PART III

A METHOD TOOLBOX

25

DISCONTINUOUS GROWTH MODELS: ILLUSTRATIONS, RECOMMENDATIONS, AND AN R FUNCTION FOR GENERATING THE DESIGN MATRIX

PAUL D. BLIESE, JASON KAUTZ, AND JONAS W. B. LANG

ABSRACT

The discontinuous growth model is a variant of a growth model that uses a block of time-related covariates to capture (1) immediate change and (2) changes in growth trajectories associated with one or more discrete events. Events might be planned such as adding a trust violation to a longitudinal study of trust or unplanned such as examining the effects of the great recession on firm performance. In this chapter, we describe the discontinuous growth model and provide examples of research questions that can be tested using the models. We also provide detailed code in R to help researchers estimate the models. Throughout the chapter we give practical advice based on our experience estimating these models in different research contexts. Finally, we introduce an R function to help researchers set up the design matrix in situations where events occur at different points for the higher-level entities.

KEYWORDS: Longitudinal, Mixed-Effects Models, Events, Discontinuous

The discontinuous growth model or DGM[1] is a variant of a mixed-effects growth model that provides considerable opportunity for researchers to propose and test hypotheses involving change associated with a discrete event or multiple events. For instance, a researcher interested in examining how trust develops over 10 trials might include an "event" of a trust violation on the third trial to see whether the event (a) impacts trust immediately and/or (b) alters future trajectories of trust development. Furthermore, the researcher might investigate individual differences to see whether participants react uniformly to the event and then potentially model characteristics of the individual that predict the nature of the reactions.

In this chapter, we provide a brief overview of mixed-effect models, growth models and the DGM using small, simulated datasets involving trust. After this introductory material, we provide a detailed step-by-step analysis of how the event of the great recession in 2008 was related to gross domestic product (GDP) trajectory changes for US States. In addition to illustrating the DGM, we also introduce and describe an R-based tool designed to help researchers set up the design matrix for both simple and complex DGMs. Finally, throughout the chapter, we provide practical advice for using the DGM based on our experiences applying the approach to different types of research questions.

BACKGROUND ON MIXED-EFFECTS MODELS AND GROWTH MODELS

The DGM is a form of a growth model estimated in a linear mixed-effect modeling program such as the R packages nlme (Pinheiro & Bates, 2000), lme4 (Bates, Mächler, Bolker, & Walker, 2015) and MCMCglmm (Hadfield, 2010). Other programs include Proc Mixed in SAS (Littell, Milliken, Stroup, Wolfinger & Schabenberger, 2006; Singer, 1998) or mixed in STATA (Rabe-

---

[1] The DGM is also referred to as the piecewise hierarchical linear model (Hernández-Lloreda et al. 2004, Raudenbush & Bryk 2002) and multiphase mixed-effects model (Cudeck & Klebe, 2002).

Hesketh & Skrondal, 2005). Nomenclature tends to differ across fields. For instance, mixed effects models are sometimes referred to as random coefficient models (RCM), random effects models (RE) or hierarchical linear models (HLM), but the term "mixed-effect models" is broadly recognized within the statistical community. We therefore use the term "mixed-effect models" to refer to the larger class of models.

Growth models are a specific type of mixed-effect model applied to panel data. In panel data, higher-level entities (typically persons) provide numerous responses over time. Data with this structure are also referred to as longitudinal or multilevel data. Mixed-effects models are well-suited for use with panel data because the models allow users to specify and model two specific forms of variance. First, mixed-effects models estimate a variance term associated with a random intercept for each of the higher-level entities (τ_{00}). To understand the nature of τ_{00}, it is helpful to consider the most basic form of the mixed-effects model: the intercept-only model. The intercept-only model – also called the null model – does not yet include a time variable. Rather, the intercept-only model reflects a unit- or person-specific deviation from the average intercept. Specifically, this model provides estimates of the random intercept (u_{0j}) and its associated variance (τ_{00}) along with the residuals (e_{ij}) and the residual variance (σ^2). The resulting model formula is $Y_{tj} = \gamma_{00} + u_{0j} + e_{tj}$, where $u_{0j} \sim N(0, \tau_{00})$ and $e_{tj} \sim N(0, \sigma^2)$.

In practical terms, by estimating τ_{00}, the models allow each higher-level entity to differ in terms of average responses across time. In most panel data, the assumption of a random intercept is warranted. For example, in a study of trusting behavior over time, we would expect some

respondents to display relatively high average levels of trusting behavior while other respondents would be expected to display relatively low levels of trusting behavior.[2]

The intercept variance from a null model can be used to estimate the ICC(1) (Bliese, 2000; Raudenbush & Bryk, 2002). The ICC(1) provides an index of how much of the total variance is associated with the higher-level entity. The ICC(1) can be estimated from a null mixed-effect model using the formula $\tau_{00}/(\tau_{00}+\sigma^2)$ (Raudenbush & Bryk, 2002). In our experience, ICC(1) values routinely exceed .40 in panel data from persons. In other words, 40% or more of the total variance can be attributed to between-person differences (e.g., individual-level factors such as personality traits), meaning that the remaining variance is due to within-person differences. In contexts involving outcomes from different types of higher-level entities (e.g., teams or firms) ICC(1) values may be much higher or much lower. Regardless of the specific value, ICC(1) is informative because it summarizes whether most of the variance occurs within higher-level entities over time or between higher-level entities.

The second form of variance that can be estimated in mixed-effects models of panel data is the variance associated with the slopes of level-1 or time-varying predictors (τ_{11}).[3] This variance term is included after the addition of the time-varying predictor to the null model and is optional in the sense that models can "fix" the time slope to be consistent across higher-level entities (i.e., not estimate τ_{11}) or allow variation in slopes across higher-level entities (i.e., estimate τ_{11}). The resulting basic model formula allowing slope variation takes the form $Y_{tj} = \gamma_{00}$

---

[2] The logic underlying a random intercept model is captured in various ways by a number of alternative statistical approaches. For instance, a random-effects one-way ANOVA also partitions variance into a within and between component. Likewise, a generalized estimation equation (GEE) model with an exchangeable error structure estimates a model similar to the random intercept mixed-effects model. McNeish, Stapleton, & Silverman (2017) provide details on alternative ways other than mixed-effects models to statistically account for differences across higher-level entities in panel and other multilevel data.

[3] For clarity, we are simplifying by using the subscripts $_{11}$ here. These subscripts only apply if we had one predictor and one random slope which is fine for illustrative purposes but would change with more complex models.

$+ u_{0j} + (\gamma_{10} + u_{1j})\text{TIME}_t + e_{tj}$. In the education and management literature, this model is frequently written in a somewhat more elaborate form that assigns the model components to their respective level of analysis as follows:

Level-1: $Y_{tij} = \beta_{00j} + \beta_{10j}\text{TIME}_t + e_{tij}$

Level-2: $\beta_{0j} = \gamma_{00} + u_{0j}$

$\beta_{1j} = \gamma_{10} + u_{1j}$

The model makes the following assumptions on the distribution and nature of the variance components.

$$e_{tj} \sim N(0, \sigma^2)$$

$$\begin{bmatrix} u_{0j} \\ \\ u_{1j} \end{bmatrix} \sim N \left(\begin{bmatrix} 0 \\ \\ 0 \end{bmatrix}, \begin{bmatrix} \tau_{00} & \tau_{01} \\ \\ \tau_{01} & \tau_{11} \end{bmatrix} \right)$$

To a large degree, the ability to feasibly estimate the τ_{11} variance term is what differentiates mixed-effect models from other analytic approaches used with panel data such as econometric fixed-effects models (Bailey, 2016) or generalized estimation equations (Zeger & Liang, 1986). That is, relatively few other analytic approaches allow one to estimate and subsequently model variance associated with the slopes of the level-1 predictors.

In the context of growth models, being able to estimate variance associated with a time-varying predictor plays a key role in testing theory and understanding how level-1 entities change over time. In a growth model, the primary time-varying predictor is TIME (in the most basic form, TIME is represented as a vector of sequentially increasing numbers starting from 0). Thus, the variance associated with the time-varying predictor of TIME provides important

information about whether higher-level entities show different trajectories of change over time. That is, some higher-level entities might increase, others might decrease and still others might remain relatively flat. These differences are captured by τ_{11}. In typical growth models, the time-varying predictors focus on the change trajectory while characteristics of the units/persons/entities are used as predictors to explain differences in trajectories. More complex forms of the growth model are possible, but they are built upon this basic concept.

<div align="center">Basic Example Data and Illustration</div>

To illustrate the basic structure of the growth model and the logic behind estimating both intercept variance and slope variance, we provide example data in Table 25.1. While somewhat subtle, notice that male subjects tend to decrease, and female subjects tend to increase. With respect to intercept variance, each of the Subjects differ in terms of average trust over the four measurement occasions. The mean ratings for Subjects 1 through 4 are 2.650, 2.275, 2.700, and 2.975, respectively. The variance of the vector [2.650, 2.275, 2.700, 2.975] is 0.083. If we estimate a mixed-effects model using the lme function in the nlme package in R, the intercept variance (τ_{00}) based on a restricted maximum likelihood (REML) algorithm is 0.077 which is similar to the estimate obtained simply by calculating the variance of the three means.[4]

In our example data in Table 25.1, most of the variance is between the three subjects, so the residual variance, or σ^2, is 0.025 and the resulting ICC(1) estimate is 0.75 [$\tau_{00}/(\tau_{00}+\sigma^2)$]. The

---

[4] Note that this is an extreme example so the raw estimate of the variance from the means is close to the estimate from the mixed-effect model. In almost all situations, the estimate based on the variance of the means (0.083 here) will be upwardly biased relative to the estimate from the mixed-effects model. The mixed-effects model estimate "removes" the part of the raw group-mean variance that can be attributed to the variance of the raw variable. In other words, some part of the 0.92 variance has been "inherited" from the combination of the raw variance and the group size. The mixed-effect model removes this inherited variance. This core idea of group-mean inheriting variance from the raw variable is a foundation of the ANOVA model where an F-Value of 1 indicates that the group-mean variance equals the expected inherited variance, suggesting that nothing meaningful is going on within groups.

ICC(1) value indicates that three-fourths of the total variance is a function of between-entity differences. Or stated otherwise, one-quarter of the total variance is associated with changes over time.

Table 25.1. Example panel data

SUBJECT	GENDER	TIME	TRUST
1	F	0	2.5
1	F	1	2.6
1	F	2	2.7
1	F	3	2.8
2	M	0	2.5
2	M	1	2.3
2	M	2	2.2
2	M	3	2.1
3	F	0	2.6
3	F	1	2.7
3	F	2	2.7
3	F	3	2.8
4	M	0	3.2
4	M	1	3.1
4	M	2	2.9
4	M	3	2.7

In Table 25.1, TIME represents the only time-varying covariate and would capture whether subjects (as a whole) showed linear change in TRUST. In these data, the estimated slope for TIME in a mixed-effects model is -0.035 (p=0.35) which is not surprising given that the trust of the two male subjects is decreasing while trust for the two female subjects is increasing.

Where mixed-effects models are particularly unique, though, is that we can formally test whether the variance associated with the TIME slope (τ_{11}) is significant. Significant variation in τ_{11} is an indication that individual linear trajectories differ more than would be expected by chance. In our example, the estimate of τ_{11} is .018 which is significant in a log likelihood test contrasting (a) the model with a random intercept with (b) the model with a random intercept and random slope (likelihood ratio = 22.10, p<.001).

Significant slope variability can be modeled in a cross-level interaction by examining whether person-level characteristics are related to differences in slopes. For example, we can formally test whether females display a positive linear increase in trust and males show a negative linear relationship. The cross-level interaction term between GENDER and TIME is significant (t-value = -8.132, p<.001). Importantly, while finding significant slope variability is helpful for understanding patterns of variability in data, authors such as Snijders and Bosker (1999) and others discourage using tests of slope variability as a reason to stop subsequent tests of cross-level interactions. In other words, even if our estimate of τ_{11} had not been significant, Snijders and Bosker recommend proceeding with an examination of the cross-level interaction if there is a strong theoretical reason to examine the moderation (it is possible to include a cross-level interaction term in the model even if the level-1 slope is fixed). Our key point is that thinking about and modeling slope variability associated with TIME provides the foundation for

understanding differences in how higher-level entities change over time and provide the basis for answering important questions in growth models.

Our short overview of growth models provided here is meant to lay the foundation for discussing the DGM, as the DGM is a growth model that simultaneously includes several variants of TIME in ways that allow the model to capture discontinuities. Further details about estimating basic growth models are provided in Bliese and Ployhart (2002), Singer and Willett (2003), Raudenbush and Bryk (2002) and Snijders and Bosker (1999) among other sources.

DISCONTINUOUS GROWTH MODEL OVERVIEW

DGM Time Trajectories

In terms of theory development, the ability to introduce a discontinuity in a growth model can be highly informative because the discontinuity provides an opportunity to examine the impact of a discrete event (Morgeson, Mitchell, & Liu, 2015) on an outcome of interest. Morgeson et al., (2015) define and describe "Event System Theory" where they argue that discrete events play a key role in organizational processes. Analytically, the DGM provides a way to examine the impact of events by modeling trajectories before and after events. Events associated with discontinuities may be planned or unplanned. For instance, designing a research study that unexpectedly changes the performance task mid-way through a series of trials represents a planned discontinuity. Such a design might be used to test hypotheses about individual adaptation to unexpected change (e.g., Lang & Bliese, 2009). Other areas where events are planned within a longitudinal design is within trust research to examine how participants react to trust violations (Korsgaard, Kautz, Bliese, Samson, & Kostyszyn, 2018; Fulmer & Gelfand,

2015) and sleep research where sleep restriction is followed by recovery sleep (e.g., Rupp, Wesensten, Bliese, & Balkin, 2009).

In other cases, events associated with the discontinuity are un-planned. For instance, Kim and Ployhart (2014) examined the impact of the great recession on firm performance specifically focusing on how firm differences in human capital practices were related to the immediate response and subsequent recovery associated with the great recession. Other examples include individual responses to events such as unemployment (Lucas, Clark, Georgellis, & Diener, 2004), long-term disability (Lucas, 2007), or change in marital status (Lucas, Clark, Georgellis, & Diener, 2003).

The DGM provides a way to specify precise theoretical hypotheses about how a dependent variable of interest changes in response to one or more events. What we refer to as the basic DGM is a growth model where TIME is sub-divided into three different components: TIME, TRANS and POST.

Level-1: $Y_{tij} = \beta_{0j} + \beta_{1j}\text{TIME}_t + \beta_{2j}\text{TRANS}_t + \beta_{3j}\text{POST}_t + e_{tij}$

Level-2: $\beta_{0j} = \gamma_{00} + u_{0j}$

$\beta_{1j} = \gamma_{10}$

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

$\beta_{3j} = \gamma_{30}$

The TIME parameter captures the linear trajectory prior to the event. The TRANS parameter captures the immediate reaction to the event, and the POST parameter captures the linear trajectory following the event. In other sources, the "POST" trajectory is referred to as "RECOV" (e.g., Bliese, Adler & Flynn, 2016; Bliese & Lang, 2016; Lang & Bliese, 2009);

however, "RECOV" implies a response to a negative event with subsequent recovery to pre-event levels. Because the post-event trajectory can take numerous forms (some of which would not in any way resemble recovery) we use the more neutral term "POST". Likewise, TIME is occasionally referred to as SA for skill acquisition (Lang & Bliese, 2009) in the specific situation that the dependent variable refers to performance in a newly acquired task. The basic form of the model shown above does not yet include random slopes parameters for TIME, TRANS, and POST (u_{1j}, u_{2j}, and u_{3j}) but a key aspect of specifying these models centers on determining whether the variance components should be added to the model.

Table 25.2 provides example data and variable set-up for TIME, TRANS, and POST. The table also includes a variable called TIME.A which we will discuss in more detail later. The dependent variable is TRUST and is presented in a way that illustrates discontinuities and changes in trajectories for SUBJECT 1. Notice, specifically, that ratings of trust increase by 1-point increments from TIME 0 to TIME 2 (first 3 occasions). Between TIME 2 and TIME 3 (the third and fourth occasions), trust increases by 3 points going from a TRUST value of 3 at TIME 2 to a TRUST value of 6 at TIME 3. Following this increase, TRUST increases at a rate of 2-point increments per time. The other two subjects in Table 25.2 follow a similar pattern with some random error added to facilitate model estimation (mixed-effect models have difficulty converging when variance terms are zero which would occur if we directly copied the same pattern for all subjects).

Table 25.2. Basic set-up for discontinuous growth model

SUBJECT	TIME.A	TIME	TRANS	POST	TRUST
1	0	0	0	0	1.00

1	1	1	0	0	2.00
1	2	2	0	0	3.00
1	2	3	1	0	6.00
1	2	4	1	1	8.00
1	2	5	1	2	10.00
2	0	0	0	0	0.95
2	1	1	0	0	1.95
2	2	2	0	0	2.85
2	2	3	1	0	5.92
2	2	4	1	1	7.96
2	2	5	1	2	10.12
3	0	0	0	0	1.08
3	1	1	0	0	1.97
3	2	2	0	0	3.03
3	2	3	1	0	6.01
3	2	4	1	1	8.00
3	2	5	1	2	10.05

Table 25.3 presents the results from a basic DGM applied to the data in Table 25.2. Notice that in this specification of the model, all subsequent terms (TRANS and POST) are interpreted relative to the TIME trajectory. More specifically, the estimate for TIME in Table 25.3 shows that TRUST increased by a 0.97-point increment prior to TRANS (also clear from raw Table 25.2 data). The relative nature of the results is apparent in the value of 2.04 for TRANS – recall

from Table 25.2 that TRUST increased 3 points, not 2 as implied by Table 25.3. The Table 25.3

value of 2.04 reflects the fact that TIME is already increasing by 1, so the total increase is, in

fact, 3 points. The value of 2.04 is testing whether the increase is *different* from what would be

expected based on the TIME trajectory. The same relative interpretation is clear in the POST

term in Table 25.3. Based on Table 25.2 we expect to see TRUST increasing by 2-point

increments following the transition event, but Table 25.3 returns a value of 1.07, which

represents the *difference* between the TIME and POST slopes. The total slope associated with

POST is, thus, 0.97 + 1.07 (approximately 2).

Table 25.3. Results from basic discontinuous growth model (relative)

	Value	Std.Error	DF	t-value	p-value
(Intercept)	1.01	0.03	12	32.98	0.00
TIME	0.97	0.02	12	47.20	0.00
TRANS	2.04	0.05	12	39.15	0.00
POST	1.07	0.03	12	36.45	0.00

Understanding the relative nature of the tests is important for theoretical clarity and

hypothesis formulation. That is, when specifying hypotheses, researchers should be clear that

they are testing whether the event produces change relative to pre-event trajectories. For

instance, a hypothesis should state that "a trust violation will lead to a post-event trajectory that

significantly differs from the pre-violation trajectory". A hypothesis that stated "a trust violation

will lead to a significant positive or negative post-event trajectory" implies a different model and

is not actually tested in Table 25.3.

In our experience, understanding *absolute* change in trajectories often helps make sense of non-simulated data. That is, in practice we have found that it can be challenging to understand what is occurring with respect to both TRANS and POST parameters in response to events using just the relative coding. Therefore, we provide an alternative model specification using TIME.A (for absolute) in Table 25.2. Notice how TIME.A takes the TIME value immediately before TRANS (in this case, 2) and then holds the effect constant over the remaining observations for that subject. Substituting TIME.A for TIME in the DGM allows us to assess if the transition (TRANS) and the post-event slope (POST) are significantly different from zero. Table 25.4 provides the estimates replacing TIME with TIME.A. In Table 25.4, the parameter estimates represent absolute effects. TRANS now has a value of 3.01, reflecting the increase of trust by 3 points, and POST now has a value of 2.04, which reflects the post-event 2-point increment trajectory of the post-event occasions.

Table 25.4. Results from basic discontinuous growth model (absolute)

	Value	Std.Error	DF	t-value	p-value
(Intercept)	1.01	0.03	12	32.98	0.00
TIME.A	0.98	0.02	12	47.20	0.00
TRANS	3.01	0.04	12	79.82	0.00
POST	2.04	0.02	12	98.75	0.00

DGM slope differences and interactions

As with the basic growth model, the DGM often relies heavily on being able to test for slope variance associated with the time-varying covariates (u_{1j}, u_{2j}, and u_{3j}). In the DGM, it is possible

to ask whether level-2 entities have different trajectories prior to the event – an examination of whether the TIME trajectory varies across higher-level entities. The DGM also allows a test of whether higher-level entities have different immediate responses to the event. In our example in Table 25.2, all subjects increased trust by 3 points at TRANS, but in practice we would expect some subjects to have more extreme reactions and other subjects to have fewer extreme reactions. Finally, the DGM allows testing whether post-event trajectories (POST) vary across higher-level entities. Again, in our example all subjects had a rate of change in trust of 2 increments per occasion, but in practice higher-level entities may differ in trajectories.

Finally, the main strength of the DGM involves being able to add entity-level predictors of the time-based covariates to answer questions such as "do males have a larger immediate reaction to a trust violation than females?" or "following a trust violation, do the trajectories of males and females differ?". For instance, a DGM with gender as a level-2 predictor would look like this.

Level-1: $Y_{tij} = \beta_{0j} + \beta_{1j}\text{TIME}_t + \beta_{2j}\text{TRANS}_t + \beta_{3j}\text{POST}_t + e_{tij}$

Level-2: $\beta_{0j} = \gamma_{00} + \gamma_{01}\text{GENDER}_j + u_{0j}$

$\beta_{1j} = \gamma_{10} + \gamma_{11}\text{GENDER}_j + u_{1j}$

$\beta_{2j} = \gamma_{20} + \gamma_{21}\text{GENDER}_j + u_{2j}$

$\beta_{3j} = \gamma_{30} + \gamma_{31}\text{GENDER}_j + u_{3j}$

In an experimental design, the condition also becomes an entity-level predictor. So, for instance, if subjects were randomly assigned to a training condition (primed towards aggression) versus a control condition (primed towards a neutral response), we could examine how the experimental condition impacted the transition and/or the post-event trajectory.

As highlighted in Bliese and Lang (2016), researchers have to carefully consider the absolute versus relative nature of the model specification when testing interactions. For example, to most cleanly test the hypothesis "following a trust violation, the trajectories of males and females will differ", we would recommend using TIME.A (absolute) coding because the hypothesis is only considering the post-event trajectory differences between males and females. Using TIME (relative coding) is testing a hypothesis of "are the differences between pre and post event trajectories the same for males and females?" In other words, the relative coding (TIME) is effectively asking a three-way interaction (see Bliese & Lang, 2016). We provide more elaboration of this point in our detailed example.

A challenge in the analyses of the DGM is the fact that a fully specified DGM can include a lot of "moving parts" and numerous variants. For instance, variants of the model can include absolute and relative change, and extensions with quadratic change (see Bliese & Lang, 2016; Lang & Bliese, 2009). A successful analytical strategy typically requires researchers to make informed decision about which model components to include. In so doing, it is possible to err by either including too much complexity or not enough. In other words, an overly complex model will burry any meaningful real-life effect or may even discover effects that do not really exist (see Gelman & Imbens, 2018). Conversely, an overly simplistic model may not adequately describe the patterns in the dataset. There is also a continuum between entirely *a-priori* defined analytical strategies and so-called model building approaches. The mixed-effects modeling literature typically recommends performing a decent amount of model building to ensure that the model is not fundamentally misspecified (Bates et al., 2015; Matuschek, Kliegl, Vasishth, Baayen & Bates, 2017; Pinheiro & Bates, 2000, however also see Barr, Levy, Scheepers & Tily, 2013). However, it is important to avoid model building that goes beyond the basic nature of the

design. In experimental studies, in particular, we recommend using a relatively fixed design relying on *a-priori* theoretical assumptions about where variance should be occurring.

A basic model examination strategy (Bliese & Lang, 2016) involves first studying the overall nature of the fixed change effects TIME, TRANS, and POST with a model that only included a random intercept. Step 1 can be performed by studying the size of these effects and examining, graphically, whether these effects adequately capture the overall pattern. Note that it is crucial to leave components in the model for which there are potential hypotheses of group differences even when they are not substantial in this first basic step.

In a second step (Step 2), we recommend examining the slope random effects. A challenge in doing so is that including all random effects and their interactions may lead to an overly complex model that is over-specified, overly conservative and cannot meaningfully be estimated by software. For instance, a researcher adding all possible random effects and their correlations would end up with the following 4x4 random effects matrix after adding random slopes parameters for TIME, TRANS, and POST (u_{1j}, u_{2j}, and u_{3j}) and all random effects correlations:

$$
\begin{bmatrix} u_{0j} \\ u_{1j} \\ u_{2j} \\ u_{3j} \end{bmatrix} \sim N \left(\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \tau_{00} & \tau_{01} & \tau_{02} & \tau_{03} \\ \tau_{01} & \tau_{11} & \tau_{12} & \tau_{13} \\ \tau_{02} & \tau_{12} & \tau_{22} & \tau_{23} \\ \tau_{03} & \tau_{13} & \tau_{23} & \tau_{33} \end{bmatrix} \right)
$$

One strategy we recommend is to use a model with an orthogonal matrix in which all random effects are uncorrelated (τ_{01}, τ_{02}, τ_{03}, τ_{12}, τ_{13}, and τ_{23} set to 0).[5] An alternative (Bates et

---

[5] In lme this can be achieved by specifying `random=list(UNIT=pdDiag(~TIME+TRANS+POST))` instead of specifying `random=~TIME+TRANS+POST|UNIT` which is equivalent to specifying

al., 2015) is to start with such a matrix and to then only retain those random effects that differ from zero in this model. On the basis of those random effects that exist in an orthogonal model the researcher can then test for meaningful random effects covariances. Both strategies will typically ensure that the random effects matrix can meaningfully be estimated. Another recommended strategy, especially for experimental studies with a clear a-prior design, is to only include random effects that are theoretically assumed or important. In many cases it may be feasible to assume that participants in a laboratory study all react in the same way to manipulation and stimuli after accounting for basic individual differences with a random intercept so that a basic DGM without random slopes may be acceptable.

The final step in studying DGMs is to add level-2 predictors. While the typical strategy is to add each predictor to the intercept and all change slopes in correlational data, the strategy for experimental design is usually more specific and targeted. In experimental designs, we recommend adding the predictor to the *a-priori* targeted or manipulated effects because adding too many effects can easily conceal a true underlying experimental effect. For instance, when a researcher uses a DGM where there is an intervention after an initial skill acquisition period, the researcher would commonly only include a dichotomous predictor of TRANS as the groups did not theoretically differ before the intervention and because there is not necessarily theory suggesting that the intervention would change both the level and the trajectory after the change. Likewise, an intervention that only starts at a defined point in time would require one to predict the POST trajectory with a predictor but not the TRANS change term. While it is important to account for potential *a-priori* differences in DGMs, it is also important to not statistically

---

`random=list(UNIT=pdLogChol(~TIME+TRANS+POST))`. In lmer, it is possible by simply using `(TIME+TRANS+POST||UNIT)` instead of `(TIME+TRANS+POST|UNIT)`.

distribute the effect of one predictor among several competing predictors as the later strategy may conceal an effect that is expected by design through separating it into several non-significant effects that are not all expected by design.

In conclusion, it is important to carefully weigh model complexity and model parsimony and *a-prior* and *post-hoc* assumptions in modeling DGMs. In our experience, and also in the experience of other authors, organizational researchers and reviewers have a tendency to recommend additional model complexity out of a desire to be conservative. It is important for researchers to be aware that this approach is not a cure for the complexity of DGMs. In many cases, a more parsimonious model with fewer random effects may be more solid and lead to more replicable results and findings.

With this background, we now provide a step-by-step analysis and R based code for a dataset containing a discontinuity involving US State-level GDP and the great recession. Following the illustration, we expand our coverage of the DGM by describe an R-based tool designed to help researchers set up the design matrix for both simple and complex DGMs. Throughout our detailed illustrating we provide practical advice for using the DGM based on our experiences.

DETAILED DGM EXAMPLE

Illustrative Data[6]

To provide a detailed illustration of the DGM, we use data obtained from the US Department of Commerce, Bureau of Economic Analysis (http://www.bea.gov). The data represent Gross Domestic Product (GPD) for the 50 US States and the District of Columbia for the years 2004 to

---

[6] The data and the R Code are available on https://osf.io/d4xmw/

2011. The event on which we focus is the "great recession" associated with the economic downturn in 2008. Table 25.5 presents a subset of the data containing the unemployment rate for 2003 (eventually used as a level-2 or state-level predictor) and annual GDP data for 10 US States (the complete data is available from the lead author). The data in Table 25.5 have been modified to facilitate the subsequent interpretation of model parameters. Specifically, each GDP value has been divided by 1,000 to avoid having very small parameter estimates and standard errors that require rounding tables to five or six decimal places to see differences.

Table 25.5. Yearly GDP data from first ten states for detailed example

STATE	UNEMP03	Y2004	Y2005	Y2006	Y2007	Y2008	Y2009	Y2010	Y2011
Alabama	5.8	142	151	159	166	170	165	170	173
Alaska	7.7	34	38	42	45	50	45	48	51
Arizona	5.7	201	223	246	259	261	246	250	258
Arkansas	5.9	84	89	94	97	100	99	102	106
California	6.8	1570	1689	1798	1871	1900	1829	1878	1959
Colorado	6.2	202	217	230	243	252	244	253	264
Connecticut	5.5	188	196	209	221	219	214	221	230
Delaware	4.0	51	54	56	60	58	60	64	66
District of Columbia	7.2	78	82	87	92	97	98	104	108
Florida	5.3	621	681	731	761	748	726	736	754

To guide our illustrative example, we provide a set of specific hypotheses on which to focus the analyses and overlay the results. Our specific hypotheses are:

1. As a whole, GDP will increase between 2004 and 2008.

2. The great recession as indexed by the year 2008 will be associated with a significant discontinuity in GDP. The specific form of the discontinuity will be a drop in absolute terms.

3. The rate of GDP growth post 2008 will be lower than the rate of GDP growth prior to 2008.

4. The level-2 units (US States) will significantly differ in terms of all model parameters. That is, States will significantly differ in terms of (a) pre-2008 growth trajectories, (b) the drop associated with 2008 and (c) growth trajectories following 2008.

5. Rates of unemployment in 2003 will be related to average levels of GDP.

6. In the post 2008 trajectories, States with lower levels of unemployment in 2003 will show a stronger increase in GDP than States with higher levels of unemployment in 2003.

Note that both our second and our sixth hypotheses are presented in <u>absolute</u> terms. With respect to the sixth hypothesis, we expect differences in the post 2008 trajectories as a function of 2003 unemployment rates, but we are not taking into consideration the pattern that existed prior to 2008. That is, we are <u>*not*</u> hypothesizing that interactive patterns involving unemployment observed prior to 2008 will be significantly different from interactive patterns observed post 2008 (a hypothesis involving a <u>*relative*</u> test).

Data Set-Up: Create Time-Related Covariates

Estimating growth models requires transforming data from wide or multivariate form to long or univariate form. Here we use the make.univ function from the multilevel library (Bliese, 2016). Help files provide details on using the function, but briefly two arguments are provided – one argument to specify the full wide or univariate dataset (Table 25.5 data, but for all States) and another argument to specify the columns representing the time-varying outcome in the correct chronological order (columns 3:10 in our example). An optional parameter allows us to name the new DV (named "GDP" in our example). The code and the output for the first state in the sample, Alabama, are presented below.

```
> library(multilevel)
> names(gdp)
 [1] "STATE"   "UNEMP03" "Y2004"   "Y2005"   "Y2006"   "Y2007"
 [7] "Y2008"   "Y2009"   "Y2010"   "Y2011"
> gdp.univ<-make.univ(gdp,gdp[,3:10],outname="GDP")
> gdp.univ[1:8,c(1,2,11:12)]
    STATE UNEMP03 TIME GDP
1   Alabama   5.8   0 142
1.1 Alabama   5.8   1 151
1.2 Alabama   5.8   2 159
1.3 Alabama   5.8   3 166
1.4 Alabama   5.8   4 170
1.5 Alabama   5.8   5 165
1.6 Alabama   5.8   6 170
1.7 Alabama   5.8   7 173
```

In the transformed data set (gdp.univ) TIME is indexed from 0 (2004) to 7 (2011) with a GDP

value for each year. In this specification, 2008 (the year of the discontinuity) occurs at TIME 4.

The level-2 variable of State-level unemployment in 2003 is consistent across values of TIME,

but varies across States. To estimate the DGM, we create the TRANS, POST and TIME.A

variables using the following code and again present the data for Alabama.

> gdp.univ$TRANS<-ifelse(gdp.univ$TIME>3,1,0)

> gdp.univ$POST<-ifelse(gdp.univ$TIME>4,gdp.univ$TIME-4,0)

> gdp.univ$TIME.A<-ifelse(gdp.univ$TIME<4,gdp.univ$TIME,3)

> gdp.univ[1:8,c(1,2,11:15)]

	STATE	UNEMP03	TIME	GDP	TRANS	POST	TIME.A
1	Alabama	5.8	0	142	0	0	0
1.1	Alabama	5.8	1	151	0	0	1
1.2	Alabama	5.8	2	159	0	0	2
1.3	Alabama	5.8	3	166	0	0	3
1.4	Alabama	5.8	4	170	1	0	3
1.5	Alabama	5.8	5	165	1	1	3
1.6	Alabama	5.8	6	170	1	2	3
1.7	Alabama	5.8	7	173	1	3	3

Notice with the data from Alabama that there is no immediately clear discontinuity when

TRANS changes from 0 to 1 nor is there any obviously clear change in trajectory associated with

POST. There does appear to be an overall increase in GDP. Other States would show more

pronounced changes associated with the great recession, which goes back to the importance of being able to model differences across higher-level entities.

<div align="center">Step 1: Estimate the ICC</div>

The first step in analyzing discontinuous growth data is to estimate the ICC from the null model. In the nlme package, the lme (linear mixed-effects) function is used to estimate the mixed-effects model and the corresponding ICC(1) value. The term, control=list(opt="optim"), is optional and changes the optimization program from the default nlimb optimizer to the optim optimizer. The optim optimizer is the standard optimization procedure in R. In its default setting, it uses the popular Nelder-Mead method (Nelder & Mead, 1965) and was initially used in the Pinheiro and Bates (2000) publication of the nlme library. More recent versions of nlme use another optimizer – the nlminb optimizer based on the port library by Gay (1990). Typically, both optimization procedures converge and give comparable results. In our experience, the optim optimizer tends to be somewhat more likely to converge on complex models with the standard parametrization of the random effects matrix in lme. When nlminb has convergence problems, a simple fix is therefore to change the optimizer to opt="optim".[7]

> null.mod<-lme(GDP~1,random=~1|STATE,gdp.univ,control=list(opt="optim"))

> VarCorr(null.mod)

STATE = pdLogChol(1)

---

[7] Another alternative is to specify a general-positive random effects matrix instead of the standard general-positive matrix with Log-Cholesky parametrization for the random effects matrix in lme when the model has random slopes (see subsequent examples). The Log-Cholesky parametrization ensures that there are no correlations of 1 in the random effects matrix and is thus preferred by some authors even though newer programs tend to use symmetric matrices explicitly allowing for random effects correlations of 1. This is borne out of the philosophy that authors should use this information to reduce model complexity (e.g. Bates at al., 2015). In lme, it is possible to specify a general positive matrix by using `random=list(STATE=pdSymm(~TIME+TRANS+POST))` instead of specifying `random=~TIME+TRANS+POST|STATE` which is equivalent to specifying `random=list(STATE=pdLogChol(~TIME+TRANS+POST))`.

Variance   StdDev

(Intercept) 104309.498 322.96981

Residual   1156.165 34.00243

> 104309.498/(104309.498+1156.165)

[1] 0.9890375

The ICC(1) of .99 associated with GDP within States is much larger than would typically be

observed with repeated measures data from persons. With these types of economic data,

however, the value is not that surprising given the large differences among states. For instance,

the GDP of California is roughly 46 times the GDP of Alaska (see Table 25.5). As we show,

however, the DGM can still identify discontinuities in trajectories when the ICC(1) value is this

large. In practice, a value this large would be problematic as it suggests that only 1% of the total

variance can be explained by the time-related covariates in the DGM. We strongly suggest that

researchers report the ICC(1) value because doing so informs readers where the major sources of

variance reside (Bliese, Schepker, Essman, & Ployhart, in review). That is, a balanced discussion

of the within-entity findings (e.g., the DGM) should mention that most of the differences in the

panel data exist between higher-level entities and therefore the results of the DGM may be

significant but at the same time explain a small amount of the total variance.

In terms of providing practical advice, if we observed an ICC(1) value this high on a

substantive research variable (rather than an illustration of an approach) we would *not* proceed

with subsequent analyses. We would also be unwilling to use the DGM on a variable with a

distribution that is far from being normally distributed as is raw State-level GDP. Nonetheless,

even with these flaws, these data provide interesting results that illustrate the utility of the DGM

so we proceed with analyzing the data in the DGM framework.

Step 2: Model Relative and Absolute Time Effects

As previously noted, the core predictors in the DGM are TIME, TRANS and POST. These

predictors are entered as a block and the results describe the overall trajectory patterns. We begin

by estimating a model with a random intercept for State. This model allows States to vary in

terms of overall levels of GDP, but fixes parameter estimates across States for the pre-

discontinuity slope (TIME), the transition (TRANS), and the recovery slope (POST). In other

words, the model assumes that the parameter estimates for TIME, TRANS and POST describe

every State (an unlikely scenario, but a logical starting place).

```
> mod.base<-lme(GDP~TIME+TRANS+POST,random=~1|STATE,gdp.univ,
+          control=list(opt="optim"))
> round(summary(mod.base)$tTable,dig=3)
```

```
            Value Std.Error  DF t-value p-value
(Intercept) 231.402    45.350 354   5.103   0.000
TIME         14.167     1.723 354   8.220   0.000
TRANS       -15.276     5.716 354  -2.673   0.008
POST         -8.365     2.437 354  -3.432   0.001
```

In this model, all three variables are significant. The model indicates that in 2004, the

average level of GDP across the 51 States was 231,402 (231.402*1000) and that GDP increased

by 14,167 each year. Results show that GDP dropped by 15,276 in 2008 _relative_ to the expected

14,167 yearly increase (an absolute drop of -1,109). Finally, the negative parameter for POST

indicates a decline in GDP slope of 8,365 _relative_ to the pre-recession increase of 14,167, for an

absolute slope of 5,802. We can recover the absolute values by using TIME.A instead of TIME

(with a slight rounding difference of -1,110 instead of -1,109 for TRANS).

> mod.absol<-lme(GDP~TIME.A+TRANS+POST,random=~1|STATE,gdp.univ,

+   control=list(opt="optim"))

> round(summary(mod.absol)$tTable,dig=3)

	Value	Std.Error	DF	t-value	p-value
(Intercept)	231.402	45.350	354	5.103	0.000
TIME.A	14.167	1.723	354	8.220	0.000
TRANS	-1.110	4.560	354	-0.243	0.808
POST	5.802	1.723	354	3.366	0.001

Notice that the t-tests for TRANS and POST now determine whether the absolute change

is significantly different from zero rather than different from the expected values associated with

TIME. We see, for instance, that the decrease of -1,110 in GDP associated with the great

recession is not large enough to statistically differ from zero; however, an absolute slope of

5,802 for the post-event phase is significantly different from zero.

In describing these data, we had hypothesized the great recession as indexed by the year

2008 would produce a significant discontinuity in GDP in absolute terms. Given the specificity

of our hypothesis, we would fail to reject the null because the drop in absolute terms was not

significant even though it was significant in relative terms. We remain agnostic as to whether

theories to which the DGM will be applied will support relative change or absolute change for

hypotheses, but we want to emphasize that hypotheses should be precise. We also offer that if

theory is not sufficiently precise, it may be informative to present DGM results from both an

absolute and relative model specification.

With respect to the post-event trajectory (POST), we had hypothesized that the rate of GDP growth post 2008 will be lower than the rate of GDP growth prior to 2008. Thus, the test for this hypothesis is based on a model using TIME (relative effects) rather than a model using TIME.A (absolute effects). Our hypothesis would be supported, as a 8,365 decline _relative_ to the pre-recession increase of 14,167 was statistically significant. Note that if we had failed to specify the hypotheses as we had (an absolute effect for TRANS and a relative effect for POST), we would have had to run and interpret two different DGM models (one with TIME and one with TIME.A).

<div align="center">Step 3: Model Variability in Time Effects</div>

Often a key reason to conduct a DGM is because theory suggests that the time-based covariates will differ among higher-level entities and that these differences can be explained by attributes of the higher-level entity. To test whether the time-based covariates differ, we contrast models with and without random slopes for the level-1 time covariates. For clarity we repeat the base model and then update the model before conducting -2log likelihood tests of the models. When adding complexity to the model in terms of random effects we are interested in determining at what point the complexity no longer increases model fit (as indexed by the -2log likelihood test). When complexity no longer increases model fit, we select the last best fitting model.

```
> mod.base<-lme(GDP~TIME+TRANS+POST,random=~1|STATE,gdp.univ,
+       control=list(opt="optim"))
> mod.t<-update(mod.base,random=~TIME|STATE)
> mod.tt<-update(mod.base,random=~TIME+TRANS|STATE)
> mod.ttp<-update(mod.base,random=~TIME+TRANS+POST|STATE)
> anova(mod.base,mod.t,mod.tt,mod.ttp)
```

Model df   AIC   BIC   logLik   Test L.Ratio p-value

mod.base   1   6 4209.018 4233.027 -2098.509

mod.t   2   8 3754.137 3786.148 -1869.069 1 vs 2 458.8811 <.0001

mod.tt   3 11 3721.302 3765.318 -1849.651 2 vs 3 38.8347 <.0001

mod.ttp   4 15 3545.679 3605.700 -1757.839 3 vs 4 183.6237 <.0001

The results of the model comparisons indicate that allowing each higher-level entity to have a different trajectory for TIME, TRANS and POST improves the overall fit. To make explicit what is being contrasted, we provide the empirical Bayes estimate for the base model for the first 5 states and the estimates of the variance components for the base model followed by the empirical Bayes estimates and variance components for the final model.

> coef(mod.base)[1:5,] #Empirical Bayes estimates for base model

        (Intercept)   TIME   TRANS   POST

Alabama   126.423673 14.16667 -15.27647 -8.364706

Alaska   8.655518 14.16667 -15.27647 -8.364706

Arizona   207.350252 14.16667 -15.27647 -8.364706

Arkansas   60.858157 14.16667 -15.27647 -8.364706

California 1774.678296 14.16667 -15.27647 -8.364706

> var(coef(mod.base)[,1]) #Illustrative estimate of variance

[1] 104264.7

> VarCorr(mod.base) #Variance components for base model

STATE = pdLogChol(1)

        Variance   StdDev

(Intercept) 104359.339 323.04696

Residual   757.439 27.52161

> coef(mod.ttp)[1:5,] #Empirical Bayes for model with 3 random slopes

         (Intercept)   TIME   TRANS   POST

Alabama   143.41757 7.431306   -8.5211039 -4.2051836

Alaska   37.35557 1.750592   0.7534417   0.4637427

Arizona   213.18077 12.240521 -14.2976316 -7.6055818

Arkansas   84.84063 4.130082   -3.4486017 -1.5661432

California 1582.12679 98.855295 -128.0656844 -70.0628491

> apply(coef(mod.ttp),2,var) #Illustrative estimate of variance

(Intercept)   TIME   TRANS   POST

 77468.5200   364.3460   515.2555   166.2434

> VarCorr(mod.ttp) #Variance components for model with 3 random slope

STATE = pdLogChol(TIME + TRANS + POST)

           Variance   StdDev   Corr

(Intercept) 77499.63414 278.387561 (Intr) TIME   TRANS

TIME   365.78272 19.125447 0.978

TRANS   517.05656 22.738878 -0.996 -0.958

POST   166.78921 12.914690 -0.997 -0.986 0.990

Residual   82.25159   9.069266

   In the base model, each State has a different intercept value, but the slopes for each

parameter are identical and match the values provided in the initial t-table and previously

discussed. In the base model, there is only one parameter randomly varying (the intercept), so we

obtain a variance estimate of 104359.339 for the intercept. The code above shows that we obtain

a similar estimate if we simply take the variance of the empirical Bayes estimates of the intercepts (the value then is 104264.7). We do *not* recommend estimating or reporting variance terms from the empirical Bayes estimates as these estimates often differ from the correct estimates returned by VarCorr. We provide the variance estimates based on the empirical Bayes to conceptually illustrate how variance estimates are calculated.

The model with random slopes for TIME, TRANS and POST (mod.ttp) provides the best fit to the data and in this model each State has a different value for all intercept and all the slope parameters. Again, because all the parameters differ, we could calculate the variance associated with each column of empirical Bayes estimates and obtain rough estimates of the variance for each term (see the apply command above). In practice, though we use the estimates from the VarCorr command. Notice that VarCorr command also estimates correlations between random effects, which is why we progressively use one more DF in the model contrasts. For instance, mod.ttp uses 4 DF relative to mod.tt – one variance term of 166.8 for POST and three for correlations. As noted, models often fail to converge when numerous random effects are included in part because an increasing number of covariance terms are also estimated, so while we were able to include random effects for TIME, TRANS and POST here, users should expect to run a series of models to identify exactly where the variance differs across higher-level entities.

In sum, the model contrasting results and the variance estimates indicate that States significantly vary in terms of (a) the change in GDP between 2004 and 2008, (b) the degree to which GDP changed in 2008, and (c) the degree to which GDP increased after 2008. Parenthetically, while there is no ambiguity about the significance of the -2log likelihood tests in this example, it is important to keep in mind that these tests are widely known to be conservative

(Pinheiro & Bates, 2000), so a generally acceptable practice is to halve the p-values as though the tests were one-tailed (LaHuis & Ferguson, 2009).

<p style="text-align:center">Step 4: Level-2 Predictors</p>

The final step (and generally the most interesting) is to include level-2 predictors of the parameters previously identified to randomly vary among higher-level entities. In our example, we explore whether unemployment levels in 2003 are related to (a) overall levels of GDP, and (b) GDP growth following the 2008 decline. Our two specific hypotheses were:

5. Rates of unemployment in 2003 will be related to average levels of GDP.

6. In the post 2008 trajectories, States with lower levels of unemployment in 2003 will show a stronger increase in GDP than States with high unemployment in 2003.

The test of hypothesis 5 involves adding unemployment in 2003 as a main effect. The results are provided below. Notice that the DF associated with UNEMP03 reflect the number of level-2 entities indicating that the model is testing whether unemployment in 2003 is related to *average* levels of GDP across the entire time period (see Bliese, Maltarich, & Hendricks, 2018; Preacher, Zyphur, & Zhang, 2010). The results indicate that State-level unemployment rates in 2003 are unrelated to average levels of GDP across the subsequent eight years.

```
> mod.main<-lme(GDP~UNEMP03+TIME+TRANS+POST,
+        random=~TIME+TRANS+POST|STATE,gdp.univ,
+        control=list(opt="optim"))
> round(summary(mod.main)$tTable,dig=3)
```

	Value	Std.Error	DF	t-value	p-value
(Intercept)	171.725	74.327	354	2.310	0.021
UNEMP03	10.672	11.364	49	0.939	0.352

TIME   14.167   2.741 354   5.169   0.000

TRANS   -15.276   3.697 354   -4.133   0.000

POST   -8.365   1.983 354   -4.219   0.000

To better illustrate how to interpret the level-2 variable of unemployment, we contrast a mixed-effects model that only includes a level-2 predictor of the unemployment rate in 2003 (UNEMP03) of GDP with a simple OLS regression model where GDP has been averaged across the eight measurement occasions. The OLS model is based on average GDP for each of the 51 higher-level entities (States) in the sample.

```
> mod.temp<-lme(GDP~UNEMP03,

+          random=~1|STATE,gdp.univ,

+          control=list(opt="optim"))

> round(summary(mod.temp)$tTable,dig=3)
```

Value Std.Error DF t-value p-value

(Intercept) -304.227   232.026 357 -1.311   0.191

UNEMP03   102.161   40.771 49 2.506   0.016

```
> TDAT<-aggregate(gdp.univ[,c("UNEMP03","GDP")],

+          list(gdp.univ$STATE),mean)

> names(TDAT)
```

[1] "Group.1" "UNEMP03" "GDP"

```
> nrow(TDAT)
```

[1] 51

```
> mod.ols<-lm(GDP~UNEMP03,TDAT)

> round(summary(mod.ols)$coef,dig=3)
```

Estimate Std. Error t value Pr(>|t|)

(Intercept) -304.227   232.026 -1.311   0.196

UNEMP03   102.161   40.771 2.506   0.016

When we omit TIME, TRANS and POST, we see that States with higher unemployment in 2003 also had higher average GDP when we collapse across all eight measurement occasions. The important point in interpreting the mixed-effect model, however, is that the estimates and standard errors are identical to a simple OLS regression model based on 51 observations where we explicitly averaged GDP over time using the aggregate function.

Hypothesis 6 states that unemployment in 2003 will be related to the post-2008 slopes. In the DGM, this phase of the analysis tends to be most complex and to require estimating multiple comparative models. We provide several alternative ideas for how to approach this phase.

Given the nature of our hypothesis, a logical starting point would be to estimate a simple growth model using only POST and restricting the model to the last four observations.

```
> mod.post<-lme(GDP~POST*UNEMP03,
+       random=~POST|STATE,gdp.univ,
+       control=list(opt="optim"),subset=TIME>=4)
> round(summary(mod.post)$tTable,dig=3)
```

Value Std.Error DF t-value p-value

(Intercept) -310.690   237.936 151 -1.306   0.194

POST   -7.149   5.969 151 -1.198   0.233

UNEMP03   104.339   41.810 49 2.496 0.016

POST:UNEMP03   2.316   1.049 151 2.208   0.029

Notice the use of the subset option and the significant interaction between the POST

trajectory and unemployment. Figure 1 suggests a finding opposite our hypothesis showing that

States with high unemployment in 2003 (one SD above the mean at 6.66) actually increase

slightly faster in terms of GDP post 2008 than states with low unemployment (one SD below the

mean at 4.53). Figure 1 also shows a surprisingly large difference between States 1 SD below the

mean relative to States 1 SD above the mean on unemployment (over 200). This large difference,

given the non-significant main effect parameter estimate of 10.7 (from a main effect model), is

an artifact of the non-normally distributed dependent variable and is further indication that these

raw data violate assumptions of the mixed-effects model.

Figure 25.1. Interaction with unemployment post-event only
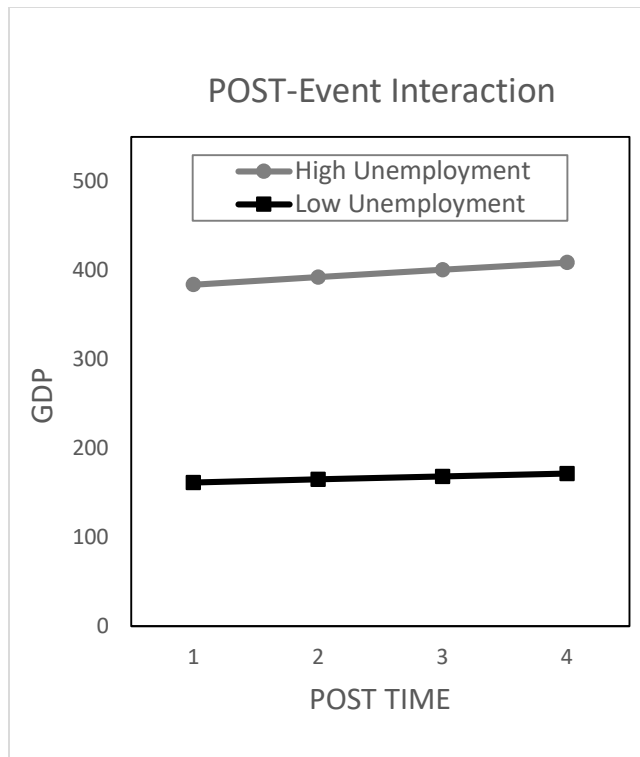
The advantage to estimating a partial model based on only the data associated with the POST event period is that the results provide a way to evaluate the findings using all the data and the full specification of the DGM. Below we estimate a model based on all time points, but which includes only one two-way interaction term involving POST:UNEMP03. Notice that the interaction is non-significant. The fact that the interaction is significant when based only on the subset of the last four time periods, but is non-significant when based on the full data suggests that the model with the full data is miss-specified. Note that we use TIME.A to estimate absolute effects because our hypothesis centers on whether the unemployment and post-event trajectory interaction is significant not whether it differs from the pre-event estimates in any relative way.

```
> mod.post.2<-lme(GDP~UNEMP03+TIME.A+TRANS+POST*UNEMP03,
+       random=~TIME.A+TRANS+POST|STATE,gdp.univ,
+       control=list(opt="optim"))
> round(summary(mod.post.2)$tTable,dig=3)
```

```
        Value Std.Error  DF t-value p-value
```

(Intercept) 177.267   74.641 353 2.375 0.018

UNEMP03   9.680   11.426 49 0.847 0.401

TIME.A   14.167   2.739 353 5.171 0.000

TRANS   -1.110   1.789 353 -0.620 0.535

POST   3.003   3.700 353 0.812 0.418

UNEMP03:POST 0.501   0.633 353 0.791 0.429

An alternative model includes all two-way interactions involving time-varying covariates and unemployment. In this model we recover the estimate for the interaction slope (2.316) identified when we focused only on the post-event slope.

```
> mod.full.a<-lme(GDP~UNEMP03+TIME.A*UNEMP03+TRANS*UNEMP03+POST*UNEMP03,
+       random=~TIME.A+TRANS+POST|STATE,gdp.univ,
+       control=list(opt="optim"))
> round(summary(mod.full.a)$tTable,dig=3)
```

```
          Value Std.Error  DF t-value p-value
```

(Intercept)   -259.271   200.092 351 -1.296 0.196

UNEMP03   87.743   35.160 49 2.496 0.016

TIME.A   -18.513   14.136 351 -1.310 0.191

TRANS   4.120   9.708 351 0.424 0.672

POST   -7.149   5.687 351 -1.257 0.210

UNEMP03:TIME.A 5.844   2.484 351 2.353 0.019

UNEMP03:TRANS -0.935   1.706 351 -0.548 0.584

UNEMP03:POST   2.316   0.999 351   2.318 0.021

      We might consider dropping the interaction between TRANS and unemployment to simplify the model leaving two 2-way interactions in the model. As shown in a variant of the model that omits the interaction involving TRANS, the other two interactions remain significant with slightly altered estimates.

> tmod<-lme(GDP~TIME.A*UNEMP03+TRANS+POST*UNEMP03,

+   random=~TIME.A+TRANS+POST|STATE,gdp.univ,

+   control=list(opt="optim"))

> round(summary(tmod)$tTable,dig=3)

          Value Std.Error DF t-value p-value

(Intercept)   -227.512   190.965 352 -1.191   0.234

TIME.A   -16.280   13.511 352 -1.205   0.229

UNEMP03   82.064   33.502 49   2.449   0.018

TRANS   -1.110   1.778 352 -0.624   0.533

POST   -6.314   5.494 352 -1.149   0.251

TIME.A:UNEMP03   5.445   2.370 352   2.297   0.022

UNEMP03:POST   2.167   0.964 352   2.247   0.025

      Finally, it is worth reviewing the model based on TIME.A versus the models based on TIME in the context of the two-way interactions. A model based on TIME provides a test of whether the form of the interaction involving POST:UNEMP03 differs from the form of the interaction involving TIME:UNEMP03. In the output below, we repeat the model based on TIME.A followed by the model based on TIME.

```
> mod.full.a<-lme(GDP~UNEMP03+TIME.A*UNEMP03+TRANS*UNEMP03+POST*UNEM
P03,
+        random=~TIME.A+TRANS+POST|STATE,gdp.univ,
+        control=list(opt="optim"))
> round(summary(mod.full.a)$tTable,dig=3)
```

|	Value	Std.Error	DF	t-value	p-value
(Intercept)	-259.271	200.092	351	-1.296	0.196
UNEMP03	87.743	35.160	49	2.496	0.016
TIME.A	-18.513	14.136	351	-1.310	0.191
TRANS	4.120	9.708	351	0.424	0.672
POST	-7.149	5.687	351	-1.257	0.210
UNEMP03:TIME.A	5.844	2.484	351	2.353	0.019
UNEMP03:TRANS	-0.935	1.706	351	-0.548	0.584
UNEMP03:POST	2.316	0.999	351	2.318	0.021

```
> mod.full<-lme(GDP~UNEMP03+TIME*UNEMP03+TRANS*UNEMP03+POST*UNEMP03
, +        random=~TIME+TRANS+POST|STATE,gdp.univ,
+        control=list(opt="optim"))
> round(summary(mod.full)$tTable,dig=3)
```

	Value	Std.Error	DF	t-value	p-value
(Intercept)	-259.271	199.955	351	-1.297	0.196
UNEMP03	87.743	35.136	49	2.497	0.016
TIME	-18.513	14.133	351	-1.310	0.191
TRANS	22.633	19.343	351	1.170	0.243

POST   11.364   10.386 351   1.094   0.275

UNEMP03:TIME   5.844   2.483 351   2.353   0.019

UNEMP03:TRANS   -6.779   3.399 351 -1.995   0.047

UNEMP03:POST   -3.528   1.825 351 -1.933   0.054

Notice that the interaction value for UNEMP03:TIME is the same in both models at 5.844. In contrast, the estimate for UNEMP03:POST is -3.528 in the model based on TIME instead of TIME.A, representing the difference between 5.844 and 2.316 from the absolute model. When using TIME, we are testing whether the form of the two-way interaction differs pre and post event (see also Bliese & Lang, 2016). In Figure 2 we can see that the interaction between TIME and UNEMP03 is more pronounced than the interaction between POST and UNEMP03 and this difference between the interactions is captured in the relative test.

Figure 25.2. Predicted values and interactions full model

Interactions Full Model
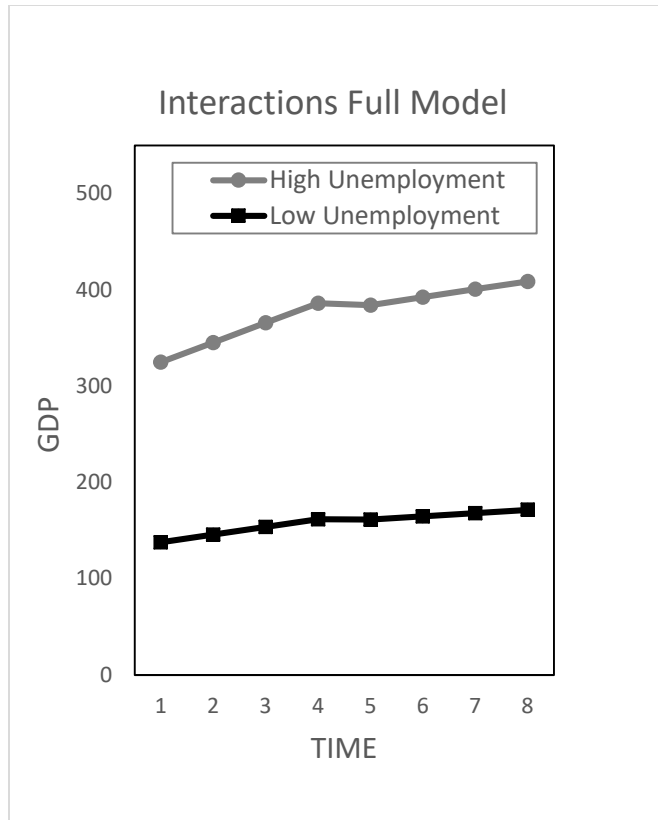
Alternative Models and Practical Advice

We conclude this section by discussing some alternative models and offering practical advice on identifying models. First, as noted in Bliese and Lang (2016) and Singer and Willett (2003) different forms of discontinuity may be captured by using combinations of two time-varying covariates (TIME and TRANS or TIME and POST) rather than all three time-varying covariates. Using just TIME and TRANS estimates a model where the pre and post event slopes are held constant but where a distinct increase or decrease occurs at the transition point. Using just TIME and POST specifies a model where the slope changes at the transition point, but where no immediate increase or decrease occurs. Often one of these two alternatives may be as good as or better than the model with all three time-varying covariates. We have found, for instance, that models based on TIME and POST often fit the data well and are conceptually logical (an event starts changing a trajectory but does not immediately cause a change). Bliese and Lang (2016)

provide some suggestions on how to compare models with and without the full set of time-varying covariates, but the decision ultimately comes down to whether one of the alternatives can be supported on theoretical grounds and knowledge about the underlying change process.

Second, as previously discussed we have found that it is often difficult to determine the exact nature of the random variability in the time-varying covariates. In our example, the comparison of -2log likelihood tests clearly suggested that data contained random variability in TIME, TRANS and POST. In our experience with other data sets, however, this type of clarity is rare because terms tend to be confounded with each other. For instance, a large variance estimate for TRANS suggests that some entities responded quite differently to the event than did other entities. A large change at TRANS, though, is then likely to cause the initial value of POST to begin at different places for each higher-level entity which can then lead to large differences in the POST slope.

We want to again emphasize that we often estimate several alternative models to determine which time-related covariates need to be specified as random effects. There are multiple options for specifying random effects to include or omit in the correlations among random terms that can help uncover the nature of the random variability. We reiterate that a guiding principal is to attempt to minimize the number of terms in the random statement. That is, it can be tempting to include as many random effects as possible, but our experience is that fewer is better. At a practical level, we also sometimes simply revert back to the random intercept model even when running interactions involving TIME, TRANS or POST with the idea that effects should be robust to various model specifications.

Third, we find that identifying cross-level interactions is often challenging. As we showed, even if a research hypothesis only centers on an interaction involving POST, it may be

DGM Models 45

necessary to include other interactions to avoid a miss-specified model. One solution we offer is to break the sample into segments and test smaller models to ensure that the final model is congruent with the subset model. Sub-setting the full dataset works fine for trends and could be used in our example for both TIME and POST. We do not have a strategy on how to estimate a subset model for TRANS as this term is highly influenced by TIME and POST trends and therefore hard to isolate.

Fourth, we encourage researchers to plot model-based estimates. Plots may be particularly informative with respect to TRANS to show whether a clear change occurs at the time of the event. Plots can be created in R, but we often simply copy predicted values to EXCEL. For example, the code below was used to obtain high and low values for unemployment and plot values for Figure 2. The last line of code writes the predicted values to the Windows clipboard so they can be pasted into EXCEL.

```
> mod.full<-lme(GDP~UNEMP03+TIME*UNEMP03+TRANS*UNEMP03+POST*UNEMP03
, +         random=~TIME+TRANS+POST|STATE,gdp.univ,
+         control=list(opt="optim"))
> #Create dataframe with high unemployment
> TDAT.H<-gdp.univ[1:8,c("TIME","TIME.A","TRANS","POST")]
> TDAT.H$UNEMP03<-mean(gdp$UNEMP03)+sd(gdp$UNEMP03)
> TDAT.H
```

	TIME	TIME.A	TRANS	POST	UNEMP03
1	0	0	0	0	6.658336
1.1	1	1	0	0	6.658336
1.2	2	2	0	0	6.658336

1.3   3   3   0   0 6.658336

1.4   4   3   1   0 6.658336

1.5   5   3   1   1 6.658336

1.6   6   3   1   2 6.658336

1.7   7   3   1   3 6.658336

> #Create dataframe with low unemployment

> TDAT.L<-gdp.univ[1:8,c("TIME","TIME.A","TRANS","POST")]

> TDAT.L$UNEMP03<-mean(gdp$UNEMP03)-sd(gdp$UNEMP03)

> TDAT.L

    TIME TIME.A TRANS POST UNEMP03

1   0   0   0   0 4.525978

1.1   1   1   0   0 4.525978

1.2   2   2   0   0 4.525978

1.3   3   3   0   0 4.525978

1.4   4   3   1   0 4.525978

1.5   5   3   1   1 4.525978

1.6   6   3   1   2 4.525978

1.7   7   3   1   3 4.525978

> PRED.VALS<-data.frame(TIME=1:8,HI.UNEMP=predict(mod.full,TDAT.H,level=0),

+   LO.UNEMP=predict(mod.full,TDAT.L,level=0))

> PRED.VALS

  TIME HI.UNEMP LO.UNEMP

1   1 324.9518 137.8522

2   2 345.3490 145.7882

3   3 365.7463 153.7243

4   4 386.1436 161.6604

5   5 384.0367 161.5476

6   6 392.3078 164.8804

7   7 400.5790 168.2132

8   8 408.8502 171.5459

> write.table(PRED.VALS,file="clipboard",sep="\t",col.names=NA)

Finally, with respect to practical advice, we emphasize the integration of theory with analytics. That is, carefully consider whether an included term or interaction term is theoretically necessary. For instance, if a randomized trial collects four waves of baseline data prior to randomly assigning participants to condition, then it may not make theoretical sense to include a TIME*Condition interaction because TIME precedes the random assignment and, thus, the initial slope cannot be caused by differences in condition. While it may seem prudent and conservative to include this interaction term (as a way to help ensure equivalence in random assignment), we have tended to find that the inclusion of this term produces a misspecified model. In short, our experience is that models provide the most clarity when random slopes and interactions involving time-varying covariates are kept to a minimum.

AN R BASED TOOL FOR CREATING THE MODEL MATRIX

In the examples provide up to this point, and in many situation where the DGM is used, setting up the design matrix of TIME, TRANS and POST is a simple programming exercise: we showed that researchers can start with TIME and program conditional statements to obtain TRANS and

POST. Unfortunately, there are many situations where the DGM might be useful but where the data is not uniformly organized around an event. For instance, if we are examining turnover and team-level data over 6 years, we would fully expect that the event (turnover) would not occur at the same time for each team. In this case, we need to set up a design matrix where TIME, TRANS and POST potentially differ for each higher-level entity. An example of such data is presented in Table 25.6. Notice that turnover occurs at TIME 3 for TEAM 1, at TIME 1 for TEAM 2, and at TIME 4 for TEAM 3.

Table 25.6. Discontinuous events at different times

TEAM	TIME	TURNOVER	PERF
1	0	0	1.00
1	1	0	2.00
1	2	0	3.00
1	3	1	2.00
1	4	0	4.00
1	5	0	5.00
2	0	0	0.95
2	1	1	0.00
2	2	0	2.05
2	3	0	2.92
2	4	0	3.96
2	5	0	5.06
3	0	0	1.08

3	1	0	1.97
3	2	0	3.03
3	3	0	4.01
3	4	1	3.00
3	5	0	5.01

The Appendix provides an R code function (dgm.code) to automate creating the design matrix. At a minimum, three parameters are provided to the function: the higher-level entity (grp), the time variable (time) and a vector of 0's and 1's for when the event occurs. In the data provided in Table 25.6, the call to dgm.code would be:

```
> dgm.matrix<-with(TABLE6.DAT,dgm.code(TEAM, TIME, TURNOVER))
> dgm.matrix
```

	grp	time	event	trans1	post1	time.a	tot.events	event.first
1	1	0	0	0	0	0	1	0
2	1	1	0	0	0	1	1	0
3	1	2	0	0	0	2	1	0
4	1	3	1	1	0	2	1	0
5	1	4	0	1	1	2	1	0
6	1	5	0	1	2	2	1	0
7	2	0	0	0	0	0	1	0
8	2	1	1	1	0	0	1	0
9	2	2	0	1	1	0	1	0
10	2	3	0	1	2	0	1	0
11	2	4	0	1	3	0	1	0

DGM Models 50

12	2	5	0	1	4	0	1	0
13	3	0	0	0	0	0	1	0
14	3	1	0	0	0	1	1	0
15	3	2	0	0	0	2	1	0
16	3	3	0	0	0	3	1	0
17	3	4	1	1	0	3	1	0
18	3	5	0	1	1	3	1	0

In the current version of the dgm.code the variable names will always be those listed above (e.g., "grp", "time", "event", etc.). Notice that trans1 and post1 differ for each of the groups. Also notice that the code provides a column for tot.events (total count of events) and event.first (which returns a 1 if the event at the first time point). Additional options to the dgm.code allow users to restrict the total number of events to the first set number of occurrences and/or to ignore the event if it occurs at the first time period. If multiple events occur and are of interest, the dgm.code will create transN and postN matrices to correspond to the multiple events based on the specifications presented in Bliese and Lang (2016). Interested readers can refer to the Open Science Foundation site (https://osf.io/d4xmw/) for the function and additional examples of how to use the dgm.code options.

For the purposes of creating a dataset to run substantive modeling, the dgm.matrix object needs to be merged back with the original dataset (TABLE6.DAT) in this case. The merge code is:

```
> TABLE6.FINAL<-merge(TABLE6.DAT,dgm.matrix,by.x=c("TEAM","TIME"),
+          by.y=c("grp","time"))
```

In this example, as long as the dgm.matrix is the second object in the merge command, the by.y=c("grp","time")option will remain the same and users only need to change the names of the grouping variable and the time variable in the original dataset.

Ultimately, no automated tool should be used without a thorough understanding of the data. Even with the use of this tool, users will have to check the data. That said, users may decide on a specific set of principles (e.g., only examine patterns in higher level entities where the at least 3 time periods proceeded the event and where at least 3 time periods follow an event and/or only model up to 2 event occurrences). If these types of principles are established, then the subset of data meeting the criteria can be provided to the dgm.code function and the automated creation of the necessary time variants should still facilitate data analyses.

CONCLUSIONS

In this chapter we have provided details and practical advice related to estimating the DGM. We believe that the DGM represents a valuable tool for researchers interested in understanding change over time when some type of event potentially disrupts the change process. In our minds, one of the strengths of the DGM approach is that the components of the model allow for clear and precise theoretical specification about how events impact change over time (immediate impact associated with the transition or a change in slope following the event or both). The models also provide a way to specify whether change is considered relative to pre-event trajectories or absolute. As noted elsewhere, the precision associated with the DGM has the potential to help advance theory on many important topics such as adaptability and resilience (Bliese, et al., 2017).

While the basic form of the DGM is relatively simple to specify in terms of adding more time-varying covariates, we have also shown that researchers need to carefully approach model building and model testing. The ability to add random slopes for numerous time-varying covariates is a strength of the approach but can quickly produce models that are over-fit and miss-specified. At the same time, including unnecessary interaction terms or conversely omitting terms that should be included can lead to model misspecification. The issues we raise with respect to fitting models are not specifically unique to the DGM because all relatively complex statistical models risk being over-fit or miss-specified. With respect to the DGM, however, we have attempted to provide numerous alternatives and recommendations for helping users fit models. Our overall recommendations are to (a) specify models that use fewer (rather than more) random terms, (b) test alternative models by potentially sub-setting parts of the temporal sequence to determine the robustness of findings, and (c) to plot out predicted values from the models to help understand the nature of change over time.

REFERENCES

Bailey, M. 2016. *Real Econometrics: The Right Tools to Answer Questions*. Oxford: Oxford University Press.

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language, 68(3)*, 255-278.

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

Bliese, P. (2016). Package 'multilevel'. Available at http://cranr-

project.org/web/packages/multilevel/multilevel.pdf (accessed date March 26th, 2019)

Bliese, P. D. (2000). Within-group agreement, non-independence, and reliability: Implications

for data aggregation and analyses. In K. J. Klein & S. W. J. Kozlowski (Eds.), *Multilevel*

Theory, Research, and Methods in Organizations: Foundations, Extensions, and New

Directions (pp. 349– 381). San Francisco: Jossey-Bass.

Bliese, P. D., Adler, A. B., & Flynn, P. J. (2017). Transition processes: A review and synthesis

integrating methods and theory. *Annual Review of Organizational Psychology and*

Organizational Behavior, *4*, 263-286.

Bliese, P. D., & Lang, J. W. (2016). Understanding relative and absolute change in discontinuous

growth models: Coding alternatives and implications for hypothesis

testing. *Organizational Research Methods*, *19*(4), 562-592.

Bliese, P. D., Maltarich, M. A., & Hendricks, J. L. (2018). Back to basics with mixed-effects

models: Nine take-away points. *Journal of Business and Psychology*, *33*(1), 1-23.

Bliese, P. D., & Ployhart, R. E. (2002). Growth modeling using random coefficient models:

Model building, testing, and illustrations. *Organizational Research Methods*, 5, 362-387.

Bliese, P. D., Schepker, D. J. Essman, S., & Ployhart, R. E. (in review). Bridging

Methodological Divides between Macro and Micro Research: Methods and Endogeneity

in Panel Data.

Cudeck, R., & Klebe, K. J. (2002). Multiphase mixed-effects models for repeated measures data.

Psychological Methods, 7, 41-63.

Fulmer, C. A., & Gelfand, M. J. (2015). Trust after violations: Are collectivists more or less

forgiving?. *Journal of Trust Research*, *5*(2), 109-131.

Gay, D.M. (1990): Usage Summary for Selected Optimization Routines, Computing Science
Technical Report No. 153, AT&T Bell Laboratories, Murray Hill NJ.

Gelman, A., & Imbens, G. (2018). Why high-order polynomials should not be used in regression
discontinuity designs. *Journal of Business & Economic Statistics*, 1-10.

Hadfield, J (2010). MCMC Methods for Multi-Response Generalized Linear Mixed Models: The
MCMCglmm R Package. *Journal of Statistical Software*, 33(2), 1-22

Hernández-Lloreta, M. V., Colmenares, F., & Martı´nez-Arias, R. (2004). Application of
piecewise hierarchical linear growth modeling to the study of continuity in behavioral
development of baboons (Papio hamadryas). *Journal of Comparative Psychology*, 118,
316-324.

Kim, Y., & Ployhart, R. E., (2014). The effects of staffing and training on firm productivity and
profit growth before, during, and after the great recession. *Journal of Applied
Psychology*, 99, 361-389.

Korsgaard, M. A., Kautz, J., Bliese, P., Samson, K., & Kostyszyn, P. (2018). Conceptualising
time as a level of analysis: New directions in the analysis of trust dynamics. *Journal of
Trust Research*, *8*(2), 142-165.

LaHuis, D. M., & Ferguson, M. W. (2009). The accuracy of significance tests for slope variance
components in multilevel random coefficient models. *Organizational Research
Methods*, *12*(3), 418-435.

Lang, J. W. B., & Bliese, P. D. (2009). General mental ability and two types of adaptation to
unforeseen change: Applying discontinuous growth models to the task-change paradigm.
Journal of Applied Psychology, 92, 411-428.

Littell, R. C., Milliken, G. A., Stroup, W. W., Wolfinger, . R. D. & Schabenberger. O (2006). *SAS for Mixed Models, Second Editio.* Cary NC, SAS Institute.

Lucas, R. E. (2007). Long-term disability is associated with lasting changes in subjective well-being: evidence from two nationally representative longitudinal studies. *Journal of Personality and Social Psychology*, *92*(4), 717-730.

Lucas, R. E., Clark, A. E., Georgellis, Y., & Diener, E. (2003). Reexamining adaptation and the set point model of happiness: reactions to changes in marital status. *Journal of Personality and Social Psychology*, *84*(3), 527-539.

Lucas, R. E., Clark, A. E., Georgellis, Y., & Diener, E. (2004). Unemployment alters the set point of life satisfaction. *Psychological Science*, 15, 8-13.

Matuschek, H., Kliegl, R., Vasishth, S., Baayen, H., & Bates, D. (2017). Balancing Type I error and power in linear mixed models. *Journal of Memory and Language, 94*, 305-315.

McNeish, D., Stapleton, L. M., & Silverman, R. D. (2017). On the unnecessary ubiquity of hierarchical linear modeling. *Psychological Methods*, *22*(1), 114-140.

Morgeson, F. P., Mitchell, T. R., & Liu, D. (2015). Event system theory: An event-oriented approach to the organizational sciences. *Academy of Management Review*, *40*(4), 515-537.

Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *The Computer Journal*, *7*(4), 308-313.

Preacher, K. J., Zyphur, M. J., & Zhang, Z. (2010). A general multilevel SEM framework for assessing multilevel mediation. *Psychological Methods*, *15*(3), 209-233.

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

Rabe-Hesketh, S. & Skrondal A. (2005). *Multilevel and Longitudinal Modeling using Stata*.
College Station, TX: StataCorp LP

Raudenbush, S. W., & Bryk, A. S. (2002). *Hierarchical Linear Models: Applications and Data
Analysis Methods*. Newbury Park, CA: Sage.

Rupp, T. L., Wesensten, N. J., Bliese, P. D., & Balkin, T. J. (2009). Banking sleep: Realization
of benefits during subsequent sleep restriction and recovery. *Sleep*, 32, 311-321.

Singer, J. D. (1998). Using SAS PROC MIXED to Fit Multilevel Models, Hierarchical Models,
and Individual Growth Models. *Journal of Educational and Behavioral Statistics, 23*,
323–355.

Singer, J. D., & Willett, J. B. (2003). *Applied Longitudinal Data Analysis: Modeling Change and
Event Occurrence*. New York, NY: Oxford University Press.

Snijders, T. A. B., & Bosker, R. J. (1999). *An Introduction to Basic and Advanced Multilevel
Modeling*. Thousand Oaks, CA: Sage.

Zeger, S. L., & Liang, K.-Y. (1986). Longitudinal data analysis for discrete and continuous
outcomes. *Biometrics*, 42, 121-130.

APPENDIX: R CODE FOR DGM DESIGN MATRIX

```
dgm.code<-function(grp,time,event,n.events=FALSE,first.obs=FALSE){

#ensure data structure is correct

  newdata<-data.frame(grp=grp,time=time,event=event)

  newdata<-na.exclude(newdata)

  newdata<-newdata[order(newdata$grp,newdata$time),]


#If first observation is an event and first.obs is TRUE

#change the first observation to a non-event

  if(first.obs){

  fobs.grps<-newdata$grp[!duplicated(newdata$grp)&(newdata$event==1)]

  #print(fobs.grps)

  newdata$event[!duplicated(newdata$grp)&(newdata$event==1)]<-0

  }


#Check to see if first observation is an event

  s.event<-nrow(newdata[!duplicated(newdata$grp)&(newdata$event==1),])

  if(s.event>0){

  print("The following groups start with an event")

  print(newdata[!duplicated(newdata$grp)&(newdata$event==1),])

  print("Drop the groups or use the first.obs=TRUE option")

  stop()
```

```
}


#count the maximum number of events for any group

  max.events<-max(with(newdata,tapply(event,grp,sum)))

  n.grps<-length(unique(newdata$grp))


#adjust the maximum number of events to a specified level

if(n.events){

  max.events=n.events

  }


# Set up the structure for the output

  ANS<-matrix(0,nrow(newdata),ncol=max.events*2)

  ANS<-data.frame(ANS)

  names(ANS)<-c(paste0("trans",c(1:max.events)),paste0("post",c(1:max.events)))

  ANS<-data.frame(ANS,grp=newdata$grp,time=newdata$time,event=newdata$event)

  g.size<-tapply(newdata$grp,newdata$grp,length)
#  ANS$cum.event<-unlist(tapply(ANS$event,ANS$grp,cumsum))

  ANS$time.a<-ANS$time

  ANS$cum.event<-do.call(c, tapply(ANS$event, ANS$grp, FUN=cumsum))

  ANS$num.grp<-rep(1:n.grps,times=g.size) #create a numeric group for loops


# Add two check variables for total events and whether an event
```

```
# occured on the first occasion

  ANS$tot.events<-NA

  ANS$event.first<-0


if(first.obs){

  ANS$event.first[ANS$grp%in%fobs.grps]<-1

  }


#collapse the number of events to a specified level

if(n.events){

  ANS$cum.event<-ifelse(ANS$cum.event>n.events,n.events,ANS$cum.event)

}


# Set up a factor outside of the loop to get all levels

  ANS$cum.event.f<-factor(ANS$cum.event,levels=c(0:max.events))


# Set up a loop to put values in trans and post variables

# First skip groups with no events

for(i in 1:n.grps){

  if(sum(ANS$event[ANS$num.grp==i])==0){

  ANS$tot.events[ANS$num.grp==i]<-0

  next(i)

  }
```

```
# Use model.matrix to set up dummy codes for trans and post

  ANS[ANS$num.grp==i,1:max.events]<-

model.matrix(~C(cum.event.f,contr.treatment),data=ANS[ANS$num.grp==i,])[,2:(max.events+1

)]

  ANS[ANS$num.grp==i,(max.events+1):(2*max.events)]<-

model.matrix(~C(cum.event.f,contr.treatment),data=ANS[ANS$num.grp==i,])[,2:(max.events+1

)]

if(max.events==1){

  ANS[ANS$num.grp==i,2]<-cumsum(ANS[ANS$num.grp==i,2])-1

}

if(max.events>1){

  ANS[ANS$num.grp==i,(max.events+1):(2*max.events)]<-

apply(ANS[ANS$num.grp==i,(max.events+1):(2*max.events)],2,cumsum)-1

}

  ANS$time.a[ANS$num.grp==i]<-

ifelse(ANS$cum.event[ANS$num.grp==i]==0,ANS$time[ANS$num.grp==i],NA)

  ANS$time.a[is.na(ANS$time.a)&ANS$num.grp==i]<-

max(ANS$time.a[!is.na(ANS$time.a)&ANS$num.grp==i])

  ANS$tot.events[ANS$num.grp==i]<-sum(ANS$event[ANS$num.grp==i])

next(i)

}

# Clean up the ANS matrix column by column

# print(ANS) to see the structure of the previous loop
```

```
for(j in 1:max.events){

  ANS[,max.events+j]<-ifelse(ANS[,j]==0,0,ANS[,max.events+j])

  next(j)

  }


#rearrange the ANS matrix for output

#print(ANS[1:30,]) to see the first 30 rows of complete data

ANS[,c((max.events*2)+1:3,1:(max.events*2),(max.events*2)+c(4,7,8))]

}
```