British Journal of Mathematical and Statistical Psychology (2011), 64, 134–160

The British Psychological Society
www.wileyonlinelibrary.com

The linear mixed model and the hierarchical Ornstein–Uhlenbeck model: Some equivalences and differences

Zita Oravecz* and Francis Tuerlinckx
Department of Psychology, University of Leuven, Belgium

We focus on comparing different modelling approaches for intensive longitudinal designs. Two methods are scrutinized, namely the widely used linear mixed model (LMM) and the relatively unexplored Ornstein–Uhlenbeck (OU) process based state-space model. On the one hand, we show that given certain conditions they result in equivalent outcomes. On the other hand, we consider it important to emphasize that their perspectives are different and that one framework might better address certain types of research questions than the other. We show that, compared to a LMM, an OU process based approach can cope with modelling inter-individual differences in aspects that are more substantively interesting. However, the estimation of the LMM is faster and the model is more straightforward to implement. The models are illustrated through an experience sampling study.

1. Introduction

1.1. Intensive longitudinal data designs

New data collection strategies can provide a solid basis to achieve a greater ecological validity in the measurement of psychological phenomena (Walls, Jung, & Schwartz, 2006). With modern technological devices, such as wearable computers or mobile phones, one can easily collect large volumes of longitudinal data from several individuals, in this way jointly measuring intra- and inter-individual variation. A typical example is the experience sampling method (Bolger, Davis, & Rafaeli, 2003; Csikszentmihalyi & Larson, 1987; Russell & Feldman-Barrett, 1999), in which measurements from several individuals are acquired frequently in a natural setting. Typically, the participants of an experience sampling study carry a handheld computer, which beeps them several times a day according to a certain time scheme (preferably random or semi-random) and upon the beep the individual has to answer to a set of questions. The

*Correspondence should be addressed to Zita Oravecz, Department of Psychology, University of Leuven, Tiensestraat 102, Box 3713, B-3000 Leuven, Belgium (e-mail: zita.oravecz@psy.kuleuven.be).

DOI:10.1348/000711010X498621

outcome of such a study consists of numerous waves of measurements for several individuals, and such data have been called intensive longitudinal data (ILD; Walls *et al.*, 2006).

ILD have some particular characteristics. First, the number of waves does not need to be (and often is not) equal for the participants, leading to unbalanced data. Second, the spacing of the data waves is very often unequal. Third, the participants might not share the same data collection schedule (i.e., a time unstructured design).

Besides the formal properties of the design, ILD are useful because they allow us to investigate different sources of variation both within and between individuals. Sources of variation that can be localized within individuals are measurement error, serial dependency, and variability due to the heterogeneity of within-person situations. Apart from intra-individual variation, individuals tend to differ on a variety of aspects (e.g., with respect to mean level and sizes of intra-individual variation).

As an example of such ILD collected through an experience sampling study, we will analyse a subset of a data set from the field of emotion research (Kuppens, Oravecz, & Tuerlinckx, 2010). The study focused on exploring the changes in individuals core affect (Russell, 2003), which is a consciously accessible momentary feeling of one's pleasure and activation level. In our application, we will analyse 50 measurements collected for a sample of 60 persons during a single day. In such a design, we definitely have to take into account different sources of variation, such as intra-individual variation and serial correlation structure. Moreover, it is of considerable interest to find out whether individuals differ with respect to these key terms.

1.2. Methods for ILD analysis

Different data analytic frameworks can exhibit different merits when they are applied to analyse ILD. We consider it useful to give a short overview of the possible approaches before we set the actual focus of the paper. Note that we restrict our focus to models for continuous outcomes.

A first method to consider is classical time-series analysis (TSA; e.g., Shumway & Stoffer, 2006). In many intensive longitudinal designs, the length of each individual's measurement chain would permit a TSA. However, it is important to note that the great bulk of the classical TSA models (i.e., the ARIMA models) have been developed for equally spaced measurements. Moreover, while an interesting aspect of ILD lies in the opportunity to investigate and explain individual differences in terms of covariates, these models do not focus on these substantively meaningful problems. Although this framework is a possible candidate for ILD analysis, it will not directly be in the focus of the current paper.

A second method is structural equation modelling (SEM; e.g., Bollen, 1989; McArdle, 2009, for a recent overview). The SEM approach is well developed and widely used in social sciences for exploring inter-individual variation, using latent variables to represent individual differences. The latent variables in the majority of the SEM models are assumed to be equally spaced, although they can be fitted for unequally or equally spaced data as well (McArdle, 2009). However, in ILD designs, this approach may not be feasible. First, assuming a large number of latent variables (because there are many measurements) may lead to numerical problems. Second, it may be important to take into account the time lag when modelling the latent states if one believes that there is a truly continuously changing latent process. An additional problem may be that in ILD one often encounters a larger number of observations than individuals, and this may result in a singular covariance matrix (for a summary of the problem and a possible solution, see Hamaker, Dolan, & Molenaar, 2003). Although this framework provides a large variety of models, some of them are actually very close to the linear mixed model (LMM) framework which we discuss in this paper (see Skrondal & Rabe-Hesketh, 2004, for more information). For these reasons, we do not focus on SEM in the current paper.

A third widely used method to analyse longitudinal data is the LMM framework (Laird & Ware, 1982; Verbeke & Molenberghs, 2000), developed primarily in biostatistics (see Verbeke & Molenberghs, 2000) and educational research (Raudenbush & Bryk, 2002). Depending on the area, the LMM is also called a multi-level and hierarchical model. LMMs can handle unequally spaced measurements and provide possibilities to analyse the different sources of variation listed above. This approach will be analysed further in this paper.

Finally, state-space models can also be used to analyse ILD. The core of this approach is to express the change by two sets of equations: a dynamic equation describing the changes in the true score of the variable and an observation equation which links the latent variables to the observed data. Note that the general set-up bears some resemblance to the SEM framework, in which a similar distinction is made between the structural and observation part. Therefore, some state-space models – especially the discrete time variants – can be translated into a corresponding SEM model (McCallum & Ashby, 1986). In addition, the transition equation in a discrete time state-space model is related to a (possibly multivariate) time-series model for the latent data. Besides, there are close links between the state-space framework and the LMM, as has been demonstrated by Jones (1993) and Jones and Boadi-Boateng (1991).

Despite the communalities between the state-space modelling framework and other models, it is treated separately since it comprises a broad class of models. In this paper, we will consider a particular instance of this class, the hierarchical Ornstein–Uhlenbeck (OU) process (or HOU model) recently proposed by Oravecz and Tuerlinckx (2008) and Oravecz, Tuerlinckx, and Vandekerckhove (2009a) based on earlier work by Oravecz, Tuerlinckx, and Vandekerckhove (2009b) and which can be seen as a state-space model in continuous time. It is different from typical state-space models because it allows for individual differences in all levels of the model, not only in the mean structure but also in the variances and serial dependencies. The HOU model will be explained in detail in Section 2.2.

1.3. Aim of the study and overview

The main goal of the paper is theoretical. We aim to compare the LMM with the HOU model to find the conditions under which they are equivalent and when they yield different results. Such an exploration should help in clarifying what the models can do and how they should be interpreted. It may also help to understand when one model is preferred over the other. For example, as we will see, the LMM provides a flexible and straightforward way of exploring inter-individual differences in the mean trajectory over time, while the main benefit of the HOU model based approach is to explore individual differences with respect to the dynamical properties of change, such as the serial correlation.

We also have a more pragmatic subgoal. In recent decades, estimation routines have been developed for the LMM using direct optimization of the likelihood. These routines have been implemented in many general purpose software packages (e.g., SAS, SPSS, R). The hallmark of these programs is that they work rapidly for a large amount of data. On the other hand, the inference for the relatively new HOU model is done in a Bayesian framework relying on Markov chain Monte Carlo techniques (see Robert & Casella, 2004), which have been applied to other stochastic differential equation models as well (De la Cruz-Mesía & Marshall, 2003, 2006; Golightly & Wilkinson, 2006). An advantage of the Bayesian inference method in the case of the current model is that it uses exact conditional densities of the latent states in setting up the likelihood, so there is no approximation step involved. We will provide software codes for fitting both types of models in Appendices A–F. Although inference for the HOU model works well (see Oravecz & Tuerlinckx, 2008; Oravecz *et al.*, 2009a), it is much more computer and time intensive than fitting an LMM. For this reason, if under some circumstances both models tend to be equivalent or nearly equivalent, one may make use of existing LMM software to fit the HOU model.

The remainder of the paper has the following structure. We start by providing a general description of the two frameworks. Some exact or approximately equivalent representations are derived. Then an experience sampling data set is used to illustrate the derived results. Finally, we close the paper with some conclusions and a discussion.

2. Introducing the LMM and the HOU models

2.1. The LMM model

In the LMM framework, the repeated measurements are modelled using a linear regression model, in which some of the parameters are constant over individuals and others are subject-specific. The constant parameters are called fixed effects and the subject-specific effects are assumed to be drawn from some distribution (mostly the normal) and are called random effects. The random effects reflect the natural heterogeneity in the population and represent a deviation from the mean in the population. The LMM for the outcome Y_{ps} of person p at measurement occasion s can be formulated as

$$Y_{ps} = \beta_0^* + \beta_1^* x_{ps} + b_{p0}^* + b_{p1}^* z_{ps} + \epsilon_{ps}^*, \tag{1}$$

where x_{ps} and z_{ps} are (possibly time-varying) covariates (they can be the same two variables). The fixed regression coefficients are β_0^* and β_1^* and the subject-specific effects are denoted by b_{p0}^* and b_{p1}^*. The subject-specific effects are assumed to follow a bivariate normal distribution $(\beta_{p0}^*, \beta_{p1}^*)^{\mathrm{T}} \sim MVN(\mathbf{0}, \mathbf{D}^*)$ with mean $\mathbf{0}$ and $\mathbf{D}^*$ as a 2×2 covariance matrix. The residual ϵ_{ps}^* follows a normal distribution $\epsilon_{ps}^* \sim N(0, \sigma_{\epsilon^*}^2)$ with $\sigma_{\epsilon^*}^2$ as its variance. The residuals and the random effects are assumed to be independent of each other (in addition, the residuals are independent of one another). Note that in order to distinguish the parameters of the LMM and the HOU model, we add an extra asterisk ($*$) to every LMM parameter.

With respect to the mean structure of the model and the subject-specific deviations from it, the model as presented in equation (1) is only a very simple version of the more general LMM. According to the traditional representation, the mean structure is written as $\boldsymbol{x}_{ps}^{\mathrm{T}}\boldsymbol{\beta}^* + \boldsymbol{z}_{ps}^{\mathrm{T}}\mathbf{b}_p^*$. However, we did not use the most general formulation of the LMM, but a simpler version, in order to provide maximal comparability with the HOU model.

Some authors (e.g., Diggle, Liang, & Zeger, 1994; Verbeke & Molenberghs, 2000) propose a further decomposition of the residuals, namely into true measurement error and a component that exhibits serial correlation over time. The residual ϵ^*_{ps} can be decomposed as follows:

$$\epsilon^*_{ps} = \epsilon^*_{(1)ps} + \epsilon^*_{(2)ps}, \tag{2}$$

where $\epsilon^*_{(1)ps} \sim N(0, \sigma^2_{\epsilon^*(1)})$ represents measurement error, and $\epsilon^*_{(2)ps}$ reflects a component of serial correlation for which a normal distribution is assumed: $\epsilon^*_{(2)ps} \sim N(0, \tau^{2*} g^*(\Delta_{ps}))$. There are two components in the variance of the latter distribution. The autocorrelation function is represented by $g^*(\Delta_{ps})$, where g^* is a monotonic decreasing function with $g^*(0) = 1$ (perfect autocorrelation when there is no time difference) and Δ_{ps} represents the elapsed time between measurement $Y_{p,s-1}$ and Y_{ps}. A frequently used function for $g^*(\Delta)$ is the exponential correlation function, which is specified as $e^{-\phi^* \Delta}$ (with $\phi^* > 0$) and is illustrated in Figure 1. The steepness of the decay in the autocorrelation is determined by the parameter ϕ^*. It is important to note that in the LMM framework it is customary to assume that the serial correlation effect is a population phenomenon, and independent of the individual (Verbeke & Molenberghs, 1997).

2.2. The HOU model

The HOU model builds further on the recent work of Oravecz *et al.* (2009b) who introduced an OU stochastic process based hierarchical model that is directly observable (hence without any measurement error). The model of Oravecz *et al.* (2009b) is based on the work of Blackwell (2003) and Brillinger, Preisler, Ager, and Kie (2004), who applied a (single, non-hierarchical) OU process for modelling animal movements.

A straightforward way to present the HOU model would be to write the transition and observation equation, with the former containing the OU stochastic differential equation (OUSDE) representing the continuous time random change on the state level. However, because the OUSDE is one of the few stochastic differential equations with a closed-form solution, it is clearer to present the solution of the equation directly at the state level.

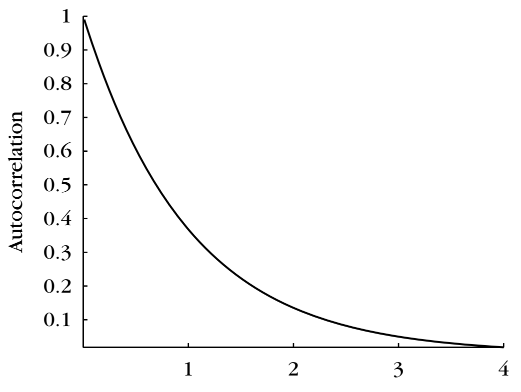

Figure 1. The exponential ($e^{-\phi^* \Delta}$) autocorrelation function.

The model then becomes

$$\theta_{ps} = \mu_p + e^{-\lambda_p \Delta_{ps}}(\theta_{p,s-1} - \mu_p) = \eta_{ps}, \tag{3}$$

$$Y_{ps} = \theta_{ps} + \epsilon_{ps}, \tag{4}$$

with Y_{ps} standing for the observation for person p at occasion s and Δ_{ps} denoting the time difference between two consecutive measurements (see Section 2.1 for its definition). Equation 3 represents the transition equation of the state-space model by describing the changes in the 'true score' θ_{ps}. The parameter μ_p represents the mean level of the latent process (sometimes also called the attractor level or home base). The random innovation η_{ps} in equation (3) is normally distributed as

$$\eta_{ps} \sim N(0, \gamma(1 - e^{-2\lambda_p \Delta_{ps}})). \tag{5}$$

The somewhat complex variance term appears as such in the solution of the OUSDE; we will give a more detailed description later. Equation (4) of the HOU model adds measurement error to the 'true position' in order to get the observed data. This observation equation relates this part to the observed data by adding a measurement error ϵ_{ps} sampled from $N(0, \sigma_\epsilon^2)$.

In some cases, it is helpful to rewrite the dynamic equation part of the HOU model by subtracting the mean level from both sides, arriving at a zero mean process $\tilde{\theta}_{ps}$:

$$\begin{aligned} \theta_{ps} &= \mu_p + e^{-\lambda_p \Delta_{ps}}(\theta_{p,s-1} - \mu_p) + \eta_{ps}, \\ \theta_{ps} - \mu_p &= e^{-\lambda_p \Delta_{ps}}(\theta_{p,s-1} - \mu_p) + \eta_{ps} \\ \tilde{\theta}_{ps} &= e^{-\lambda_p \Delta_{ps}}\tilde{\theta}_{p,s-1} + \eta_{ps}. \end{aligned} \tag{6}$$

The mean can then be added to the observation equation:

$$Y_{ps} = \mu_p + \tilde{\theta}_{ps} + \epsilon_{ps}. \tag{7}$$

With this representation, it is more straightforward to derive the autocorrelation function of the process (see the exact derivation in Appendix A), which equals $e^{-\lambda_p \Delta_{ps}}$. The same autocorrelation function holds for the original non-zero centred process θ_{ps}. The autocorrelation is determined by λ_p (with $\lambda_p > 0$). If we keep the elapsed time Δ_{ps} fixed, and increase λ_p, we can see that a higher value of λ_p causes the autocorrelation function to approach 0, hence the position of the process $\tilde{\theta}_{ps}$ will be very close to 0 (or θ_{ps} close to μ_p). On the other hand, a low value of λ_p causes the autocorrelation to approach 1, hence θ_{ps} stays close to $\theta_{p,s-1}$. Therefore, the parameter λ_p can be interpreted as an adjusting force or centralizing tendency, controlling how strong the adjustment is to the mean level.

If we let the process run for a very long time (i.e., let $\Delta_{ps} \rightarrow \infty$), it converges to a normal limiting distribution with μ_p as its mean and γ as its variance. Hence, γ can be defined as the variance of the equilibrium distribution and can be interpreted as the true intra-individual variance.

Two remarks are in order. First, note that in this general formulation of the OU process we allow the parameters μ_p and λ_p to be person-specific but we have not yet assigned distributions to them. The reason is that we will first examine the equivalences with the LMM and then conclude what we may derive about the distributions for the random effects in the HOU if we assume that the random effects in the LMM are normally distributed. The reason for taking the LMM as a starting-point and concluding what the consequences are for the HOU is that in practice the LMM with normally distributed random effects will be used as a proxy for the HOU, and then it is important to know what the implications are for the HOU model. Second, we do not consider the variance parameter γ to be person-specific right now since variances cannot be modelled as a random effect in a LMM and it is not possible to transform the parameter out of the variance part of the model into the mean level.

To conclude these observations on the HOU model, we remark that substantive covariates x_{ps} and z_{ps} can be included as well. Howevers in the first instance, we will try to examine the equivalence between the LMM from equation (1) and the HOU from equations (3) and (4), without adding additional covariates in the HOU model.

3. Equivalences between the LMM and HOU for equally spaced data

3.1. No individual differences

To give a general idea of our working method, we start out from a very simple regression model with equally spaced observations and no individual differences (i.e., no random effects). Also, let us assume for now that there is no observation error in the HOU (so that $Y_{ps} = \theta_{ps}$). The equal time differences will be denoted by Δ. We can then formulate the following simplified version of the HOU model:

$$Y_{ps} = \mu + e^{-\lambda\Delta}(Y_{p,s-1} - \mu) + \eta_{ps}. \tag{8}$$

Now, we rearrange its terms and equate it to a simple fixed effects model (hence with no random effects) in which we use the previous observation as a predictor:

$$\begin{aligned} Y_{ps} &= \mu(1 - e^{-\lambda\Delta}) + e^{-\lambda\Delta} Y_{p,s-1} + \eta_{ps} \\ &= \beta_0^* + \beta_1^* Y_{p,s-1} + \epsilon_{ps}^*, \end{aligned}$$

where the distribution of ϵ_{ps}^* is similar to the already introduced formulation in equation (1) (hence, we do not yet consider a split in two residual terms as in equation (2)). The covariate x_{ps} from equation (1) equals $Y_{p,s-1}$, that is, the previous measurement for person p. An easily solvable system of nonlinear equations can then be derived:

$$\beta_0^* = \mu(1 - e^{-\lambda\Delta}), \qquad \beta_1^* = e^{-\lambda\Delta}, \qquad \sigma_{\epsilon^*}^2 = \gamma(1 - e^{-2\lambda\Delta}).$$

The solution of this system shows how to transform the parameters of the regression model into the parameters of the HOU:

$$\mu = \frac{\beta_0^*}{1 - \beta_1^*}, \tag{9}$$

$$\lambda = -\frac{1}{\Delta}\log \beta_1^*, \tag{10}$$

$$\gamma = \frac{\sigma_{\epsilon^*}^2}{1-\beta_1^{*2}}. \tag{11}$$

As we can see, the HOU parameters can be recovered exactly based on the regression estimates, apart from the positivity restriction on λ in the HOU model. In an application, one replaces the parameters by their estimates and the standard errors of the transformed parameters can be approximated using the delta method (see Agresti, 2002).

Often, there is good reason to add (possibly time-varying) covariates. If we include such a covariate in the HOU model, denoted as x_{ps}, we obtain

$$\begin{aligned} Y_{ps} &= (\mu + \omega x_{ps}) + e^{-\lambda\Delta}(Y_{p,s-1} - \mu - \omega x_{ps}) + \eta_{ps} \\ &= \mu(1 - e^{-\lambda\Delta}) + e^{-\lambda\Delta} Y_{p,s-1} + \omega(1 - e^{-\lambda\Delta})x_{ps} + \eta_{ps} \\ &= \beta_0^* + \beta_1^* Y_{p,s-1} + \beta_2^* x_{ps} + \epsilon_{ps}^*. \end{aligned}$$

Adding the covariate evidently only affects the mean structure. Then we can write out the equation system in the same manner as above. After solving it, we can see that nothing changes for the parameters already introduced (see equations (9)–(11)) and the regression coefficient of x_{ps} in the HOU can be obtained by

$$\omega = \frac{\beta_2^*}{1-\beta_1^*}.$$

The linear regression model thus provides a straightforward way to estimate the parameters of the HOU model and the two models can again be considered equivalent. Because adding covariates does not pose additional problems, we will not consider the issue further in order not to complicate matters unnecessarily.

We note here that in some cases, it is also customary to centre a predictor in a LMM. If in the LMM one centres the predictor (i.e., grand mean centring: see Raudenbush & Bryk, 2002), then the regression model has the following form:

$$Y_{ps} = \beta_0^* + \beta_1^*(Y_{p,s-1} - \bar{Y}) + \epsilon_{ps}^*.$$

The centred LMM can be obtained by approximating the second μ parameter in equation (8) by the sample mean $\hat{Y}$, and the corresponding HOU model has an easily interpretable equivalence, namely:

$$\begin{aligned} Y_{ps} &\approx \mu + e^{-\lambda\Delta}(Y_{p,s-1} - \bar{Y}) + \eta_{ps} \\ &= \beta_0^* + \beta_1^*(Y_{p,s-1} - \bar{Y}) + \epsilon_{ps}^*. \end{aligned}$$

After grand mean centring, the LMM intercept β_0 can be interpreted as μ, the mean level of the HOU model.

3.2. Individual differences in the mean level

Although we have been calling the models HOU and LMM, there has been no hierarchical or random part in the equations up to this point, since we have only looked at examples without differences between the units. Let us now add individual differences in the HOU model. First, let us assume that the μ parameter of the HOU model can differ among individuals, such that we can write that $\mu_p = \mu + m_p$. At this point, we do not specify a distribution for the person-specific parameter m_p yet. Instead, we will approximate the model with the LMM in which all random effects have a normal distribution and then deduce what this implies for the HOU parameters. We write out the HOU model, rearrange the terms, and equate it with a corresponding LMM (i.e., a random intercept model) the following way:

$$\begin{aligned}
Y_{ps} &= (\mu + m_p) + \mathrm{e}^{-\lambda\Delta}(Y_{p,s-1} - \mu - m_p) + \eta_{ps} \\
&= \mu(1 - \mathrm{e}^{-\lambda\Delta}) + m_p(1 - \mathrm{e}^{-\lambda\Delta}) + \mathrm{e}^{-\lambda\Delta} Y_{p,s-1} + \eta_{ps} \\
&= \beta_0^* + b_{p1}^* + \beta_1^* Y_{p,s-1} + \epsilon_{ps}^*,
\end{aligned}$$

with $b_{p1}^* \sim N(0, d_{11}^*)$, as is standard for a LMM.

Again, we can derive a system of equations, solve it and find that there is an exact equivalence between the two models. For the fixed effect parameters λ, μ, and γ, the same formulas as in equations (9)–(11) are valid. The random effect m_p is

$$m_p = \frac{b_{p1}^*}{1 - \beta_1^*}.$$

Since m_p is the rescaled version of the normally distributed random effect b_{p1}, the distribution of m_p is also normal with mean zero and variance $d_{11}^*/(1 - \beta_1^*)^2$. This makes perfect sense as a distribution on m_p.

3.3. Individual differences in the mean level and in the autocorrelation

Next, we make the autocorrelation parameter λ person-specific as follows: $\lambda_p = \lambda + a_p$. Because λ is restricted to be larger than zero, this assumption is not the most suitable one, but it is pragmatically motivated since it simplifies the derivations considerably. Unfortunately, the LMM is now no longer exactly equal to the HOU. But we can see that the equivalence approximately holds between the two models (by using a first-order Taylor series expansion):

$$\begin{aligned}
Y_{ps} &= (\mu + m_p) + \mathrm{e}^{-(\lambda + a_p)\Delta}(Y_{p,s-1} - \mu - m_p) + \eta_{ps} \qquad (12)\\
&\approx \mu(1 - \mathrm{e}^{-\lambda\Delta}) + m_p(1 - \mathrm{e}^{-\lambda\Delta}) + a_p\Delta\mathrm{e}^{-\lambda\Delta}(m_p + \mu) \\
&\quad + \mathrm{e}^{-\lambda\Delta} Y_{p,s-1} - a_p \mathrm{e}^{-\lambda\Delta}\Delta Y_{p,s-1} + \eta_{ps}
\end{aligned}$$

$$= \beta_0^* + b_{p1}^* + \beta_1^* Y_{p,s-1} + b_{p2}^* Y_{p,s-1} + \epsilon_{ps}^*, \qquad (13)$$

where we have made use of the following approximation (based on the first-order Taylor series expansion $\mathrm{e}^{-x} \approx 1 - x$): $\mathrm{e}^{-(\lambda + a_p)\Delta} = \mathrm{e}^{-\lambda\Delta}\mathrm{e}^{-a_p\Delta} \approx \mathrm{e}^{-\lambda\Delta}(1 - a_p\Delta)$. This is justified for $a_p\Delta$ small enough, say $a_p\Delta < 0.5$. The magnitude of the product will depend on the size of a_p and also on the time difference Δ. For example, if we take a time difference of 2 hours, due to the characteristics of the exponential autocorrelation function, the autocorrelation values of 0.9 and 0.1 correspond to λ values $\sim$0.05 and $\sim$1.15, respectively. Then for a_p values smaller than 0.25, our method provides a good approximation. For values outside this range, the LMM approach will lead to an underestimation of the HOU parameters (because the linear approximation is smaller than the exponential target). As we will see, this might prove to be a serious disadvantage when there is substantial inter-individual variation in the autocorrelation parameter.

With respect to the distribution of the random effects, in the LMM framework it is assumed that the random intercept and slope are jointly normally distributed:

$$\begin{pmatrix} b_{p1}^* \\ b_{p2}^* \end{pmatrix} \sim N\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} d_{11}^* & d_{12}^* \\ d_{12}^* & d_{22}^* \end{pmatrix}\right].$$

As before, λ and μ can be equated to the LMM parameters as is specified in equations (9) and (10). For the random effect a_p, we can see that it is equal to the following function of LMM parameters and random effects:

$$a_p = -\frac{b_{p2}^*}{\beta_1^*\Delta}.$$

Since this is a simple rescaling of the normally distributed random effect b_{p2}, the distribution of a_p is also normal:

$$a_p \sim N\left(0, \frac{d_{22}^*}{\left(\beta_1^*\right)^2\Delta^2}\right). \tag{14}$$

For the other random effect m_p, the situation is somewhat more complicated:

$$m_p = \frac{b_{p1}^* + \frac{\beta_0^*}{1-\beta_1^*} b_{p2}^*}{\left(1-\beta_1^*\right) - b_{p2}^*}. \tag{15}$$

This shows that m_p is expressed in terms of a ratio of two random variables, so that the distribution of m_p is no longer exactly normal. However, we can approximate the moments of m_p and can simulate its distribution; see the corresponding derivations and simulations in Appendix B.

Finally, we encounter a problem with the variance of η_{ps} for person p, since it depends on a_p (see equation (5) and assume λ in that equation is person-specific). In the LMM framework, such person variability in the error term cannot be captured for the corresponding parameter $\text{Var}(\epsilon_{ps}^*)$, so we need to find out what $\text{Var}(\epsilon_{ps}^*)$ represents approximately. It turns out that

$$\gamma = \frac{\sigma_\epsilon^{2*}}{\left[1 - \exp\left(-2\Delta - \frac{1}{\Delta}\log\beta_1^*\right)\right]\exp\left(2\Delta^2\frac{d_{22}^*}{\beta_1^{*2}}\right)}.$$

The corresponding derivations can be found in Appendix C.

To summarize, we have used the foregoing derivation to demonstrate the relation between the LMM and the HOU: first we linearized the random effect on λ, next we approximated the distribution of m_p, and finally we set the person-specific variability of the error term equal to its expectation. Because of these approximations, it is good to check the overall equivalence of both models. For this purpose, we generated data according to the HOU, fitted the LMM, and used the conversion formulas to transform the results back to the HOU parameter scales.

We simulated data under the HOU model for 50 individuals, each having a time series of 50 equally spaced measurements (Δ was arbitrarily set equal to 0.3, but any other value could have been chosen). In the simulations, the mean level and the autocorrelation were allowed to vary randomly across persons as follows: $\mu_p = \mu + m_p$ and $\lambda_p = \lambda + a_p$, where we assumed normal distributions for the random effects (see Table 1 for the precise values). The parameter estimates were found based on an LMM with random intercept and slope (estimation was done with the SAS procedure PROC MIXED; see the program script in Appendix D). Having obtained the LMM estimates, they were converted into the HOU estimates based on the equations presented above. Table 1 shows the averaged recovered estimates of 10 replications. It can be seen that the HOU parameters were recovered reasonably well. However, as the value Var(a_p) (i.e., the variance of the random effect on λ) increases, the approximation for this parameter is negatively biased, due to the linear approximation for $e^{a_p\Delta}$, as described above.

Table 1. Simulation results with equally spaced data, random effect on μ and λ

	$\mu^O(\hat{\mu}^L)$	$\mathrm{Var}(m_p)^O$ $(\widehat{\mathrm{Var}(m_p)}^L)$	$\lambda^O(\hat{\lambda}^L)$	$\mathrm{Var}(a_p)^O$ $(\widehat{\mathrm{Var}(a_p)}^L)$	γ^O $(\hat{\gamma}^L)$
Simulation 1	9.00 (8.99)	1.00 (1.06)	1.00 (0.99)	0.10 (0.00)	2.00 (1.98)
Simulation 2	6.00 (5.96)	2.00 (1.74)	2.00 (2.03)	0.50 (0.51)	2.00 (1.99)
Simulation 3	6.00 (5.97)	1.00 (0.91)	2.00 (1.95)	1.00 (0.63)	4.00 (4.03)
Simulation 4	9.00 (8.92)	2.00 (1.97)	2.00 (2.11)	2.00 (1.07)	4.00 (4.01)

Note. The first values in each column are the true values (denoted by a superscript O in the column heading) under the HOU model and the values in parentheses are the averaged estimates based on the transformed LMM parameter estimates (denoted by a superscript L in the column heading).

As a conclusion in certain cases, use of the LMM as a way of roughly estimating the HOU parameters seems to be justified. On a more fundamental level, the two models can be considered equivalent when there is no random effect or only a single random effect on the μ (mean level) parameter. The two models are nearly equivalent when there are random effects on both μ (mean level) and λ (autocorrelation). When the variability in a_p is not too large, the LMM can be used as a proxy for the HOU.

Two remarks should be made with regard to our analysis. First, we have not discussed the case of only random effects in the autocorrelation separately, but from the analysis above, it can be deduced that this is another case in which the relation is only rough. Second, we did not turn the third HOU parameter, γ (i.e., the intra-individual variation), into a random effect. The reason is that it is impossible to estimate the population variation regarding this aspect of the model with the LMM because it will only affect the variation of σ_ϵ^{2*} which is assumed to be constant. It shows, however, that the LMM is less well equipped to handle sources of person-specific variability.

In the next section, we will allow for unequal time differences between the measurements and additional measurement error.

4. Equivalences between the LMM and HOU for unequally spaced data

4.1. Individual differences in the mean level

If the measurements are not sampled at equal time intervals, using the previous position as a predictor to account for the autocorrelation structure of the data can lead to inaccurate results. As a demonstration, a simulation study similar to the previous one was set up, the only difference being that now the simulated data were not equally spaced, but the time difference varied uniformly between 0.01 and 1 h (the computer script used was similar to that in Appendix D). To be able to use the conversion formulas and express the estimates in terms of OU parameters, we used the average time difference for Δ, namely 0.5 h. Table 2 displays the results. We can see that the autocorrelation parameter and its random effect were both heavily underestimated.

Table 2. Simulation results with unequally spaced data, random effect on μ and λ, and the estimates are calculated by using the average time difference

	μ^{O} ($\hat{\mu}^{L}$)	$\mathrm{Var}(m_p)^{O}$ ($\widehat{\mathrm{Var}(m_p)})^{L}$	λ^{O} ($\hat{\lambda}^{L}$)	$\mathrm{Var}(a_p)^{O}$ ($\widehat{\mathrm{Var}(a_p)})^{L}$	γ^{O} ($\hat{\gamma}^{L}$)
Simulation 1	9.00 (9.07)	1.00 (0.92)	1.00 (0.91)	0.10 (0.06)	2.00 (2.01)
Simulation 2	6.00 (5.98)	2.00 (1.97)	2.00 (1.66)	0.50 (0.33)	2.00 (1.99)
Simulation 3	6.00 (5.94)	1.00 (1.20)	2.00 (1.70)	1.00 (0.50)	4.00 (3.98)
Simulation 4	9.00 (9.08)	2.00 (2.21)	2.00 (1.64)	2.00 (0.56)	4.00 (4.01)

Note. The first values in each column are the true values (denoted by a superscript O in the column heading) under the HOU model and the values in parentheses are the averaged estimates based on the transformed LMM parameter estimates (denoted by a superscript L in the column heading).

However, the LMM framework is able to include serial correlation even when the measurements are unequally spaced, by means of an adequately defined residual term (see equation (2)). An additional advantage of such a model is that we can separate the variation due to the measurement error from the variation due to the serial correlation component.

Let us start from the HOU model as specified in equation (7) (with person-specific mean level, but not yet a person-specific autocorrelation):[1]

$$Y_{ps} = \mu + m_p + \epsilon_{ps} + \tilde{\theta}_{ps}, \tag{16}$$

where $\tilde{\theta}_{ps}$ is a zero-mean OU process and ϵ_{ps} is measurement error. As derived in Appendix A, the autocovariance function of this model is exponential: $\gamma e^{-\lambda\Delta_{ps}}$. On the other hand, the LMM with two residual terms is defined as

$$Y_{ps} = \beta^* + b^*_{p1} + \epsilon^*_{(1)ps} + \epsilon^*_{(2)ps}, \tag{17}$$

with autocovariance function $\tau^{*2}e^{-\phi^*\Delta_{ps}}$. It is not difficult to see that the models in equations (16) and (17) can be equated term by term.

To verify our claim that both models are equivalent, we ran simulations to check how well the LMM model can recover the parameters from the HOU model that was used to simulate the data. The size of the data set was set to 50 individuals with 50 observations each for every simulation and with no missing data. We set the time scale to hours and the time difference uniformly distributed between 0.01 and 1 h, thus mimicking a typical experience sampling design. Again, SAS PROC MIXED was used to estimate the LMM parameters; see the fitted computer script in Appendix E. Different parameter values were chosen and 10 simulations were run with each setting. Table 3 displays the simulated values, with the averaged recovered estimates of the 10 simulations in parentheses. In the four settings, the level of the measurement error was increased step by step, keeping the rest of the parameters fixed. As we can see, even with a rather increased noise level, the parameters were recovered sufficiently. Our final conclusion is that the LMM approach was able to estimate the parameters of an HOU with a fixed autocorrelation structure (across persons).

Table 3. Simulation results with unequally spaced data with random effects on the mean parameter (μ)

	μ^O ($\widehat{\beta^{*L}}$)	$\mathrm{Var}(m_p)^O$ ($\widehat{\mathrm{Var}(b^*_{p1})^L}$)	λ^O ($\widehat{1/\phi^{*L}}$)	γ^O ($\widehat{\tau^{*2L}}$)	σ^2_ϵ ($\widehat{\sigma^{2*}_{\epsilon(1)}}^L$)
Simulation 1	9.00 (9.08)	1.00 (1.03)	1.00 (1.05)	2.00 (1.94)	0.10 (0.09)
Simulation 2	6.00 (5.99)	2.00 (2.03)	2.00 (1.98)	2.00 (2.02)	0.50 (0.49)
Simulation 3	6.00 (6.05)	1.00 (0.99)	2.00 (2.06)	4.00 (4.01)	1.00 (0.97)
Simulation 4	9.00 (9.04)	2.00 (2.03)	2.00 (2.07)	4.00 (4.17)	2.00 (1.87)

Note. The first values in each column are the true values (denoted by a superscript O in the column heading) under the HOU model and the values in parentheses are the averaged LMM parameter estimates (denoted by a superscript L in the column heading).

4.2. Individual differences in the mean level and in the autocorrelation

By analogy with the previous section where the time differences were equal, we may wonder what happens with unequal time differences and a model with both random mean level and autocorrelation. The problem is that the LMM cannot incorporate both unequal time differences and random autocorrelation and thus we may ask how the

[1]Note that we do not include other predictors because they do not pose particular difficulties with respect to the equivalence of both models.

parameter estimates are affected. We have run some simulations to investigate this problem. The simulations followed the same settings as before: the time difference was uniformly distributed between 0.01 and 1 h, SAS PROC MIXED was used to estimate the LMM parameters (the computer script in Appendix E was used again) and 10 simulations were run for each condition. However, now the only difference among the conditions was in the value of $\mathrm{Var}(a_p)$ (i.e., the variance of the random effect on λ). Thus the stochastic contribution had a person-specific component, a_p:

$$\tilde{\theta}_{ps} = e^{-(\lambda + a_p)\Delta_{ps}} \tilde{\theta}_{p,s-1} + \eta_{ps},$$

and Y_{ps} was defined as in equation (16). The results of the simulations are displayed in Table 4. We can observe that as the variance of the random effect of a_p increases, the estimates of λ become slightly worse. Also, the estimated level of the measurement error increases as we raise the value of $\mathrm{Var}(a_p)$, although we did not change this parameter over the simulations. We can thus conclude that having a random person-specific effect in the autoregressive structure adds an extra bias to the autoregression parameter estimates, but the LMM method still does well in recovering the true values for most of the other parameters. Thus the model misspecification only has a local effect since the parameters not related to the autocorrelation are hardly affected. However, the disadvantage of this approach is that it gives no idea of the presence and size of the inter-individual differences in the autocorrelation structure.

Table 4. Simulation results with unequally spaced data with random effects on the mean parameter (μ) and on the autoregression (λ) parameter

	μ^{O} ($\widehat{\beta^{*\mathrm{L}}}$)	$\mathrm{Var}(m_p)^{\mathrm{O}}$ ($\widehat{\mathrm{Var}(b^*_{p1})}^{\mathrm{L}}$)	λ^{O} ($\widehat{1/\phi^{*\mathrm{L}}}$)	γ^{O} ($\widehat{\tau^{*2\mathrm{L}}}$)	$\sigma_\epsilon^{2\mathrm{O}}$ ($\widehat{\sigma^{2*}_{\epsilon(1)}}^{\mathrm{L}}$)	$\mathrm{Var}(a_p)^{\mathrm{O}}$
Simulation 1	6.00 (6.00)	0.50 (0.53)	2.00 (2.05)	4.00 (3.93)	0.10 (0.10)	0.01
Simulation 2	6.00 (5.98)	0.50 (0.49)	2.00 (1.79)	4.00 (4.02)	0.10 (0.11)	1.00
Simulation 3	6.00 (5.98)	0.50 (0.60)	2.00 (1.84)	4.00 (3.79)	0.10 (0.15)	2.00
Simulation 4	6.00 (6.04)	0.50 (0.61)	2.00 (2.12)	4.00 (3.80)	0.10 (0.18)	3.00

Note. The first values in each column are the true values (denoted by a superscript O in the column heading) under the HOU model and the values in parentheses are the averaged LMM parameter estimates (denoted by a superscript L in the column heading).

4.3. Individual differences in the mean level and in the intra-individual variance

Additionally, in applications such as experience sampling studies, by calculating sample-based intra-individual variance (which can be considered a rough estimate of the γ parameter of the OU process) for the different persons, people are shown to exhibit variety. Moreover, some studies have shown (e.g., Kuppens, Van Mechelen, Nezlek, Dossche, & Timmermans, 2007) that if the sample-based person-specific intra-individual variances are calculated and correlated to some covariates, interesting associations might be revealed. Therefore, when the focus of the research is on inter-individual differences, we can take advantage of the OU approach in which the intra-individual variance parameter (γ) can be turned into a random effect (such as $\gamma_p = \gamma + g_p$, where $g_p \sim N(0, \sigma_g^2)$). Such an extension is not customary in the LMM framework: the corresponding LMM parameter τ^{*2} is specified through the elements of the serial covariance matrix and τ^{*2} is typically regarded as a population phenomenon.

To get some insight into the effect of inter-individual variation in this parameter, we again ran some simulations with the usual basic settings. SAS PROC MIXED was used again to estimate the parameters (the script used can be found in Appendix E), and the results are displayed in Table 5. We found that even if there is some difference among individuals with respect to their intra-individual variation (Var(g_p) ranged from 0.01 to 6), the LMM recovers all parameters sufficiently, and with respect to γ it recovers the population mean. Again, the only disadvantage here remains that in the LMM approach we neglect possible individual differences in the intra-individual variation parameter.

Table 5. Simulation results with unequally spaced data with random effects on the mean parameter (μ) and on the intra-individual variance (γ) parameter

	μ^{O} ($\widehat{\beta^{*\mathrm{L}}}$)	Var(m_p)$^{\mathrm{O}}$ ($\widehat{\mathrm{Var}(b^*_{p1})^{\mathrm{L}}}$)	λ^{O} ($\widehat{1/\phi^{*\mathrm{L}}}$)	γ^{O} ($\widehat{\tau^{*2\mathrm{L}}}$)	σ^2_ϵ ($\widehat{\sigma^{2*}_{\epsilon(1)}}^{\mathrm{L}}$)	Var(g_p)$^{\mathrm{O}}$
Simulation 1	6.00 (6.01)	0.50 (0.54)	2.00 (2.02)	4.00 (4.08)	0.10 (0.07)	0.01
Simulation 2	6.00 (5.97)	0.50 (0.57)	2.00 (1.96)	4.00 (3.92)	0.10 (0.13)	2.00
Simulation 3	6.00 (6.02)	0.50 (0.48)	2.00 (2.01)	4.00 (3.84)	0.10 (0.09)	4.00
Simulation 4	6.00 (6.02)	0.50 (0.48)	2.00 (1.95)	4.00 (4.09)	0.10 (0.10)	6.00

Note. The first values in each column are the true values (denoted by a superscript O in the column heading) under the HOU and the values in parentheses are those based on the LMM parameter estimates (denoted by a superscript L in the column heading).

4.4. Individual differences in the mean level, autocorrelation, and intra-individual variance

Finally, we can let all HOU parameters be random at the same time. Table 6 displays the results of simulations with random effects on the mean parameter (μ), on the autoregression (λ) parameter, and on the intra-individual variance (γ) parameter. The settings of the simulation study are similar to those of the studies reported above and SAS PROC MIXED was used to recover the parameters (see the computer script in Appendix E). As we can see, when the two random effects are not modelled by the LMM (a_p and g_p), it is mainly the variability of a_p that influences the recovery of the autoregression parameter λ. The other parameters are to a large extent not affected and the manipulation of Var(g_p) does not have a large influence (which was already the case when g_p was the only unmodelled random effect).

Table 6. Simulation results with unequally spaced data with random effects on the mean parameter (μ), on the autoregression (λ) parameter, and on the intra-individual variance (γ) parameter

	μ^{O} ($\widehat{\beta^{*\mathrm{L}}}$)	Var(m_p)$^{\mathrm{O}}$ ($\widehat{\mathrm{Var}(b^*_{p1})^{\mathrm{L}}}$)	λ^{O} ($\widehat{1/\phi^{*\mathrm{L}}}$)	γ^{O} ($\widehat{\tau^{*2\mathrm{L}}}$)	σ^2_ϵ ($\widehat{\sigma^{2*}_{\epsilon(1)}}^{\mathrm{L}}$)	Var(a_p)$^{\mathrm{O}}$	Var(g_p)$^{\mathrm{O}}$
Simulation 1	6.00 (5.97)	0.50 (0.49)	2.00 (2.05)	4.00 (3.95)	0.10 (0.09)	0.01	0.01
Simulation 2	6.00 (5.99)	0.50 (0.53)	2.00 (1.81)	4.00 (3.65)	0.10 (0.11)	1.00	2.00
Simulation 3	6.00 (5.96)	0.50 (0.51)	2.00 (1.62)	4.00 (3.80)	0.10 (0.17)	2.00	4.00
Simulation 4	6.00 (5.98)	0.50 (0.52)	2.00 (1.56)	4.00 (3.78)	0.10 (0.16)	3.00	6.00

Note. The first values in each column are the true values (denoted by a superscript O in the column heading) under the HOU model and the values in parentheses are the averaged LMM parameter estimates (denoted by a superscript L in the column heading).

5. Illustration of the relationship between the LMM and the HOU model in an experience sampling study

To illustrate the findings and the types of research questions which these models can address, we present an application of some of the models discussed to an experience sampling study. In the study, students carried a palmtop and were asked 50 times (at semi-random time points) during a single day to report their core affect. Core affect consists of a hedonic (pleasure–displeasure) and an arousal (deactivated–activated) dimension and is an important theoretical concept in the psychology of emotion (Barrett, Mesquita, Ochsner, & Gross, 2007; Russell, 2003). To measure it, a person has to report their current core affect by locating it in a two-dimensional space defined by the axes pleasure and activation (the resulting measurement was a score from 1 to 99 on each axis but then rescaled by dividing the score by 10). In this paper, we analyse the activation measurements taken from 60 students.

A major difference between the previously introduced simulation studies and this application is that here we do not know the true model behind the data. Thus we cannot observe the effect of allowing for random effects in such a systematic way. However, we can still fit several possible models to the data set and compare the results. Given that the data have unequally spaced measurements, models taking this aspect into account are to be preferred. However, we will also discuss models that are based on equal time intervals for the sake of comparison. The unequally spaced time points were determined in a semi-random way: the gap between consecutive measurements is between 0.90 and 401.15 min. On average, 17.27 min elapsed between consecutive time points and the standard deviation was 11.90.

In total, eight models were fitted to the data and Table 7 shows the results. The upper four rows refer to the models fitted to the unequally spaced data, and for the lower four rows the time differences are assumed to be constant (and set equal to the average time difference of 0.28 h). The LMM parameters are estimated using SAS PROC MIXED (for the first model the corresponding script can be found in Appendix E, and for the second in Appendix D), which returns the maximum likelihood estimates. The different versions of the HOU model are fitted in WinBUGS; the corresponding scripts can be found in Appendix F (altogether five of them, following the order in which they appear in Table 7, and for the last (sixth) OU model for equal time differences the same code was used as for the one with unequal time differences (the third one with all parameters person-specific), to facilitate their comparison). In order to facilitate comparison with the maximum likelihood method, we report the expected *a posteriori* estimate (the mean of the posterior distribution) for each of the parameters of interest. In addition and also for reasons of comparison, all parameter estimates were converted to HOU parameters.

In the LMM case, the model fitted with unequal time differences cannot contain a random effect for the autocorrelation parameter (therefore the entry for $\widehat{\mathrm{Var}(a_p)}$ contains 'np', for 'not possible'). If it is possible to estimate the variance of the random effect, but it is not done, the cell entry becomes 'ne' for 'not estimated'. For both the equal and unequal time interval methods, three versions of the HOU are estimated. The first HOU model is presented in equation (16) and does not allow for individual differences either in the intra-individual variance or in the autocorrelation. In the second HOU model (see equation (5)), the autocorrelation may differ across individuals (but not the intra-individual variance). The final HOU model is unrestricted and allows for individual differences at all levels. The same type of OU models are fitted by using the average time

Table 7. Results from the experience sampling data set using LMM and HOU models

Model	Δ	$\hat{\mu}$	$\widehat{\mathrm{Var}(m_p)}$	$\hat{\gamma}$	$\widehat{\mathrm{Var}(g_p)}$	$\hat{\lambda}$	$\widehat{\mathrm{Var}(a_p)}$	$\widehat{\sigma^2_\epsilon}$
LMM	Uneq	4.07 (0.12)	0.37 (0.18)	3.53 (0.24)	np	0.90 (0.12)	np	0.98 (0.07)
HOU	Uneq	4.12 (0.12)	0.38 (0.20)	3.45 (0.25)	ne	0.87 (0.10)	ne	1.01 (0.07)
HOU	Uneq	3.96 (0.13)	0.30 (0.17)	5.12 (0.28)	ne	1.56 (0.24)	2.48 (0.76)	0.23 (0.03)
HOU	Uneq	4.06 (0.13)	0.57 (0.19)	4.46 (0.42)	6.56 (3.64)	1.87 (0.25)	2.28 (0.75)	0.20 (0.03)
LMM	Eq	4.15 (0.13)	0.63 (0.14)	4.41 (0.31)	np	1.99 (0.23)	1.32 (0.38)	np
HOU	Eq	4.15 (0.12)	0.66 (0.19)	4.20 (0.17)	ne	1.98 (0.10)	ne	ne
HOU	Eq	3.95 (0.12)	0.28 (0.16)	5.30 (0.27)	ne	1.55 (0.19)	1.59 (0.45)	ne
HOU	Eq	4.06 (0.13)	0.55 (0.20)	4.46 (0.41)	6.16 (3.58)	1.65 (0.21)	1.50 (0.47)	0.17 (0.03)

Note. The column headings from the third column onwards refer to the estimated parameters (all LMM estimates are converted to HOU parameters). An 'np' entry means it is 'not possible' to estimate the given variance in this model and 'ne' denotes 'not estimated' (because the parameter is set to some value). If a parameter is followed by a variance that is estimated, the parameter is the mean of the random effects distribution; otherwise it is a fixed effect. The values in parentheses are the standard errors (for the LMM) or posterior *SD*s (for the HOU).

difference. In this way we consider the observations equally spaced. For this case, we can also estimate an LMM model which allows a random effect on the autoregression parameter. The second half of Table 7 displays the estimated results of these simplified models.

Although the statistical inference framework is different for the two models (frequentist vs. Bayesian), Table 7 indicates that the corresponding models give very similar estimates with respect to the mean level (i.e., $\hat{\mu}$ and its uncertainty). In general, we can see that the participating students on average do not feel very activated (since the neutral point is 5) and the variance of the intercept ($\hat{\sigma}^2_\mu$) indicates that there is some inter-individual variation with respect to the average activation level. Only the HOU with individual differences in the mean level and autocorrelation but not the variance has a slightly lower estimated μ (both for equal and unequal time differences), but it still is in the overall range of uncertainty. Further steps can be taken to model the mean structure more extensively (e.g., by allowing for μ to vary over time), but we will not do that here.

With respect to the other parameters and uncertainty estimates, the variation across models is much larger than for the mean and the result depends much more on the chosen model. Nevertheless, a number of general issues should be emphasized. First, when the inter-individual variability of a_p (i.e., $\mathrm{Var}(a_p)$) is estimated, it is quite large (around 2.5 for the unequal time difference models and around 1.5 for the equal time difference models). In addition, it should be noted that there is also quite some inter-individual variability in intra-individual variance (i.e., the estimate of $\mathrm{Var}(g_p)$ is around 5.9). Such large variances indicate that it is important to take these individual differences in variability parameters into account.

Second, the variability of the autocorrelation (i.e., variance of a_p) and the measurement error interact with each other. If the variance of a_p is not taken into account, the variability due to serial correlation shows up in the measurement error. This is especially notable in the unequal time difference models: without $\mathrm{Var}(a_p)$, the measurement error variance is about 1, but when estimating $\mathrm{Var}(a_p)$, the measurement error variance drops to about 0.2.

Third, the equal and unequal full HOU models give very similar results. This is due to the fact that for this data set there were very many measurements taken, such that the bulk of the inter-measurement times was concentrated around the mean. Taking the two previous issues together, we may distil the following advice: it is better to equalize the time differences than to ignore the inter-individual variability in autocorrelation.

6. Discussion

In this paper, our aim has been to provide a comparison of the traditional LMM approach and a more recent dynamical model for investigating ILD, the HOU model. In short, we can conclude that there are equivalences and approximate equivalences between both models. With equal time differences, only individual differences in the mean level and no measurement error, both models are exactly equal. The models are also equivalent under unequal time differences, only individual differences in the mean level but with measurement error. Altering any one of the conditions leads to a breakdown in equivalence. Of crucial importance is the size of inter-individual variation in the autoregression parameter (i.e., $\mathrm{Var}(a_p)$), because if it is large, the two models diverge with respect to the results.

In sum, there is a clear difference between the two model frameworks. The LMM approach focuses on the mean structure of the data and provides an efficient (and computationally fast) method for analysing it. However, if one is genuinely interested in the dynamic structure of the change over time and the inter-individual variation therein, the HOU is clearly more suitable. In such a case, the limitation of the LMM approach is apparent: it is only for equally spaced data that the autoregression parameter can be turned into a random effect, although in this case the residual variation cannot be split further into measurement error and intra-individual variation. Alternatively, applying the HOU approach offers a wider range of possibilities for modelling the dynamical properties. Random effects can be introduced with respect to the mean, autoregression, and intra-individual variation, and variation given by the measurement error can be separated from these factors. All in all, depending on the inclinations of the researcher, a model with the proper focus should be chosen.

We also note that a clear advantage of the LMM framework is the fast computation time. In the cases where equivalences and near equivalences were shown, this property of the LMM can be exploited for the faster estimation of the HOU model.

Another possible advantage of the HOU model concerns the case where more longitudinal variables are measured simultaneously. For instance, in the application section, it was outlined that at each time point, two measures were taken (pleasure and activation) and thus ideally a bivariate model would be used to analyse the data. A bivariate HOU model exists (Oravecz & Tuerlinckx, 2008; Oravecz *et al.*, 2009a) that is naturally equipped to properly model the possible cross-effects between the two variables, not only in terms of cross-correlation, but also with respect to their dependence on each other over time. Furthermore, the cross-effect parameter can also be turned into a random effect. In contrast, an LMM approach would concentrate on the mean structure in the bivariate case as well, and differences with respect to dynamics (between the two variables as well as between persons) and dynamical crossed effects could not be taken into account by applying this framework.

References

Agresti, A. (2002). *Categorical data analysis*. New York: Wiley.

Barrett, L. F., Mesquita, B., Ochsner, K. N., & Gross, J. J. (2007). The experience of emotion. *Annual Review of Psychology*, *58*, 373–403. doi:10.1146/annurev.psych.58.110405.085709

Blackwell, P. G. (2003). Bayesian inference for Markov processes with diffusion and discrete components. *Biometrika*, *90*, 613–627. doi:10.1093/biomet/90.3.613

Bolger, N., Davis, A., & Rafaeli, E. (2003). Diary methods: Capturing life as it is lived. *Annual Review of Psychology*, *54*, 579–616. doi:10.1146/annurev.psych.54.101601.145030

Bollen, K. A. (1989). *Structural equations with latent variables*. New York: Wiley.

Brillinger, D. R., Preisler, H. K., Ager, A. A., & Kie, J. G. (2004). An exploratory data analysis (EDA) of the paths of moving animals. *Journal of Statistical Planning and Inference*, *122*, 43–63. doi:10.1016/j.jspi.2003.06.016

Csikszentmihalyi, M., & Larson, R. (1987). Validity and reliability of the experience sampling method. *Journal of Nervous and Mental Disease*, *175*, 526–536.

De la Cruz-Mesía, R., & Marshall, G. (2003). A Bayesian approach for nonlinear regression models with continuous errors. *Communications in Statistics*, *32*, 1631–1646. doi:10.1081/STA-120022248

De la Cruz-Mesía, R., & Marshall, G. (2006). Non-linear random effects models with continuous time autoregressive errors: A Bayesian approach. *Statistics in Medicine*, *25*, 1471–1784. doi:10.1002/sim.2290

Diggle, P. J., Liang, K. Y., & Zeger, S. L. (1994). *Analysis of longitudinal data*. New York: Oxford University Press.

Golightly, A., & Wilkinson, D. J. (2006). Bayesian sequential inference for nonlinear multivariate diffusions. *Statistics and Computing*, *16*, 323–338. doi:10.1007/s11222-006-9392-x

Hamaker, E. L., Dolan, C. V., & Molenaar, P. C. M. (2003). ARMA-based SEM when the number of time points T exceeds the number of cases N: Raw data maximum likelihood. *Structural Equation Modeling*, *10*(3), 352–379. doi:10.1207/S15328007SEM1003_2

Jones, R. H. (1993). *Longitudinal data with serial correlation: A state-space approach*. New York: Chapman & Hall/CRC.

Jones, R. H., & Boadi-Boateng, F. (1991). Unequally spaced longitudinal data with AR(1) serial correlation. *Biometrics*, *47*, 161–175. doi:10.2307/2532504

Kuppens, P., Oravecz, Z., & Tuerlinckx, F. (2010). *Feelings change: Accounting for individual differences in the temporal dynamics of affect*. Manuscript submitted for publication.

Kuppens, P., Van Mechelen, I., Nezlek, J. B., Dossche, D., & Timmermans, T. (2007). Individual differences in core affect variability and their relationship to personality and adjustment. *Emotion*, *7*, 262–274. doi:10.1037/1528-3542.7.2.262

Laird, N. M., & Ware, J. H. (1982). Random-effects models for longitudinal data. *Biometrics*, *38*, 963–974. doi:10.2307/2529876

McArdle, J. J. (2009). Latent variable modeling of differences and changes with longitudinal data. *Annual Review of Psychology*, *60*, 577–605. doi:10.1146/annurev.psych.60.110707.163612

McCallum, R., & Ashby, F. G. (1986). Relationships between linear systems theory and covariance structure modeling. *Journal of Mathematical Psychology*, *30*, 1–27. doi:10.1016/0022-2496(86)90039-8

Mood, A. M., Graybill, F. A., & Boes, D. C. (1974). *Introduction to the theory of statistics*. New York: McGraw-Hill.

Oravecz, Z., & Tuerlinckx, F. (2008). *Bayesian estimation of an Ornstein–Uhlenbeck process based hierarchical model*. Proceedings of the Seventh International Conference on Social Science Methodology [CD–ROM], Naples: University of Naples.

Oravecz, Z., Tuerlinckx, F., & Vandekerckhove, J. (2009a). *A hierarchical latent stochastic differential equation model for affective dynamics*. Manuscript submitted for publication.

Oravecz, Z., Tuerlinckx, F., & Vandekerckhove, J. (2009b). A hierarchical Ornstein–Uhlenbeck model for continuous repeated measurement data. *Psychometrika*, *74*, 395–418. doi:10.1007/s11336-008-9106-8

Raudenbush, S. W., & Bryk, A. S. (2002). *Hierarchical linear models: Applications and data analysis methods*. Newbury Park, CA: Sage.

Robert, C. P., & Casella, G. (2004). *Monte Carlo statistical methods*. New York: Springer.

Russell, J. A. (2003). Core affect and the psychological construction of emotion. *Psychological Review*, *110*, 145–172. doi:10.1037/0033-295X.110.1.145

Russell, J. A., & Feldman-Barrett, L. (1999). Core affect, prototypical emotional episodes, and other things called emotion: Dissecting the elephant. *Journal of Personality and Social Psychology*, *76*, 805–819. doi:10.1037/0022-3514.76.5.805

Shumway, R. H., & Stoffer, D. S. (2006). *Time series analysis and its applications: With R examples* (2nd ed.). New York: Springer.

Skrondal, A., & Rabe-Hesketh, S. (2004). *Generalized latent variable modeling*. Boca Raton, FL: Chapman & Hall/CRC.

Verbeke, G., & Molenberghs, G. (1997). *Linear mixed models in practice*. New York: Springer.

Verbeke, G., & Molenberghs, G. (2000). *Linear mixed models for longitudinal data*. New York: Springer.

Walls, T. A., Jung, H., & Schwartz, J. E. (2006). Multilevel models for intensive longitudinal data. In T. A. Walls & J. L. Schafer (Eds.), *Models for intensive longitudinal data* (pp. 148–175). New York: Oxford University Press.

Received 14 July 2009; revised version received 22 January 2010

Appendix A

Derivation of the autocorrelation function of the OU process

Let us start from the autocovariance function and use the double expectation theorem:

$$
\begin{aligned}
\mathrm{E}[\tilde{\theta}_{ps}\tilde{\theta}_{p,s-1}] &= \mathrm{E}[\mathrm{E}(\tilde{\theta}_{ps}\tilde{\theta}_{p,s-1}|\tilde{\theta}_{p,s-1})] \\
&= \mathrm{E}[\tilde{\theta}_{p,s-1}\mathrm{E}(\tilde{\theta}_{p,s}|\tilde{\theta}_{p,s-1})] \\
&= \mathrm{E}\left[\tilde{\theta}^2_{p,s-1}\mathrm{e}^{-\lambda_p\Delta_{ps}}\right] \\
&= \mathrm{E}\left[\tilde{\theta}^2_{p,s-1}\right]\mathrm{e}^{-\lambda_p\Delta_{ps}} \\
&= \gamma\mathrm{e}^{-\lambda_p\Delta_{ps}}.
\end{aligned}
$$

Since $\mathrm{E}[\tilde{\theta}^2_{p,s-1}] = \mathrm{Var}[\tilde{\theta}_{ps}] = \mathrm{Var}[\tilde{\theta}_{p,s-1}] = \gamma$, the autocorrelation function $\rho_p(\Delta_{ps})$ is equal to $\mathrm{e}^{-\lambda_p\Delta_{ps}}$.

Appendix B

Approximation of the moments of m_p

We can approximate the moments of m_p based on the formulas in Mood, Graybill, and Boes (1974, p. 181) as follows:

$$
\begin{aligned}
\mathrm{E}(m_p) \approx\ & \frac{\mathrm{E}\left(b^*_{p1}+\frac{\beta^*_0}{1-\beta^*_1}b^*_{p2}\right)}{\mathrm{E}\left(\left(1-\beta^*_1\right)-b^*_{p2}\right)} - \frac{1}{\mathrm{E}\left(\left(1-\beta^*_1\right)-b^*_{p2}\right)^2}\mathrm{Cov}\left(b^*_{p1}+\frac{\beta^*_0}{1-\beta^*_1}b^*_{p2},\left(1-\beta^*_1\right)-b^*_{p2}\right) \\
& + \frac{\mathrm{E}\left(b^*_{p1}+\frac{\beta^*_0}{1-\beta^*_1}b^*_{p2}\right)}{\mathrm{E}\left(\left(1-\beta^*_1\right)-b^*_{p2}\right)^3}\mathrm{Var}\left(\left(1-\beta^*_1\right)-b^*_{p2}\right) \\
=\ & \frac{1}{\left(1-\beta^*_1\right)^2}\left(\mathrm{Cov}\left(-b^*_{p2},\frac{\beta^*_0}{1-\beta^*_1}b^*_{p2}\right)+\mathrm{Cov}\left(b^*_{p1},-b^*_{p2}\right)\right) \\
=\ & \frac{1}{\left(1-\beta^*_1\right)^2}\left(\frac{\beta^*_0}{1-\beta^*_1}d^*_{22}+d^*_{12}\right),
\end{aligned}
$$

$$
\mathrm{Var}(m_p) \approx \frac{\mathrm{Var}\left(b^*_{p1}+\frac{\beta^*_0}{1-\beta^*_1}b^*_{p2}\right)}{\left[\mathrm{E}\left(\left(1-\beta^*_1\right)-b^*_{p2}\right)^2\right]} = \frac{d^*_{11}+\left(\frac{\beta^*_0}{1-\beta^*_1}\right)^2 d^*_{22}+2\frac{\beta^*_0}{1-\beta^*_1}d^*_{12}}{\left(1-\beta^*_1\right)^2}.
$$

With respect to its distribution, it is hard to derive analytical results. For this reason (and also to get more insight into the formulas for the mean and the variance of m_p), we resort to simulations. Ideally, the random effect m_p has approximately a normal distribution with mean zero. In Figure 2, normalized histograms for simulations of m_p based on chosen values for the LMM parameters are shown; the black lines correspond to a normal distribution with the same mean and variance as the histograms. Among the plots, the covariance matrix of the random effects was changed (i.e., d^*_{11}, d^*_{22}, and d^*_{12}) while the other parameters were kept fixed, namely to values $\beta^*_0 = 3.5606$ (corresponding to $\mu = 6$ under the HOU model), $\beta^*_1 = 0.4066$ (corresponding to $\lambda = 3$ under the HOU model) and $\Delta = 0.3$. In the first two columns of the figure, the variance of the second random effect (for b^*_{p2}) was relatively small and then the normal approximation is almost perfect. The reason for this almost perfect approximation can be seen in equation (15): if the variance of b^*_{p2} goes to zero, the random effect disappears from the formula and m_p is a rescaled version of the normally distributed b^*_{p1}. Also, the covariation (or correlation) does not seem to influence the approximation. However, increasing d^*_{22} causes more substantial deviation from the reference normal distribution since in that case the histograms become somewhat skewed to the right. To explore the approximation in more detail, we simulated distributions of m_p based on large random effect variances, as shown in Figure 3. We can see that the deviation from the reference normal distribution is more substantial now, and some histograms show a bimodal pattern. However, despite the discrepancies, the mean of m_p is close to zero (such that μ is the average position in the population) and the variability is picked up quite well by the LMM parameters.

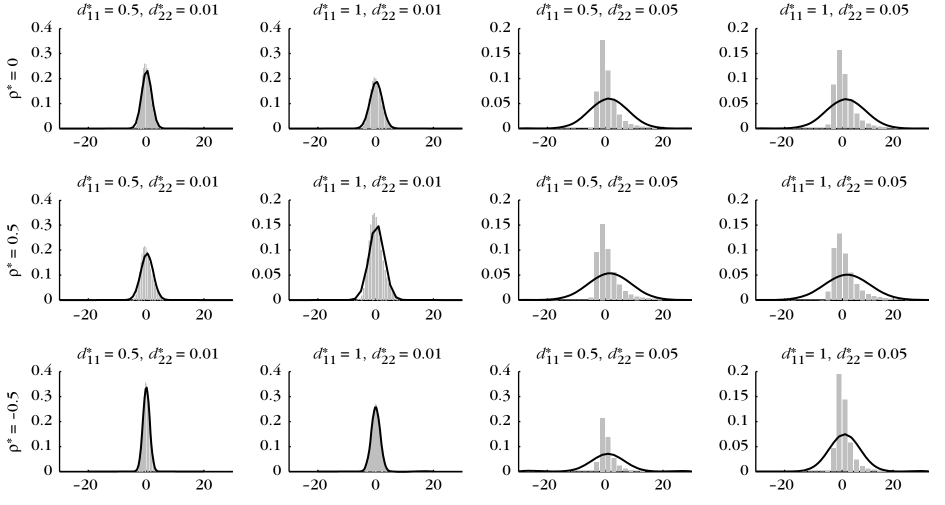

Figure 2. Simulated distributions of m_p. Each row has a different level of correlation, $\rho^* = d^*_{12}/\sqrt{d^*_{11}d^*_{22}}$, indicated on the for left of the row. The first two columns differ only with respect to the parameter value d^*_{11}, which is set to 0.5 in the first column and 1 in the second. The variance of the random effect on the autocorrelation (d^*_{22}) is 0.01 in the first two columns. In the third and fourth columns, this variance parameter is increased, namely $d^*_{22} = 0.05$, while d^*_{11} is set to 0.5 in the third column and 1 in the last.

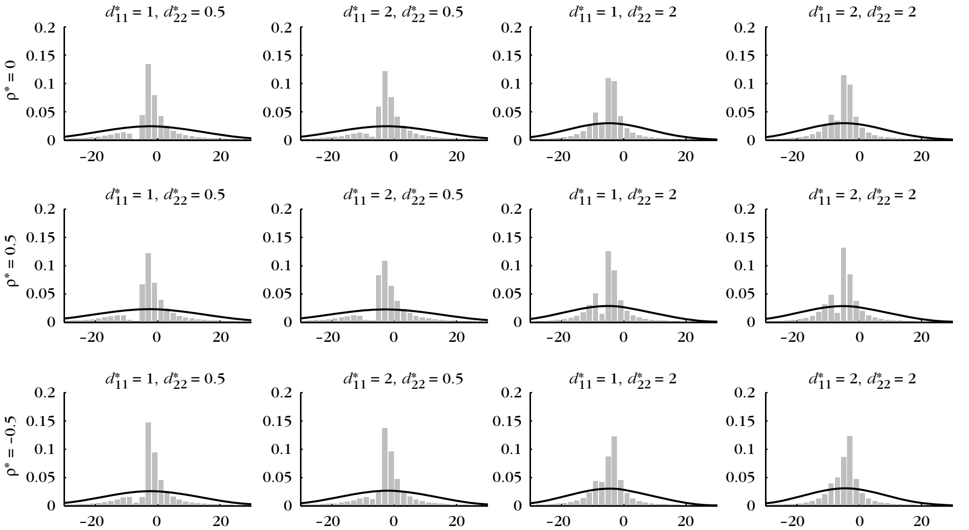

Figure 3. Simulated distributions of m_p. Each row has a different level of correlation, $\rho^* = d^*_{12}/\sqrt{d^*_{11}d^*_{22}}$, indicated on the for left of the row. The first two columns differ only with respect to the parameter value d^*_{11}, which is set to 1 in the first column and 2 in the second. The variance of the random effect on the autocorrelation (d^*_{22}) is 0.5 in the first two columns. In the third and fourth columns, this variance parameter is increased, namely $d^*_{22} = 2$, while d^*_{11} is set to 1 in the third column and 2 in the last.

Appendix C

Derivation of the equivalence between γ and the LMM parameters

If it is assumed that $\mathrm{Var}(\epsilon^*_{ps})$ approximates the unconditional variance of η_{ps}, that is, with the random effect a_p averaged out, we get the following result:

$$\begin{aligned}\mathrm{Var}\left(\epsilon^*_{ps}\right) &= \mathrm{Var}(\eta_{ps}) = \mathrm{Var}(\mathrm{E}[\eta_{ps}|a_p]) + \mathrm{E}[\mathrm{Var}(\eta_{ps}|a_p)]\\ &= 0 + \mathrm{E}[\gamma(1 - \mathrm{e}^{-2\Delta(\lambda+a_p)})]\\ &= \gamma - \gamma\mathrm{E}[\mathrm{e}^{-2\Delta(\lambda+a_p)}]\\ &= \gamma - \gamma\mathrm{e}^{-2\Delta\lambda}\mathrm{E}[\mathrm{e}^{-2\Delta a_p}].\end{aligned} \tag{A1}$$

To calculate $\mathrm{E}[\mathrm{e}^{-2\Delta a_p}]$, we note that a_p follows a normal distribution with mean zero and variance $\mathrm{Var}(a_p)$ equal to the expression given in equation (14). Therefore, $\mathrm{e}^{-2\Delta a_p}$ will be distributed log-normally with mean $\mathrm{e}^{2\Delta^2\mathrm{Var}(a_p)}$ and variance $\mathrm{e}^{8\Delta^2\mathrm{Var}(a_p)} - \mathrm{e}^{4\Delta^2\mathrm{Var}(a_p)}$ (see Mood *et al.*, 1974). Inserting this result into equation (A1) then gives

$$\sigma^{2*}_{\epsilon} = \gamma - \gamma\mathrm{e}^{-2\Delta\lambda}\mathrm{e}^{2\Delta^2\mathrm{Var}(a_p)},$$

and by using the expression for the variance of a_p (see equation (14)), we find that γ of the HOU process is defined as

$$\gamma = \frac{\sigma_\epsilon^{2*}}{\left[1 - \exp\left(-2\Delta - \frac{1}{\Delta}\log\beta_1^*\right)\right]\exp\left(2\Delta^2 \frac{d_{22}^*}{\beta_1^{*2}}\right)}.$$

Appendix D

LMM with random intercept and slope

In this LMM model, the previous measurement occasion (denoted by *yp*) is used as a predictor for the current one, personwise. We let this regression coefficient vary among persons. Moreover, a random intercept term is incorporated in the model. The covariance matrix is specified as unstructured.

```
PROC MIXED data = mylib.data;
class person;
model y = yp/solution;
random intercept yp/sub = person type = un;
run;
```

Appendix E

LMM with random intercept and split residual term (into serial correlation and measurement error)

In this LMM model, the residual variance structure is expressed in terms of exponential serial correlation and measurement error. The intercept is allowed to differ among persons. The covariance matrix is specified as unstructured.

```
PROC MIXED data = mylib.data;
class person beep_number;
model y = /solution;
random intercept/type = un subject = person;
repeated beep_number/type = sp(exp)(measurement_time) local
subject = person;
```

Appendix F

WinBUGS scripts for the different OU models

F.1. OU model for unequal time differences with only the mean level allowed to differ among persons

```
model{ # N = number of persons
for (p in 1:N){
# specifying the sampling distribution of the person-specific mean
# level
mu[p] ~ dnorm(mmu,p.mu)
# sampling the first latent state (Td = 0)
theta[p,1] ~ dnorm(mu[p],inv.gamma)
# sampling the rest of the latent states
for (i in 2:nrobs[p]){
instmean[p, i]
<- mu[p] + exp(−1*lambda*Td[p, i])*(theta[p, i − 1] − mu[p])
instprec[p, i] <-1/(gamma*(1 − exp(−2*lambda*Td[p, i])))
theta[p, i] ~ dnorm(instmean[p, i], instprec[p, i])
Y[p, i] ~ dnorm(theta[p, i], p.me) }}
# priors and parameter transformations
mmu ~ dnorm(0, 0.01)
sd.mu ~ dunif(0.01,10)
p.mu <- pow(sd.mu, − 2)
var.mu <- 1/p.mu
invgamma <- pow(sd.gamma, − 2)
sd.gamma ~ dunif(0.01,5)
gamma <- 1/invgamma
lambda ~ dnorm(0,0.01)I(0.001,)
sd.me ~ dunif(0.01,10)
p.me <- pow(sd.me, − 2)
var.me <- 1/p.me
```

F.2. OU model for unequal time differences with person-specific mean level and autocorrelation

Compared to the previous script, the only difference is that we allow for person-specific autocorrelation (denoted as 'lambda'), so only those parts of the previous code which contain lambda will change. The following modification of the previous code should be implemented:

```
# In the beginning, specify the sampling distribution of
#the person-specific autocorrelation parameter
lambda[p] ~ dnorm(mlambda,p.lambda)I(0.001,)
# Subsequently, insert index [p] after each lambda parameter
# In the section of priors and parameter transformations,
#delete the line containing lambda and insert:
mlambda ~ dnorm(0, 0.01)
sd.lambda ~ dunif(0.01,5)
p.lambda <- pow(sd.lambda, − 2)
var.lambda <- 1/p.lambda
```

F.3. OU model for unequal time differences with all parameters allowed to differ among persons
Compared to the previous script, we now allow the intra-individual variance parameter (denoted as 'gamma') also to differ among persons. As before, only the parts containing this parameter have to be changed:

```
# In the beginning, specify the sampling distribution (and some
# transformations)
# of the person-specific intra-individual variance parameter
loggamma[p] ~ dnorm(mloggamma,p.loggamma)
gamma[p] <- exp(loggamma[p])
inv.gamma[p] <- 1/gamma[p]
# Subsequently, insert index [p] after each gamma and inv.gamma
# parameters
# In the section of priors and parameter transformations,
# delete all three lines containing 'gamma' and insert:
mloggamma ~ dnorm(0,0.01)
sd.loggamma ~ dunif(0.01,5)
p.loggamma <- pow(sd.loggamma, − 2)
var.loggamma <- 1/p.loggamma
Emgamma <- exp(mloggamma + 0.5*var.loggamma)
Evargamma <-
exp(2*mloggamma + 2*var.loggamma) − exp(2*mloggamma + var.loggamma)
```

F.4. OU model for equal time differences with only the mean level allowed to differ among persons

```
model{
# N = number of persons
for (p in 1:N) {
# sampling the person-specific OU parameters
mu[p] ~ dnorm(mmu,p.mu)
Y[p,1] ~ dnorm(mu[p],invgamma)
for (i in 2:nrobs[p]) {
instmean[p, i] <- mu[p]+exp(-1*lambda*Td[p, i])*(Y[p, i - 1] - mu[p])
instprec[p, i] <- 1/(gamma*(1 - exp(-2*lambda * Td[p, i])))
Y[p, i] ~ dnorm(instmean[p, i], instprec[p, i]) }}
# priors and parameter transformations
mmu ~ dnorm(0, 0.01)
sd.mu ~ dunif(0.01,10)
p.mu <- pow(sd.mu, - 2)
var.mu <- 1/p.mu
invgamma <- pow(sd.gamma, - 2)
sd.gamma ~ dunif(0.01,5)
gamma <- 1/invgamma
lambda ~ dnorm(0,0.01)I(0.001,)
```

F.5. OU model for equal time differences with person-specific mean level and autocorrelation
To allow for person-specific autocorrelation ('lambda') for equal time differences, the same changes should be made in the previous code as in Section F.2.