Example 5b: Time-Invariant and Time-Varying Predictors of Within-Person Change in General Multilevel Models for Two-Level Longitudinal Outcomes (complete syntax, data, and output available for R electronically)

The data for this example (from the Hoffman 2026 chapter, "A Rosetta Stone for Modeling Change") are simulated to resemble a four-year longitudinal study of 200 children (ages 8 to 12 at baseline), each measured on four annual occasions, for a total of 800 observations. Occasions were intended to be one year apart, but the timing was imperfect: The exact time since baseline varies by about ±0.33 years around each nominal year (i.e., unbalanced time), so a child's age at each occasion contains both between-person information (children started at different ages) and within-person information (time passing during the study). We will examine how depression (*dep*) changes over time, whether that change differs by age at baseline (*baseage*), and how it relates to the number of hours per day spent video gaming (*game*), a time-varying predictor. It uses several custom functions to process output and add additional computations, as shown below (all in the online files).

R Syntax for Packages; Importing and Preparing Data for Analysis

```
# Install and load required packages
if (!require(pacman)) install.packages("pacman")
pacman::p_load(readxl, TeachingDemos, lme4, lmerTest, nlme, ggplot2)
# readxl: import .xls or .xlsx files; TeachingDemos: save output to text file
# lme4/lmerTest: MLMs with random intercepts and slopes (Satterthwaite DDF)
# nlme: gls for the UN (saturated) covariance model, which lmer cannot estimate
# ggplot2: figures of predicted values

# Set working directory (to import and export files to)
setwd("C:/Dropbox/26_EDF9850/Unit5/Example5b")

# Load R functions for this class from R file in folder above working directory
source("../../EDF9850_Functions.R")  # Remove ../ for same folder
source("https://www.lesahoffman.com/EDF9850/EDF9850_Functions.R") # From the web

# Import Example 5b excel data file in sheet "Data"
Example5b = read_excel(path="Example5b_Data.xlsx", sheet="Data")
# Convert to data frame to use for analysis
Example5b = as.data.frame(Example5b)

# Create L2 person mean gaming
Example5b = addUnitMeans(data=Example5b, unitVariable="PersonID",
                         meanVariables="game", newNames="PMgame")

# Create L1 within-person gaming using person-mean-centering
Example5b$WPgame = Example5b$game - Example5b$PMgame

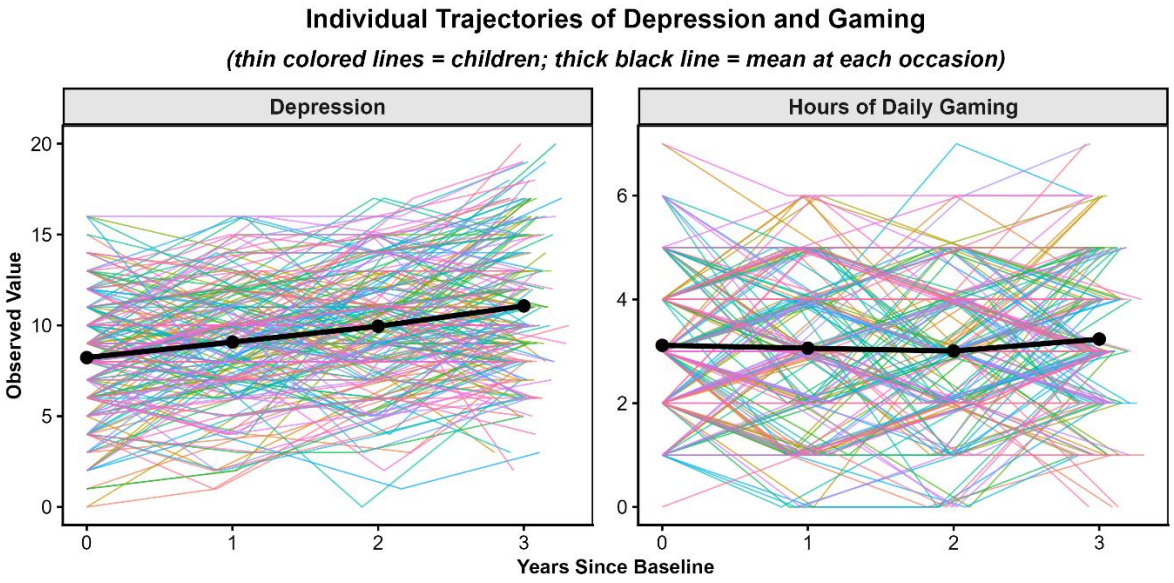
# Center L2 person mean gaming at 3 hours for the between-person effect
Example5b$PMgame3 = Example5b$PMgame - 3

# Center L2 baseline age at 10 so that the intercept is for a 10-year-old at baseline
Example5b$baseage10 = Example5b$baseage - 10

print("Descriptive Statistics from Occasion-Level-1 Data")
univSummary(data=Example5b, vars=c("baseage","time","dep","PMgame","WPgame"))
```

```
Variable    N  Mean    SD Variance    Min    Max
 baseage  800 9.899 0.646    0.418  8.200 12.120
 time     800 1.500 1.121    1.258  0.000  3.310
 dep      800 9.581 3.611   13.042  0.000 20.000
 PMgame   800 3.104 1.061    1.125  0.750  6.250
 WPgame   800 0.000 0.869    0.755 -2.500  2.500
```

See code online for how to make the "spaghetti" plots below with observed means overlayed in heavy black lines



Model 1: Empty Means, Random Intercept Model for Depression

Level 1: $Dep_{ti} = \beta_{0i} + e_{ti}$

Level 2: $\beta_{0i} = \gamma_{00} + U_{0i}$

```
print("Model 1: Empty Means, Random Intercept for Person, to Get ICC for Depression")
Model1 = lmer(data=Example5b, REML=TRUE, formula=dep~1+(1|PersonID))
lmerSummary(Model1, ICC1=TRUE, LRT=TRUE)
```

```
Model Fit for Predicting dep:
      AIC       BIC    logLik -2*logLik
 4143.157  4157.211 -2068.579  4137.157

Model for the Variance:
 Groups    Name        Variance Std.Dev
 PersonID  Intercept    5.905   2.430
 Residual               7.160   2.676

Fixed Effects:
 Term         Est    SE      DF t.value p.value
 Intercept  9.581 0.196 199.000  48.847   <.001

ICC1 intraclass correlation:
    Adjusted ICC: 0.452
  Unadjusted ICC: 0.452

LRT for random effect(s):
 Random Effect Dropped  Npar    LL -2LL  AIC  BIC     LRT DF p.value p.mix
 <none>                    3 -2069 4137 4143 4157
 Intercept | PersonID      2 -2163 4326 4330 4340 188.975  1   <.001 <.001
```

$$ICC_1 = \frac{\tau^2_{U_0}}{\tau^2_{U_0} + \sigma^2_e} = \frac{5.905}{5.905 + 7.160} = .\mathbf{452}$$

Model 2: Saturated Means over Occasion, Unstructured Variances and Covariances—*This is an approximate answer key as to what the fixed and random effects of time will be capturing!*

Composite: $Dep_{ti} = \gamma_{00} + \gamma_{10}(Occ2_{ti}) + \gamma_{20}(Occ3_{ti}) + \gamma_{30}(Occ4_{ti}) + e_{ti}$

Residuals: $e_i \sim MVN(\mathbf{0}, R_i)$, $R_i = unstructured$ $(4\ variances\ and\ 6\ covariances)$

Occ2, Occ3, and Occ4 are binary contrasts for occasions 2, 3, and 4 (occasion 1 is the reference), ignoring differences in exact time for the sake of plotting an "average" trajectory. Model 2 has no random effects: Instead, the person-level dependency is represented entirely by the unstructured R matrix (all 4 variances and all 6 covariances are estimated), creating a single-level multivariate model. The rounded-time means are close to what the fixed time slopes will be trying to reproduce, and the rounded-time variances are close to what the model for the variance will be trying to reproduce.

```
print("Model 2: Saturated Means, UN Variances and Correlations by Rounded Occasion")
Model2b = gls(data=Example5b, method="REML", model=dep~1+factor(occasion),
              correlation=corSymm(form=~occasion|PersonID), # Unstructured correlations
              weights=varIdent(form=~1|occasion))       # Separate variance by occasion
glsSummary(Model2, predMeans=TRUE)
```

```
Model Fit for Predicting dep:
      AIC      BIC    logLik -2*logLik
 3996.206 4061.721 -1984.103  3968.206

Overall F-test for All Fixed Slopes:
 Model   NumDF   DenDF F.value p.value
 Model2      3 198.858  35.875   <.001
```

I've omitted additional output related to the variances and correlations because it is unnecessarily confusing—see my replacement below.

The F-test below is for the "omnibus" main effect of time in ANOVA terms.

```
Fixed Effects:
 Term                  Est    SE      DF t.value p.value
 Intercept           8.225 0.237 206.548  34.725   <.001
 factor(occasion)2   0.855 0.207 199.035   4.126   <.001
 factor(occasion)3   1.725 0.239 199.128   7.230   <.001
 factor(occasion)4   2.845 0.277 198.292  10.257   <.001

Model-Predicted Means for factor(occasion):
 Level   Mean    SE      DF
 1      8.225 0.237 206.548
 2      9.080 0.230 194.243
 3      9.950 0.239 199.421
 4     11.070 0.271 200.661
```

```
print("R matrix (variances and covariances) and correlations for first person")
getVarCov(Model2, individual="1", type="marginal")
corMatrix(Model2$modelStruct$corStruct)[[1]]
```

```
Marginal variance covariance matrix
        [,1]    [,2]    [,3]    [,4]
[1,] 11.2200  6.6101  5.6294  5.2505
[2,]  6.6101 10.5870  6.0844  6.0597
[3,]  5.6294  6.0844 11.4250  7.9733
[4,]  5.2505  6.0597  7.9733 14.6680
```

The R matrix shows increasing variances from occasion 2 and declining correlations over time (rather than both being constant, as is implied by only a random intercept).

```
          [,1]      [,2]      [,3]      [,4]
[1,] 1.0000000 0.6064888 0.4972048 0.4092631
[2,] 0.6064888 1.0000000 0.5532491 0.4862755
[3,] 0.4972048 0.5532491 1.0000000 0.6159253
[4,] 0.4092631 0.4862755 0.6159253 1.0000000
```

This model is used only for descriptive purposes, since it ignores differences in exact years since baseline (the unbalanced timing, ±0.33 years). The remaining models will use fixed and random effects of exact time instead. Based on the pattern of means, let's try a linear model of change first.

Model 3a: Random Linear Change over Exact Time

Level 1: $Dep_{ti} = \beta_{0i} + \beta_{1i}(Time_{ti}) + e_{ti}$

Level 2: $\beta_{0i} = \gamma_{00} + U_{0i}$

$\beta_{1i} = \gamma_{10} + U_{1i}$

Note that time is created using baseline-centering (see Hoffman 2015 ch. 10) as: $Time_{ti} = Age_{ti} - BaseAge_i$

```
print("Model 3a: Random Linear Change over Exact Time")
Model3a = lmer(data=Example5b, REML=TRUE, formula=dep~1+time+(1+time|PersonID))
lmerSummary(Model3a, LRT=TRUE, randCI=TRUE)
```

```
Model Fit for Predicting dep:
      AIC       BIC    logLik -2*logLik
 3976.546  4004.653 -1982.273  3964.546

Model for the Variance:
 Groups    Name        Variance Std.Dev   Corr
 PersonID  Intercept    7.045   2.654
           time         0.793   0.890 -0.321
 Residual               4.341   2.084

Fixed Effects:
 Term         Est    SE      DF t.value p.value
 Intercept  8.156 0.225 198.742  36.327   <.001
 time       0.948 0.091 199.507  10.405   <.001

LRT for random effect(s):
 Random Effect Dropped  Npar    LL -2LL  AIC  BIC    LRT DF p.value p.mix
 <none>                    6 -1982 3965 3977 4005
 time | PersonID           4 -1999 3998 4006 4025 33.241  2   <.001 <.001

Random Effect 95% CIs (Fixed Effect +/- 1.96*SQRT(random effect variance)):
 Term                 FixedEffect Variance    SD  Lower  Upper
 Intercept | PersonID       8.156    7.045 2.654  2.953 13.358
 time | PersonID            0.948    0.793 0.890 -0.797  2.693
```

The random time slope variance (0.793, SD = 0.890) indicates that persons differ in their linear change: assuming normally distributed random slopes, 95% of children would be expected to have a rate of change between 0.948 ± 1.96(0.890) = −0.797 and 2.692 per year. The intercept and time slope were negatively correlated (r = −.321), so children who started higher in depression tended to increase less.

```
# Model-implied means: Model 2 (occasion means) and Model 3a (fixed effects only)
MeanM2  = predict(Model2, newdata=data.frame(occasion=1:4))
MeanM3a = predict(Model3a, newdata=data.frame(time=0:3), re.form=NA)

# Model-implied total variances: Model 2 (diagonal of its R matrix) and Model 3a
# (diagonal of its marginal V matrix, V = Z G Z' + sigma^2 I, given by marginalV)
VarM2 = diag(as.matrix(getVarCov(Model2, individual="1", type="marginal")))
V3a = marginalV(Model3a, times=0:3); VarM3a = diag(V3a$V)
print("Predicted means and variances: Model 2 versus Model 3a")
print(data.frame(time=0:3, MeanModel2=MeanM2, MeanModel3a=MeanM3a,
                 VarModel2=VarM2, VarModel3a=VarM3a, row.names=NULL), digits=5)
```

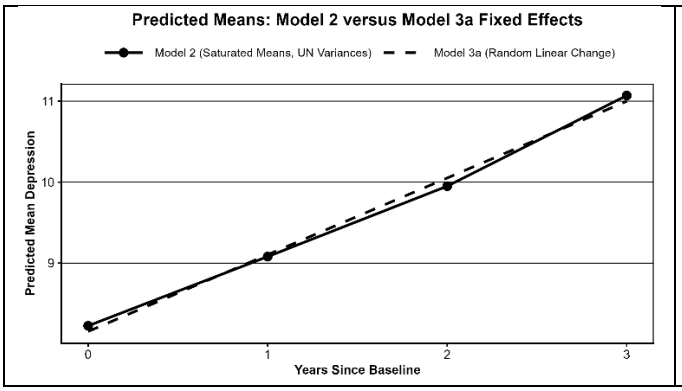
```
Model-Implied Marginal Variance-Covariance Matrix (V) for Model3a:
       0      1      2      3
0 11.386  6.287  5.530  4.772
1  6.287 10.664  6.358  6.393
2  5.530  6.358 11.527  8.014
3  4.772  6.393  8.014 13.976
Model-Implied Marginal Correlation Matrix:
      0     1     2     3
0 1.000 0.571 0.483 0.378
1 0.571 1.000 0.573 0.524
2 0.483 0.573 1.000 0.631
3 0.378 0.524 0.631 1.000
(Rows and columns are values of time)
```

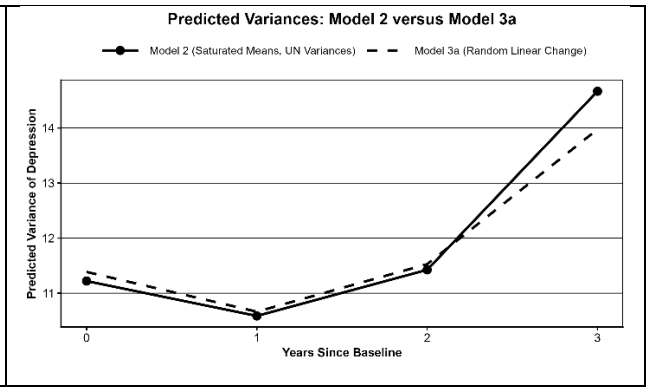
A random linear time slope predicts quadratically changing variance over time, as well as decreasing correlations over time (whereas a random intercept only would predict constant for both).

The marginalV function computes for us at any exact time we choose (here, 0, 1, 2, 3 exactly).

```
  time MeanModel2 MeanModel3a VarModel2 VarModel3a
1    0      8.225      8.1555    11.220     11.386
2    1      9.080      9.1034    10.587     10.664
3    2      9.950     10.0514    11.425     11.527
4    3     11.070     10.9993    14.668     13.976
```

See code online for plots comparing the means and variances across models





The predicted means and variances from Model 3a appear to be a pretty good match to those of the rounded-time "answer key" Model 2, but do we need any "bendiness" in the trajectory instead?

Model 3b: Random Quadratic Change over Exact Time

Level 1: $Dep_{ti} = \beta_{0i} + \beta_{1i}(Time_{ti}) + \beta_{2i}(Time_{ti})^2 + e_{ti}$

Level 2: $\beta_{0i} = \gamma_{00} + U_{0i}$, $\beta_{1i} = \gamma_{10} + U_{1i}$, $\beta_{2i} = \gamma_{20} + U_{2i}$

```
print("Model 3b: Random Quadratic Change over Exact Time")
Model3b = lmer(data=Example5b, REML=TRUE,
               formula=dep~1+time+I(time^2)+(1+time+I(time^2)|PersonID))
lmerSummary(Model3b, LRT=TRUE)
```

```
Model Fit for Predicting dep:
      AIC       BIC    logLik -2*logLik
 3983.686  4030.533 -1981.843  3963.686

Model for the Variance:
 Groups    Name        Variance Std.Dev   Corr
 PersonID  Intercept      7.321   2.706
           time           1.397   1.182 -0.328
           I(time^2)      0.143   0.378 -0.687
 Residual                 4.061   2.015
```

```
Warning message:
In checkConv(attr(opt, "derivs"),
opt$par, ctrl = control$checkConv, :
Model failed to converge with
max|grad| = 0.00633285 (tol = 0.002,
component 1) See ?lme4::convergence
and ?lme4::troubleshooting.
```

```
Fixed Effects:
 Term         Est    SE      DF t.value p.value
 Intercept  8.233 0.236 198.946  34.846   <.001
 time       0.716 0.238 202.832   3.005   0.003
 I(time^2)  0.077 0.076 202.959   1.014   0.312

LRT for random effect(s):
 Random Effect Dropped      Npar    LL -2LL  AIC  BIC   LRT DF p.value p.mix
 <none>                       10 -1982 3964 3984 4031
 time across PersonID          7 -1983 3966 3980 4013 2.416  3   0.491 0.395
 I(time^2) across PersonID     7 -1983 3967 3981 4014 3.077  3   0.380 0.297
```

```
print("Model 3b: Random Quadratic Change over Exact Time using lme")
Model3b_lme = lme(data=Example5b, method="REML", fixed=dep~1+time+I(time^2),
                  random=~1+time+I(time^2)|PersonID)
lmeSummary(Model3b_lme, LRT=TRUE, randCI=TRUE)
```

Given the warning message, let's try the same model in lme to see its solution

```
Model Fit for Predicting dep:
      AIC       BIC    logLik -2*logLik
 3983.686  4030.495 -1981.843  3963.686

Model for the Variance:
 Groups    Name       Variance Std.Dev   Corr
 PersonID  Intercept     7.321   2.706
           time          1.396   1.182 -0.328
           I(time^2)     0.143   0.378  0.107 -0.687
 Residual                4.061   2.015

Fixed Effects:
 Term         Est    SE      DF t.value p.value
 Intercept  8.233 0.236 209.137  34.846   <.001
 time       0.716 0.238 246.167   3.005   0.003
 I(time^2)  0.077 0.076 236.200   1.014   0.312

LRT for random effect(s):
 Random Effect Dropped  Npar    LL -2LL  AIC  BIC   LRT DF p.value p.mix
 <none>                   10 -1982 3964 3984 4031
 time | PersonID           7 -1983 3966 3980 4013 2.416  3   0.491 0.395
 I(time^2) | PersonID      7 -1983 3967 3981 4014 3.077  3   0.380 0.297

Random Effect 95% CIs (Fixed Effect +/- 1.96*SQRT(random effect variance)):
 Term                 FixedEffect Variance    SD  Lower  Upper
 Intercept | PersonID       8.233    7.321 2.706  2.930 13.536
 time | PersonID            0.716    1.397 1.182 -1.601  3.032
 I(time^2) | PersonID       0.077    0.143 0.378 -0.664  0.819
```

The lme results are identical to those of lmer, so they agree that neither the fixed nor random quadratic time slope is needed. However, does a random linear slope do a good enough job of reproducing the correlations across occasions? Let's see if an AR1 residual correlation helps…

Model 3c: Add an Exact Time AR(1) Residual Correlation to Random Linear Change

Level 1: $\boldsymbol{Dep_{ti} = \beta_{0i} + \beta_{1i}(Time_{ti}) + e_{ti}}$

Level 2: $\boldsymbol{\beta_{0i} = \gamma_{00} + U_{0i}}$

$\boldsymbol{\beta_{1i} = \gamma_{10} + U_{1i}}$

Residuals: $\boldsymbol{Cov(e_{ti}, e_{t-k,i}) = \sigma_e^2 \rho^k}$

In Model 3c the level-1 e_{ti} residuals are no longer assumed to be independent: Residuals k exact time apart have correlation ρ^k – an AR(1) structure.

```
print("Model 3c: Add AR1 Residual Correlation to Random Linear Change")
# corAR1 is for balanced (ordinal) time, corCAR1 is for unbalanced (continuous)
Model3c = lme(data=Example5b, method="REML", fixed=dep~1+time, random=~1+time|PersonID,
              correlation=corCAR1(form=~time|PersonID))
lmeSummary(Model3c, LRT=TRUE)
```

Note: corAR1 is for integer time (as in Example 5a), whereas corCAR1 is for continuous exact time below. Both require the lme function in the nlme package.

```
Model Fit for Predicting dep:
      AIC       BIC    logLik -2*logLik
 3976.694  4009.469 -1981.347  3962.694

Model for the Variance:
 Groups    Name      Variance Std.Dev   Corr
 PersonID  Intercept    6.345   2.519
           time         0.632   0.795 -0.256
 Residual               4.883   2.210
 Residual  CAR1 Phi                   0.104 = 1yr, .104^2 = 2yrs, .104^3 = 3yrs

Fixed Effects:
 Term         Est    SE      DF t.value p.value
 Intercept  8.163 0.224 201.943  36.391   <.001
 time       0.947 0.091 202.087  10.423   <.001

LRT for each random effect and residual correlation, tested separately (full model vs.
that piece dropped):
 Random Effect Dropped            Npar    LL -2LL  AIC  BIC    LRT DF p.value p.mix
 <none>                              7 -1981 3963 3977 4009
 time | PersonID                     5 -1986 3972 3982 4005  9.052  2   0.011 0.007
Residual correlation (CAR1)          6 -1982 3965 3977 4005  1.852  1   0.174
```

The estimated AR(1) residual correlation between adjacent occasions was ρ = 0.104, and adding it did not significantly improve fit, so Model 3a is retained. In Example 5a, an AR(1) residual correlation was needed in a model with only a random intercept. Why might it be much less needed here?

Model 4: Add Moderation of Linear Change by Baseline Age (back to lmer)

Level 1: $\boldsymbol{Dep_{ti} = \beta_{0i} + \beta_{1i}(Time_{ti}) + e_{ti}}$

Level 2: $\boldsymbol{\beta_{0i} = \gamma_{00} + \gamma_{01}(BaseAge_i - 10) + U_{0i}}$

$\boldsymbol{\beta_{1i} = \gamma_{10} + \gamma_{11}(BaseAge_i - 10) + U_{1i}}$

Composite equation to show the cross-level interaction (with all fixed effects first by convention):

$$\boldsymbol{Dep_{ti} = \gamma_{00} + \gamma_{10}(Time_{ti}) + \gamma_{01}(BaseAge_i - 10) + \gamma_{11}(Time_{ti})(BaseAge_i - 10) + U_{0i} + U_{1i}(Time_{ti}) + e_{ti}}$$

```
print("Model 4: Add Moderation of Linear Change by Baseline Age (centered at 10)")
Model4 = lmer(data=Example5b, REML=TRUE,
              formula=dep~1+time*baseage10+(1+time|PersonID))
lmerSummary(Model4)
```

```
Model Fit for Predicting dep:
      AIC       BIC    logLik -2*logLik
 3918.907  3956.384 -1951.453  3902.907

Model for the Variance:
 Groups    Name      Variance Std.Dev   Corr
 PersonID  Intercept    6.270   2.504
           time         0.342   0.585 -0.108
 Residual               4.341   2.083

Fixed Effects:
 Term               Est    SE      DF t.value p.value
 Intercept        8.018 0.218 197.359  36.727   <.001
 time             1.050 0.079 198.277  13.362   <.001
 baseage10       -1.400 0.334 196.782  -4.197   <.001
 time:baseage10   1.046 0.120 198.979   8.697   <.001
```

Interpret the Intercept (γ_{00}):

Interpret time (γ_{10}):

Interpret baseage10 (γ_{01}):

Interpret time:baseage10 (γ_{11}):

```
PseudoR2(baseModel=Model3a, altModel1=Model4)
TotalR2(altModel1=Model3a, altModel2=Model4)
```

```
Pseudo-R2
 Term       Model   Variance Pseudo.R2
 Intercept  Model3a    7.045
 Intercept  Model4     6.270     0.110
 time       Model3a    0.793
 time       Model4     0.342     0.568
 Residual   Model3a    4.341
 Residual   Model4     4.341     0.000

Total-R2 and Change in Total-R2
 Model   Total.R2 Change.R2
 Model3a    0.089
 Model4     0.136     0.047
```

Which new slope accounted for which variance?

Before continuing to examine the role of daily hours spent gaming as a time-varying predictor, we need to understand its sources of between-person differences by predicting it as an outcome on its own. In particular, **person-mean-centering should only be used for time-varying predictors that do not contain individual differences in change over time** (see Hoffman, 2015, ch. 9 or Hoffman, 2025). However, using person-mean-centering for fixed-only change over time in a time-varying predictor would still be ok, so long as those fixed time effects are included in predicting the outcome along with the time-varying predictor (whose "unique" contribution then controls for time).

Empty Means, Random Intercept Model for Gaming, the Time-Varying Predictor

```
print("Empty Means, Random Intercept to Get ICC for the Gaming Predictor")
Gaming = lmer(data=Example5b, REML=TRUE, formula=game~1+(1|PersonID))
lmerSummary(Gaming, ICC1=TRUE)
```

```
Model for the Variance:
 Groups    Name      Variance Std.Dev
 PersonID  Intercept    0.878   0.937
 Residual               1.005   1.003
Number of obs: 800 , groups: PersonID, 200

Fixed Effects:
 Term         Est    SE      DF t.value p.value
 Intercept  3.104 0.075 199.000  41.300   <.001

ICC1 intraclass correlation:
    Adjusted ICC: 0.466
  Unadjusted ICC: 0.466
```

So what does this result tell us so far about how to specify time-varying gaming in predicting time-varying depression?

Gaming Models 2 and 3a: Does Gaming Change Over Time?

Person-mean-centering gaming into a within-person part (WPgame) and a between-person part (PMgame) is only appropriate if gaming has no individually-varying trend over time (which would remain buried in WPgame, smushing its WP slope). We check by fitting the same models used for depression to gaming as an outcome: saturated means with an unstructured R matrix (Gaming2), and random linear change over exact time (Gaming3a).

```
print("Gaming Model 2: Saturated Means, UN Variances and Correlations by Occasion")
Gaming2 = gls(data=Example5b, method="REML", model=game~1+factor(occasion),
              correlation=corSymm(form=~occasion|PersonID), # Unstructured correlations
              weights=varIdent(form=~1|occasion))       # Separate variance by occasion
glsSummary(Gaming2, predMeans=TRUE)

Model Fit for Predicting game:
      AIC      BIC    logLik -2*logLik
 2597.187 2662.701 -1284.594  2569.187

Overall F-test for All Fixed Slopes:
 Model    NumDF   DenDF F.value p.value
 Gaming2      3 200.430   1.784   0.151

Fixed Effects:
 Term                  Est    SE      DF t.value p.value
 Intercept           3.115 0.094 223.454  33.296   <.001
 factor(occasion)2  -0.055 0.100 199.951  -0.553   0.581
 factor(occasion)3  -0.110 0.103 198.878  -1.064   0.289
 factor(occasion)4   0.120 0.090 199.512   1.328   0.186

Model-Predicted Means for factor(occasion):
 Level  Mean    SE      DF
 1     3.115 0.094 223.454
 2     3.060 0.101 199.325
 3     3.005 0.097 200.509
 4     3.235 0.096 196.903

print("R matrix (variances and covariances) and correlations for first person")
getVarCov(Gaming2, individual="1", type="marginal")
corMatrix(Gaming2$modelStruct$corStruct)[[1]]
```

What does this F-test result tell us?

```
Marginal variance covariance matrix
        [,1]    [,2]    [,3]    [,4]
[1,] 1.75050 0.90261 0.74314 0.98791
[2,] 0.90261 2.03660 0.99970 0.82000
[3,] 0.74314 0.99970 1.87430 0.82294
[4,] 0.98791 0.82000 0.82294 1.85910
  Standard Deviations: 1.3231 1.4271 1.3691 1.3635

          [,1]      [,2]      [,3]      [,4]
[1,] 1.0000000 0.4780411 0.4102613 0.5476291
[2,] 0.4780411 1.0000000 0.5116732 0.4214193
[3,] 0.4102613 0.5116732 1.0000000 0.4408529
[4,] 0.5476291 0.4214193 0.4408529 1.0000000
```

```
print("Gaming Model 3a: Random Linear Change over Exact Time")
Gaming3a = lmer(data=Example5b, REML=TRUE, formula=game~1+time+(1+time|PersonID))
lmerSummary(Gaming3a, LRT=TRUE)
```

```
Model Fit for Predicting game:
      AIC       BIC    logLik -2*logLik
 2593.563  2621.671 -1290.782  2581.563

Model for the Variance:
 Groups    Name      Variance Std.Dev   Corr
 PersonID  Intercept    0.884   0.940
           time         0.000   0.002 -1.000
 Residual               1.005   1.003
Number of obs: 800 , groups: PersonID, 200

Overall F-test for All Fixed Slopes:
 Model    NumDF   DenDF F.value p.value
 Gaming3a     1 599.362   1.010   0.315

Fixed Effects:
 Term         Est    SE      DF t.value p.value
 Intercept  3.056 0.089 214.106  34.321   <.001
 time       0.032 0.032 599.362   1.005   0.315

LRT for random effect(s):
 Random Effect Dropped  Npar    LL -2LL  AIC  BIC   LRT DF p.value p.mix
 <none>                    6 -1291 2582 2594 2622
 time | PersonID           4 -1291 2582 2590 2608 0.003  2   0.999 0.978
```

The fixed linear effect of time was not significant, and the random time slope was not needed. Gaming therefore appears to have only stable between-person differences in the mean over time only, plus occasion-to-occasion fluctuation, so person-mean-centering to create WPgame and PMgame (to be centered at 3, near the sample mean) is appropriate.

Model 5: Add Within-Person and Between-Person Fixed Slopes of Gaming, with Moderation of Linear Change by Between-Person Gaming

Level 1: $Dep_{ti} = \beta_{0i} + \beta_{1i}(Time_{ti}) + \beta_{2i}(Game_{ti} - \overline{Game}_i) + e_{ti}$

Level 2: $\beta_{0i} = \gamma_{00} + \gamma_{01}(BaseAge_i - 10) + \gamma_{02}(\overline{Game}_i - 3) + U_{0i}$

$\beta_{1i} = \gamma_{10} + \gamma_{11}(BaseAge_i - 10) + \gamma_{12}(\overline{Game}_i - 3) + U_{1i}$

$\beta_{2i} = \gamma_{20}$

Composite equation to show the cross-level interactions (with all fixed effects first by convention):

$$Dep_{ti} = \gamma_{00} + \gamma_{10}(Time_{ti}) + \gamma_{01}(BaseAge_i - 10) + \gamma_{02}(\overline{Game}_i - 3) + \gamma_{11}(Time_{ti})(BaseAge_i - 10) + \gamma_{12}(Time_{ti})(\overline{Game}_i - 3) + U_{0i} + U_{1i}(Time_{ti}) + e_{ti}$$

```
print("Model 5: Separate WPgame and PMgame3 Fixed Slopes Predicting Depression")
Model5 = lmer(data=Example5b, REML=TRUE,
              formula=dep~1+time*baseage10+WPgame+time*PMgame3+(1+time|PersonID))
lmerSummary(Model5)
```

```
Model Fit for Predicting dep:
      AIC       BIC    logLik -2*logLik
 3671.165  3722.696 -1824.583  3649.165

Model for the Variance:
 Groups    Name      Variance Std.Dev   Corr
 PersonID  Intercept    6.739   2.596         Var(U_0i)
           time         0.295   0.543 -0.101  Var(U_1i)
 Residual               2.720   1.649         Var(e_ti)

Overall F-test for All Fixed Slopes:
 Model  NumDF   DenDF F.value p.value
 Model5     6 227.520 108.809   <.001

Fixed Effects:
 Term               Est    SE      DF t.value p.value
 Intercept        8.068 0.211 196.559  38.233   <.001  gamma00
 time             1.006 0.066 198.477  15.313   <.001  gamma10
 baseage10       -1.373 0.324 196.143  -4.231   <.001  gamma01
 WPgame           1.283 0.071 561.181  18.111   <.001  gamma20
 PMgame3          0.164 0.198 196.018   0.827   0.409  gamma02
 time:baseage10   1.055 0.101 199.395  10.431   <.001  gamma11
 time:PMgame3     0.016 0.062 199.802   0.252   0.801  gamma12
```

Interpret the Intercept (γ_{00}):

Interpret WPgame (γ_{20}):

Interpret PMgame3 (γ_{02}):

Interpret time:PMgame3 (γ_{12}):

```
PseudoR2(baseModel=Model3a, altModel1=Model4, altModel2=Model5)
TotalR2(altModel1=Model3a, altModel2=Model4, altModel3=Model5)
```

```
Pseudo-R2 and Change in Pseudo-R2
 Term       Model   Variance Pseudo.R2 Change.R2
 Intercept  Model3a    7.045
 Intercept  Model4     6.270     0.110
 Intercept  Model5     6.739     0.043    -0.066
 time       Model3a    0.793
 time       Model4     0.342     0.568
 time       Model5     0.295     0.628     0.060
 Residual   Model3a    4.341
 Residual   Model4     4.341     0.000
 Residual   Model5     2.720     0.373     0.373

Total-R2 and Change in Total-R2
 Model   Total.R2 Change.R2
 Model3a    0.089
 Model4     0.136     0.047
 Model5     0.238     0.101
```

Which new slope accounted for which variance?

This was the end of my chapter example, as my goal was to link together how these models can be specified equivalently in univariate MLM, single-level SEM, and multilevel SEM. However, one could keep going by testing random effects of WPgame to see how the effect of gaming more than usual varies between persons (as well as cross-level interactions to see what predicts those individual differences if so). Such extensions can be where the equivalence across these three analysis frameworks breaks down, depending on what software is being used. In this case, however, since I simulated the data, I know those additional terms won't be likely be found 😊

Sample Results Section starts here—see the last page for figures that could be included!
[indicates notes about what to customize or also include; note that SE and p-values are not needed if you provide tables for the model solutions; these are wordy on purpose to help show how to interpret each effect]

The extent to which depression changed over time, how that change differed by age at baseline, and how it was predicted by daily hours of video gaming was examined in a series of multilevel models. A total of 800 observations were modeled as nested within 200 children, each observed on four annual occasions beginning at ages 8–12 years (M = 9.90, SD = 0.65). Time was the exact years since baseline (0 to 3.31) given that the occasions were not equally spaced. Restricted maximum likelihood (REML) within the lmer function in the lme4 R package (v. 2.0.6; Bates et al., 2015) was used in estimating and reporting all model parameters, except for models with an unstructured covariance matrix or adding an AR(1) residual correlation, which instead required the gls and lme functions for REML estimation within the nlme R package (v. 3.1-169; Pinheiro & Bates, 2000). The significance of fixed effects was evaluated with univariate and multivariate Wald tests using Satterthwaite denominator degrees of freedom (computed using the lmerTest R package, v. 3.2.1; Kuznetsova et al., 2017, for lmer models, and the emmeans R package, v. 2.0.4; Lenth & Piaskowski, 2026, for gls models), whereas additions to the model for the variance were evaluated via likelihood ratio tests (i.e., $-2\Delta LL$ with degrees of freedom equal to the number of new parameters). Alpha was chosen as .05. Fixed effect sizes were evaluated via pseudo-R^2 values for the proportion reduction in each variance component relative to the model without the predictors in question, as well as with total-R^2, the squared correlation between the actual depression outcomes and those predicted by the model fixed effects.

As derived from an empty means, random intercept model, depression had an intraclass correlation of ICC1 = .452, indicating that 45.2% of its variance reflected mean differences between children, with the remaining 54.8% within children across occasions. A model with saturated means and unstructured variances and correlations by rounded occasion was estimated for descriptive purposes. The mean trajectory increased across the four annual occasions (means = 8.23, 9.08, 9.95, and 11.07), residual variances increased after the second occasion, and correlations declined from about .61 for adjacent occasions to .41 for occasions three years apart. Given these trends, change over time was modeled using fixed and random linear effects for the exact time since baseline.

A model with a fixed and random linear effect of time showed significant mean increase in depression of 0.948 per year (SE = 0.091, $p < .001$), with significant individual differences in the rate of change (random slope variance = 0.793, SD = 0.890), $-2\Delta LL(2) = 33.24$. The additions of fixed or random quadratic time slopes were not significant, nor was an AR(1) correlation among the level-1 residuals, and so the random linear change model with independent residuals was retained. Cross-sectional differences in age (as baseline age centered at 10 years) were then examined as a moderator of the intercept and linear change. For every year older children were at baseline, their predicted initial depression was significantly lower by 1.400 (SE = 0.334, $p < .001$) but their rate of linear change was significantly more positive by 1.046 (SE = 0.120, $p < .001$). These age slopes accounted for 11.0% of the level-2 random intercept variance, 56.8% of the level-2 random linear time slope variance, as well as 13.6% of the total variance when combined with the main effect of linear time.

Daily hours of gaming was then considered as a time-varying predictor. We first examined the same sequence of models for gaming as an outcome as used for depression. An empty means, random inbtercept model indicated that 46.6% of the variance in gaming reflected mean differences between children. A model with saturated means and unstructured variance and correlation by rounded occasion indicated no clear pattern of change over time, and a random linear model of change showed no significant fixed or random linear change over time. Consequently, using person-mean-centering, gaming was separated into a level-1 within-person deviation from the child's own mean, along with their level-2 person mean over time (centered at 3 hours, near the sample mean). Fixed main effects for gaming were included at each level, and the gaming person mean also moderated the linear time slope. Neither the main effect nor the interaction of time of person mean gaming was significant: For additional hour of person mean gaming, depression at baseline was nonsignificantly higher by 0.164 (SE = 0.198, $p = .409$) and the rate of increase in depression over time was nonsignificantly greater by 0.016 (SE = 0.062, $p = .801$). However, the level-1 slope was significant: For every additional hour spent gaming more than usual, depression at that occasion was significantly higher by 1.283 (SE = 0.071, $p < .001$). Together these accounted for 37.3% of the level-1 residual variance (due to the main effect of within-person gaming), another 6.0% of the level-2 random linear time slope variance (due to person mean gaming by time), and a net increase of 6.6% of the level-2 random intercept variance (due to the main effect of between-person gaming along with the lesser level-1 residual variance relative to the previous model), yielding a total $R^2 = .238$.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software, 67*(1), 1–48.

Hoffman, L. (2026). A Rosetta Stone for modeling change: Connections among multilevel models, structural equation models, and multilevel structural equation models. In D. M. Teti, H. H. Cleveland, & K. L. Rulison (Eds.), *Handbook of research methods in developmental science* (2nd ed.). Wiley.

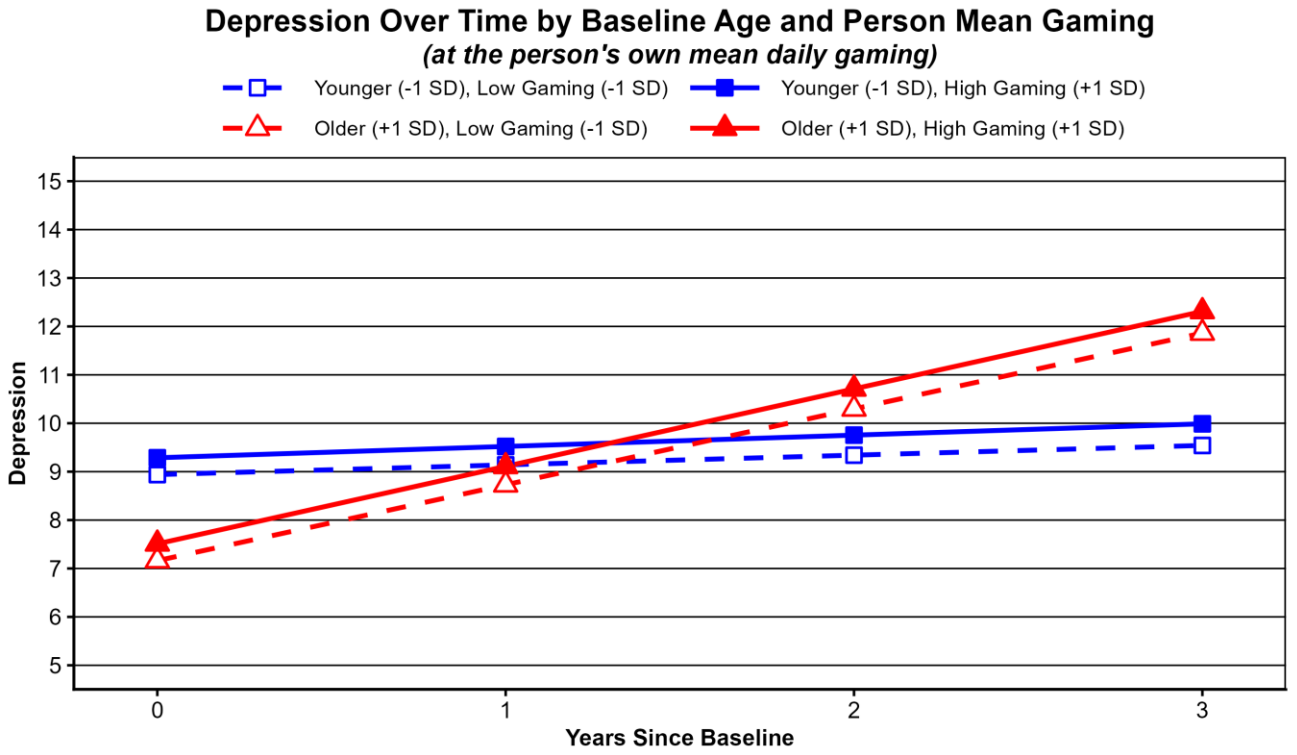
Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software, 82*(13), 1–26.

Lenth, R. V., & Piaskowski, J. (2026). *emmeans: Estimated marginal means, aka least-squares means* (R package version 2.0.4) [Computer software]. https://CRAN.R-project.org/package=emmeans

Pinheiro, J. C., & Bates, D. M. (2000). Mixed-effects models in S and S-PLUS. Springer.

Figures to Illustrate the Model 5 Results—see Example 5b files online R syntax!

Predicted depression is based on the fixed effects of Model 5 only, for children who were younger (baseline age = −1 SD, or 9.25 years) or older (+1 SD, or 10.55 years; SD = 0.65) at baseline, with low (−1 SD, or 2.04 hours) or high (+1 SD, or 4.17 hours; SD = 1.06) person mean gaming, at their own mean gaming on each occasion (WPgame = 0). The significant time × baseline age interaction is the difference in steepness of the blue (younger = 0.20 per year) and red (older = 1.57 per year) lines, which cross over because older children started lower. The nonsignificant between-person gaming main effect is shown by the small vertical distance between the dashed (low) and solid (high) lines of the same color at year 0, and its nonsignificant interaction with time is shown by the nearly parallel dashed and solid lines of the same color. The significant within-person gaming effect (1.283 per hour) is not shown here, but is shown in the next plot.



Predicted depression is based on the fixed effects of Model 5 only, at baseline (time = 0) for a 10-year-old (baseline age centered at 10 = 0), for children with low (−1 SD, or 2.04 hours), average (3.10 hours), and high (+1 SD, or 4.17 hours; SD = 1.06) person mean gaming, at within-person gaming of −1 SD (lower than usual), the child's own mean (usual), and +1 SD (higher than usual; SD = 0.87 hours). The within-person effect is the slope of each line (1.283 per hour, or about 2.2 points between lower than usual and higher than usual), which is the same for every line given the lack of WP × BP. The nonsignificant between-person effect is the vertical distance between the lines at the same within-person gaming: children whose person mean gaming differed by 2 SDs (low versus high) differed by only 0.35 points.

