term in the equation is a random disturbance term that, like most regression models, is assumed to be independent of the X and has a mean of zero.

Theoretically, I am interested in estimating the various γ in Equation 1 because they provide insight into the causal impact of X on the propensity to engage in Y. When we conduct a traditional logistic regression analysis, we apply the model

$$\ln[\text{Odds}(Y)] = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \ldots + \beta_k X_k \quad [2]$$

That is, we estimate the logistic coefficients rather than the γs in Equation 2. We ultimately want to link the various β to the γ so we can understand and estimate the values of γ. As well, we want to identify the value of τ. For reasons that will become apparent shortly, I prefer to express the disturbance term in Equation 1 using a notation scheme by Mood (2010) and Karlson (2015) that introduces a scaling parameter, *s*, that allows the variance of ε to vary from an a priori assumed value as a function of this scaling parameter, as follows:

[2] Do not confuse my use of the term "latent" here with the concept of a latent variable in structural equation modeling (SEM). Although there are similarities, the literature and meanings surrounding the term "latent" is distinct for logistic/probit regression compared to SEM.

$$y_i^* = \gamma_0 + \gamma_1 X_1 + \gamma_2 X_2 + \ldots + \gamma_k X_k + (s)(\varepsilon_i) \quad [3]$$

Note that when s = 1.0, Equation 3 is no different from Equation 1; as s becomes larger than 1.0, then the variance of the disturbances increases; as s becomes less than 1.0, the variance decreases. The need for s will be apparent shortly.

One challenge of the latent response formulation in Equation 3 is that y* is not directly measured and consequently, it has no metric. Is it scored from 1 to 10, from 1 to 100, from -5 to +5, or what? In classic multiple regression, the regression coefficients and error variances take on values based on the metrics of Y and the metrics of the Xs. However, because y* has no metric, the propensity model is under-identified. By under-identified, I mean that the model has too many unknowns relative to the data at hand, so it cannot be estimated. This is because there are an infinite number of solutions for the values of γ, τ, and ε.

It turns out, that if I fix the mean and variance of $(s)(\varepsilon)$ to take on certain values, then this reduces the degree of under-identification because it reduces the number of unknowns. For example, if I set the mean of $(s)(\varepsilon)$ to be zero (a common assumption in regression models), s = 1, and the variance of ε to 3.290, then I do not need to estimate any of these values. This is the strategy that the latent response model uses to reduce under-identification; it fixes these parameters at the values just mentioned. It also assumes the ε follows a standard logistic distribution, which is a bell-shaped curve, much like a normal distribution. For probit models, s also is assumed to equal 1.0, the variance of ε is fixed at 1.0 instead of 3.290, and the ε are assumed to be normally distributed. Why choose the value 3.290 for the variance of ε in the logistic model? This is done because the variance of scores in a standard logistic distribution is $\pi^2/3$ (which equals 3.290) and because it therefore has several statistical properties that assist statistical inference.

The act of fixing the error variance, ε, at 3.290 in logistic regression or 1.0 in probit regression is important. In traditional regression modeling, the variance of ε is impacted by the metric of the outcome and can take on any non-negative value based on that metric. If the outcome is measured on a 1 to 100 scale, then the variance of ε will be different than if the outcome is measured on a 1 to 10 scale. One goal of traditional regression analysis is to estimate the magnitude of the variance of ε. By contrast, in the latent response logit/probit regression framework, the variance of the ε is fixed and it never changes. This is, to say the least, unorthodox. Fixing the error variance at 3.290 or 1.00 actually is a form of "standardization," but it is a form of standardization that is different from what we are familiar with.

So, to summarize to this point, for Equation 3, the mean, variance and distribution of ε per se never changes; for logit models, the mean is zero, the variance is 3.209, and it has a standard logit distribution; for probit models, the mean is zero, the variance is 1.0, and it

has a normal distribution. The value of *s* can vary, but these other features of ε are set in stone. The parameter *s* is essentially an adjustment factor that modifies the disturbance variance to reflect its true variance. In theory, it equals the ratio of the true standard deviation of the errors divided by the assumed standard deviation of the errors (the latter of which is 3.290 in logit regression and 1.00 in probit regression). In practice, the value of *s* is not knowable. However, we need to make the latent response formulation "work."

Even with these assumptions, there still remain sources of under-identification in the model. The source of this additional under-identification is the threshold value, τ, and the intercept, γ_0. These parameters cannot be simultaneously estimated. One of them has to be fixed at an *a priori* value. In logistic and probit models, some statistical software fixes the threshold value at zero and the intercept is estimated. Other software fixes the intercept at 0 and estimates the threshold value.[3] Once this potential source of under-identification is rectified, the model is estimable and, coupled with the other statistical assumptions, the γ in Equation 3 will equal the β in Equation 2 (but with some qualifications).

Given the assumptions that must be made to deal with under-identification as well as some of the contorted statistical properties of the latent response framework, some methodologists prefer not to think of logistic or probit regression in such terms. They instead work with logistic-derived or probit-derived probabilities and average marginal effects. The approach to latent growth modeling with binary or ordinal outcomes relies heavily on the latent propensity framework. The version I emphasize uses probit rather than logit frameworks because doing so has many statistical advantages.

The Latent Propensity Framework as Applied to Latent Growth Curves

The latent propensity framework as applied to a traditional growth curve model with a single binary variable Y measured at 6 time points is show in Figure 8. This figure is taken from Lee et. al (2018). Each of the observed Y binary variables is assumed to be a function of an underlying latent propensity signified by Y*. Between each Y and Y* is a threshold value, τ, that specifies at what point the continuous Y* translate from a Y score of 0 to a Y score of 1. The latent intercept and latent slope factors focus on the mean and variance of Y* at the first time period and the mean and variance of the slope of the growth function for the Y* across time. The model I use in the example RET is more complex than this because it includes a continuous mediator at each of the time points that predict Y in accord with a parallel process design.

[3] It turns out that if one fixes the intercept to zero, the threshold value will equal the intercept value for the case where the threshold is fixed at zero, but they will be opposite in sign.

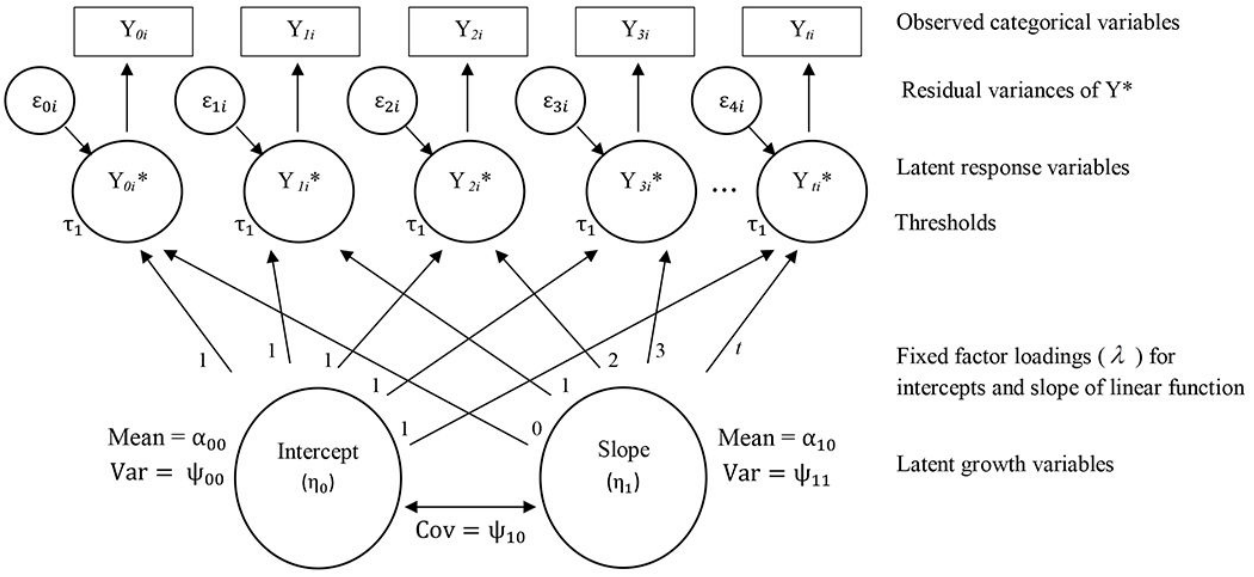

FIGURE 8. Binary latent growth curve model with latent propensities

As noted, I use a probit formulation of the binary growth model coupled with maximum likelihood estimation. Some researchers use variants of weighted least squares (WLS or WLSMV in Mplus) instead of maximum likelihood but the presence of a continuous mediator in the model complicates the use of these methods.

Mplus offers two versions of the probit model, one known as a **theta parameterization** and another known as a **delta parameterization**. These parameterizations focus on the $(s)(\varepsilon_i)$ term of Equation 3, with one emphasizing estimation of the scale factor and the other emphasizing estimation of the error variance. There is debate about which parameterization is best (see Kim & Whittaker, 2025; Lim & Kim, 2024; Grimm & Liu, 2016) and the answer is "it depends." For the probit-based multi-group maximum likelihood approach that I use in the RET example, the residual variances are fixed to 1.0 in all groups and at all time points. The latent slope factor in the growth model then represents the rate of change in the latent, underlying continuous propensity Y* given a one unit increase in time. The advantage of focusing on Y* is that it is construed as being linearly related to time, which simplifies the growth modeling analytics. The mean of the slope parameter represents the average rate of change in Y* as a function of time. Given the somewhat ethereal nature of the Y* metric in binary latent growth curve models, it is challenging to give precise meaning to the value of the Time→Y* slope coefficient other than (a) it signals whether the change in the latent propensity across time is non-zero and (b) it indicates the direction (positive or negative) of that change. Researchers supplement

the analysis of latent slope means and variances for Y* by examining how the probability that the observed Y = 1 changes across the different time points, but this usually is non-linear in form. .

For the example RET, I use a multiple group strategy that separates the treatment and control conditions into analytic groups. When implementing the binary-outcome-based probit latent growth curve model, it is necessary to introduce equality constraints and fixed coefficients to avoid model under-identification. Traditionally, researchers force the thresholds at each of the time points to be equal and they fix the mean of the latent intercept to equal 0. I adopt these conventions. The latter constraint ties my analytic hands in some respects, but I provide work arounds later for purposes of answering the three fundamental questions for an RET.

One other technical point is worth noting. When working with binary and ordinal outcomes you sometimes will encounter references to **sparse data**. This typically refers to cases where the two way contingency tables between variables have small frequencies in one or more of the individual cells. In SEM, when there are a large number of contingency tables with sparse data, such as when you have many time points in a growth model with binary outcomes, statistical tests can be adversely affected. You need to be cautious of scenarios were sparse data conditions are prevalent.

The Design and Analysis of the RET

The RET focused on a digitized program for individuals with anger management issues. The program consisted of web-based self-administered interactive exercises that individuals were to practice on a daily basis for a period of one month. The intervention focused on anger-based cognitive restructuring with an emphasis on identifying anger-inducing, irrational, or negative thoughts and replacing them with realistic, calm alternatives. During the four weeks over which the intervention was administered, a binary outcome measure was obtained at the end of each week. It assessed whether the individual had experienced an anger-related outburst in the past seven days (1 = experienced at least one outburst, 0 = did not experience an outburst). This was used as the Y1, Y2, Y3 and Y4 measures. A weekly measure of positive thinking also was obtained at the beginning of each week (M1, M2, M3 and M4). The measures served as an indicator of people's status on the mediator. The measure ranged from -5 to +5 with the typical standard deviation being 1 to 1.5, give or take. The research team was interested in documenting the improvement trajectory in anger outbursts as experience with the digital exercises accumulated over the course of the month. This was contrasted with naturally occurring trajectories of improvement for a control group. The hypothesis was that participation in the program would create a more favorable trajectory than what occurs naturally in the control group. I used a parallel process

growth model per my other examples in this document. I do not include covariates to keep things simple. I discuss later methods for including them. Random assignment to the treatment condition maintains the fundamental advantages of an RET even without baseline covariates.

Mplus does not support multiple group analyses of the type I need but I can trick Mplus into doing the analysis using a method called **known groups mixture modeling**. This approach closely parallels the structure of traditional multi-group analyses but uses slightly different syntax and nomenclature. Table 8 presents the relevant Mplus syntax.

Table 8. Mplus Syntax for Multi-Group Binary Outcome Analysis

```
1.  TITLE: LGCM with binary outcomes  ;
2.  FILE = LGAbinaryM.dat;
3.  VARIABLE: NAMES ARE Y1 Y2 Y3 Y4 M1 M2 M3 M4 TREAT;
4.  USEVARIABLES ARE Y1 Y2 Y3 Y4 M1 M2 M3 M4 ;
5.  CATEGORICAL ARE Y1 Y2 Y3 Y4 ;
6.  CLASSES = c(2);   !specify class label and number of classes
7.  KNOWNCLASS = c (TREAT=0 TREAT=1);
8.  ANALYSIS: TYPE = MIXTURE;
9.  ALGORITHM=INTEGRATION; LINK=PROBIT;
10. MODEL:
11.   %OVERALL%
12.     i s | y1@0 y2@1 y3@2 y4@3;          !Specify growth model for y
13.     [y1$1] ; [y2$1] ;  [y3$1] ;[y4$1] ; !Estimate thresholds
14.     [i@0] ; i ;   !Fix y latent intercept to zero; estimate its var
15.     [s] ; s ;      !Estimate mean/intercept and var of latent slope for y
16.     i with s ;     !Correlate latent slope and latent intercept for y
17.     mi ms | m1@0 m2@1 m3@2 m4@3 ;       !Specify growth model for m
18.     [mi]; [ms] ;   !Estimate mean of latent intercept and slope for m
19.     mi; ms ;       !Estimate var of latent intercept and slope for m
20.     mi WITH ms ;   !Correlate latent slope and mean for m
21.     m1-m4 ;        !Estimate disturbance var for observed m
22.     i ON mi ;      !Regress y latent intercept on m latent intercept
23.     s ON ms ;      !Regress y latent slope on m latent slope
24.   %c#1%            !Add labels to some of the above for control grp
25.     [y1$1] (gp1t); [y2$1] (gp1t); [y3$1] (gp1t); [y4$1] (gp1t);
26.     [s] (gp1s);
27.     [mi] (gp1mi); [ms] (gp1ms);
28.     i ON mi (gp1p1);
29.     s ON ms (gp1p2);
30.  %c#2%             !Add labels to some of the above for intervention grp
31.   [y1$1] (gp2t); [y2$1] (gp2t); [y3$1] (gp2t); [y4$1] (gp2t);
32.   [s] (gp2s);
33.   [mi] (gp2mi); [ms](gp2ms);
34.   i ON mi (gp2p1);
35.   s ON ms (gp2p2);
```

```
36. MODEL CONSTRAINT:
37.   NEW (yst ysc ysdiff ) ;
38.   yst = gp2s ;          !define y latent slope for treated grp
39.   ysc = gp1s ;          !define y latent slope for control grp
40.   ysdiff = yst-ysc ; !  !calculate difference between them
41.   NEW (mst msc msdiff m1t m2t m3t m4t m1c
42.   m2c m3c m4c mdiff1 mdiff2 mdiff3 mdiff4) ;
43.   mst = gp2ms ;         !define m latent slope for treated grp
44.   msc = gp1ms ;         !define m latent slope for treated grp
45.   msdiff = mst-msc ;    !calculate difference between them
46.   m1t = gp2mi ;         !calculate pred m1 for treated grp
47.   m2t = m1t + mst*1 ;   !calculate pred m2 for treated grp
48.   m3t = m1t + mst*2  ;   !calculate pred m3 for treated grp
49.   m4t = m1t + mst*3  ;   !calculate pred m4 for treated grp
50.   m1c = gp1mi ;          !calculate pred m1 for control grp
51.   m2c = m1c + msc*1 ;    !calculate pred m2 for control grp
52.   m3c = m1c + msc*2 ;    !calculate pred m3 for control grp
53.   m4c = m1c + msc*3 ;    !calculate pred m4 for control grp
54.   mdiff1 = m1t - m1c ;   !calculate the differences
55.   mdiff2 = m2t - m2c ;
56.   mdiff3 = m3t - m3c ;
57.   mdiff4 = m4t - m4c ;
58. OUTPUT: Samp StdYX Residual Tech4  ;
```

Much of the syntax should be familiar. On line 5, I use the CATEGORICAL command to declare y1-y4 as being binary. Instead of the term "groups" to denote the two groups, mixture models refer to them as "classes." On Line 6, I tell Mplus I have 2 classes and I am going to label each of the `c`. Mplus will internally assign a separate integer to the classes preceded by a #, in this case `c#1` and `c#2`. On Line 7, I tell Mplus I want to use the known class strategy with the variable `TREAT` used to define the two classes; participants with a score of 0 on `TREAT` are assigned to the first class and participants with a score of 1 on `TREAT` are assigned to the second class. On Line 8 I tell Mplus I want to do a mixture model. On Line 9, I tell Mplus to invoke the special mathematical process of numeric integration and that I want to use a probit link.

Recall that in multiple group SEM, it is understood that all of the lines under the word `MODEL` are applicable to all of the groups being analyzed unless you specifically tell Mplus to override some of them. The same is true for mixture models, but I signify the start of the generalized commands with the keyword `%OVERALL%` on Line 11. Lines 23 provide the generalized commands all of which should be familiar when coupled with the within-syntax comment remarks.

Line 24 tells Mplus you are going to add class specific commands that augment or override the generalized commands the first class `%c#1%`. Line 30 does the same but for the second class `%c#1%`. All of the commands in each class appeared in the generalized

commands but I have added unique group/class labels to the parameters. I make use of these labels in the MODEL CONSTRAINT section that begins on Line 36.

In prior examples, I did not include the MODEL CONSTRAINT commands in the initial analysis so that I could obtain modification indices to evaluate model fit. Modification indices are not allowed in mixture models more generally, so I do not request them on the OUTPUT line and I simply conduct a single analysis that includes the MODEL CONSTRAINT commands. Coupled with the within syntax comments I provide, all of the commands within labels I provide, all of the commands within MODEL CONSTRAINT should be straightforward, based on prior examples in this document. I am not able to calculate the predicted Y* for Y1 to Y4 because I had to fix the mean latent intercept for Y* to zero to avoid model under identification. Later, I describe my workaround for sacrificing this information for the sake of model identification.

I execute the above syntax and turn first to evaluation of the overall model fit. Mixture models provide different global fit indices than what you are accustomed to. Here are the core fit statistics of interest from the output:

```
MODEL FIT INFORMATION

     Chi-Square Test of Model Fit for the Binary and Ordered Categorical
     (Ordinal) Outcomes

               Pearson Chi-Square

               Value                             28.039
               Degrees of Freedom                    18
               P-Value                           0.0615

               Likelihood Ratio Chi-Square

               Value                             26.479
               Degrees of Freedom                    18
               P-Value                           0.0893
```

The Pearson chi-square test and the likelihood ratio chi-square test document the discrepancy between the predicted frequency patterns of the observed Y with the observed frequency patterns taking into account both classes of the mixture model. These two indices calculate a chi square statistic but using different approaches. The Pearson chi square statistic focuses on the squared differences between observed and expected cell frequencies. It is sensitive to small expected frequencies in sparse contingency tables, which is a potential shortcoming of the index. The likelihood ratio chi square statistic is based on the ratio of the likelihood of a saturated model to the likelihood of the estimated model (see Chapter 7).

For both indices, a statistically non-significant p value for the respective chi square statistic suggests there is unlikely to be notable disparities relative to model-data correspondence. Both indices should be in more or less agreement with one another. If they are widely discrepant, then you likely should not trust either of them. However, with small sample sizes and in the presence of sparse data, the likelihood ratio chi square statistic often is viewed as more reliable. In the present case, both approaches yielded statistically non-significant p values but the borderline nature of the p values makes me a bit nervous.

For localized fit, Mplus does not provide modification indices nor cell-by cell z tests of predicted versus observed covariances when analyzing mixture models. In the current case, for the Y portion of the model, Mplus reports the model predicted proportion of Y=1 cases for each time point in each class and these can be compared visually to the corresponding observed proportions. Here are the relevant model predicted probabilities/proportions (note: Category 2 refers to the estimated probabilities/proportions for the case of Y=1. Class 1 is the control group and class 2 is the intervention group. I put in bold the observed proportions/probabilities that I calculated external to Mplus):

```
RESULTS IN PROBABILITY SCALE

                  Estimate

Latent Class 1

 Y1
    Category 2         0.696     (0.71)
 Y2
    Category 2         0.672     (0.65)
 Y3
    Category 2         0.638     (0.64)
 Y4
    Category 2         0.609     (0.63)

Latent Class 2

 Y1
    Category 2         0.531     (0.54)
 Y2
    Category 2         0.284     (0.28)
 Y3
    Category 2         0.164     (0.18)
 Y4
    Category 2         0.116     (0.14)
```

The correspondence between the predicted and observed values seems reasonable.

For the continuous mediator variables, Mplus reports the model estimated means for M1 through M4 and the difference between the predicted and observed means for them in the RESIDUAL OUTPUT section,. The differences (called Residuals for Means on the output) should all be zero or close to zero; keep in mind that in the current case M1 to M4 have metrics ranging from -5 to +5 with standard deviations near 1 to 1.5. Here is the relevant output:

```
ESTIMATED MODEL AND RESIDUALS (OBSERVED - ESTIMATED) FOR CLASS 1

      Model Estimated Means
         M1             M2             M3             M4
         ________       ________       ________       ________
           0.003         -0.006         -0.014         -0.023

      Residuals for Means
         M1             M2             M3             M4
         ________       ________       ________       ________
          -0.021          0.028          0.002         -0.014

ESTIMATED MODEL AND RESIDUALS (OBSERVED - ESTIMATED) FOR CLASS 2

      Model Estimated Means
         M1             M2             M3             M4
         ________       ________       ________       ________
           0.588          1.178          1.767          2.357

      Residuals for Means
         M1             M2             M3             M4
         ________       ________       ________       ________
          -0.007          0.022         -0.028          0.010
```

The correspondence between the predicted and observed M1 to M4 means seems reasonable.

Mplus also reports differences between the predicted and observed covariances for each group for M1, M2, M3 and M4, but these diagnostics are more challenging to interpret because they are metric dependent and a bit non-intuitive.

When all was said and done and taking into account a weight-of-the-evidence perspective elaborated in Chapter 7, I ultimately decided to move forward with the model and interpret the parameter estimates. I now turn to the standard questions for an RET, namely, (1) does the intervention have a meaningful effect on the outcome, (2) does the intervention have a meaningful effect on the mediator(s), and (3) do the mediators have a meaningful effect on the outcome.

The Effect of the Intervention on the Outcome

The use of the multiple group (mixture) strategy does not yield coefficients that directly address the effect of the intervention on the binary outcomes. Evaluation of the effect is further complicated by the fact that the intervention is expected to be more impactful as time passes and participants get more practice with the on-line exercises. One simple approach to addressing the total effect question is to use a limited information estimation strategy outside the context of growth modeling that compares the proportion of people reporting an anger outburst in the past 7 days for the intervention versus control conditions at each of the four time points, respectively. I first used the *Temporal line plot* program on my website to plot the relevant proportions for the two groups as calculated from the observed data. Here is the plot:

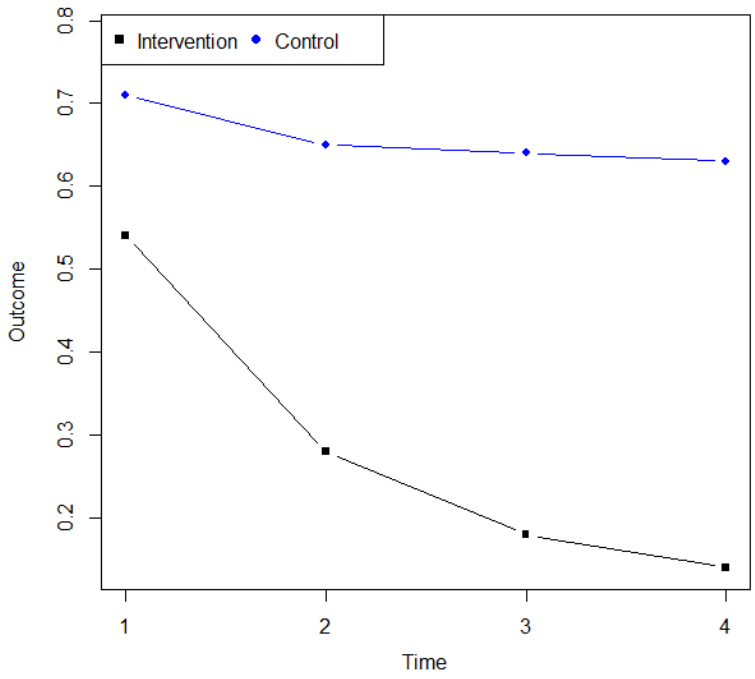

In the control condition, the proportion of participants reporting anger outbursts is relatively stable across time but with a trending slight decline. The intervention condition shows decays in the proportion of people exhibiting anger outbursts over time with the rate of change decelerating as the learning of positive thinking skills likely reaches an intervention-based asymptote. I tested for differences in the proportion of people who exhibited no anger at each time point separately for the intervention versus control conditions using the program *Test of Proportions* on my website. The program calculates confidence intervals for each separate proportion using the Agresti-Croul method as well as a more computationally intense method by Wilson. The two methods usually produce

similar results. Confidence intervals for the proportion differences use methods by Mee and by Miettinen–Nurminen. They tend to produce similar results but the Miettinen–Nurminen method is the more popular of the two. For these latter tests, if the confidence interval does not contain zero, the proportion difference is statistically significant at $p < .05$ based on a 95% confidence interval. For details about the methods, see Agresti (1996) and Wilcox (2021). Here is a summary of what I found when I subtracted the intervention proportion from the control proportion (note: the 95% confidence interval based margins of error (MOE) are asymmetrical so I report the absolute value of the larger MOE for the lower limit MOE versus the upper limit MOE; the amount of asymmetry was minor):

	Control Proportion	Treatment Proportion	Difference
Time 1	0.71 ± 0.04	0.54 ± 0.04	0.17 ± 0.06
Time 2	0.65 ± 0.04	0.28 ± 0.04	0.37 ± 0.06
Time 3	0.64 ± 0.04	0.18 ± 0.04	0.46 ± 0.05
Time 4	0.63 ± 0.04	0.14 ± 0.02	0.49 ± 0.05

All of the proportion differences were statistically significant ($p < 0.05$). Suppose the researchers set a meaningfulness standard proportion difference of 0.10 or greater. The lower limit of the confidence interval for the proportion difference was larger than 0.10 for all four contrasts, so I conclude the intervention had a meaningful effect on the outcome.

In addition to the above contrasts, I am interested in whether the trajectory of change in anger outbursts differs for intervention versus control groups. The relevant Mplus output focuses not on the proportions per se, but does provide information on the slope of the trajectory for the *propensity* to have anger outbursts, namely Y*, for the two groups. Here is the relevant output:

```
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

New/Additional Parameters
    YST               -0.445      0.062     -7.222      0.000
    YSC               -0.029      0.042     -0.677      0.499
    YSDIFF            -0.417      0.077     -5.381      0.000
```

The slope for Y* across the four time points for the intervention group was -0.445 ±0.12 (CR = 7.22, $p < 0.05$). This suggests that the propensity for anger outbursts statistically significantly decreased with time. The slope for Y* for the control group was -0.029 ±0.08 (CR = 0.68, *ns*). The results do not support that Y* decreased across time for them.

For the intervention condition, the above analysis characterizes changes across time for outburst propensities, but these changes do not necessarily translate into actual outbursts

because this also depends on the thresholds. If I want to evaluate within a condition (e.g., the intervention condition) the proportion of individuals who reported an anger outburst at one point in time versus the proportion who reported an anger outburst at another point in time, I can do so using the program *Dependent proportions* on my website. For example, The proportion of individuals in the intervention condition who reported anger outbursts at Time 1 was 0.54 ± 0.04 and the proportion who reported such outbursts at Time 2 was 0.28 ± 0.04. Here is the output from the program:

```
Proportion Difference Analysis
                Result     Lower CI     Upper CI
y1 proportion    0.536    0.4921784    0.5792727
y2 proportion    0.282    0.2443016    0.3230226
Difference       0.254    0.2017469    0.3055692
```

In the intervention condition, there was a statistically significant decline in the proportion of individuals who reported anger outbursts at Time 2 relative to Time 1 (difference = 0.25 ±0.05, $p < 0.05$). The declaration of statistical significance is based on the fact that the confidence interval for the proportion difference did not contain the value of zero. I can repeat this analysis for any other pairwise difference that is of theoretical interest. This allows me to explore the Y trajectory rather than the Y* trajectory in more depth, as desired.

In sum, there is evidence that the intervention had meaningful effects on Y and Y*.

The Effect of the Intervention on the Mediator

The measure for the mediator was continuous, so this portion of the Mplus output follows traditional two-condition multiple group formats discussed earlier in this document. I first examined treatment condition differences for the mediator at the immediate posttest. I obtain this contrast and the separate means for the intervention and control groups from the `MODEL CONSTRAINT` commands. Here is the relevant output:

```
                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

New/Additional Parameters
    MDIFF1           0.585      0.059      9.966      0.000
    M1T              0.588      0.041     14.204      0.000
    M1C              0.003      0.042      0.074      0.941
```

The adjusted mean difference is 0.59 ±0.12, which is statistically significant (CR = 9.97, $p < 0.05$). The 95% confidence interval is 0.47 to 0.71. Suppose the standard for a meaningful mean difference for the mediator at the immediate posttest is a population absolute difference of 0.20 or greater. Because the lower limit of the mean difference

confidence interval was greater than this standard, I conclude that the intervention effect for the immediate posttest is meaningful. The mediator means for the intervention and control conditions at the immediate posttest were 0.59 ±0.08 and 0.003 ±0.08, respectively.

Next, I examined a second facet of the intervention effect on the mediator, namely whether there was meaningful decay in the effect of the intervention on the mediator over time. As a first step, I isolate the coefficient for mediator decay for the intervention group and that for the control group. I then examine a contrast comparing the two coefficients. The relevant values are from the `MODEL CONSTRAINT` command. Here is the edited output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value
    MST                0.589      0.032     18.349      0.000
    MSC               -0.009      0.019     -0.450      0.653
    MSDIFF             0.598      0.038     15.927      0.000
```

The model implied linear slope for the mediator of positive thinking in the intervention condition is positive and statistically significant (slope = 0.59 ±0.06, CR = 18.35, $p < 0.05$). The predicted mean increases as time from the immediate posttest increases. For the control condition, the slope for the mean of M is flat (slope = -0.01 ±0.04, CR = 0.45, *ns*). The difference between the two coefficients (0.598 ±0.08) was statistically significant.

I can plot the two curves using the information in the output section from the `MODEL CONSTRAINT` commands for the predicted means at M1 though M4 for the intervention group (labeled `M1T`, `M2T`, `M3T`, and `M4T`) and for the control group (labeled `M1C`, `M2C`, `M3C`, and `M4C`). Using the program on my website called *Temporal line plot*, here is the plot:

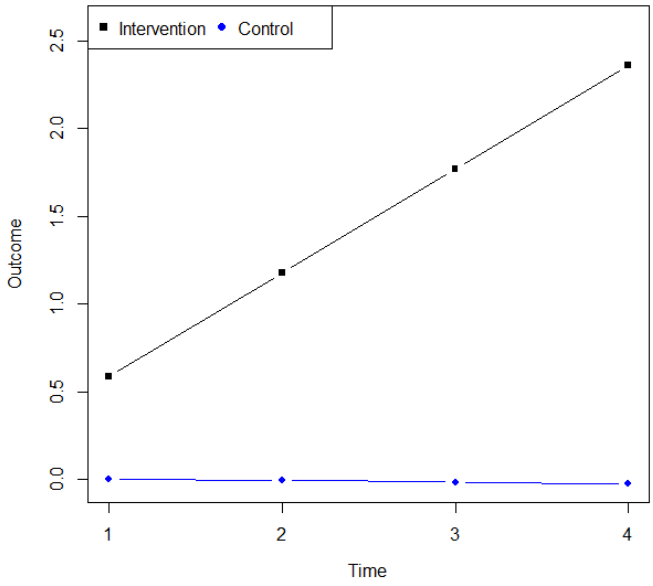

The positive slope for the intervention group is evident as is the flat slope for the control group. The separation between lines at a given time point reflects the predicted mean difference between the intervention and control groups. The mean difference between the treatment and control conditions increases as time passes.

The contrasts that I specified in the `MODEL CONSTRAINT` commands also test the mediator mean difference at each of the four time points using the model-informed predicted means. The contrasts are labeled `MDIFF1`, `MDIFF2`, `MDIFF3`, `MDIFF4` for times 1, 2, 3 and 4, respectively. Here is the relevant output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    MDIFF1             0.585      0.059      9.966      0.000
    MDIFF2             1.184      0.063     18.838      0.000
    MDIFF3             1.782      0.085     20.903      0.000
    MDIFF4             2.380      0.116     20.556      0.000
```

The intervention group predicted mediator mean is statistically significantly larger than the control group predicted mean at all of the time points. Suppose a meaningfulness standard is set by the research team to be a population absolute difference of 0.20 or greater. The 95% confidence interval each contrast exceeds this standard, so all of the above differences are deemed as being meaningful.

Effect of the Mediator on the Outcome

For the multi-group mixture model, the analysis estimates the effect of the mediator on the outcome separately for the intervention and control groups. The relevant coefficients are taken from the `MODEL RESULLTS` section and appear as follows (edited; note that Latent Class 1 is the control group and Latent Class 2 is the intervention group):

```
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

Latent Class 1 (0)

I ON MI               -0.675      0.107     -6.304      0.000
S ON MS               -0.786      0.123     -6.378      0.000

Latent Class 2 (1)

I ON MI               -0.710      0.100     -7.085      0.000
S ON MS               -0.886      0.099     -8.939      0.000
```

The coefficient for MI→I is the path coefficient for Y* at the immediate posttest regressed onto the mediator at the immediate posttest. It was statistically significant in both conditions and suggests that as positive thinking skills increase, the propensity for angry outbursts decreases. The coefficient for MS→S is the path coefficient for the slope of the Y* regressed onto the slope for the mediator. It showed the same trend and magnitude as that of the intercepts. The path coefficient was statistically significant for both the intervention and control conditions.

Omnibus Mediation Analysis

As noted, omnibus mediation tests are of lower priority for me. I instead prefer to focus on the individual links in mediational chains, evaluating their effects sizes. If I want an overall omnibus test of statistical significance, I rely on the joint significance test which in this case implied non-zero mediation effects, all things considered. In LGC models, evaluations of mediation take two forms. The first focuses on treatment differences in the latent intercept for the mediator and whether these differences, in turn, translate into treatment differences on the outcome latent intercept. The second mediation form focuses on the effect of the latent slopes of the mediator on the latent slopes of the outcome. The pattern of results in the current example is consistent with omnibus mediation.

SECOND ORDER LGCM AND MATTERS OF MEASUREMENT

An important issue in LGC modeling is how to deal with bias that measurement error can introduce. A strength of SEM is that it offers strategies to address or gain perspectives on measurement error bias. For models with all single indicator constructs, one strategy is to pursue sensitivity checks outlined in the document on my web page on the *Resources* tab for Chapter 3 for the link *measurement error for single indicator SEM models*. This strategy converts observed measures to latent constructs with single indicators that have fixed error variances mapping onto different reliability scenarios. Changes in parameter estimates for different reliability scenarios are then explored to give the analyst a sense of how sensitive conclusions are to different levels of measurement error. Here, I focus instead on a second strategy, namely the use of multiple indicators of constructs, often referred to as **second order growth models**.

Second order growth models have the same structure as the prior growth curve models I have considered in this document but now the M and Y are represented by latent variables with multiple indicators. Figure 9 provides an example in which the measures of M and Y at the different time points all are continuous. To reduce clutter, I omit the circled disturbance and error terms from the diagram, instead representing their presence with just

an arrow pointing to the appropriate endogenous variable. I also omit correlated disturbances and correlated errors in the diagram to reduce clutter but you will see that I chose to model some of them in the Mplus syntax. Finally, I also omit covariates from this example, but it is straightforward to include them as described in a later section.

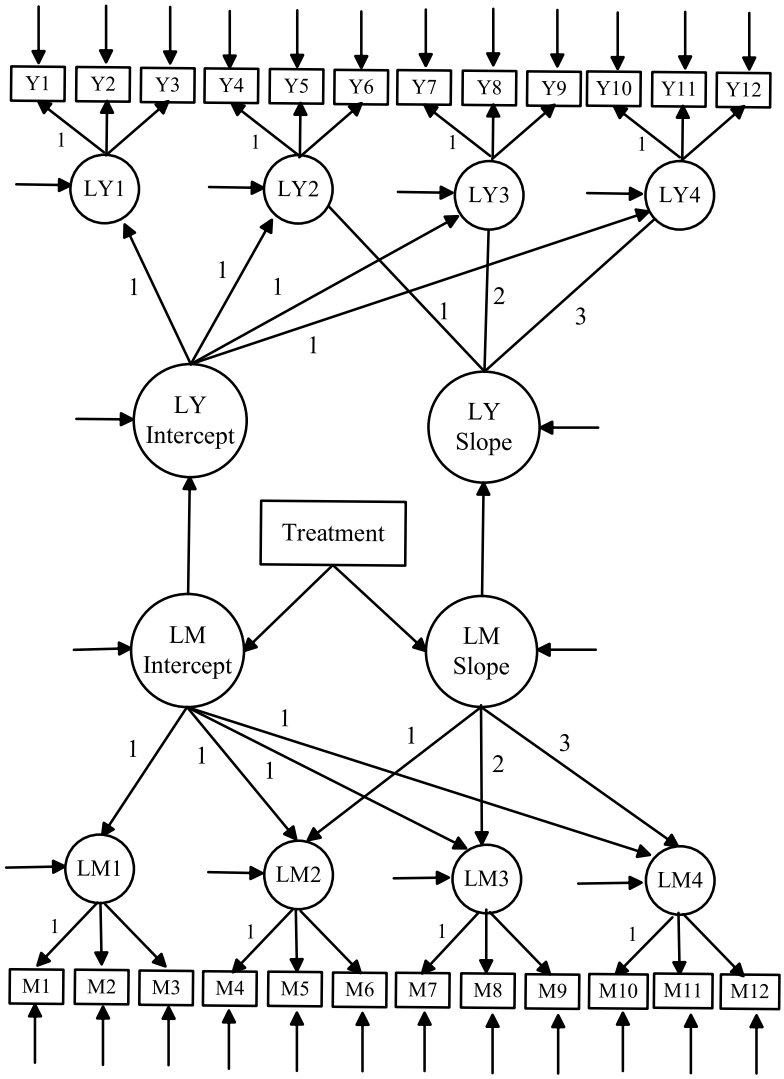

FIGURE 9. Second order latent growth curve model

The most common strategy for implementing second order growth curve models is to choose one of the indicators of the first order latent variables as a reference indicator and then fix its loading to 1.0 and its measurement intercept to zero in order to pass the

indicator's (error corrected) metric to the latent variable as well as its mean value. In this sense, the factor mean in the model becomes the same as the observed mean for the reference indicator, unless equality constraints are imposed on the loadings and/or measurement intercepts. This means that the analysis of the growth curve portion of the second order growth curve model primarily reflects the growth curve structure of the reference indicator. A consequence of this property is that the estimated curve can change somewhat depending on the variable chosen as the reference indicator, a property some methodologists find disconcerting. An alternative is to create a single composite from the three indicators of the latent variable in question and then perform a traditional growth curve analysis that reflects information from all the indicators relative to the underlying mean structure. Or, one might invoke effect coding as described by Little (2023) which defines latent means based on the average of the indicators of the latent variable. The single composite approach as well as the effect coding approach are challenging, however, when the indicators have decidedly different metrics There is debate about the best course of action to take for defining the factor means in second order growth curve models; see Newsom (2023) for a good discussion of the issues. I discuss below criteria to help you select a reference indicator. Once the reference indicator is chosen, it is used as the reference indicator at each of the different time points.

I now walk you through the Mplus syntax for an RET-based second order latent curve model. I use a parallel process modeling strategy in the context of a multigroup analysis with the two groups defined as the intervention group and the control group. I add `MODEL CONSTRAINT` commands later to address different questions but refrain from doing so here in order to focus on the basic syntax structure. The syntax appears in Table 9.

Table 9. Mplus Syntax for Second Order Analysis

```
1.  TITLE: LGCM with binary outcomes  ;
2.  DATA: FILE = LGAlatentM.dat;
3.  VARIABLE:
4.    NAMES = y1-y12 m1-m12 dtreat;
5.    USEVARIABLES = y1-y12 m1-m12;
6.    GROUPING IS DTREAT (0=control, 1=treat) ;
7.  ANALYSIS: ESTIMATOR=MLR ;
8.  MODEL:
9.   yeta1 by y1@1 y2 y3;          !define first order y factors
10.   yeta2 by y4@1 y5 y6;
11.  yeta3 by y7@1 y8 y9;
12.  yeta4 by y10@1 y11 y12;
13.  [yeta1-yeta4@0];              !fix intercepts of y factors to 0
14.  yeta1 yeta2 yeta3 yeta4  ;    !estimate var of y factors
15.  [y1@0 y2 y3];                 !define y measurement intercepts
```

```
16.  [y4@0 y5 y6];
17.  [y7@0 y8 y9];
18.  [y10@0 y11 y12];
19.  y1-y12 ;                                    !define y error variances
20.  yi by yeta1@1 yeta2@1 yeta3@1 yeta4@1;   !define y growth intercept
21.  ys by yeta1@0 yeta2@1 yeta3@2 yeta4@3;   !define y growth slope
22.  yi ys;           !estimate var of y latent intercept and latent slope
23.  [yi ys];         !estimate mean y latent intercept and y latent slope
24.  yi with ys;      !correlate y latent intercept with latent slope
25.  meta1 by m1@1 m2 m3;        !define first order m factors
26.  meta2 by m4@1 m5 m6;
27.  meta3 by m7@1 m8 m9;
28.  meta4 by m10@1 m11 m12;
29.  [meta1-meta4@0];            !fix intercepts of m factors to 0
30.  meta1 meta2 meta3 meta4  ;  !estimate var of m factors
31.  [m1@0 m2 m3];               !define m measurement intercepts
32.  [m4@0 m5 m6];
33.  [m7@0 m8 m9];
34.  [m10@0 m11 m12];
35.  m1-m12 ;                                    !define m error variances
36.  mi by meta1@1 meta2@1 meta3@1 meta4@1 ;  !define m growth intercept
37.  ms by meta1@0 meta2@1 meta3@2 meta4@3 ;  !define m growth slope
38.  mi ms;        !estimate var of m latent intercept and latent slope
39.  [mi ms];      !estimate mean y latent intercept and y latent slope
40.  mi with ms;   !correlate y latent intercept with latent slope
41.  yi ON mi ;    !regress latent intercepts
42.  ys ON ms ;    !regress latent slopes
43. MODEL control:  !assign group specific labels to some of above
44.  yeta1 by y1@1 y2 y3 (ycf1 ycf2 ycf3 );
45.  yeta2 by y4@1 y5 y6  (ycf4 ycf5 ycf6 );
46.  yeta3 by y7@1 y8 y9  (ycf7 ycf8 ycf9 );
47.  yeta4 by y10@1 y11 y12 (ycf10 ycf11 ycf12 );
48.  [y1@0 y2* y3*] (yci1 yci2 yci3);
49.  [y4@0 y5* y6*] (yci4 yci5 yci6);
50.  [y7@0 y8* y9*] (yci7 yci8 yci9);
51.  [y10@0 y11 y12](yci10 yci11 yci12);
52.  [yi] (ya1gp0) ;
53.  [ys] (ya2gp0);
54.  meta1 by m1@1 m2 m3 (mcf1 mcf2 mcf3 );
55.  meta2 by m4@1 m5 m6  (mcf4 mcf5 mcf6 );
56.  meta3 by m7@1 m8 m9  (mcf7 mcf8 mcf9 );
57.  meta4 by m10@1 m11 m12 (mcf10 mcf11 mcf12 );
58.  [m1@0 m2* m3*] (mci1 mci2 mci3);
59.  [m4@0 m5* m6*] (mci4 mci5 mci6);
60.  [m7@0 m8* m9*] (mci7 mci8 mci9);
61.  [m10@0 m11 m12](mci10 mci11 mci12);
62.  [mi] (ma1gp0) ;
63.  [ms] (ma2gp0) ;
64.  yi ON mi (p1gp0 ) ;
```

```
65.  ys ON ms (p2gp0);
66. MODEL treat: !assign group specific labels to some of above
67.  yeta1 by y1@1 y2 y3 (ytf1 ytf2 ytf3 );
68.  yeta2 by y4@1 y5 y6  (ytf4 ytf5 ytf6 );
69.  yeta3 by y7@1 y8 y9  (ytf7 ytf8 ytf9 );
70.  yeta4 by y10@1 y11 y12 (ytf10 ytf11 ytf12 );
71.  [y1@0 y2* y3*] (yti1 yti2 yti3 );
72.  [y4@0 y5* y6*] (yti4 yti5 yti6 );
73.  [y7@0 y8* y9*] (yti7 yti8 yti9 );
74.  [y10@0 y11 y12](yti10 yti11 yti12 );
75.  [yi] (ya1gp1) ;
76.  [ys] (ya2gp1);
77.  meta1 by m1@1 m2 m3 (mtf1 mtf2 mtf3 );
78.  meta2 by m4@1 m5 m6  (mtf4 mtf5 mtf6 );
79.  meta3 by m7@1 m8 m9  (mtf7 mtf8 mct9 );
80.  meta4 by m10@1 m11 m12 (mtf10 mtf11 mtf12 );
81.  [m1@0 m2* m3*] (mti1 mti2 mti3 );
82.  [m4@0 m5* m6*] (mti4 mti5 mti6 );
83.  [m7@0 m8* m9*] (mti7 mti8 mti9 );
84.  [m10@0 m11 m12](mti10 mti11 mti12 );
85.  [mi] (ma1gp1) ;
86.  [ms] (ma2gp1) ;
87.  yi ON mi (p1gp1 ) ;
88.  ys ON ms (p2gp1);
89. OUTPUT: Samp StdYX Mod(4) Residual Tech4  ;
```

Most of the syntax should be self-explanatory if you have read the prior section of this document on multiple group analysis of LGCMs for RETs. I assume you have done so. In the current syntax, I use the term `eta` to refer to a latent variable either for Y (`yeta`) or for M (`meta`) - see lines 9 to 14 and 25 to 30.

When I executed the syntax, the model fit indices suggested good model fit. The chi square test of perfect model fit in the population was statistically non-significant (chi square = 459.43 with 496 degrees of freedom, $p < 0.88$), which is consistent with the data being in accord with the model. The separate group contributions to the overall chi square were nearly equal (224.54 and 234.89, respectively). The RMSEA was <0.001. The upper limit of the 90% confidence interval for it was <0.006. The p value for close fit was not statistically significant ($p < 1.00$). The CFI is 1.00 and the standardized RMR was 0.023. For localized fit, there were no theoretically meaningful modification indices greater than 4 and no meaningful residuals between the predicted and observed covariances on a cell-by-cell basis.

After concluding in favor of reasonable model fit, I re-ran the syntax but added `MODEL CONSTRAINT` commands to address different issues. The first set of commands are to evaluate measurement invariance based on the document in the *Resources* tab of my

webpage for Chapter 3 in the link *testing measurement invariance*. Before adding the commands, I eliminate the request for modification indices from OUTPUT Line 89 because such indices are not allowed in conjunction with the MODEL CONSTRAINT command.

Measurement Invariance

My discussion of measurement invariance assumes you have read the relevant material on it in Chapter 3 and in the document on my webpage. The first set of commands I added address loading invariance using a forward analysis format. This strategy assumes that the reference indicator for the respective latent variables is approximately both loading and measurement intercept invariant across time and across the intervention and control groups. Here are the first set of commands I add just before the OUTPUT line, making use of the labels I assigned to parameters in the main program:

```
MODEL CONSTRAINT:
  NEW (ytf2f5 ytf2f8 ytf2f11 ytf5f8 ytf5f11 ytf8f11
  ycf2f5 ycf2f8 ycf2f11 ycf5f8 ycf5f11 ycf8f11
  mtf2f5 mtf2f8 mtf2f11 mtf5f8 mtf5f11 mtf8f11
  mcf2f5 mcf2f8 mcf2f11 mcf5f8 mcf5f11 mcf8f11 ) ;
  ytf2f5=ytf2-ytf5 ;
  ytf2f8=ytf2-ytf8 ;
  ytf2f11=ytf2-ytf11 ;
  ytf5f8=ytf5-ytf8 ;
  ytf5f11=ytf5-ytf11 ;
  ytf8f11=ytf8-ytf11 ;
  ycf2f5=ycf2-ycf5 ;
  ycf2f8=ycf2-ycf8 ;
  ycf2f11=ycf2-ycf11 ;
  ycf5f8=ycf5-ycf8 ;
  ycf5f11=ycf5-ycf11 ;
  ycf8f11=ycf8-ycf11 ;
  mtf2f5=mtf2-mtf5 ;
  mtf2f8=mtf2-mtf8 ;
  mtf2f11=mtf2-mtf11 ;
  mtf5f8=mtf5-mtf8 ;
  mtf5f11=mtf5-mtf11 ;
  mtf8f11=mtf8-mtf11 ;
  mcf2f5=mcf2-mcf5 ;
  mcf2f8=mcf2-mcf8 ;
  mcf2f11=mcf2-mcf11 ;
  mcf5f8=mcf5-mcf8 ;
  mcf5f11=mcf5-mcf11 ;
  mcf8f11=mcf8-mcf11 ;
```

These commands conduct significance tests on pairwise differences between factor

loadings for the same indicator at the four time periods, first for the Y variables and then for the M variables. The reference indicators are omitted from these contrasts, For example, the loading from the latent Y variable at time 1, LY1, to the observed indicator y2 is contrasted with the loadings LY2→y5, LY3→y8, and LY4→y11 in Figure 9. Each of the y represent the same scale but at different time points, e.g., y2, y5, y8, and y11 are the same scale. The loading LY1→y3 is contrasted against the loadings LY2→y6, LY3→y9, and LY4→y12 and the loading LY2→y5 is contrasted with the loadings LY3→y8 and LY4→y11. And so on. Each of the contrasts were conducted for the intervention condition and then repeated for the control condition. None of the contrasts should produce meaningful loading differences.

The notation for the labels is as follows, using `ytf2f5` as an example: The first character indicates if the loading is for the outcome (y) or the mediator (m). The second character is if it the loading is for the treatment group (t) or the control group (c). The final characters indicate the target contrasts, namely the first factor loading listed (f2) minus the second factor loading listed (f5), which is LY1→y2 minus LY2→y5. Here are the results:

```
MODEL RESULTS

                                                    Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

New/Additional Parameters
    YTF2F5          -0.004      0.042     -0.097      0.923
    YTF2F8          -0.029      0.047     -0.622      0.534
    YTF2F11          0.028      0.046      0.596      0.551
    YTF5F8          -0.025      0.043     -0.589      0.556
    YTF5F11          0.032      0.043      0.728      0.467
    YTF8F11          0.057      0.042      1.369      0.171
    YCF2F5          -0.024      0.043     -0.559      0.576
    YCF2F8           0.085      0.042      2.013      0.044
    YCF2F11          0.043      0.046      0.936      0.349
    YCF5F8           0.109      0.041      2.654      0.008
    YCF5F11          0.067      0.044      1.513      0.130
    YCF8F11         -0.042      0.043     -0.965      0.334
    MTF2F5          -0.059      0.051     -1.143      0.253
    MTF2F8          -0.030      0.054     -0.558      0.577
    MTF2F11          0.016      0.058      0.272      0.786
    MTF5F8           0.029      0.054      0.533      0.594
    MTF5F11          0.075      0.053      1.401      0.161
    MTF8F11          0.046      0.054      0.857      0.391
    MCF2F5          -0.020      0.062     -0.315      0.753
    MCF2F8           0.012      0.064      0.195      0.846
    MCF2F11         -0.045      0.056     -0.801      0.423
    MCF5F8           0.032      0.059      0.546      0.585
```

```
MCF5F11           -0.026      0.063     -0.408      0.684
MCF8F11           -0.058      0.061     -0.948      0.343
```

Of the 22 contrasts, 2 were statistically significant. I used the program on my webpage called *FDR p values* to adjust for familywise error rates given the large number of contrasts performed. None of the contrasts were statistically significant when this correction was applied. The lowest p value (.008) occurred for the contrast `YCF5F8`. I used the program on my website called *Loading non-invariance* to convert the raw difference for this contrast to a standardized effect size difference analogous to a Cohen's *d* statistic. The absolute standardized effect size was 0.12. This is modest (a small, medium, and large effect size in Cohen's framework are values of 0.20, 0.50 and 0.80, respectively).

Here are the statements I add to evaluate across-group loading invariances:

```
NEW (ytcf2 ytcf3 ytcf5 ytcf6 ytcf8 ytcf9 ytcf11 ytcf12
  mtcf2 mtcf3 mtcf5 mtcf6 mtcf8 mtcf9 mtcf11 mtcf12 ) ;
  ytcf2=ytf2-ycf2 ;
  ytcf3=ytf3-ycf3 ;
  ytcf5=ytf5-ycf5 ;
  ytcf6=ytf6-ycf6 ;
  ytcf8=ytf8-ycf8 ;
  ytcf9=ytf9-ycf9 ;
  ytcf11=ytf11-ycf11 ;
  ytcf12=ytf12-ycf12 ;
  mtcf2=mtf2-mcf2 ;
  mtcf3=mtf3-mcf3 ;
  mtcf5=mtf5-mcf5 ;
  mtcf6=mtf6-mcf6 ;
  mtcf8=mtf8-mcf8 ;
  mtcf9=mtf9-mcf9 ;
  mtcf11=mtf11-mcf11 ;
  mtcf12=mtf12-mcf12 ;
```

The notation for the labels is as follows using `ytcf2` as an example: The first character indicates if the loading is for the outcome (y) or the mediator (m). The second two characters indicate the two groups that are being compared, the treatment group (t) and the control group (c). The final two characters indicate the contrast targets, in this case f2 or LY1→y2. Here are the results:

```
                                                  Two-Tailed
                  Estimate       S.E.  Est./S.E.    P-Value

YTCF2               -0.048      0.045     -1.066      0.287
YTCF3               -0.047      0.068     -0.689      0.491
YTCF5               -0.068      0.044     -1.537      0.124
```

```
    YTCF6              0.000      0.066     -0.002      0.999
    YTCF8              0.067      0.044      1.525      0.127
    YTCF9              0.031      0.068      0.459      0.646
    YTCF11            -0.032      0.045     -0.718      0.473
    YTCF12             0.024      0.063      0.381      0.703
    MTCF2             -0.007      0.056     -0.118      0.906
    MTCF3             -0.036      0.079     -0.451      0.652
    MTCF5              0.033      0.058      0.562      0.574
    MTCF6             -0.252      0.079     -3.207      0.001
    MTCF8              0.036      0.059      0.617      0.537
    MTCF9             -0.079      0.080     -0.982      0.326
    MTCF11            -0.067      0.058     -1.159      0.246
    MTCF12             0.058      0.082      0.712      0.476
```

Of the 16 contrasts, one was statistically significant. When I adjusted the p values for multiple contrasts using the *FDR p values* program on my website, none of the contrasts were statistically significant.

The results of the two sets of analyses, taken as a whole, are consistent with loading invariance across time and across groups.

To test for measurement intercept non-invariance across-time the tradition is to force across-time loading invariance onto the model as part of the intercept non-invariance testing process. This is because, as discussed in Chapter 3, artifactual mean differences across time can be created either by loading non-invariance and/or measurement intercept non-invariance. Forcing loading invariance onto the model, in theory, removes it as a source of bias thereby permitting a cleaner test of intercept non-invariance. As I discuss below, I do not find this logic entirely convincing but it is nevertheless ingrained in traditional practice.

When I conducted the tests for intercept non-invariance across time, none of the contrasts were statistically significant after I applied the FDR corrections and this also was true when I pursued tests for across group non-invariance of measurement intercepts. Here is the syntax I used for the former tests, with red highlights of the syntax that forced loading invariance through the introduction of common labels onto the tests and all of the `MODEL CONSTRAINT` commands adapted to focus on measurement intercept difference across time:

```
TITLE: LGCM with latent outcomes  ;
DATA: FILE = LGAlatentM.dat;
  VARIABLE:
    NAMES = y1-y12 m1-m12 dtreat;
    USEVARIABLES = y1-y12 m1-m12;
    GROUPING IS DTREAT (0=control, 1=treat) ;
  ANALYSIS: ESTIMATOR=MLR ;
  MODEL:
   yeta1 by y1@1 y2 y3;        !define first order y factors
   yeta2 by y4@1 y5 y6;
```

```
  yeta3 by y7@1 y8 y9;
  yeta4 by y10@1 y11 y12;
  [yeta1-yeta4@0];            !fix intercepts of y factors to 0
  yeta1 yeta2 yeta3 yeta4  ;  !estimate var of y factors
  [y1@0 y2* y3*];             !define y measurement intercepts
  [y4@0 y5* y6*];
  [y7@0 y8* y9*];
  [y10@0 y11 y12];
  y1-y12 ;                                  !define y error variances
  yi by yeta1@1 yeta2@1 yeta3@1 yeta4@1;    !define y growth intercept
  ys by yeta1@0 yeta2@1 yeta3@2 yeta4@3;    !define y growth slope
  yi ys;           !estimate var of y latent intercept and latent slope
  [yi ys];         !estimate mean y latent intercept and y latent slope
  yi with ys;      !correlate y latent intercept with latent slope
  meta1 by m1@1 m2 m3;        !define first order m factors
  meta2 by m4@1 m5 m6;
  meta3 by m7@1 m8 m9;
  meta4 by m10@1 m11 m12;
  [meta1-meta4@0];            !fix intercepts of m factors to 0
  meta1 meta2 meta3 meta4  ;  !estimate var of m factors
  [m1@0 m2 m3];               !define m measurement intercepts
  [m4@0 m5 m6];
  [m7@0 m8 m9];
  [m10@0 m11 m12];
  m1-m12 ;                                  !define m error variances
  mi by meta1@1 meta2@1 meta3@1 meta4@1 ;   !define m growth intercept
  ms by meta1@0 meta2@1 meta3@2 meta4@3 ;   !define m growth slope
  mi ms;        !estimate var of m latent intercept and latent slope
  [mi ms];      !estimate mean y latent intercept and y latent slope
  mi with ms;   !correlate y latent intercept with latent slope
  yi ON mi ;    !regress latent intercepts
  ys ON ms ;    !regress latent slopes
MODEL control:  !assign group specific labels to some of above
  yeta1 by y1@1 y2 y3 (ycf1 ycf2 ycf3 );
  yeta2 by y4@1 y5 y6  (ycf4 ycf2 ycf3 );
  yeta3 by y7@1 y8 y9  (ycf7 ycf2 ycf3 );
  yeta4 by y10@1 y11 y12 (ycf10 ycf2 ycf3 );
  [y1@0 y2* y3*] (yci1 yci2 yci3 );
  [y4@0 y5* y6*] (yci4 yci5 yci6 );
  [y7@0 y8* y9*] (yci7 yci8 yci9 );
  [y10@0 y11 y12](yci10 yci11 yci12 );
  [yi] (ya1gp0) ;
  [ys] (ya2gp0);
  meta1 by m1@1 m2 m3 (mcf1 mcf2 mcf3 );
  meta2 by m4@1 m5 m6  (mcf4 mcf2 mcf3 );
  meta3 by m7@1 m8 m9  (mcf7 mcf2 mcf3 );
  meta4 by m10@1 m11 m12 (mcf10 mcf2 mcf3 );
  [m1@0 m2* m3*] (mci1 mci2 mci3 );
  [m4@0 m5* m6*] (mci4 mci5 mci6 );
```

```
   [m7@0 m8* m9*] (mci7 mci8 mci9 );
   [m10@0 m11 m12](mci10 mci11 mci12 );
   [mi] (ma1gp0) ;
   [ms] (ma2gp0) ;
   yi ON mi (p1gp0 ) ;
   ys ON ms (p2gp0);
 MODEL treat:
   yeta1 by y1@1 y2 y3 (ytf1 ytf2 ytf3 );
   yeta2 by y4@1 y5 y6  (ytf4 ytf2 ytf3 );
   yeta3 by y7@1 y8 y9  (ytf7 ytf2 ytf3 );
   yeta4 by y10@1 y11 y12 (ytf10 ytf2 ytf3 );
   [y1@0 y2* y3*] (yti1 yti2 yti3 );
   [y4@0 y5* y6*] (yti4 yti5 yti6 );
   [y7@0 y8* y9*] (yti7 yti8 yti9 );
   [y10@0 y11 y12](yti10 yti11 yti12 );
   [yi] (ya1gp1) ;
   [ys] (ya2gp1);
   meta1 by m1@1 m2 m3 (mtf1 mtf2 mtf3 );
   meta2 by m4@1 m5 m6  (mtf4 mtf2 mtf3 );
   meta3 by m7@1 m8 m9  (mtf7 mtf2 mtf3 );
   meta4 by m10@1 m11 m12 (mtf10 mtf2 mtf3 );
   [m1@0 m2* m3*] (mti1 mti2 mti3 );
   [m4@0 m5* m6*] (mti4 mti5 mti6 );
   [m7@0 m8* m9*] (mti7 mti8 mti9 );
   [m10@0 m11 m12](mti10 mti11 mti12 );
   [mi] (ma1gp1) ;
   [ms] (ma2gp1) ;
   yi ON mi (p1gp1 ) ;
   ys ON ms (p2gp1);
MODEL CONSTRAINT:
   NEW (yti2i5 yti2i8 yti2i11 yti5i8 yti5i11 yti8i11
   yci2i5 yci2i8 yci2i11 yci5i8 yci5i11 yci8i11
   mti2i5 mti2i8 mti2i11 mti5i8 mti5i11 mti8i11
   mci2i5 mci2i8 mci2i11 mci5i8 mci5i11 mci8i11 ) ;
   yti2i5=yti2-yti5 ;
   yti2i8=yti2-yti8 ;
   yti2i11=yti2-yti11 ;
   yti5i8=yti5-yti8 ;
   yti5i11=yti5-yti11 ;
   yti8i11=yti8-yti11 ;
   yci2i5=yci2-yci5 ;
   yci2i8=yci2-yci8 ;
   yci2i11=yci2-yci11 ;
   yci5i8=yci5-yci8 ;
   yci5i11=yci5-yci11 ;
   yci8i11=yci8-yci11 ;
   mti2i5=mti2-mti5 ;
   mti2i8=mti2-mti8 ;
   mti2i11=mti2-mti11 ;
```

```
   mti5i8=mti5-mti8 ;
   mti5i11=mti5-mti11 ;
   mti8i11=mti8-mti11 ;
   mci2i5=mci2-mci5 ;
   mci2i8=mci2-mci8 ;
   mci2i11=mci2-mci11 ;
   mci5i8=mci5-mci8 ;
   mci5i11=mci5-mci11 ;
   mci8i11=mci8-mci11 ;
OUTPUT: Samp StdYX Residual Tech4  ;
```

My overall conclusion is that for the current example, loading and measurement non-invariance are not problematic and do not require corrective actions. However, I offer more commentary to help put additional matters in perspective.

Some texts and methodologists argue that the imposition of measurement invariance based equality constraints for second order latent growth curve models is important enough that it should be enacted *a priori* without formally exploring the presence of non-invariance or its implications. As long as one obtains satisfactory global model fit in conjunction with the constraints, the argument goes, one moves forward with model evaluation accordingly.

There are counter-arguments to this approach. First, strict equality of factor loadings and measurement intercepts across groups and time is a very stringent condition that rarely holds in the real-world. *A priori* forcing of invariance constraints onto a model does not make non-invariance go away. Instead it often just redistributes the non-invariance to other model parameters that then can create bias in them. It also can lead to poor global model fit that results in the rejection of a potentially viable structural model. In my opinion, rather than trying to force measurement non-invariance into submission, it is best instead to determine if non-invariance is a problem using the methods in my measurement invariance document for Chapter 3. Research on partial invariance analytic strategies often offer viable remedies should non-invariance be found. These solutions can be invoked accordingly.

Second, it is well-known that evaluations of model fit should not rely strictly on global fit indices. It is the responsibility of the analyst to perform localized tests of ill fill and to probe fit at deeper levels than what global indices reflect. Measurement non-invariance can raise its ugly head in the context of localized explorations that are not detected by global fit analyses. Again, if non-invariance appears at the local level, then it often will be addressable using partial invariance strategies.

If statistically significant non-invariance occurs in a model, this does not necessarily mean it is consequential for estimating the key model parameters that are of substantive interest. In the measurement invariance document for Chapter 3, I describe methods by Oberski (2014) and Pornprasertmanit (2026) that provide perspectives on whether the observed non-invariance can reasonably be ignored because it has trivial consequences.

In the current case where no meaningful measurement non-invariance was observed, the question becomes whether the primary analysis should impose the traditional non-invariance equality constraints on the model in the form of loading invariance across time and across groups as well as measurement intercept invariance across time and across groups. The answer to this question is controversial.

An argument against introducing the equality constraints is that assuming strict loading and measurement intercept invariance in which values of loadings and intercepts are assumed to be exactly equal at, for example, all time points is unrealistic. To be sure, the loading values might be close in value but to say they are exactly equal lacks credulity. A fully unconstrained model is more realistic in that it allows for small variation in population loading values and intercept values that is likely to exist. As noted earlier, introducing the equality constraints in the model does not make the (small) levels of non-invariance disappear. Instead, it typically diverts it to other model parameters which can introduce bias into them.

An argument in favor of imposing constraints is that the unconstrained approach might overfit the data by being sensitive to sample-specific data idiosyncrasies. As well, the constrained model results in more degrees of freedom and fewer parameter estimates, which can increase statistical power and solution stability.

For the current example, I report the unconstrained analysis. Of course, one can always pursue sensitivity analyses and analyze the data both ways. Here is the Mplus syntax I used:

```
TITLE: LGCM with latent outcomes  ;
DATA: FILE = LGAlatentM.dat;
VARIABLE:
  NAMES = y1-y12 m1-m12 dtreat;
  USEVARIABLES = y1-y12 m1-m12;
  GROUPING IS DTREAT (0=control, 1=treat) ;
ANALYSIS: ESTIMATOR=MLR ;
MODEL:
  yeta1 by y1@1 y2 y3;        !define first order y factors
  yeta2 by y4@1 y5 y6;
  yeta3 by y7@1 y8 y9;
  yeta4 by y10@1 y11 y12;
  [yeta1-yeta4@0];            !fix intercepts of y factors to 0
  yeta1 yeta2 yeta3 yeta4  ;  !estimate var of y factors
  [y1@0 y2* y3*];             !define y measurement intercepts
  [y4@0 y5* y6*];
  [y7@0 y8* y9*];
  [y10@0 y11 y12];
  y1-y12 ;                                  !define y error variances
  yi by yeta1@1 yeta2@1 yeta3@1 yeta4@1;   !define y growth intercept
  ys by yeta1@0 yeta2@1 yeta3@2 yeta4@3;   !define y growth slope
  yi ys;           !estimate var of y latent intercept and latent slope
```

```
 [yi ys];         !estimate mean y latent intercept and y latent slope
 yi with ys;      !correlate y latent intercept with latent slope
 meta1 by m1@1 m2 m3;        !define first order m factors
 meta2 by m4@1 m5 m6;
 meta3 by m7@1 m8 m9;
 meta4 by m10@1 m11 m12;
 [meta1-meta4@0];            !fix intercepts of m factors to 0
 meta1 meta2 meta3 meta4  ;  !estimate var of m factors
 [m1@0 m2 m3];               !define m measurement intercepts
 [m4@0 m5 m6];
 [m7@0 m8 m9];
 [m10@0 m11 m12];
 m1-m12 ;                                  !define m error variances
 mi by meta1@1 meta2@1 meta3@1 meta4@1 ;  !define m growth intercept
 ms by meta1@0 meta2@1 meta3@2 meta4@3 ;  !define m growth slope
 mi ms;        !estimate var of m latent intercept and latent slope
 [mi ms];      !estimate mean y latent intercept and y latent slope
 mi with ms;   !correlate y latent intercept with latent slope
 yi ON mi ;    !regress latent intercepts
 ys ON ms ;    !regress latent slopes
MODEL control:  !assign group specific labels to some of above
 yeta1 by y1@1 y2 y3 (ycf1 ycf2 ycf3 );
 yeta2 by y4@1 y5 y6  (ycf4 ycf5 ycf6 );
 yeta3 by y7@1 y8 y9  (ycf7 ycf8 ycf9 );
 yeta4 by y10@1 y11 y12 (ycf10 ycf11 ycf12 );
 [y1@0 y2* y3*] (yci1 yci2 yci3 );
 [y4@0 y5* y6*] (yci4 yci5 yci6 );
 [y7@0 y8* y9*] (yci7 yci8 yci9 );
 [y10@0 y11 y12](yci10 yci11 yci12 );
 [yi] (ya1gp0) ;
 [ys] (ya2gp0);
 meta1 by m1@1 m2 m3 (mcf1 mcf2 mcf3 );
 meta2 by m4@1 m5 m6  (mcf4 mcf5 mcf6 );
 meta3 by m7@1 m8 m9  (mcf7 mcf8 mcf9 );
 meta4 by m10@1 m11 m12 (mcf10 mcf11 mcf12 );
 [m1@0 m2* m3*] (mci1 mci2 mci3 );
 [m4@0 m5* m6*] (mci4 mci5 mci6 );
 [m7@0 m8* m9*] (mci7 mci8 mci9 );
 [m10@0 m11 m12](mci10 mci11 mci12 );
 [mi] (ma1gp0) ;
 [ms] (ma2gp0) ;
 yi ON mi (p1gp0 ) ;
 ys ON ms (p2gp0);
MODEL treat:
 yeta1 by y1@1 y2 y3 (ytf1 ytf2 ytf3);
 yeta2 by y4@1 y5 y6  (ytf4 ytf5 ytf6);
 yeta3 by y7@1 y8 y9  (ytf7 ytf8 ytf9);
 yeta4 by y10@1 y11 y12 (ytf10 ytf11 ytf12);
 [y1@0 y2* y3*] (yti1 yti2 yti3);
```

```
  [y4@0 y5* y6*] (yti4 yti5 yti6);
  [y7@0 y8* y9*] (yti7 yti8 yti9);
  [y10@0 y11 y12](yti10 yti11 yti12);
  [yi] (ya1gp1) ;
  [ys] (ya2gp1);
  meta1 by m1@1 m2 m3 (mtf1 mtf2 mtf3);
  meta2 by m4@1 m5 m6  (mtf4 mtf5 mtf6);
  meta3 by m7@1 m8 m9  (mtf7 mtf8 mct9);
  meta4 by m10@1 m11 m12 (mtf10 mtf11 mtf12);
  [m1@0 m2* m3*] (mti1 mti2 mti3);
  [m4@0 m5* m6*] (mti4 mti5 mti6);
  [m7@0 m8* m9*] (mti7 mti8 mti9);
  [m10@0 m11 m12](mti10 mti11 mti12);
  [mi] (ma1gp1) ;
  [ms] (ma2gp1) ;
  yi ON mi (p1gp1) ;
  ys ON ms (p2gp1);
MODEL CONSTRAINT:
  NEW (mst msc msdiff meta1t meta2t meta3t meta4t
  meta1c meta2c meta3c meta4c mdiff1 mdiff2
  mdiff3 mdiff4) ;
  mst = ma2gp1 ;
  msc = ma2gp0 ;
  msdiff = mst-msc ;
  meta1t = ma1gp1 ;
  meta2t = meta1t + mst*1  ;
  meta3t = meta1t + mst*2  ;
  meta4t = meta1t + mst*3  ;
  meta1c = ma1gp0 ;
  meta2c = meta1c + msc*1 ;
  meta3c = meta1c + msc*2 ;
  meta4c = meta1c + msc*3 ;
  mdiff1 = meta1t - meta1c ;
  mdiff2 = meta2t - meta2c ;
  mdiff3 = meta3t - meta3c ;
  mdiff4 = meta4t - meta4c ;
  NEW (yst ysc ysdiff yeta1t yeta2t yeta3t yeta4t
  yeta1c yeta2c yeta3c yeta4c ydiff1 ydiff2
  ydiff3 ydiff4) ;
  yst = ya2gp1 + (ma2gp1)*p2gp1 ;
  ysc = ya2gp0 + ma2gp0*p2gp0 ;
  ysdiff = yst-ysc ;
  yeta1t = ya1gp1 + ma1gp1*p1gp1;
  yeta2t = yeta1t + yst*1;
  yeta3t = yeta1t + yst*2;
  yeta4t = yeta1t + yst*3;
  yeta1c = ya1gp0 + ma1gp0*p1gp0;
  yeta2c = yeta1c + ysc*1;
  yeta3c = yeta1c + ysc*2;
```

```
  yeta4c = yeta1c + ysc*3;
  ydiff1 = yeta1t - yeta1c ;
  ydiff2 = yeta2t - yeta2c ;
  ydiff3 = yeta3t - yeta3c ;
  ydiff4 = yeta4t - yeta4c ;
OUTPUT: Samp StdYX Residual Tech4  ;
```

All of the syntax should be familiar. The `MODEL CONSTRAINT` commands follow the logic outlined in the prior example on multiple group analysis of LGM in RETs.

Choosing a Reference Indicator

The choice of a reference indicator in second order latent growth curve models is critical because it ultimately dominates the analysis of the mean structure of the latent factors. One criterion that is important when selecting a reference indicator is that it should be reasonably measurement invariant as reflected in the analyses described in the prior section. A second criterion is that the measure should have solid reliability and validity data backing it up. A third criterion is that the metric of the measure should be inherently meaningful or it should have acquired meaning from prior research, i.e., the metric should be non-arbitrary (Blanton & Jaccard, 2006). A given measure may not maximize all three of these desiderata, but we seek to do the best we can, all things considered.

The RET Analysis

I now turn to the standard questions we ask of mediation analyses in RETs, namely (1) does the intervention meaningfully impact the outcome, (2) does the intervention meaningfully affect the mediators or mechanisms through which the intervention is presumed to impact the outcome, and (3) do the mediators meaningfully impact the outcome.

Total Effect of the Treatment on the Outcome

To test the effects of the intervention on the outcome, I first examine treatment condition differences for Y at the immediate posttest. I obtain this contrast and the separate means for the intervention and control groups from the `MODEL CONSTRAINT` commands. Here is the relevant (edited) output:

```
                                                    Two-Tailed
                   Estimate       S.E.  Est./S.E.    P-Value

New/Additional Parameters
    YDIFF1             0.494      0.033     14.947      0.000
    YETA1T             1.033      0.023     44.089      0.000
    YETA1C             0.539      0.023     23.113      0.000
```

Note that my focus is on the latent variables not the observed Y *per se*. The adjusted mean difference is 0.49 ±0.07, which is statistically significant (CR = 14.95, p < 0.05). The 95% confidence interval is 0.42 to 0.56. Suppose that the standard for a meaningful population mean difference at the immediate posttest is a population absolute difference of 0.20 or greater. Because the lower limit of the mean difference confidence interval was greater than this standard, I conclude that the intervention effect for the immediate posttest is meaningful. The estimated Y latent means for the intervention and control conditions at the immediate posttest were 1.03 ± 0.05 and 0.54 ± 0.05, respectively.

Next, I examined a second facet of the effect of the intervention on the outcome, namely whether there was meaningful decay in the intervention effect over time. As a first step, I isolate the coefficient characterizing the decay curve for the intervention group and for the control group. I then examine a contrast comparing the two coefficients. The relevant values are obtained in the `MODEL CONSTRAINT` command. Here is the edited output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    YST               -0.030      0.008     -3.637      0.000
    YSC                0.018      0.008      2.178      0.029
    YSDIFF            -0.048      0.012     -4.117      0.000
```

The model implied linear decay in the intervention condition is negative and statistically significant (slope = -0.03 ± 0.02, CR = 3.64, p < 0.05) indicating that the predicted latent mean of Y decreases as time from the immediate posttest increases. For the control condition, the change in the latent mean of Y across time is relatively flat although there is a slight positive increase (slope = 0.018 ±0.016, CR = 2.18, p < 0.05). The difference between the coefficients (-0.048 ±0.02) was statistically significant (CR = 4.12, p < 0.05).

I can plot the two curves using the information in the output section from the `MODEL CONSTRAINT` commands for the predicted latent means `YETA1T` through `YETA4T` for the intervention group and the corresponding means for the control group, `YETA1C` through `YETA4C`. Using the program on my website called *Temporal line plot*, here is the plot:

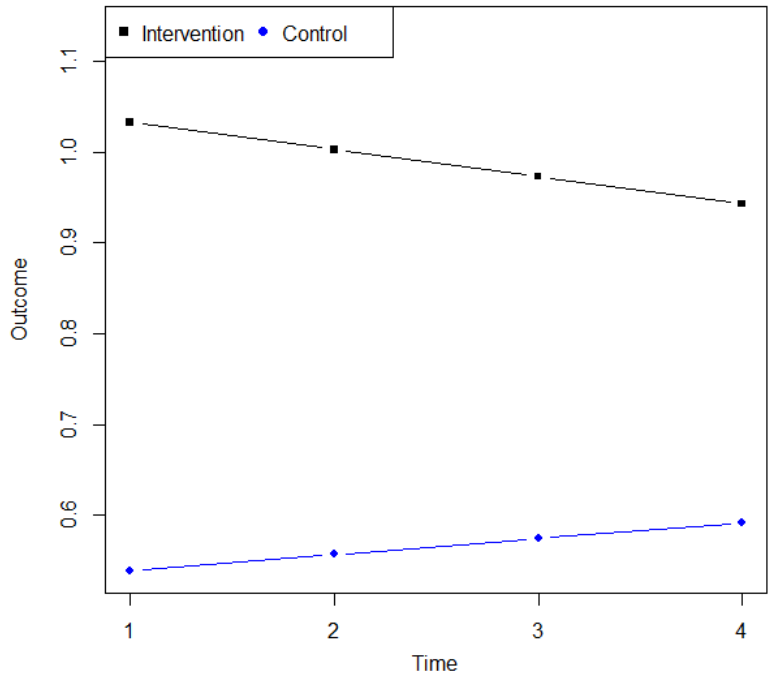

The downward decay in the predicted latent Y means for the intervention group is evident as is the slight upward trend in the means for the control group. The separation between lines at a given time reflects the predicted mean difference between the intervention and control groups. It can be seen from the plot that the mean difference between the treatment and control conditions decreases with time.

The contrasts that I specified in the `MODEL CONSTRAINT` commands include tests of the outcome mean difference at each of the four time points using the model-informed predicted latent means. The contrasts are labeled `YDIFF1`, `YDIFF2`, `YDIFF3`, `YDIFF4` for times 1, 2, 3 and 4, respectively (keep in mind that time 1 is the immediate posttest). Here is the output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    YDIFF1             0.494      0.033     14.947      0.000
    YDIFF2             0.446      0.029     15.580      0.000
    YDIFF3             0.399      0.029     13.937      0.000
    YDIFF4             0.351      0.033     10.654      0.000
```

The intervention group predicted mean is statistically significantly larger than the control group mean at all four time points. Suppose a meaningfulness standard is set by the research team to be a population absolute difference of 0.20 or greater. For all contrasts, the

lower limit of the 95% confidence interval (roughly, the `Estimate` minus double the `S.E.`) exceeds the meaningfulness standard. Given this, I conclude these Y predicted mean differences are meaningful.

Effect of the Treatment on the Mediator

To test the effects of the intervention on the mediator, I follow the same structure as that for the total effect on the outcome. I first examine treatment condition differences for the mediator at the immediate posttest. I obtain this contrast from the `MODEL CONSTRAINT` commands. Here is the relevant output:

```
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

New/Additional Parameters
    MDIFF1             0.318      0.029     10.846      0.000
    META1T             0.596      0.021     28.912      0.000
    META1C             0.278      0.021     13.323      0.000
```

The adjusted mean difference is 0.32 ±0.06, which is statistically significant (CR = 10.85, $p < 0.05$). The 95% confidence interval is 0.20 to 0.44. Suppose the standard for a meaningful mean difference for the mediator at the immediate posttest is a population absolute difference of 0.20 or greater. Because the lower limit of the mean difference confidence interval meets this standard, I conclude that the intervention effect for the immediate posttest is meaningful. The mediator means for the intervention and control conditions at the immediate posttest were 0.60 ±0.04 and 0.28 ±0.04, respectively.

Next, I examined a second facet of the intervention effect on the mediator, namely whether there was meaningful decay in the effect of the intervention on the mediator over time. As a first step, I isolate the coefficient for mediator decay for the intervention group and that for the control group. I then examine a contrast comparing the two coefficients. The relevant values are from the `MODEL CONSTRAINT` command. Here is the edited output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    MST               -0.025      0.008     -3.239      0.001
    MSC                0.026      0.009      3.031      0.002
    MSDIFF            -0.051      0.011     -4.419      0.000
```

The model implied linear decay for the mediator in the intervention condition is negative and statistically significant (slope = -0.025 ±0.02, CR = 3.24, $p < 0.05$) indicating

that the predicted latent mean for the mediator only slightly as time from the immediate posttest increases. For the control condition, the slope for the latent mean of M also was relatively flat but there was a statistically significant tendency for the latent mean to increase somewhat (slope = 0.026 ±0.02, CR = 3.03, $p < 0.05$). The difference between the two coefficients (-0.05 ±0.02) was statistically significant (CR = 4.42, $p < 0.05$).

I can plot the two curves using the information in the output section from the `MODEL CONSTRAINT` commands for the predicted latent means `META1` though `META4` for the intervention group (labeled `MESTA1T`, `META2T`, `META3T`, and `META4T`) and for the control group (labeled `META1C`, `META2C`, `META3C`, and `META4C`). Using the program on my website called *Temporal line plot*, here is the plot:

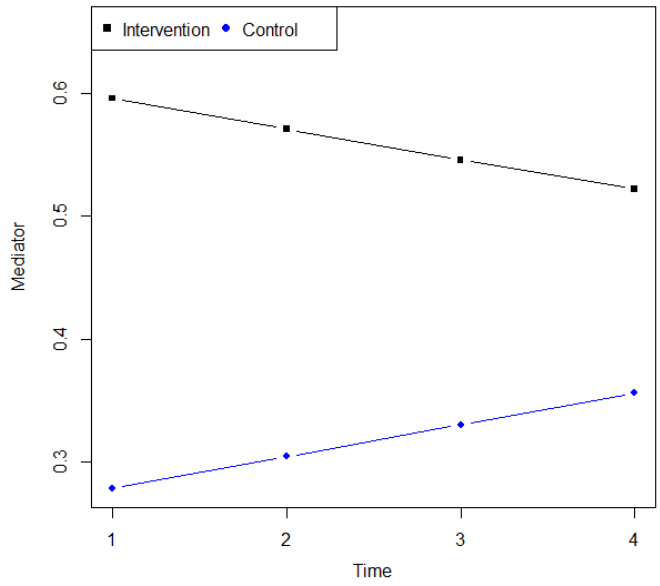

The slight downward decay in the predicted mediator latent means for the intervention group is evident as is the upward trend in the means for the control group. The separation between lines at a given time point reflects the predicted latent mean difference between the intervention and control groups. The latent mean difference between the treatment and control conditions decreases as time passes.

The contrasts that I specified in the `MODEL CONSTRAINT` commands also test the mediator latent mean difference at each of the four time points using the model-informed predicted latent means. The contrasts are labeled `MDIFF1`, `MDIFF2`, `MDIFF3`, `MDIFF4` for times 1, 2, 3 and 4, respectively (keep in mind that time 1 is the immediate posttest). Here is the relevant output:

```
New/Additional Parameters
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

    MDIFF1             0.318      0.029     10.846      0.000
    MDIFF2             0.267      0.024     10.994      0.000
    MDIFF3             0.217      0.024      8.941      0.000
    MDIFF4             0.166      0.029      5.698      0.000
```

The intervention group predicted mediator mean is statistically significantly larger than the control group predicted mean at each time point. Suppose a meaningfulness standard is set by the research team to be a population absolute difference of 0.20 or greater. The 95% confidence interval for `MDIFF1` is 0.25 to 0.38. Because the lower limit exceeds the meaningfulness standard, this effect is deemed meaningful. This is also true for the `MDIFF2` contrast. However, this cannot be said for the remaining contrasts.

Effect of the Mediator on the Outcome

For the multi-group model, the analysis estimates the effect of the mediator on the outcome separately for the intervention and control groups. The relevant coefficients are taken from the `MODEL RESULLTS` section and appear as follows (edited):

```
MODEL RESULTS
                                                    Two-Tailed
                    Estimate       S.E.  Est./S.E.    P-Value

Group CONTROL

  YI ON MI             0.737      0.056     13.174      0.000
  YS ON MS             0.782      0.454      1.723      0.085

Group TREAT

  YI ON MI             0.783      0.050     15.587      0.000
  YS ON MS             0.865      0.429      2.018      0.044
```

The coefficient for MI→YI is the path coefficient for the outcome at the immediate posttest regressed onto the mediator at the immediate posttest controlling for the baseline outcome. It was statistically significant and meaningful (using a meaningfulness standard of 0.50) in both the intervention and control conditions. The coefficient for MS→YS is the path coefficient for the decay slope of the outcome regressed onto the decay slope for the mediator. A positive coefficient implies that people with more positive slopes on the mediator tend to have more positive slopes on the outcome, i.e., that there is an association

between the two. The path coefficient was about the same magnitude as for the immediate posttest but only one of them was statistically significant, the other marginally so. Overall, I would conclude there is evidence for a link between the mediator and the outcome but it is somewhat mixed.

Omnibus Mediation Analysis

Omnibus mediation tests are of lower priority for me. I instead prefer to focus on the individual links in mediational chains, evaluating their effects sizes. If I want an overall omnibus test of statistical significance, I rely on the joint significance test. In LGC models, evaluations of mediation take two forms. The first focuses on treatment differences in the latent intercept for the mediator and whether these differences, in turn, translate into treatment differences on the outcome latent intercept. The second mediation form focuses on the effect of the latent slopes of the mediator on the latent slopes of the outcome. The RET showed evidence towards such omnibus mediation.

COVARIATE INCLUSION IN LGCMs

Inclusion of covariates in LGCMs is straightforward but decisions about them need to be made based on the questions you seek to answer. Generally speaking, there are two types of covariates, time-invariant and time-varying. A **time-invariant covariate** is one that does not change over the course of the growth period being studied, such as race, biological sex, place of birth, certain genetic markers, and the like. The treatment condition one is assigned to is a time-invariant variable in all of the growth curve examples in this document but it is a focal independent variable, not a nuisance covariate. Time-invariant covariates basically have the same value for an individual across all the waves of the target growth period for all individuals. A **time-varying covariate** is one that takes on two or more different values over the course of the target growth period. Examples include health status, stress, employment status, and emotional states. The definition of time-invariant and time-varying covariates can vary as a function of substantive theory and/or the characteristics of the empirical data. If a potentially time-varying covariate happens to remain constant across the time periods for all individuals in a study, it essentially takes on the property of a time invariant covariate.

To control for time-invariant confounds, one typically includes them as predictors of the latent intercept or latent slope, much like I did in the examples in this document when I controlled for baseline mediators and outcomes. In theory, one could instead include them as common cause predictors of each observed variable in the growth model directly, but this is done less often. Doing so serves to reduce or eliminate correlated disturbances between

the observed variables.

Time varying covariates can be introduced via the logic of parallel process modeling, as I did with mediators in the examples of this document. An alternative strategy is shown in Figure 10 which, for pedagogy, focuses only on a single growth outcome. To reduce clutter in the figure, I omit the circled disturbance/error terms, instead representing their presence with just an arrow pointing to the appropriate endogenous variable. I also omit correlations among all exogenous variables (including the latent slope and intercept in this case) to reduce clutter but they would take them into account in Mplus modeling. In this approach, the Y at a given time point is directly regressed onto the Z at the same point. This means that the observed Y are now conceptualized as a *joint* function of the latent intercept and slope factors as well as the time-specific influences of Z. A way of thinking about this conceptualization is that one is examining the growth parameters of Y net the effect of Z. Alternatively, one can conceptualize the model as examining the effect of Z on Y at a given time point net the latent growth parameters (Bollen & Curran, 2006).

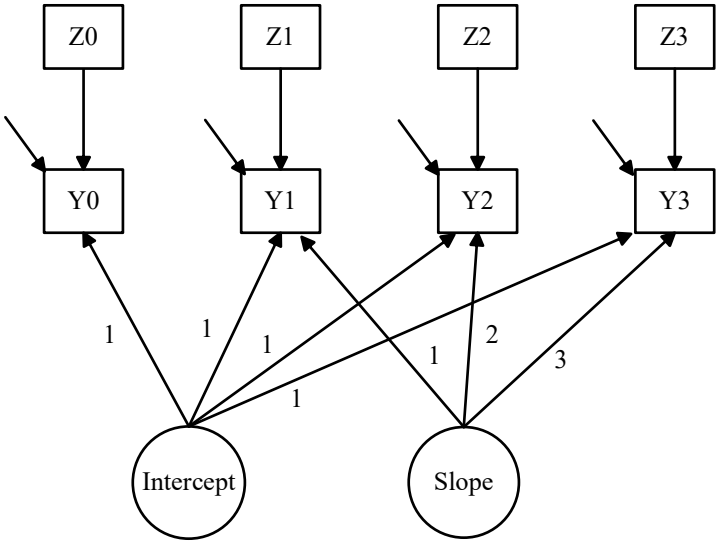

FIGURE 10. Second order latent growth curve model

The approach you use to model and interpret covariates depends on the substantive questions you seek to address. Latent growth curve models offer flexibility in this regard.

.

ALTERNATIVES TO TRADITIONAL LGCM

There are several variants of the traditional latent growth curve model. One that has received considerable attention is the **autoregressive latent trajectory (ALT) model** by Bollen and Curran (2006). The model combines latent growth modeling with traditional panel modeling and requires at least 4 time points to avoid under-identification. Figure 11 presents a simple example that uses only a single variable Y measured at four time points. Note that I added first order autoregressive paths to the observed Y. The idea is that the covariation among the various Y can be explained by the joint impact of (a) the developmental process captured by the latent intercept and slope factors, and (b) the first autoregressive effects of Y at time *t-1* on Y at time *t*. This integrative mindset opens numerous modeling possibilities in RETs. Consideration of them is beyond the scope of my book. For details on autoregressive latent trajectory models, see Bollen and Curran (2004, 2006), Bollen and Zimmer (2010), Bauldry and Bollen (2018) and Newsom (2023).

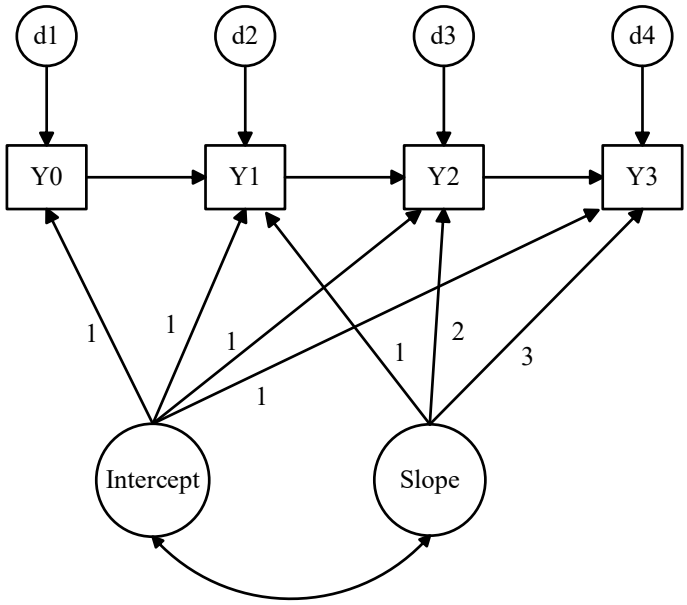

FIGURE 11. Autoregressive latent trajectory model

CONCLUDING COMMENTS

Latent growth curve modeling is a viable approach to analyzing longitudinal data in RETs when your interest is in making statements about trajectories. This document is an

introduction to the basics of such modeling. For an excellent treatment of latent growth modeling more generally, see Newsom (2023).

My examples in this document are simplistic in that they used only a single mediator, they downplayed covariate control, and they largely ignore correlated disturbances. I encourage you to carefully think through the possible role of correlated disturbances in your LGCMs as they tend to be downplayed in most tutorials. Often the assumption is made that correlations between disturbances among the observed variables in a trajectory are minimal due to lengthy lags between them. But this is not always the case. If the theoretical data generating process is akin to an ALT model, for example, and the autoregressive effects are omitted from the model, then the omitted effects will manifest themselves in correlated disturbances. Most of my examples did not articulate rationales for correlated disturbances between endogenous latent slopes and latent intercepts (or lack of such correlations) and this is an issue you will want to consider carefully in your modeling efforts.

My examples tended to use large sample sizes to ensure convergence and well behaved parameter estimates for the hypothetical data I generated. I did so for pedagogical reasons. However, you will want to carefully explore issues surrounding sample size for your LGCMs, ideally using the localized simulation strategies outlined in Chapter 28. See the discussions by Rast and Hofer (2014) for tentative guidelines.

All of the examples I consider in this document used designs where the time intervals between measures were functionally the same for all individuals. There is a literature on how to handle cases where the time intervals vary from one individual to the next. For a summary of these methods, see Sterba (2014). Mplus offers a procedure based on its TSCORE option that accommodates these types of models.

My examples relied strictly on maximum likelihood methods, usually the robust variant of them defined by the MLR option in Mplus. Almost all of the analyses can be pursued using Bayesian estimation instead. For an introduction to Bayesian SEM, see Chapter 8 and 25 of the main text.

My examples also tended to use straightforward quantitative coding to link the latent slope factor to the observed Y. Methodologists have discussed a wider range of coding strategies for LGCMs. See, for example, Biesanz et al. (2004) and Hancock and Choi (2006).

Latent growth curve modeling has interesting connections to multilevel modeling strategies. For a comparison of the approaches, see McNeish and Matta (2018).

Finally for an interesting description of the history of growth curve modelling, see Bollen (2007).

REFERENCES

Agresti, A. (1996). *An introduction to categorical data analysis*. Wiley.

Atkinson, A. C., Riani, M., & Corbellini, A. (2021). The Box–Cox transformation: Review and extensions. *Statistical Science*, 36, 239-255.

Bendler, J. & Reinecke, J. (2025). A tutorial on Bayesian multiple-group comparisons of latent growth curve models with count distributed variables. *Behavioral Research Methods*, 57, 112-122.

Bollen, K. & Curran, P. (2004). Autoregressive latent trajectory (ALT) models: A synthesis of two traditions. *Sociological Methods and Research*, 32, 336–383.

Bollen, K. & Curran, P. (2006). *Latent curve models : A structural equation perspective*. Wiley.

Bauldry, S. & Bollen, K. (2018). Non-linear autoregressive latent trajectory models. *Sociological Methodology,* 48, 269–302

Biesanz, J., Deeb-Sossa, N., Papadakis, A., Bollen, K. & Curran, P. (2004). The role of coding time in estimating and interpreting growth curve models. *Psychological Methods*, 9, 30–52.

Bollen, K. (2007). On the origins of latent curve models. In R. Cudeck & R. MacCallum (Eds.), *Factor analysis at 100* (pp. 79–98). Erlbaum

Bollen, K. & Zimmer, C. (2010). An overview of the autoregressive latent trajectory (ALT) model. In K. van Montfort, J. Oud & A. Satorra (Eds.), *Longitudinal research with latent variables* (pp. 153–176). Springer-Verlag.

Box, G. & Cox, D. (1964). An analysis of transformation. *Journal of the Royal Statistical Society: Series B (Methodological)*, 26, 211–243.

Grimm, K. & Liu, Y. (2016),=. Residual structures in growth models with ordinal outcomes. *Structural Equation Modeling*, 23, 466-475.

Grimm, K. & Ram, N. (2009). Nonlinear growth models in Mplus and SAS. *Structural Equation Modeling*, 16, 676–701,

Harring, J. Strazzeri, M. & Blozis, S. (2021). Piecewise latent growth models: Beyond modeling linear-linear processes. *Behavior Research Methods*, 53, 593–608.

Hancock, G. R., & Choi, J. (2006). A vernacular for linear latent growth models. *Structural Equation Modeling*, 13, 352–377.

Kim, J. & Whittaker, T. (2025) Latent growth model with categorical indicators: Comparison of the different estimators. *Structural Equation Modeling*, 32, 400-414.

Lee, T., Wickrama, K., & O'Neal, C. (2018). Application of latent growth curve analysis with categorical responses in social behavioral research. *Structural Equation Modeling, 25*, 294–306.

Lim, K. & Kim, S. (2024). Scaling and estimation of latent growth models with categorical indicator variables. *Psychological Methods.* Advance online publication.

McNeish, D. (2020). Relaxing the proportionality assumption in latent basis models for nonlinear growth. *Structural Equation Modeling,* 27, 817–824.

McNeish, D., & Matta, T. (2018). Differentiating between mixed-effects and latent-curve approaches to growth modeling. *Behavior Research Methods*, 50, 1398–1414.

Newsom, J. (2023). *Longitudinal structural equation modeling*. Taylor and Francis.

Oberski, D. (2014). Evaluating sensitivity of parameters of interest to measurement invariance in latent variable models. *Political Analysis*, 22, 45-60.

Pornprasertmanit, S. (2026). A note on effect size for measurement invariance. Downloaded from https://cran.r-project.org/web/packages/semTools/vignettes/partialInvariance.pdf

Preacher, K. J. (2019). Latent growth curve models. In G. R. Hancock, L. M. Stapleton, & R. O. Mueller (Eds.), *The reviewer's guide to quantitative methods in the social sciences* (2nd ed., pp. 178–192). Routledge/Taylor & Francis Group.

Preacher, K., Wichman, A., MacCallum, R. & Briggs, N. (2008). *Latent growth curve modeling*. SAGE.

Rast, P., & Hofer, S. M. (2014). Longitudinal design considerations to optimize power to detect variances and covariances among rates of change: Simulation results based on actual longitudinal studies. *Psychological Methods*, 19, 133–154.

Singer, J & Willett, J. (2003). *Applied longitudinal data analysis: Modeling change and event occurrence. Oxford University Press*.

Sterba, S. K. (2014). Fitting nonlinear latent growth curve models with individually varying time points. *Structural Equation Modeling, 21*, 630–647.

Wilcox, R. (2021). *Introduction to robust estimation and hypothesis testing*. Academic Press (Fifth edition).

Willitt, J. & Bub, K. (2025), Structural equation modeling: Latent growth curve analysis**.** In the *Encyclopedia of statistics in behavioral science*. Wiley. Downloaded from https://userweb.ucs.louisiana.edu/~rmm2440/Willett_and_Bub_2004.pdf

Wu, W. & Lang, K. (2016) Proportionality assumption in latent basis curve models: A cautionary note. *Structural Equation Modeling*, 23, 140-154,