

Multivariate inhomogeneous diffusion models with covariates and mixed effects

Mareile Große Ruse

University of Copenhagen, Copenhagen, Denmark

E-mail: mareile@math.ku.dk

Adeline Samson

Laboratoire Jean Kutzmann, Université Grenoble-Alpes, Grenoble, France

Susanne Ditlevsen

University of Copenhagen, Copenhagen, Denmark.

Summary. Modeling of longitudinal data often requires diffusion models that incorporate overall time-dependent, nonlinear dynamics of multiple components and provide sufficient flexibility for subject-specific modeling. This complexity challenges parameter inference and approximations are inevitable. We propose a method for approximate maximum-likelihood parameter estimation in multivariate time-inhomogeneous diffusions, where subject-specific flexibility is accounted for by incorporation of multidimensional mixed effects and covariates. We consider N multidimensional independent diffusions $X^i = (X_t^i)_{0 \leq t \leq T^i}$, $1 \leq i \leq N$, with common overall model structure and unknown fixed-effects parameter μ . Their dynamics differ by the subject-specific random effect ϕ^i in the drift and possibly by (known) covariate information, different initial conditions and observation times and duration. The distribution of ϕ^i is parametrized by an unknown ϑ and $\theta = (\mu, \vartheta)$ is the target of statistical inference. Its maximum likelihood estimator is derived from the continuous-time likelihood. We prove consistency and asymptotic normality of $\hat{\theta}_N$ when the number N of subjects goes to infinity using standard techniques and consider the more general concept of local asymptotic normality for less regular models. The bias induced by time-discretization of sufficient statistics is investigated. We discuss verification of conditions and investigate parameter estimation and hypothesis testing in simulations.

Keywords: Approximate maximum likelihood, asymptotic normality, consistency, covariates, LAN, mixed effects, non-homogeneous observations, random effects, stochastic differential equations

1. Introduction

Many physical and biological processes recorded over time exhibit time-varying and non-linear dynamics. There are two key demands for reliable statistical investigation of such data: First, the model should enable a sufficiently comprehensive description of the dynamics, which translates into a minimum requirement on the model complexity. Often, this is well captured by ordinary differential equations (ODEs) with a suitable degree of non-linearity and dimensionality. Second, the model should be sufficiently parsimonious to facilitate robust statistical estimation. This parsimony can to some degree be achieved by including mixed effects (Pinheiro and Bates, 2006; Davidian and Giltinan, 2003): While assuming an overall model structure for all experimental units, some parameters are allowed to vary across the population to capture also individual-specific characteristics. Hence, it

is not surprising that ODE models with mixed effects have become a popular tool for inference on longitudinal data, (Lindstrom and Bates, 1990; Wolfinger, 1993; Tornøe et al., 2004; Guedj et al., 2007; Wang, 2007; Ribba et al., 2014; Lavielle, 2014). This framework has, however, one important deficiency: The deterministic nature of ODE models does not capture uncertainties in the model structure and this can lead to biased estimates and false inference. For example, if the data have a periodic component with fluctuations in the phase, least squares estimation from a deterministic ODE will lead to predicted dynamics that are close to constant, and other estimation methods have to be sought for (Ditlevsen et al., 2005). This shortcoming can be addressed by replacing ODEs with stochastic differential equations (SDEs), thereby facilitating a more robust estimation (Donnet et al., 2010; Møller et al., 2010; Leander et al., 2014). Furthermore, SDE models with mixed effects use data more efficiently. Consider an SDE model based on one longitudinal measurement. It is well-known that for fixed time horizon the drift estimator is inconsistent (Kessler et al., 2012). In many applications one has observations of several experimental units at hand, but dynamics seem too subject-specific to assume that all individuals share the same parameter values. This prohibits the otherwise natural approach to reduce the bias by pooling the data. However, when dynamics are individual-specific but structurally similar, the unknown drift parameter can be modeled as a mixed effect, which may reduce bias considerably. In combination with mixed effects, nonlinear, time-inhomogeneous and multidimensional SDEs thus become a highly versatile tool for the intuitive and comprehensive modeling of complex longitudinal data, allowing for more robust statistical inference. This framework of stochastic differential mixed-effects models (SDMEM) lends itself to numerous applications and thereby opens up for new insights in various scientific areas.

Combining the benefits that are specific to mixed-effects and SDE models, however, entails particular challenges in terms of statistical inference. The key challenge lies in the intractability of the data likelihood, which now has two sources: The likelihood for (nonlinear) SDE models (given fixed parameter values) is analytically not available, rendering parameter inference for standard SDE models a nontrivial problem in itself. This intractable quantity has then to be integrated over the distribution of the random effects and one realizes that numerical or analytical approximations are inevitable. The likelihood in SDE models can be approximated in various ways. Given discrete-time observations, the likelihood is expressed in terms of the transition density. Approximation methods for the latter reach from solving the Fokker-Planck equation numerically (Lo, 1988), over standard first-order (Euler-Maruyama) or higher-order approximation schemes and simulation-based approaches (Pedersen, 1995; Durham and Gallant, 2002) to a closed-form approximation via Hermite polynomial expansion (Aït-Sahalia, 2002). If continuous-time observations are assumed (e.g., if high-frequency data is available), transition densities are not needed and the likelihood can be obtained from the Girsanov formula (Phillips and Yu, 2009). Popular analytical approximation techniques for nonlinear mixed-effects models are first-order conditional estimation (FOCE) (Beal and Sheiner, 1981) and Laplace approximation (Wolfinger, 1993). A computational alternative to analytical approximation is the expectation-maximization (EM) algorithm, or stochastic versions thereof (Delyon et al., 1999).

In the context of SDMEMs, the above mentioned approximation methods have been combined in various ways, depending on whether observations are modeled in discrete or continuous time and with or without measurement noise. Models for discrete-time data including measurement noise require marginalization of the likelihood over both the state and the random effects distribution. Approximation via FOCE in combination with the extended Kalman filter has been pursued by several authors (Tornøe et al., 2005; Overgaard et al., 2005; Mortensen et al., 2007; Klim et al., 2009; Leander et al., 2014, 2015). The most general

setting was considered by Leander et al. (2015), who allowed drift and diffusion function of the multivariate state SDE to depend on time, state and the individual parameters. However, theoretical convergence properties are not available in this setting. Donnet et al. (2008) approximate a one-dimensional SDE by the Euler-Maruyama scheme, and employ a stochastic approximation EM algorithm to avoid marginalization entirely. To circumvent the computationally expensive simulation of the SDE solution, Delattre and Lavielle (2013) consider only the random effects as latent and use a Metropolis-Hastings algorithm to simulate these conditioned on the observation. Finally, a Bayesian setting for a one-dimensional homogeneous diffusion is considered by Donnet et al. (2010). In models for discrete-time observations without observation noise, Ditlevsen and De Gaetano (2005), consider one-dimensional linear SDEs with linear mixed effects, where the likelihood is available in closed form. Picchini et al. (2010) and Picchini and Ditlevsen (2011) approximate the transition density by Hermite expansion and explore Gaussian quadrature algorithms and Laplace's approximation to compute the integral over the mixed effect. Mixed effects that enter the diffusion coefficient are investigated by Delattre et al. (2015). Continuous-time observations are the starting point of investigations in Delattre et al. (2013). They consider a univariate SDMEM without measurement noise for Gaussian mixed effects, which enter the drift linearly.

None of the previously mentioned works provide theoretical investigations of the estimators, when the state process is modeled by a multivariate, time-inhomogeneous and nonlinear SDE. However, many biological and physical processes are time-varying and require a certain degree of model complexity, such as non-linearity and multidimensional states. Furthermore, none of the cited works include covariates. Especially for practitioners, being used to regression analyses, not including covariate information in a model seems highly restrictive.

The purpose of this article is two-fold. On the one hand, we extend the setting in Delattre et al. (2013) to a multidimensional state process with non-linear, time-inhomogeneous dynamics. We obtain an integral expression for the likelihood. If the drift function for the SDE is linear in the random effect and if the random effects are independent and identically $\mathcal{N}(0, \Omega)$ -distributed with unknown covariance matrix Ω , the integral expression for the likelihood can be solved explicitly. If also the fixed effect enters the drift linearly, the likelihood turns into a neat expression, in which all remaining model complexity (multidimensionality of the state, nonlinearity, covariates) is conveniently hidden in the sufficient statistics. Under standard, but rather strict regularity conditions, we derive in subsection 3.1 the consistency and asymptotic normality of the MLE of the fixed effect and Ω . We also investigate the discretization error that arises when replacing the continuous-time statistics by their discrete-time versions. All results here can be shown using the same techniques as in Delattre et al. (2013), however, they are tedious to write down due to the more general and multidimensional setup. Therefore, these proofs are omitted here. Nevertheless, the approach has two main drawbacks. The first one is model-related: It is assumed that observations are identically distributed, not allowing for subject-specific covariates. The other is proof-related: The imposed regularity assumptions are rather restrictive, for instance, the density of the random effects may not be smooth. The second part of this work addresses these two issues. In subsection 3.2, we allow the inclusion of deterministic covariates. Moreover, we switch from the standard verification of asymptotic properties of the MLE to a more general strategy, which builds upon the concept of local asymptotic normality (LAN) of statistical experiments as introduced by Le Cam (2012), and further studied by Ibragimov and Has'minskii (2013). It allows statements on asymptotic normality of estimators even when the density functions have some degree of "roughness". The general conditions for consistency and asymptotic normality of the MLE that were formulated by Ibragimov

and Has'minskii (2013) are adapted to our setting. We then discuss intuitive conditions (on, for instance, the N -sample Fisher information), which are familiar from regression settings (such as $\frac{1}{N}I_N(\theta) \rightarrow I(\theta)$) and which help to verify the assumptions for MLE asymptotics. We point out the difficulties that arise with these conditions when observations are generated by not identically distributed SDMEMs, and propose a way to remedy these issues. Section 4 covers simulations for multivariate models that are linear in the fixed and random effects, the latter being Gaussian distributed. In this setting the likelihood is explicitly available. The first example is linear in state and covariates. More specifically, we consider a multidimensional Ornstein-Uhlenbeck model with one covariate having two levels (e.g., *treatment* and *placebo*). This type of model is common in, e.g., pharmacokinetics, and is motivated by a recent study on Selenium metabolism in humans (Große Ruse et al., 2015). Moreover, we perform hypothesis testing on the treatment effect and investigate the performance of the Wald test. The second example is the stochastic Fitzhugh-Nagumo model, used to model electrical activity in neurons, which, after parametrization, is still linear in the parameter, but non-linear in the state. We explore the quality of estimation for different sample sizes and sampling frequencies.

2. Preliminaries

This section introduces the general statistical setup and derives the likelihood function.

2.1. The setting

We consider N r -dimensional stochastic processes $X^i = (X_t^i)_{0 \leq t \leq T^i}$ whose dynamics are governed by the stochastic differential equations

$$dX_t^i = F(t, X_t^i, \mu, \phi^i)dt + \Sigma(t, X_t^i)dW_t^i, \quad 0 \leq t \leq T^i, \quad X_0^i = x_0^i, \quad i = 1, \dots, N. \quad (1)$$

The r -dimensional Wiener processes $W^i = (W_t^i)_{t \geq 0}$ and the d -dimensional random vectors ϕ^i are defined on some filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$, which is rich enough to ensure independence of all random objects $W^i, \phi^i, i = 1, \dots, N$. The d -dimensional vectors $\phi^i, i = 1, \dots, N$, are the so-called *random effects*. They are assumed to be \mathcal{F}_0 -measurable and have a common (usually centered) distribution which is specified by a (parametrized) Lebesgue density $g(\varphi; \vartheta)d\varphi$. The parameter $\vartheta \in \mathbb{R}^{q-p}$ is unknown as well as the *fixed effect* $\mu \in \mathbb{R}^p$. Together, these two quantities are gathered in the parameter $\theta = (\mu, \vartheta)$. This is the object of statistical inference and is assumed to lie in the parameter space Θ , which is a subset of \mathbb{R}^q . The functions $F : [0, T] \times \mathbb{R}^r \times \mathbb{R}^p \times \mathbb{R}^d \rightarrow \mathbb{R}^r, \Sigma : [0, T] \times \mathbb{R}^r \rightarrow \mathbb{R}^{r \times r}, T = \max_{1 \leq i \leq N} T^i$, are deterministic and known and the initial conditions x_0^i are independent and identically distributed (i.i.d.) r -dimensional random vectors. We assume that we observe X^i at time points $0 \leq t_0^i < t_1^i < \dots < t_{n_i}^i = T^i$ and the inference task consists in recovering the "true" underlying θ based on the observations $X_{t_0^i}^i, \dots, X_{t_{n_i}^i}^i$ of $X^i, i = 1, \dots, N$. To this end, we first suppose to have the entire paths $(X_t^i)_{0 \leq t \leq T^i}, i = 1, \dots, N$, at our disposal and derive the continuous-time MLE. Later on we will investigate the error that arises when it is approximated by its discrete-time analogue.

STANDING ASSUMPTION (SA1). *To assure that the inference problem is well-defined, we assume that the coefficient functions F, Σ and the distributions of initial conditions x_0^i and random effects ϕ^i are such that (1) has unique, continuous solutions $X^i, i = 1, \dots, N$, satisfying $\sup_{0 \leq t \leq T^i} \mathbb{E} \left(\|X_t^i\|^k \right) < \infty$ for all $k \in \mathbb{N}$. If not further specified, it is assumed*

that μ, φ lie in relevant subspaces of \mathbb{R}^p and \mathbb{R}^d , respectively. Natural choices would be the projection of Θ on the μ -coordinate and the support of the $g(\cdot; \vartheta)$ – or the largest support of those, if $g(\cdot; \vartheta)$ and $g(\cdot; \tilde{\vartheta})$ have different supports for different $\vartheta, \tilde{\vartheta}$. We assume that for all μ, φ (in relevant subspaces) there is a continuous, adapted solution $X^{i, \mu, \varphi}$ to

$$dX_t^{i, \mu, \varphi} = F(t, X_t^{i, \mu, \varphi}, \mu, \varphi)dt + \Sigma(t, X_t^{i, \mu, \varphi})dW_t^i, \quad 0 \leq t \leq T^i, \quad X_0^{i, \mu, \varphi} = x_0^i, \quad (2)$$

with existing moments of any order (see above). We moreover assume that there are μ_0, φ_0 such that for all μ, φ the $X^{i, \mu, \varphi}$ satisfy, with $\Gamma = \Sigma \Sigma'$,

$$\mathbb{P} \left(\int_0^T F(s, X_s^{i, \mu, \varphi}, \mu_0, \varphi_0)' \Gamma(s, X_s^{i, \mu, \varphi})^{-1} F(s, X_s^{i, \mu, \varphi}, \mu_0, \varphi_0) ds < \infty \right) = 1.$$

REMARK 1. Sufficient assumptions for the above are the standard Lipschitz and sublinear growth conditions on F and Σ .

In all what follows, an integral of a matrix will be understood component-wise, $\|\cdot\|$ denotes the Euclidean vector norm, $\|\cdot\|$ a sub-multiplicative matrix norm, A' the transpose of a matrix A and A^{-1} its inverse. Further, all statements given for subject i are implied to hold for all $i = 1, \dots, N$. To ease notation, we will write $F_{\mu, \varphi}(s, x)$ instead of $F(s, x, \mu, \varphi)$ and simply $F(s, x)$ for $F_{\mu_0, \varphi_0}(s, x)$.

We denote by C_T the space of continuous \mathbb{R}^r -valued functions defined on $[0, T]$, which is endowed with the Borel- σ -algebra \mathcal{C}_T , the latter being associated with the topology of uniform convergence. A generic vector in \mathbb{R}^r is denoted by x , while we use μ for one such in \mathbb{R}^p and φ for one in \mathbb{R}^d . If $X^{i, \mu, \varphi}$ is the unique continuous solution to (2), we denote the measure induced by $X^{i, \mu, \varphi}$ on the space $(C_{T^i}, \mathcal{C}_{T^i})$ by $\mathbb{Q}_{\mu, \varphi}^i$ and for $\phi^i \sim g(\cdot; \vartheta)$ and X^i as the unique continuous solution to (1) under \mathbb{P}_θ , we write \mathbb{Q}_θ^i for the distribution of X^i on $(C_{T^i}, \mathcal{C}_{T^i})$. On the product space $\prod_{i=1}^N C_{T^i}$ we introduce the product measures $\mathbb{Q}_{N, \mu, \varphi} = \bigotimes_{i=1}^N \mathbb{Q}_{\mu, \varphi}^i$ and $\mathbb{Q}_{N, \theta} = \bigotimes_{i=1}^N \mathbb{Q}_\theta^i$. Expectations with respect to $\mathbb{Q}_{\mu, \varphi}^i$ and \mathbb{Q}_θ^i are written as $\mathbb{E}_{\mu, \varphi}$ and \mathbb{E}_θ , respectively, (for convenience, omitting the index i) and expectation with respect to the joint distribution of $(\phi^i, X^i)_{1 \leq i \leq N}$ on the product space $\prod_{i=1}^N (\mathbb{R}^d \times C_{T^i})$ will be denoted by \mathbb{E}_θ . From now on, we let (ϕ^i, X^i) be the canonical process on $\mathbb{R}^d \times C_{T^i}$.

2.2. Derivation of the likelihood

One can show that the conditional distribution of X^i conditioned on $\phi^i = \varphi$ coincides with the distribution of $X^{i, \mu, \varphi}$. This, combined with Fubini's theorem, implies that the likelihood of X^i , denoted by $p^i(\theta)$, is the integral over the likelihood $q^i(\mu, \varphi)$ of $X^{i, \mu, \varphi}$, weighted by $g(\varphi; \vartheta)$, that is, $p^i(\theta) = \int q^i(\mu, \varphi) g(\varphi; \vartheta) d\varphi$. The following theorem, which is a standard result from the theory on SDEs, specifies the likelihood $q^i(\mu, \varphi)$.

THEOREM 1 (CONDITIONAL LIKELIHOOD). *The distribution $\mathbb{Q}_{\mu, \varphi}^i$ is dominated by $\nu^i := \mathbb{Q}_{\mu_0, \varphi_0}^i$ and the Radon-Nikodym derivative is a.s. given by*

$$q^i(\mu, \varphi; X^i) = \exp \left(\int_0^{T^i} [F_{\mu, \varphi}(s, X_s^i) - F(s, X_s^i)]' \Gamma(s, X_s^i)^{-1} dX_s^i - \frac{1}{2} \int_0^{T^i} [F_{\mu, \varphi}(s, X_s^i) - F(s, X_s^i)]' \Gamma(s, X_s^i)^{-1} [F_{\mu, \varphi}(s, X_s^i) + F(s, X_s^i)] ds \right).$$

STANDING ASSUMPTION (SA2). *To assure that p^i is well-defined and measurable, the integrand $(\mu, \varphi, \vartheta, x) \mapsto q^i(\mu, \varphi; x)g(\varphi; \vartheta)$ should be measurable (w.r.t. the product Borel- σ -algebra). Therefore, we assume that the model is such that g and q^i are both product-measurable. From above, we see that q^i is measurable in the x -component, such that $q^i(\mu, \varphi; x)$ is surely product-measurable if it is continuous in its remaining two components. For instance, a sufficient condition for continuity in φ is that for any μ there is $\kappa = \kappa(\mu) > 0$ such that for all φ_0, φ, x (in relevant sets) and $0 \leq s \leq T^i$, one has $\| [F_{\mu, \varphi}(s, x) - F_{\mu, \varphi_0}(s, x)]' \Gamma(s, x)^{-1} \| \leq K(1 + \|x\|^\kappa) \|\varphi - \varphi_0\|$. This, together with a linear growth condition on Σ , implies that there is a $\tilde{\kappa}$ such that*

$$\| ([F_{\mu, \varphi} - F_{\mu, \varphi_0}]' \Gamma^{-1} [F_{\mu, \varphi} - F_{\mu, \varphi_0}])(s, x) \| \leq C(1 + \|x\|^{\tilde{\kappa}}) \|\varphi - \varphi_0\|^2.$$

One can then apply Kolmogorov's continuity criterion, which will yield the continuity (rather, existence of an in φ continuous version) of q^i in φ .

The likelihood of the sample $X^{(N)} = (X^1, \dots, X^N)$ is now an immediate consequence of Fubini's theorem.

THEOREM 2 (UNCONDITIONAL LIKELIHOOD). *The distribution \mathbb{Q}_θ^i admits the ν^i -density $p^i(\theta) := p^i(\theta; X^i) = \int_{\mathbb{R}^d} q^i(\mu, \varphi) \cdot g(\varphi; \vartheta) d\varphi$ and the corresponding product measure $\mathbb{Q}_{N, \theta}$ has the ν_N -density $p_N(\theta) := p_N(\theta; X^{(N)}) = \prod_{i=1}^N p^i(\theta)$ (with $\nu_N = \bigotimes_{i=1}^N \nu^i$).*

We remark that the absolute continuity of \mathbb{Q}_θ^i w.r.t. ν^i implies that all ν^i -a.s. statements made in the sequel also hold \mathbb{Q}_θ^i -a.s., for all $\theta \in \Theta$.

3. Asymptotic results for the MLE

The present section deals with asymptotic properties of the MLE (consistency, asymptotic normality and time discretization) and is divided into two parts. The first part, which assumes identically distributed observations, is in spirit close to the work of Delattre et al. (2013) and extends their results to a multidimensional state process with time-inhomogeneous dynamics. Given a drift function that is linear in the fixed and random effects (but possibly nonlinear in the state variable), we derive results on consistency and asymptotic normality of the MLE, following the traditional road of proof. Moreover, we leave the theoretical setting of continuous-time observations and switch to the practical situation, in which data are only available at discrete time points. We bound the discretization bias that arises when the continuous-time statistics are replaced by their discrete-time analogues. The second part is more general in two aspects: On the one hand, we allow the drift to be subject-specific by inclusion of covariate information and on the other hand, we suggest an alternative road to verification of consistency and asymptotic normality of the MLE, which poses less regularity assumptions and is based on the LAN property of the statistical model(s).

3.1. Independent and identically distributed observations

In this first subsection, we state results on consistency and asymptotic normality of the MLE when observations are independent and identically distributed. The proofs, which closely follow those in Delattre et al. (2013), while, however, getting more tedious to work out due to the multidimensional setup, are omitted here, but are available upon request.

The drift function in (1) is assumed to be linear in the fixed and random effects and random effects have a centered d -dimensional Gaussian distribution with unknown covariance matrix Ω , such that the likelihood is explicitly available. More specifically, we consider the r -dimensional processes X^i , whose dynamics are given by

$$dX_t^i = [A(t, X_t^i) + B(t, X_t^i)\mu + C(t, X_t^i)\phi] dt + \Sigma(t, X_t^i)dW_t^i, \quad X_0^i = x_0^i, \quad 0 \leq t \leq T, \quad (3)$$

and the parameter to be estimated based on the (continuous-time) observations $(X_t^i)_{0 \leq t \leq T}$, $i = 1, \dots, N$, is $\theta = (\mu, \Omega)$. The parameter space Θ is a bounded subset of $\mathbb{R}^p \times \mathfrak{S}_d(\mathbb{R})$, where $\mathfrak{S}_d(\mathbb{R})$ is the set of symmetric, positive definite $(d \times d)$ -matrices. The conditional likelihood of subject i is (cf. Theorem 1) $q^i(\mu, \varphi) = e^{\mu' U_{1i} - \frac{1}{2} \mu' V_{1i} \mu + \varphi' U_{2i} - \frac{1}{2} \varphi' V_{2i} \varphi - \varphi' S_i \mu}$ with the sufficient statistics

$$\begin{aligned} U_{1i} &= \int_0^T (B' \Gamma^{-1})(s, X_s^i) [dX_s^i - A(s, X_s^i) ds], & V_{1i} &= \int_0^T (B' \Gamma^{-1} B)(s, X_s^i) ds, \\ U_{2i} &= \int_0^T (C' \Gamma^{-1})(s, X_s^i) [dX_s^i - A(s, X_s^i) ds], & V_{2i} &= \int_0^T (C' \Gamma^{-1} C)(s, X_s^i) ds, \\ S_i &= \int_0^T (C' \Gamma^{-1} B)(s, X_s^i) ds. \end{aligned}$$

According to Theorem 2, the N -sample log-likelihood then turns into the explicit expression $l_N(\theta) = \log \left(\prod_{i=1}^N p^i(\theta) \right)$, where

$$p^i(\theta) = \frac{1}{\sqrt{\det(I + V_{2i}\Omega)}} \exp \left([U'_{1i} - U'_{2i} R^i(\Omega) S_i] \mu - \frac{1}{2} \mu' [V_{1i} - S'_i R^i(\Omega) S_i] \mu + \frac{1}{2} U'_{2i} R^i(\Omega) U_{2i} \right)$$

and with $R^i(\Omega) = (V_{2i} + \Omega^{-1})^{-1}$. We set $G^i(\Omega) = (I + V_{2i}\Omega)^{-1} V_{2i}$ and assuming that V_{1i}, V_{2i} are strictly positive definite, we can write $P^i(\theta) = G^i(\Omega) V_{2i}^{-1} (U_{2i} - S_i \mu)$, and the (vectorized) N -sample score function is given by $S_N(\theta) = \sum_{i=1}^N S^i(\theta) = \left[\frac{d}{d\mu} l_N(\theta), \frac{d}{d\Omega} l_N(\Omega)' \right]$, with

$$\begin{aligned} \frac{d}{d\mu} l_N(\theta) &= \sum_{i=1}^N [U'_{1i} - U'_{2i} R^i(\Omega) S_i] - \mu' \sum_{i=1}^N [V_{1i} - S'_i R^i(\Omega) S_i], \\ \frac{d}{d\Omega} l_N(\theta) &= \frac{1}{2} \sum_{i=1}^N \left[-G^i(\Omega) + P^i(\theta) P^i(\theta)' \right]. \end{aligned}$$

The MLE $\hat{\theta}_N = (\hat{\mu}_N, \hat{\Omega}_N)$ solves the equations

$$\begin{aligned} \hat{\mu}_N &= \left[\sum_{i=1}^N [V_{1i} - S'_i R^i(\hat{\Omega}) S_i] \right]^{-1} \left[\sum_{i=1}^N [U_{1i} - S'_i R^i(\hat{\Omega}) U_{2i}] \right] \\ \sum_{i=1}^N \left[-G^i(\hat{\Omega}_N) + P^i(\hat{\theta}_N) P^i(\hat{\theta}_N)' \right] &= 0. \end{aligned} \quad (4)$$

Note that the likelihood p^i is explicit even if the fixed effect enters the drift nonlinearly. However, only a linear fixed effect μ leads to an explicit expression for its ML estimator $\hat{\mu}$. If one wants to impose a random variation on all components of μ , one simply sets $B(t, x) = C(t, x)$. In that case, $U_{1i} = U_{2i} =: U_i$ and $V_{1i} = V_{2i} = S_i =: V_i$. The conditional likelihood simplifies to $q^i(\mu, \varphi) = e^{(\mu + \varphi)' U_i - \frac{1}{2} (\mu + \varphi)' V_i (\mu + \varphi)}$ and the unconditional likelihood to

$$p^i(\theta) = \frac{1}{\sqrt{\det(I + V_i \Omega)}} \exp \left(-\frac{1}{2} (\mu - V_i^{-1} U_i)' G^i(\Omega) (\mu - V_i^{-1} U_i) \right) \exp \left(\frac{1}{2} U_i' V_i^{-1} U_i \right). \quad (5)$$

With $\gamma^i(\theta) = G^i(\Omega)(V_i^{-1}U_i - \mu)$ (and clearly, $\gamma^i(\theta) = P^i(\theta)$), the MLE $\hat{\theta}_N = (\hat{\mu}_N, \hat{\Omega}_N)$ then solves

$$\begin{aligned} \hat{\mu}_N &= \left[\sum_{i=1}^N G^i(\hat{\Omega}_N) \right]^{-1} \left[\sum_{i=1}^N (I + V_i \hat{\Omega}_N)^{-1} U_i \right] \\ \sum_{i=1}^N \left[-G^i(\hat{\Omega}_N) + \gamma^i(\hat{\theta}_N) \gamma^i(\hat{\theta}_N)' \right] &= 0. \end{aligned} \quad (6)$$

To simplify the subsequent outline, we impose random effects on all components of μ , such that the dynamics are $dX_t^i = [A(t, X_t^i) + C(t, X_t^i)(\mu + \phi^i)] dt + \Sigma(t, X_t^i) dW_t^i$ and $p^i(\theta)$ is given by (5).

LEMMA 1 (MOMENT PROPERTIES). *For all $\theta \in \Theta$ the following statements hold.*

- (i) *For all $u \in \mathbb{R}^d$, $\mathbb{E}_\theta \left(\exp \left(u' (I + V_i \Omega)^{-1} U_i \right) \right) < \infty$.*
- (ii) *$\mathbb{E}_\theta (\gamma^i(\theta)) = 0$, $\mathbb{E}_\theta^i \left(\gamma^i(\theta) \gamma^i(\theta)' - G^i(\Omega) \right) = 0$.*

The asymptotic normality of the normalized score function is an immediate consequence of the law of large numbers together with the standard multivariate central limit theorem.

THEOREM 3 (ASYMPTOTIC NORMALITY OF THE NORMALIZED SCORE FUNCTION). *For all $\theta \in \Theta$, under \mathbb{Q}_θ and as N tends to infinity, the (vectorized version of) the normalized score function $\sqrt{N}S_N(\theta)$ converges in distribution to $\mathcal{N}(0, \mathcal{I}(\theta))$, where $\mathcal{I}(\theta)$ is the covariance matrix of $\text{vec}(S^i(\theta))$, with $S^i(\theta) = [\gamma^i(\theta), \frac{1}{2}(\gamma^i(\theta)\gamma^i(\theta)' - G^i(\Omega))]$.*

STANDING ASSUMPTION (SA3).

- (a) *The function $x \mapsto (C'\Gamma^{-1}C)(s, x) : \mathbb{R}^r \rightarrow \mathbb{R}^{d \times d}$ is not constant and under $\mathbb{Q}_{\varphi_0}^i$ the $\mathbb{R}^{d \times (d+1)}$ -valued random variable (U_i, V_i) admits a continuous density function (w.r.t. the Lebesgue measure), which is positive on an open ball of $\Theta \subset \mathbb{R}^d \times \mathbb{R}^{d \times d}$.*
- (b) *Θ is convex and compact. In particular, we assume that there are positive constants $M_1, M_{2,1}, M_{2,2}$ such that for all $\theta \in \Theta$: $\|\mu\| \leq M_1$ and $M_{2,1} \leq \|\Omega\| \leq M_{2,2}$.*
- (c) *The true value θ_0 belongs to $\text{int}(\Theta)$ and the matrix $\mathcal{I}(\theta_0)$ is invertible.*

THEOREM 4 (CONTINUITY OF KL INFORMATION AND UNIQUENESS OF ITS MINIMUM).

Let $K(\mathbb{Q}_{\theta_0}^i, \mathbb{Q}_\theta^i) = \mathbb{E}_{\theta_0}(\log p^i(\theta_0) - \log p^i(\theta))$ be the Kullback-Leibler information of $\mathbb{Q}_{\theta_0}^i$ w.r.t. \mathbb{Q}_θ^i . Then the function $\theta \mapsto K(\mathbb{Q}_{\theta_0}^i, \mathbb{Q}_\theta^i)$ is continuous and has a unique minimum at $\theta = \theta_0$.

THEOREM 5 (WEAK CONSISTENCY AND ASYMPTOTIC NORMALITY OF THE MLE). *Let $\hat{\theta}_N$ be an ML estimator defined as any solution of $l_N(\hat{\theta}_N) = \sup_{\theta \in \Theta} l_N(\theta)$. Then, as $N \rightarrow \infty$, $\hat{\theta}_N$ converges to θ_0 in \mathbb{Q}_{θ_0} -probability, and $\sqrt{N}(\hat{\theta}_N - \theta_0) \rightarrow \mathcal{N}(0, \mathcal{I}^{-1}(\theta_0))$ under \mathbb{Q}_{θ_0} .*

3.1.1. Discrete data

So far, we have assumed that we observe the entire paths $(X_t^i)_{0 \leq t \leq T}$ of the processes generated by (1). This is a severe restriction as in practice, observations are only available at discrete time points t_0, \dots, t_n . A natural approach is to replace the continuous-time integrals in $q^i(\theta)$ by discrete-time approximations and to derive an approximate MLE based on the resulting approximate likelihood. For instance, the stochastic integral term in $q^i(\theta)$, which is an expression of the form $\int_{t_k}^{t_{k+1}} h(s, X_s^i) dX_s^i$, may be replaced by a first-order approximation $\int_{t_k}^{t_{k+1}} h(s, X_s^i) dX_s^i \approx h(t_k, X_k^i) \Delta X_k^i$ or, by a higher-order approximation using Ito's formula, which gives $\int_{t_k}^{t_{k+1}} h(s, X_s^i) dX_s^i \approx H(t_{k+1}, X_{k+1}^i) - H(t_k, X_k^i) - \frac{\Delta t}{2} \sum_{j,l=1}^r (H_{x_j, x_l} \Sigma_j \Sigma_l')(t_k, X_k^i)$, where $h(t, x) = \nabla_x H(t, x)$. Note that this requires h to be of gradient-type, i.e., it requires the existence of a differentiable function

H such that h can be obtained as $h(t, x) = \nabla_x H(t, x)$. A higher-order approximation scheme is preferable, if the time step is not sufficiently small (non-high-frequency data) and/or the dynamics are highly non-linear. In the linear model (3), the first-order approximation of the continuous-time likelihood (which breaks down to the discretization of the sufficient statistics U_i, V_i) corresponds to the exact likelihood of its Euler scheme approximation. In particular, if we assume for simplicity that we observe all individuals at time points $t_i = \frac{i}{n}$ and denote by U_i^n, V_i^n the first-order discrete-time approximations to the continuous-time statistics U_i, V_i , one has the following result:

THEOREM 6. *Assume model (3) and suppose that $A(t, x), (C'\Gamma^{-1}C)(t, x)$ and $(C'\Gamma^{-1})(t, x)$ are globally Lipschitz-continuous in t and x and that in addition to $A(t, x), C(t, x)$ and $\Sigma(t, x)$ also $(C'\Gamma^{-1})(t, x)$ is of sublinear growth in x , uniformly in t . Then for any $p \geq 1$, the error behaves like $\mathbb{E}_\theta (\|V_i - V_i^n\|^p + \|U_i - U_i^n\|^p) = O(n^{-p/2})$.*

3.2. Non-homogeneous observations and covariates

In this section, we consider the asymptotic behavior of the MLE when the observations $X^i = (X_t^i)_{0 \leq t \leq T^i}$, $i = 1, \dots, N$, are independent between subjects i , but not necessarily identically distributed. This occurs, for instance, if the drift contains subject-specific covariate information D^i and these covariates are not i.i.d. If they are assumed to be deterministic, as in standard regression, the drift function F varies to a certain degree across subjects, $F^i(t, x, \mu, \phi^i) = F(t, x, D^i, \mu, \phi^i)$. As in the i.i.d. case, one would naturally wonder, under which conditions on the degree of variation among the F^i the derived MLEs still satisfy standard asymptotic results and, equally important, how to verify conditions that assure a regular asymptotic behavior.

Results on asymptotic normality of MLEs for independent, not identically distributed (i.n.i.d.) random variables are well-known (Bradley and Gart, 1962; Hoadley, 1971). They commonly built upon regularity conditions on the density functions p^i , such as third-order differentiability and boundedness of the derivatives, to ensure that integration and differentiation can be interchanged (as in the i.i.d. case, see also subsection 3.1). To achieve a limiting behavior when the observations do not share a common distribution, the variation across these non-homogeneous distributions has to be controlled. This is usually achieved by, on the one hand, imposing that the family of score functions $\{S^i(\theta); i \in \mathbb{N}\}$ satisfies the Lindeberg condition (a condition that bounds the variation of each $S^i(\theta)$ in relation to the total variation of the N -sample score function $\sum_{i=1}^N S^i(\theta)$). On the other hand, by requiring that the sample average $\frac{1}{N} \sum_{i=1}^N I^i(\theta)$ of the Fisher information matrices $I^i(\theta)$ converges to a positive definite limiting matrix $I(\theta)$. Under these conditions, the Lindeberg-Feller central limit theorem assures that the scaled N -sample score function is asymptotically $\mathcal{N}(0, I(\theta))$ distributed and a Taylor expansion gives the asymptotic normality of the MLE (Bradley and Gart, 1962; Hoadley, 1971; Gabbay et al., 2011).

The regularity conditions imposed on the densities as mentioned above are as standard as restrictive. If, for instance, random effects are supposed to have a double exponential distribution, i.e., a distribution whose density is not differentiable at $x = \theta$, those regularity conditions can not be met. The Laplace density is, however, "almost" regular. In fact, it satisfies a particular type of first-order differentiability and can perfectly be treated by a less standard, but more general road to verification of consistency and asymptotic normality. It dispenses with the previously mentioned strong regularity conditions on the density functions and instead builds upon L_2 -differentiability and the LAN property of a sequence of statistical models (Le Cam, 2012; Ibragimov and Has'minskii, 2013).

3.2.1. The convergence of the averaged Fisher informations

When studying the asymptotic behavior of the MLE in the setting of independent, but not identically distributed observations, a common - and intuitive - assumption is to require that the sample average $\frac{1}{N} \sum_{i=1}^N I^i(\theta) = \frac{1}{N} I_N(\theta)$ of the individual Fisher information matrices converges to a deterministic, symmetric, positive definite (SPD) limit matrix $I(\theta)$ as the sample size grows to infinity, (see, e.g., Bradley and Gart (1962, condition N7), or Hoadley (1971, equation (13))). This not only

simplifies verification of the assumptions considerably, it is also natural when compared to the i.i.d. case, where $I_N(\theta) = NI(\theta)$ and $[\frac{1}{N}I_N(\theta)]^{-1} = I(\theta)^{-1}$ is the asymptotic variance of the scaled MLE $\sqrt{N}\hat{\theta}_N$. However, it would be convenient to break the requirement down to the level of the model structure. If, for instance, the only structural difference between the distributions of the X^i is caused by inclusion of covariates, it is natural to ask whether one can formulate conditions on the average behavior of the covariates, as it is done in standard linear regression. This, however, is not possible for SDMEM, not even if we assume the simplest case where the drift function F is linear in state, covariates, fixed and random effects and if the latter are Gaussian distributed with known covariance matrix. And here is why. Assume a standard linear regression model $y_N = \mathcal{X}_N\mu + \epsilon_N$ with N observations collected in the response vector y_N , deterministic $N \times p$ design matrix \mathcal{X}_N containing the covariate information for all subjects, unknown parameter vector μ and N -dimensional vector ϵ_N of uncorrelated noise. The Fisher information matrix is given by $I_N(\mu) = \mathcal{X}_N'\mathcal{X}_N$ and the standard assumption is $\frac{1}{N}(\mathcal{X}_N'\mathcal{X}_N) \rightarrow V$ for some SPD matrix V . The matrix $I_N(\mu)$ has elements $\sum_{i=1}^N x_k^i x_l^i$, therefore the convergence requirement translates to assuming that second order sample averages of the covariates converge. If the linear model additionally includes random effects $\phi \sim \mathcal{N}(0, \Omega)$, that is, $y_N = \mathcal{X}_N\mu + \mathcal{Z}_N\phi + \epsilon_N$ and \mathcal{Z}_N is a deterministic $N \times d$ design matrix, a standard assumption (see, e.g., Pinheiro and Bates (2006)) for asymptotic normality is (among others) the convergence of $\frac{1}{N}\mathcal{X}_N'R_N(\Omega)^{-1}\mathcal{X}_N$ to a SPD V , where $R_N(\Omega) = I + \mathcal{Z}_N\Omega\mathcal{Z}_N'$ is the covariance matrix of y_N . So also for linear mixed effects models one can break the convergence of the average Fisher information down to conditions on a second-order average behavior of the covariates. In the SDE case the situation is more difficult. In fact, assuming moment statistics of the covariates to converge is *not* enough to ensure convergence of the Fisher information matrix, which we illustrate now in the simplest possible example that includes covariates. Assume $r = 1$ and $d = 2$ and consider the dynamics $dX_t^i = [X_t^i(\mu^1 + \phi^{i,1}) + D_t^i(\mu^2 + \phi^{i,2})]dt + dW_t^i$ for $0 \leq t \leq T$ with $X_0^i = x_0$. The vector $\mu = (\mu^1, \mu^2)'$ is the unknown fixed effect and the $\phi^i = (\phi^{i,1}, \phi^{i,2})'$, $i = 1, \dots, N$, are i.i.d. two-dimensional random effects with $\mathcal{N}(0, \Omega)$ distribution. Assume that the covariance matrix Ω is known, such that $\theta = \mu$ is the only unknown parameter. This setup is a special case of the example in subsection 3.2.4. The sufficient statistics U_i, V_i are given by

$$U_i = \begin{pmatrix} \int_0^T X_t^i dX_t^i \\ \int_0^T D_t^i dX_t^i \end{pmatrix} \quad \text{and} \quad V_i = \begin{pmatrix} \int_0^T (X_t^i)^2 dt & \int_0^T X_t^i D_t^i dt \\ \int_0^T X_t^i D_t^i dt & \int_0^T (D_t^i)^2 dt \end{pmatrix}.$$

The Fisher information is by definition $I^i(\mu) = \mathbb{E}_\mu \left(-\frac{d}{d\mu} S^i(\mu) \right) = \mathbb{E}_\mu (f(U_i, V_i))$, where the function f is the negative second derivative of the log-likelihood function. Since the log-likelihood is quadratic in μ , f will in fact not depend on the parameter (since Ω is known). We immediately conclude from the expression of p^i in eq. (5) that $f(U_i, V_i) = G^i(\Omega) = (I + V_i\Omega)^{-1}V_i$, such that $I^i(\mu) = \mathbb{E}_\mu (G^i(\Omega))$. The 2×2 matrix $G^i(\Omega) = (I + V_i\Omega)^{-1}V_i$ is, however, a non-linear function of V_i and thus finding an explicit expression for $I^i(\mu)$ is generally impossible - even in the simple linear case, where X^i is nothing but a Gaussian process. For comparison, in the linear mixed effects model, the log-likelihood for observation y^i with covariate vectors x^i, z^i is proportional to $-\frac{1}{2}(y^i - (x^i)'\mu)'R^i(\Omega)(y^i - (x^i)'\mu)$, with $R^i(\Omega) = (I + z^i\Omega z^i)^{-1}$ as inverse covariance matrix of y^i . The Fisher information is $\mathbb{E}_\mu (f(x^i, z^i))$ with $f(x^i, z^i) = x^i R^i(\Omega)(x^i)'$. The crucial difference, as compared to the SDE case, is that the matrix $R_i(\Omega)$ is deterministic. This renders calculation of the expectation unnecessary, such that $I^i(\mu) = \mathbb{E}_\mu (f(x^i, z^i)) = f(x^i, z^i)$. Therefore, requiring convergence of $\frac{1}{N} \sum_{i=1}^N I^i(\theta)$ is nothing but asking for a limiting behavior of covariate averages $\frac{1}{N} \sum_{i=1}^N f(x^i, z^i)$. This is particularly attractive as one can often design the experiment in such a way that the required limiting behavior holds. In the SDE case, however, it will not - not even in the simple linear case - be possible to break the condition $\frac{1}{N}I_N(\theta) \rightarrow I(\theta)$ down to the level of covariates, by requiring that an expression of the form $\frac{1}{N} \sum_{i=1}^N f(D^i, \theta)$, with f being some suitable function, converges. Therefore, it will generally not be possible to determine from an analytical expression of $I_N(\theta)$, whether the condition $\frac{1}{N}I_N(\theta) \rightarrow I(\theta)$ holds! Of course, this is not the end of the day, as the direct way via specific expressions for $I_N(\theta)$ is not the only possible road to show convergence. Averages of the form $\frac{1}{N} \sum_{i=1}^N a_i$ converge, for in-

stance, if the sequence $\{a_i\}$ converges to a limit a as $i \rightarrow \infty$. In this spirit, an alternative way would be to, heuristically speaking, assume that everything which is deterministic and individual-specific, or random but not with the same distribution across all individuals, converges as i goes to infinity to a limit (for instance, $x_0^i \rightarrow x_0, T^i \rightarrow T, D^i \rightarrow D$) in a suitable sense. Such an assumption corresponds to requiring that for large i the observations are, in fact, identically distributed. Exemplified in the linear example, one could proceed as follows. The Fisher information based on observation X^i is $I^i(\theta) = \mathbb{E}_\theta(f(U_i, V_i, \theta))$ and $U_i = U(x_0^i, T^i, D^i), V_i = V(x_0^i, D^i, T^i)$ with suitable functions f, U, V . If one can show that U, V and f are continuous in all arguments, a.s. convergence of x_0^i, T^i, D^i to limits x_0, T, D (as $i \rightarrow \infty$) implies convergence of $f(U^i, V^i, \theta)$. If the family $\{f(U^i, V^i, \theta); i \in \mathbb{N}\}$ is uniformly integrable, the a.s. convergence implies the convergence of moments and thus $I^i(\theta) \rightarrow I(\theta)$ and therefore also the average converges, i.e. $\frac{1}{N}I_N(\theta) = \frac{1}{N}\sum_{i=1}^N I^i(\theta) \rightarrow I(\theta)$. Note, however, that due to the nonlinear dependence of the function V on D^i , convergence of covariate averages of, for instance, the form $\frac{1}{N}\sum_{i=1}^N D_t^i \rightarrow D_t$ is not enough to ensure that the averaged Fisher informations converge - even if the other quantities x_0^i, T^i are the same for all individuals.

3.2.2. Asymptotics of the MLE with generalized conditions

A framework that also captures less regular models is provided by Ibragimov and Has'minskii (2013) and will here be adapted to the present setting. Those results that are not included in Ibragimov and Has'minskii (2013) are adaptations of ones therein and proofs will be omitted. We make the following assumptions.

STANDING ASSUMPTION (SA4).

- $\theta \mapsto p^i(\theta)$ is ν^i -a.s. continuous and $\theta \mapsto \sqrt{p^i(\theta)}$ is $L_2(\nu^i)$ -differentiable, i.e., $p^i(\theta)$ is Hellinger differentiable with $L_2(\nu^i)$ -derivative $\psi^i(\theta)$ (a row vector). That is, for each θ , $\int \|\psi^i(\theta; x)\|^2 d\nu^i(x) < \infty$ and $\lim_{\|h\| \rightarrow 0} \|h\|^{-2} \int \|\sqrt{p^i(\theta+h; x)} - \sqrt{p^i(\theta; x)} - \psi^i(\theta; x)h\|^2 d\nu^i(x) = 0$.
- $\psi^i(\theta)$ is continuous in $L_2(\nu^i)$. Consequently, the matrix $I^i(\theta) = 4 \int \psi^i(\theta; x)' \psi^i(\theta; x) d\nu^i(x)$ exists, is continuous and the N -sample Fisher information matrix can be defined as $I_N(\theta) = \sum_{i=1}^N I^i(\theta)$.
- $0 < \inf_{\theta \in \Theta} \|\frac{1}{N}I_N(\theta)\| \leq \sup_{\theta \in \Theta} \|\frac{1}{N}I_N(\theta)\| < \infty$.
- There is a SPD matrix $I(\theta)$ such that $\lim_{N \rightarrow \infty} \sup_{\theta \in K} \|\frac{1}{N}I_N(\theta) - I(\theta)\| \rightarrow 0$ and $\lim_{N \rightarrow \infty} \sup_{\theta \in K} \|\left(\frac{1}{N}I_N(\theta)\right)^{-1/2} - I(\theta)^{-1/2}\| = 0$.

Assuming that the (norm of the) Fisher information matrix grows beyond bounds corresponds to the requirement of infinite flow of information, which is naturally connected to the consistency of estimators. The L_2 -differentiability is neither a stronger nor weaker concept than standard (point-wise) differentiability. One may think of the relation between the two differentiability concepts as of the one between L_2 -convergence and almost sure convergence - without further assumptions, in general none of them implies the other, but under certain conditions, the limits are identical. Of course, if p^i is L_2 -differentiable and differentiable in the ordinary sense, then $\psi^i(\theta; x) = \frac{d}{d\theta} [p^i(\theta; x)^{1/2}]$.

Analogously to the traditional setting, we call $S^i(\theta) = 2p^i(\theta)^{-1/2}\psi^i(\theta)$ the score function of sample i and set $S_N(\theta) = \sum_{i=1}^N S^i(\theta)$ for the N -sample score function. A result familiar from traditional theory is that the score function is centered, which under the above conditions also holds true here, (Ibragimov and Has'minskii, 2013, p. 115). The likelihood ratio process (random field), which will be defined on the local parameter space $\Theta_{N,\theta} = \{h \in \mathbb{R}^q : \theta + I_N(\theta)^{-1/2}h \in \Theta\}$, is denoted by $L_{N,\theta}(h) = p_N(\theta + I_N(\theta)^{-1/2}h)/p_N(\theta)$.

REMARK 2.

- (a) Sufficient conditions for the (a.s.) continuity of $p^i(\theta) = \int q^i(\mu, \varphi)g(\varphi; \vartheta)d\varphi$ in θ are continuity of $\mu \mapsto q^i(\mu, \varphi)$ and $\vartheta \mapsto g(\varphi; \vartheta)$, together with the existence of a dominating,

integrable function, $q^i(\mu, \varphi)g(\varphi; \vartheta) \leq H(\varphi)$. If the density g of the random effects is assumed to be Gaussian, it is naturally continuous in ϑ (provided the variance parameter is bounded away from zero). For the continuity of q^i in μ , we remark the following: Suppose F is continuous and assume for simplicity $\Sigma(t, x) \equiv I$ is the identity matrix. If F is uniformly continuous in μ (for instance differentiable with bounded derivative/Jacobian), then $\mu \mapsto \int_0^{T^i} F(s, x(s), \mu, \varphi)' F(s, x(s), \mu, \varphi) ds$ is continuous. If F moreover has the property $\|F(s, x, \mu, \varphi) - F(s, x, \mu_0, \varphi)\| \leq K(1 + \|x\|^\kappa) \|\mu - \mu_0\|$ for some $\kappa > 0$, Kolmogorov's continuity criterion yields the continuity (rather, existence of an in μ continuous version) of q^i in μ .

- (b) Suppose $\theta \mapsto \sqrt{p^i(\theta)}$ is continuously differentiable. Then (since $p^i > 0$) the quantity $\tilde{S}^i(\theta) := 2p^i(\theta)^{-1/2} \frac{d}{d\theta} p^i(\theta)$ is well-defined. If the expression $\tilde{I}^i(\theta) = \mathbb{E}_\theta(\tilde{S}^i(\theta) \tilde{S}^i(\theta)')$ is finite and moreover continuous, then $\theta \mapsto \sqrt{p^i(\theta)}$ is L_2 -differentiable (Van der Vaart, 2000, Lemma 7.6), the L_2 -derivative coincides with the point-wise derivative and $\tilde{S}^i(\theta) = S^i(\theta)$, $\tilde{I}^i(\theta) = I^i(\theta)$.

3.2.3. General results on consistency and asymptotic normality

In this part, we give conditions on the asymptotic behavior of the MLE in our present framework. For simplicity, $\Theta \subseteq \mathbb{R}^q$ is assumed to be open, bounded and convex and in all what follows, $K \subset \Theta$ is a (fixed) compact subset. Whenever we write θ_N , we mean that it is of the form $\theta_N = \theta + I_N(\theta)^{-1/2} h$ for $\theta \in K$ and $h \in \Theta_{N,\theta}$.

THEOREM 7 (CONSISTENCY). *The MLE is uniformly on K consistent, if*

- (A.1) *There is a constant $m > q$ such that $\sup_{\theta \in K} \mathbb{E}_\theta(\|S_N(\theta)\|^m) < \infty$.*
 (A.2) *There is a positive constant $a(K)$ such that for (sufficiently large N and) all $\theta \in K$ (and all $h \in \Theta_{N,\theta}$) $H_i^2(\theta, \theta_N) \geq a(K) \frac{\|\theta_N - \theta\|^2}{1 + \|\theta_N - \theta\|^2}$, where $H_i^2(\theta_1, \theta_2) = \int \left(\sqrt{p^i(\theta_1)} - \sqrt{p^i(\theta_2)} \right)^2 d\nu^i$ is the squared Hellinger distance between $\mathbb{Q}_{\theta_1}^i$ and $\mathbb{Q}_{\theta_2}^i$.*

PROOF. (A.1) is an extension of Lemma III.3.2. in Ibragimov and Has'minskii (2013) to the setting of non-homogeneous observations and (A.2) is adapted from (Ibragimov and Has'minskii, 2013, Lemma I.5.3).

REMARK 3. *If the dimension of the parameter set is 1, the first condition above can be replaced by a sub-quadratic growth condition on the Hellinger distance (for i.i.d. observations, see Ibragimov and Has'minskii (2013, Theorem I.5.3)). In that case, one can instead require that $H^2(\theta_1, \theta_2) \leq A\|\theta_2 - \theta_1\|^2$, such that for one-dimensional parameter sets, consistency here reduces to $H^2(\theta_1, \theta_2)$ behaving asymptotically as $\|\theta_2 - \theta_1\|^2$.*

The following theorem implies the so-called uniform asymptotic normality of the model, which in turn gives rise to the asymptotic normality of the MLE (cf. Theorems II.6.2. and III.1.1 in Ibragimov and Has'minskii (2013)).

THEOREM 8 (ASYMPTOTIC NORMALITY). *Assume (A.1) and (A.2) from Theorem 7 hold. If additionally*

- (B.1) *$\{S^i(\theta), i = 1, \dots, N\}$ satisfy the Lyapunov condition uniformly in K , i.e. there is $\delta > 0$ such that $\lim_{N \rightarrow \infty} \sup_{\theta \in K} \sum_{i=1}^N \mathbb{E}_\theta \left(\|I_N(\theta)^{-1/2} S^i(\theta)\|^{2+\delta} \right) = 0$.*
 (B.2) *$\forall R > 0 : \lim_{N \rightarrow \infty} \sup_{\theta \in K} \sup_{\|h\| < R} \sum_{i=1}^N \int_{C_{T^i}} ([\psi^i(\theta_N) - \psi^i(\theta)] I_N(\theta)^{-1/2} h)^2 d\nu^i = 0$,*

are satisfied, $\{\hat{\theta}_N\}_{N \in \mathbb{N}}$ is uniformly in K consistent, asymptotically Gaussian distributed with parameters $(\theta, I_N(\theta)^{-1})$ and all moments of $\{I_N(\theta)^{1/2}(\hat{\theta}_N - \theta)\}_{N \in \mathbb{N}}$ converge uniformly in K to the corresponding moments of the $\mathcal{N}(0, I)$ distribution.

Condition (B.1) can be generalized to the Lindeberg condition. If the densities $\sqrt{p^i(\theta)}$ are twice continuously differentiable with second derivative $J^i(\theta)$, (B.2) can be replaced by requiring that $\lim_{N \rightarrow \infty} \sup_{\theta \in K} \sup_{\|h\| \leq R} \mathbb{E} [I_N(\theta)^{-1/2}]^4 \sum_{i=1}^N \int_{C_{T^i}} \mathbb{E} [J^i(\theta_N)]^2 d\nu^i = 0$. In the general setting, the p^i are not explicitly available. One can, however, formulate more general conditions on the drift function F and on the random effects density g such that differentiability of $\log p^i(\theta) = \log \left(\int q^i(\mu, \varphi) g(\varphi; \vartheta) d\varphi \right)$ is guaranteed, by assuring that differentiation can be passed under the integral sign. Sufficient conditions for the differentiability of $\log p^i(\theta)$ with respect to μ would, e.g., include differentiability of q^i w.r.t. μ and a uniform in μ domination of $\frac{d}{d\mu} q^i(\mu, \varphi) d\varphi \left(\int q^i(\mu, \varphi) g(\varphi; \vartheta) d\varphi \right)^{-1}$. Explicitly formulating these conditions is not very illustrative. Instead, it is recommended to check suitable conditions in the specific application at hand. One particular case in which the $p^i(\theta)$ are explicitly available is the case of linear Gaussian random effects, which will be considered in more detail below.

3.2.4. Affine Gaussian fixed and random effects and inclusion of covariates

We revisit the example model (3), but now include for each subject i a covariate information D^i , which is a known and deterministic function $D^i : [0, T^i] \rightarrow \mathbb{R}^s$. We let $B(t, X_t^i) = C(t, X_t^i)$ in (3) and enrich the function C by the covariate, $C(t, X_t^i, D^i)$. This model, being linear in state, covariate information, fixed effect $\mu \in \mathbb{R}^d$ and in the d -dimensional random effects ϕ^i , is the simplest non-trivial models with covariates. We assume that Θ is a bounded subset of $\mathbb{R}^d \times \mathfrak{S}_d(\mathbb{R})$. The likelihood and score function are as in (5) and (6), respectively, the only difference being a possible subject-specific observation horizon T^i and that the covariate information now enters the sufficient statistics $U_i = \int_0^{T^i} C(s, X_s^i, D_s^i)' \Gamma(s, X_s^i)^{-1} [dX_s^i - A(s, X_s^i) ds]$ and $V_i = \int_0^{T^i} C(s, X_s^i, D_s^i)' \Gamma(s, X_s^i)^{-1} C(s, X_s^i, D_s^i) ds$. Again, we assume that V_i is invertible. It is clear that the model is more regular than actually required and we include this example, where we verify the conditions of Theorem 8, merely for illustration purposes. It will also be revisited in the subsequent section on simulations, where we investigate parameter estimation (and hypothesis testing) for different sample sizes and sampling frequencies.

The set $K \subset \Theta$ is compact, so there are positive constants A_K, B_K, C_K such that $\|\mu\| \leq A_K, B_K \leq \|\Omega\| \leq C_K$. One can show that $\|G^i(\Omega)\| \leq \|\Omega^{-1}\|$, which gives the upper bound $\|S^i(\theta)\| \leq (\|\gamma^i(\theta)\| + \|\Omega^{-1}\| + \|\gamma^i(\theta)\|^2)$. Moreover, the moment-generating function $\Phi_{\theta, \gamma^i(\theta)}(a)$ of $\gamma^i(\theta)$ can be bounded from above by $e^{\frac{1}{2} a' \Omega^{-1} a}$, for $a \in \mathbb{R}^d$. This can be used to find that $\mathbb{E}_\theta (\|\gamma^i(\theta)\|^m) \leq C_1$ for some constant C_1 that may depend on K, d, m . Therefore, there is another constant C_2 , which may depend on K, d, m, N , such that $\mathbb{E}_\theta (\|S_N(\theta)\|^m) \leq C_2$, proving (A.1). To verify (A.2), note that the regularity of $p_N(\theta)$ and its derivatives implies that

$$\begin{aligned} H^2(\theta, \theta_N) &= \int \left[-\psi_N(\theta)(\theta_N - \theta) + \left(\sqrt{p_N(\theta_N)} - \sqrt{p_N(\theta)} \right) + \psi_N(\theta)(\theta_N - \theta) \right]^2 d\nu \\ &= \int [-\psi_N(\theta)(\theta_N - \theta)]^2 d\nu + o(\|\theta_N - \theta\|^2) \\ &= (\theta_N - \theta)' I_N(\theta) (\theta_N - \theta) + o(\|\theta_N - \theta\|^2) - 2O(\|\theta_N - \theta\|^2) o(\|\theta_N - \theta\|^2) \\ &\geq \|(\theta_N - \theta)\|^2 \lambda_{N, \min}(\theta) + o(\|\theta_N - \theta\|^2). \end{aligned}$$

where $\lambda_{N, \min}(\theta)$ denotes the smallest eigenvalue of $I_N(\theta)$. Therefore, for N sufficiently large, there is a constant A_K such that $H^2(\theta, \theta_N) \geq A_K \|\theta_N - \theta\|^2$. Since Θ is bounded, we even have $\|(\theta_N - \theta)\|^2 \geq C \frac{\|(\theta_N - \theta)\|^2}{1 + \|(\theta_N - \theta)\|^2}$ for some positive constant C , which shows that (A.2) holds. The Lyapunov condition (B.1) follows in a straightforward way. According to the above, $\mathbb{E}_\theta (\|S^i(\theta)\|^3) \leq$

C for some C and therefore

$$\begin{aligned} \sup_{\theta \in K} \sum_{i=1}^N \mathbb{E}_{\theta} \left(\|I_N(\theta)^{-1/2} S^i(\theta)\|^3 \right) &\leq N^{-3/2} \sup_{\theta \in K} \|\sqrt{N} I_N(\theta)^{-1/2} - I(\theta)^{-1/2}\| \sum_{i=1}^N \mathbb{E}_{\theta} \left(\|S^i(\theta)\|^3 \right) \\ &\quad + N^{-3/2} \sup_{\theta \in K} \|I(\theta)^{-1/2}\| \sum_{i=1}^N \mathbb{E}_{\theta} (\|S^i(\theta)\|^3) \\ &\leq CN^{-1/2} \left[\sup_{\theta \in K} \|\sqrt{N} I_N(\theta)^{-1/2} - I(\theta)^{-1/2}\| + \sup_{\theta \in K} \|I(\theta)^{-1/2}\| \right], \end{aligned}$$

which converges to 0 as $N \rightarrow \infty$. To verify (B.2), we show that

$$\sup_{\|h\| \leq R} \frac{1}{N} \left[\frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\nu^i} \left(\left\| J^i(\theta_N) - J^i(\theta) \right\|^2 \right) \right] \quad \text{and} \quad \frac{1}{N} \left[\frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\nu^i} \left(\left\| J^i(\theta) \right\|^2 \right) \right] \quad (7)$$

converge to 0 uniformly in K . As $J^i(\theta)$ is continuous, it is uniformly continuous on compacta, such that for all $i \in \mathbb{N}$, $a_{i,N} = \sup_{\|h\| \leq R} \left\| J^i(\theta_N) - J^i(\theta) \right\|^2$ converges a.s. to 0 as $N \rightarrow \infty$.

One can show that $a_{i,N} \leq A^i(\theta, R)$ and $\mathbb{E}_{\nu^i} (A^i(\theta, R)^2) \leq D_K$. Dominated convergence therefore implies $\mathbb{E}_{\theta}(a_{i,N}) \rightarrow 0$, and the uniform (in i) bound D_K implies uniform in K convergence of the left term in (7) to 0. For the right hand side term in (7) we note that $\mathbb{E}_{\nu^i} (\|J^i(\theta)\|^2) \leq \mathbb{E}_{\theta} (\|\frac{d}{d\theta} S^i(\theta)\|^2) + \mathbb{E}_{\theta} (\|S^i(\theta)' S^i(\theta)\|^2) < C_K$, where C_K is a constant that only depends on K . We conclude uniform in $\theta \in K$ convergence of the right hand side term in (7) to 0. The right hand side follows similarly.

3.3. Hypothesis testing

It is commonly of interest to the researcher to test whether an applied treatment has a significant effect on the treated subjects, i.e., to test whether an underlying treatment effect β , a s -dimensional subparameter of the fixed effect μ , $1 \leq s \leq p$, is significantly different from 0. The asymptotic normality of the MLE in this model lends itself naturally to the application of Wald tests, which can be used to investigate two-sided null hypotheses such as $H_0 : \beta = 0$ (no treatment effect) or more generally, any k -dimensional, $1 \leq k \leq s$, linear null hypothesis $H_0 : L\beta = \eta_0$, where L is a $5 \times k$ matrix of rank k which specifies the linear hypotheses of interest and $\eta_0 \in \mathbb{R}^k$. The Wald test statistic is $\hat{W}_N = \left(L\hat{\beta}_N - \eta_0 \right)' \left(L\hat{V}_N L' \right)^{-1} \left(L\hat{\beta}_N - \eta_0 \right)'$, which is under the null hypothesis asymptotically χ^2 -distributed with k degrees of freedom (Lehmann and Romano, 2006). Here, $\hat{\beta}_N$ is the MLE of β and $\hat{V}_N = \widehat{\text{Cov}}(\hat{\beta}_N)$ denotes its estimated variance-covariance matrix of $\hat{\beta}_N$.

4. Simulations

4.1. Linear transfer model

The first example, which is inspired from a study on the selenomethionine metabolism in humans (Große Ruse et al., 2015), is a 5-dimensional linear transfer model, which finds applicability in various fields, especially in modeling population flows or in pharmacokinetics. A component in the model's state vector can be viewed to represent the concentration of a substance in a certain compartment and the model describes the (linear) flow between compartments. We consider a basic cascade-shaped transfer structure as illustrated in Figure 1. When observing N subjects, each of them following the linear transfer model in Figure 1, it is often reasonable to assume that the transfer rates are subject-specific. We moreover assume that we are given covariate information on subject i in form of a (deterministic) variable $D^i \in \{0, 1\}$. It encodes the affinity of subject i to one of two possible study groups, such as *placebo* and *treatment*. Consider the model (for simplicity assuming unit diffusion) $dX_t^i = F(t, X_t^i, D^i, \mu, \phi^i)dt + dW_t^i$, $0 \leq t \leq T$, $X_0^i = 0$, where $\mu' = (\alpha, \beta)$ is the fixed

parameter and the drift function has the specific form $F(t, X_t^i, D^i, \mu, \phi^i) = -G(\alpha + \phi^i)X_t + D^i\beta$, $i = 1, \dots, N$, with rate matrix

$$G(\alpha) = \begin{pmatrix} \alpha_1 & 0 & 0 & 0 & -\alpha_5 \\ -\alpha_1 & \alpha_2 & 0 & 0 & 0 \\ 0 & -\alpha_2 & \alpha_3 + \alpha_6 & 0 & 0 \\ 0 & 0 & -\alpha_3 & \alpha_4 & 0 \\ 0 & 0 & 0 & -\alpha_4 & \alpha_5 \end{pmatrix}.$$

The (unknown) fixed effect μ has the 6-dimensional component α , which is shared across both groups (*placebo* and *treatment*) and an additional 5-dimensional component β , which describes the effect of the covariate (treatment effect) on the subjects' dynamics. We let $\beta' = (1, 2, 3, 1, -2)$. The random effects ϕ^i are i.i.d. $\mathcal{N}(0, \Omega)$ -distributed and the covariance matrix Ω is unknown. With $\alpha' = (\alpha_1, \dots, \alpha_6) = (2, 4, 3, 2, 1, 1)$, all eigenvalues of $G(\alpha)$ have positive real parts, implying that the model has a stationary solution. The processes X^i for individuals without treatment, i.e. $D^i = 0$, are (on average) mean-reverting to 0, and the processes belonging to individuals in the treatment group, $D^i = 1$, are mean-reverting to the long-term mean $(G(\alpha + \phi^i))^{-1}\beta$. For our choice of parameters, conditional on $\phi^i = 0$, this long-term mean is $(7.50, 4.25, 5.00, 8.00, 14.00)'$, see also Figure 2. The covariance matrix Ω is taken to be a diagonal matrix with entries $\text{diag}(\Omega) = (0.5^2, 1^2, 1^2, 0.5^2, 0.3^2, 0.3^2)$. The observation horizon T is fixed to $T = 15$. A trajectory of $(X_t^1, \dots, X_t^N)_{0 \leq t \leq T}$ is simulated with the Euler-Maruyama scheme with simulation step size $\delta = 10^{-4}$. Figure 2 shows four realized (5-dimensional) trajectories of the process X^i . The upper two panels show trajectories for $D^i = 0$ and the lower two correspond to trajectories with $D^i = 1$.

4.1.1. Parameter estimation

For parameter estimation, the simulated trajectories are thinned by a factor b (taking only every b -th observation). To explore the expected time-discretization bias of the estimators, we repeated estimation for different thinning factors, $b \in \{10, 100, 1000\}$, which results in sampling intervals $\Delta t = \delta \cdot b = 0.001, 0.01, 0.1$. To also investigate the estimation performance as a function of sample size, we performed estimation on trajectories (X^1, \dots, X^N) , for sample sizes $N = 20$, $N = 50$ and $N = 100$. Estimation for all considered $(\Delta t, N)$ -combinations was repeated on $M = 500$ simulated data sets. Tables 1 and 2 report the sample estimates of relative biases and root mean squared errors (RMSE) of the fixed effects and of the variances of the random effects. The relative bias of $\hat{\alpha}_j$ is computed as $\frac{1}{M} \sum_{m=1}^M \frac{\hat{\alpha}_j^{(m)} - \alpha_j}{\alpha_j}$ and the RMSE as $\left(\frac{1}{M} \sum_{m=1}^M (\hat{\alpha}_j^{(m)} - \alpha_j)^2 \right)^{1/2}$, $j = 1, \dots, 6$, and with an analogous definition for the other parameters. Table 1 shows estimation results for a fixed sample size of $N = 50$, and different values of sampling intervals Δt , while results in Table 2 are computed for a fixed sampling interval $\Delta t = 0.001$ and different values of sample size N . In each table, the first six rows correspond to estimated biases and RMSEs of the shared fixed effects α_j , $j = 1, \dots, 6$. The subsequent five rows show the estimated biases and RMSEs of the treatment effects β_j , $j = 1, \dots, 5$ and the last six rows correspond to the estimated biases and RMSEs of the diagonal elements of Ω (i.e., the variances of the random effects). The estimation is very accurate already at sample sizes as small as $N = 20$, when the data is sampled at high frequency (here $1/0.001$), see Table 2. For a moderate sampling frequency of $1/0.01$, the results in the middle part of Table 1 reveal that estimates of the fixed effects α, β are on average biased by only about 1-2% (of the true parameter value), which is still very accurate. The variances of the random effects are estimated with an average bias of 5-9% for $N = 50$ and $\Delta t = 0.01$. When the observations are sampled at low frequency $1/0.1$, estimation gets unreliable. The bias due to the time-discretization of the continuous-time estimator is very pronounced, with values of up to 25% for the fixed effects and up to almost 50% for the variances of the random effects. The RMSEs rise - as compared to a 10 times higher frequency - by more than 100%. If only low-frequency data is available, caution is recommended and estimation should only be done on a data set that has been enlarged by imputing data in between the observation time points.

Table 1. Linear transfer model. Shown are estimated relative bias and RMSE of $\hat{\alpha}, \hat{\beta}$, and $\text{diag}(\hat{\Omega})$. The sample size is fixed to $N = 50$, but different sampling intervals are considered ($\Delta t = 0.001, 0.01, 0.1$). For each value of Δt , the estimation was repeated on $M = 500$ generated data sets.

		$\Delta t = 0.001$		$\Delta t = 0.01$		$\Delta t = 0.1$	
true value		rel. bias	RMSE	rel. bias	RMSE	rel. bias	RMSE
α	2.00	0.001	0.079	-0.018	0.086	-0.182	0.369
	4.00	-0.002	0.149	-0.024	0.172	-0.204	0.824
	3.00	0.001	0.163	-0.021	0.170	-0.203	0.624
	2.00	-0.001	0.083	-0.017	0.088	-0.162	0.332
	1.00	0.001	0.047	-0.016	0.049	-0.159	0.164
	1.00	0.002	0.091	-0.008	0.091	-0.082	0.119
β	1.00	-0.002	0.099	-0.020	0.099	-0.166	0.186
	2.00	-0.002	0.114	-0.024	0.121	-0.198	0.408
	3.00	0.002	0.152	-0.010	0.152	-0.116	0.373
	1.00	-0.001	0.148	0.014	0.146	0.140	0.188
	-2.00	0.002	0.124	-0.024	0.131	-0.255	0.522
$\text{diag}(\Omega)$	0.25	-0.037	0.062	-0.079	0.062	-0.399	0.108
	1.00	-0.035	0.208	-0.095	0.216	-0.483	0.498
	1.00	-0.035	0.215	-0.085	0.219	-0.426	0.445
	0.25	-0.026	0.061	-0.065	0.060	-0.352	0.097
	0.09	-0.009	0.022	-0.047	0.021	-0.333	0.034
	0.09	-0.040	0.036	-0.065	0.035	-0.213	0.036

Table 2. Linear transfer model. Shown are estimated relative bias and RMSE of $\hat{\alpha}, \hat{\beta}$, and $\text{diag}(\hat{\Omega})$. The sampling interval is fixed to $\Delta t = 0.001$, but different sample sizes are considered ($N = 20, 50, 100$). For each value of N , the estimation was repeated on $M = 500$ generated data sets.

true value		$N = 20$		$N = 50$		$N = 100$	
		rel. bias	RMSE	rel. bias	RMSE	rel. bias	RMSE
α	2.00	0.003	0.116	0.001	0.079	-0.001	0.058
	4.00	0.001	0.232	-0.002	0.149	0.001	0.114
	3.00	0.003	0.253	0.001	0.163	-0.001	0.106
	2.00	-0.003	0.126	-0.001	0.083	-0.000	0.052
	1.00	0.003	0.074	0.001	0.047	-0.003	0.031
	1.00	-0.003	0.146	0.002	0.091	0.000	0.068
β	1.00	0.000	0.157	-0.002	0.099	0.004	0.073
	2.00	-0.001	0.174	-0.002	0.114	0.002	0.075
	3.00	0.002	0.233	0.002	0.152	0.000	0.102
	1.00	0.010	0.231	-0.001	0.148	-0.002	0.102
	-2.00	0.006	0.203	0.002	0.124	-0.000	0.087
$\text{diag}(\Omega)$	0.25	-0.091	0.093	-0.037	0.062	-0.014	0.043
	1.00	-0.046	0.355	-0.035	0.208	-0.020	0.162
	1.00	-0.073	0.343	-0.035	0.215	-0.017	0.163
	0.25	-0.035	0.097	-0.026	0.061	-0.016	0.039
	0.09	-0.045	0.035	-0.009	0.022	-0.021	0.015
	0.09	-0.181	0.055	-0.040	0.036	-0.020	0.027




Fig. 1. Illustration of the 5-dimensional linear transfer model used in the first simulation example. The state $X_i = (X_{i,t})_{0 \leq t \leq T}$ gives the concentration (over time) of a substance in compartment $i, i = 1, \dots, 5$. The $\alpha_i, i = 1, \dots, 6$ are the unknown rates of flow between corresponding compartments (for α_6 the outflow of the system).




Fig. 2. Linear transfer model: Four realizations of the 5-dimensional state process $(X_t^i)_{0 \leq t \leq T}$. The upper two panels show realizations when the covariate is 0 ("reference group") and the lower two panels display trajectories for $D^i = 1$ ("treatment group"). Note the clearly visible difference in the long-term means between the two groups.

4.1.2. Hypothesis testing

A natural step is to test whether β , or a subparameter, is significantly different from 0. We estimate the false-positive rate of the Wald test (see subsection 3.3) in this model and investigate the test's power under different "true" (non-zero) treatment effects. The estimated variance-covariance matrix $\hat{V}_N = \widehat{\text{Cov}}(\hat{\beta}_N)$ of $\hat{\beta}_N$ is obtained from $M = 500$ (separately) computed MLEs $\hat{\beta}_N^{(m)}, m = 1, \dots, M$, where underlying data sets have been simulated under the true hypothesis (under H_0 for estimation of the false positive rate and under H_1 for power estimation). Tables 1 and 2 show that the estimation was accurate for high- and medium-frequency observations. Diagnostic plots (not shown here) reveal that the asymptotic distribution of the MLE is close to normal already for $N = 20$ subjects, such that even for a rather small data set and a medium sampling frequency, test results can be considered sufficiently reliable. The choice $(N, \Delta t) = (20, 0.01)$ provides a simulation setting that is sufficiently reliable, but at the same time not trivial and will challenge the hypothesis test, in particular for small treatment effects. The estimated false positive rate (based on M

under H_0 generated data sets) is 0.074, revealing a slightly liberal finite-sample test behavior. The power of detecting a treatment effect (rejecting $H_0 : \beta = 0$) was computed for different "true" values of β . For $\beta = (1, 2, 3, 1, -2)'$ (values as in the estimation part above), the estimated power was 1. This comes to no surprise as the long-term mean $(7.5, 4.25, 5, 8, 14)'$ of the state process in the treatment group is considerably different from the zero long-term mean of the control group. The power, estimated to 0.956, was still convincing for a much smaller treatment effect $\beta = (0.1, 0.2, 0.3, 0.1, -0.2)'$, which gives a long-term mean of $(0.75, 0.425, 0.5, 0.8, 1.4)'$. This is especially impressive as the state process' standard deviation (from its long-term mean 0) under H_0 is about $(0.66, 0.49, 0.59, 0.72, 1.21)'$. More challenging is the rejection of H_0 when the treatment has a small effect on, e.g., only one coordinate, $\beta = (0.1, 0, 0, 0, 0)'$. In this case (long-term mean $(0.2, 0.1, 0.1, 0.15, 0.3)'$), and for such a small sample size the chance of rejecting H_0 is as small as 16% and it is thus hardly possible to detect a difference between groups. However, while being only slightly conservative, the asymptotic Wald test is able to detect a treatment effect for a rather small data set, even if it causes only a little change of the long-term mean as compared to the standard deviation of the process.

4.2. Fitzhugh-Nagumo model

The deterministic Fitzhugh-Nagumo (FHN) model (FitzHugh, 1955; Nagumo et al., 1962) is a two-dimensional approximation of the well-known four-dimensional Hodgkin-Huxley neuronal model (Hodgkin and Huxley, 1952) and is typically applied to model the regenerative firing mechanism in an excitable neuron. Neural firing is a complex interplay of numerous cell processes and to account for various unexplained noise sources, a stochastic FHN model can be considered (Jensen et al., 2012),

$$\begin{aligned} dY_t &= \frac{1}{\varepsilon} (Y_t - Y_t^3 - Z_t + s) dt + \sigma_1 dW_{1,t}, \\ dZ_t &= (\gamma Y_t - Z_t + \eta) dt + \sigma_2 dW_{2,t}. \end{aligned} \tag{8}$$

The variable Y represents the membrane potential of a neuron, while the Z coordinate represents the recovery. The time scale separation ε is commonly $\ll 1$, such that Y lives on a much faster time scale than Z . The variable s is the input current. If $\gamma > 1$, the system has exactly one fixed point, which may be stable or unstable, depending on the specific parameter values. Under the reparametrization $\mu = (1/\varepsilon, s/\varepsilon, \gamma, \eta)'$, (8) may be written as in (3). We assume to study a collection of N excitable neurons and model their membrane potentials Y^i via $dX_t^i = A(X_t^i) + C(X_t^i)(\mu + \phi^i)dt + \Sigma dW_t^i$, $0 \leq t \leq T$, where $X^i = (Y^i, Z^i)'$ and the ϕ^i are the i.i.d. $\mathcal{N}(0, \Omega)$ -distributed random effects. Observe that despite being nonlinear in the state variable, the model equations are linear in the random effects and therefore an explicit likelihood is available. We assume here that both coordinates of X^i are observed. For all simulations, we let $\sigma_1 = 0.5, \sigma_2 = 0.3$ (assumed known), we fix $T = 20$ and choose the values of the unknown parameters as $\varepsilon = 0.1, s = 0.5, \gamma = 1.5$ and $\eta = 1.2$. With this choice of η the fixed point of the deterministic FHN system is stable, but small noise levels will suffice to induce large excursions through state space (spikes). The covariance matrix of the random effects is fixed as $\Omega = \text{diag}(1.5^2, 1^2, 0.2^2, 0.2^2)$. The simulation settings are as in the previous example: We simulate each trajectory $(X_t^1, \dots, X_t^N)_{0 \leq t \leq T}$ with the Euler-Maruyama scheme and a simulation time step of $\delta = 10^{-4}$. The estimation is carried out on the thinned trajectory. We conduct estimation for different values of the sample size $N \in \{20, 50, 100\}$ to investigate the finite sample behavior. To illustrate how the discrete-time bias evolves, we repeat estimation for thinning factors $b = 10, 100, 1000$, which results in sampling intervals of $\Delta t = \delta \cdot b = 0.001, 0.01, 0.1$, respectively (note that the observation horizon is always fixed to $T = 20$). For all combinations of N and Δt , the estimation is repeated on $M = 500$ generated data sets. Figure 3 shows example trace plots of four realizations, which illustrate the possible qualitatively different behaviors of the state process, depending on the realized values of the random effects. Table 3 shows, similar to the previous simulation example on the linear transfer model, the bias and the RMSE of the estimates, where estimation was based on samples with fixed sample size $N = 50$, but repeated for different sampling

intervals Δt . The estimation was done under the reparametrization $\mu = (1/\varepsilon, s/\varepsilon, \gamma, \eta)$. Estimates for the parameter ε and s on the original scale are obtained by transformation. The upper six rows show estimated bias and RSME for the fixed effects (on the original and on the transformed μ -scale) and the subsequent four rows correspond to results for the estimation of the diagonal of Ω . Despite the non-linearity (in the state) of the model, implying violation of the absolute standard assumptions on the diffusion drift for regularity of the model, the parameter estimation for high-frequency data and moderate ($N = 50$) sample size is very convincing (Table 1, first two columns), while still being satisfactory for observations sampled at medium frequency (middle two columns). If observations are sampled at low frequency (last two columns in Table 3), the bias for the estimation of s, γ, η is still rather low (with 1%, 8% and 5% bias, respectively, as compared to the true parameter value). The estimation of ε is, however, highly biased. The variances of the random effects are all estimated with an error of about 21-28%, except for the variance of the random effect that adds to ε , which has an error of as high as 69%. This comes to no surprise, since non-linearity in the state requires denser observations. For ε we estimate the inverse of a small number, making the estimator unstable. Figure 4 illustrates the distribution of the relative bias of the ν -scale estimates (i.e. of the bias divided by the true parameter value) for high-frequency observations and different sample sizes (for $N = 20$ in red, for $N = 50$ in green, for $N = 100$ in black). Here one can see that even for small sample sizes ($N = 20$), estimates are centered around the true parameter (negligible bias, though with considerable variance) and their distribution approaches a normal distribution for large N , becoming increasingly centered around the true value (zero bias).




Fig. 3. FHN model. Trace plots of four realizations of the stochastic FHN model with random effects. The corresponding (rounded) realized parameter values of $\mu + \phi^i$ are (8.86, 4.29, 1.50, 1.39), (10.16, 7.49, 1.73, 1.40), (9.98, 5.49, 1.10, 1.07), (9.26, 5.17, 1.84, 1.01).

5. Discussion

Stochastic differential equations with random effects constitute an attractive class of statistical models, especially for biological data. We extended the framework introduced in Delattre et al. (2013) to multidimensional and time-inhomogeneous state processes and proved the consistency of the maximum likelihood estimator using standard tools when the random effects enter the model linearly. However, the dynamics are allowed to be non-linear in the state. This particular setting comprises numerous well-known models, such as the predator-prey (or Lotka-Volterra) model (Murray, 2002), the Lorenz equations introduced by Lorenz (1963), which have been used to model, e.g., temperature, wind speed and humidity, the Brusselator model (Kondepudi and Prigogine, 2014,

Table 3. FHN model. Shown are estimated relative bias and RMSE of $\hat{\mu}$ and $\text{diag}(\hat{\Omega})$. The sample size is fixed to $N = 50$, but different sampling intervals are considered ($\Delta t = 0.001, 0.01, 0.1$). For each value of Δt , the estimation was repeated on $M = 500$ generated data sets.

		$\Delta t = 0.001$		$\Delta t = 0.01$		$\Delta t = 0.1$	
true value		rel. bias	RMSE	rel. bias	RMSE	rel. bias	RMSE
ε	0.10	0.003	0.022	0.030	0.037	0.356	0.356
s	0.50	0.001	0.033	0.002	0.033	0.009	0.035
γ	1.50	-0.000	0.031	-0.006	0.032	-0.079	0.121
η	1.20	-0.001	0.031	-0.005	0.032	-0.051	0.067
$1/\varepsilon$	10.00	-0.003	0.216	-0.028	0.349	-0.262	2.624
s/ε	5.00	-0.002	0.135	-0.026	0.188	-0.256	1.283
$\text{diag}(\Omega)$	2.25	-0.048	0.469	-0.155	0.539	-0.690	1.563
	1.00	-0.025	0.197	-0.044	0.197	-0.281	0.318
	0.04	-0.035	0.010	-0.044	0.010	-0.212	0.012
	0.04	-0.007	0.010	-0.028	0.009	-0.218	0.012

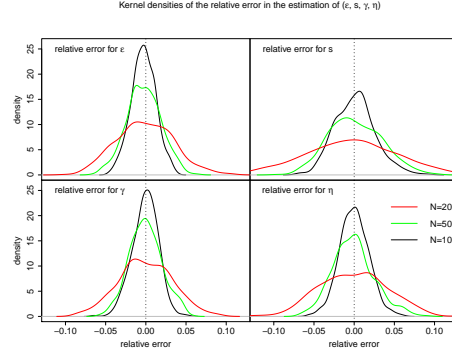


Fig. 4. FHN model. Density plots of the relative error $\text{err}_{\nu_j}(N) = \frac{\hat{\nu}_j^{(m)} - \nu_j}{\nu_j}$, $m = 1, \dots, M$, of the estimation of fixed effects, given here on the original ν -scale. The sampling interval is $\Delta t = 0.001$, but different sample sizes are considered: $N = 20$ (red), $N = 50$ (green), $N = 100$ (black). Estimation was repeated on $M = 500$ generated data sets.

19.4), the FHN model (see section 4) or the SIR (susceptible-infected-removed) model introduced by Kermack and McKendrick (1927), an epidemic model which has widely been studied and applied (Keeling and Rohani, 2008; Jiang et al., 2011; Guy et al., 2015).

We examined the extension of the i.i.d. model to independent, but not identically distributed observations, with particular emphasis on the inclusion of covariate information. We pointed out the fundamental difference to assumptions on the covariates that are standard in regression analysis, gave conditions for asymptotic normality of the MLE (or, more generally, the LAN of the models) when the observations do not come from the same distribution and illustrated their verification by means of an example.

The quality of the estimation in terms of sample size and sampling frequency was investigated in two simulation studies. In the first one, we use a model with covariates, which is linear in parameters and state. When observations are sampled at high frequency, estimation results were convincing already for small sample sizes ($N = 20$), despite the comparably large number (11 fixed effects and 6 variances) of unknown parameters. A moderate sampling interval (of $\Delta t = 0.01$) still gave good results for all considered sample sizes. However, when sampling at low frequency ($\Delta t = 0.1$), the discrete-time bias makes itself felt. Similar observations could be made in the second simulation setting. Here, we considered a stochastic FHN model, which is linear in the

parameter, but non-linear in the state. It is an example of a diffusion which violates common assumptions on the growth of the drift function of a diffusion, which ensure the validity of many results (existence of solution, absolute continuity of measures, convergence of the time-discretized estimator to its continuous-time analogue). The estimation results are accurate for high-frequency data, even at small sample sizes ($N = 20$), and also for moderate frequency data, the estimation was still convincing (except for the variance of the time-scale random effect). A considerable drop in accuracy occurred for low-frequency data (sampling interval of $\Delta t = 0.1$). If this method is to be used on such kind of data, algorithms for data imputation should be applied prior to estimation, in order to reduce the discrete-time bias.

The asymptotic normality of the MLE lends itself naturally to hypothesis testing of parameters by means of the Wald test. For the linear transfer model we estimate the false-positive rate, revealing a slight liberalism of the test procedure, and compute the test's power for different "true" values of parameters.

We have only studied the method's applicability to models with up to 17 parameters. Even in the case of an explicit likelihood, the MLE of the (unknown) covariance matrix of the random effects vector is implicit and estimation requires numerical optimization, which may hamper estimation when the parameter space has a high dimension. Another drawback is the already mentioned inherent discrete-time bias of the estimation procedure. It is negligible if observations are sampled at sufficiently high frequency, but for low-frequency observations, a severe bias occurs (cf. simulation results), which is to bear in mind in applications. A possible solution could be to impute data at time points in between observation times, and conduct the estimation on the enlarged data set (Bladt et al., 2016). Related to that is the problem of incomplete observations, where only some of the coordinates in the state space are observed, and an entire path of a completely unobserved (latent) coordinate should be inferred (Berg and Ditlevsen, 2013; Ditlevsen et al., 2014). Missing observations of one or more coordinates is not untypical for biological data. This, at a first step, prohibits application of the proposed estimation procedure, as it relies on the assumption of complete data observations. Such statistical recovery of hidden state coordinates remains a topic for future research.

Acknowledgments

The work is part of the Dynamical Systems Interdisciplinary Network, University of Copenhagen. Adeline Samson has been partially supported by the LabExPERSYVAL-Lab (ANR-11-LABX-0025-01).

References

- Aït-Sahalia, Y. (2002). Maximum likelihood estimation of discretely sampled diffusions: A closed-form approximation approach. *Econometrica* 70(1), 223–262.
- Beal, S. L. and L. B. Sheiner (1981). Estimating population kinetics. *Critical Reviews in Biomedical Engineering* 8(3), 195–222.
- Berg, R. W. and S. Ditlevsen (2013). Synaptic inhibition and excitation estimated via the time constant of membrane potential fluctuations. *Journal of Neurophysiology* 110(4), 1021–1034.
- Bladt, M., S. Finch, and M. Sørensen (2016). Simulation of multivariate diffusion bridges. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 78(2), 343–369.
- Bradley, R. A. and J. J. Gart (1962). The asymptotic properties of ML estimators when sampling from associated populations. *Biometrika* 49(1/2), 205–214.
- Davidian, M. and D. M. Giltinan (2003). Nonlinear models for repeated measurement data: an overview and update. *Journal of Agricultural, Biological, and Environmental Statistics* 8(4), 387–419.

- Delattre, M., V. Genon-Catalot, and A. Samson (2013). Maximum likelihood estimation for stochastic differential equations with random effects. *Scandinavian Journal of Statistics* 40(2), 322–343.
- Delattre, M., V. Genon-Catalot, and A. Samson (2015). Estimation of population parameters in stochastic differential equations with random effects in the diffusion coefficient. *ESAIM: Probability and Statistics* 19, 671–688.
- Delattre, M. and M. Lavielle (2013). Coupling the SAEM algorithm and the extended Kalman filter for maximum likelihood estimation in mixed-effects diffusion models. *Statistics and its Interface* 6(4), 519–532.
- Delyon, B., M. Lavielle, and E. Moulines (1999). Convergence of a stochastic approximation version of the EM algorithm. *Annals of Statistics*, 94–128.
- Ditlevsen, S. and A. De Gaetano (2005). Mixed effects in stochastic differential equation models. *REVSTAT-Statistical Journal* 3(2), 137–153.
- Ditlevsen, S., A. Samson, et al. (2014). Estimation in the partially observed stochastic morris–lecar neuronal model with particle filter and stochastic approximation methods. *The Annals of Applied Statistics* 8(2), 674–702.
- Ditlevsen, S., K.-P. Yip, and N.-H. Holstein-Rathlou (2005). Parameter estimation in a stochastic model of the tubuloglomerular feedback mechanism in a rat nephron. *Mathematical Biosciences* 194(1), 49–69.
- Donnet, S., J.-L. Foulley, and A. Samson (2010). Bayesian analysis of growth curves using mixed models defined by stochastic differential equations. *Biometrics* 66(3), 733–741.
- Donnet, S., A. Samson, et al. (2008). Parametric inference for mixed models defined by stochastic differential equations. *ESAIM P&S* 12, 196–218.
- Durham, G. B. and A. R. Gallant (2002). Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes. *Journal of Business & Economic Statistics* 20(3), 297–338.
- FitzHugh, R. (1955). Mathematical models of threshold phenomena in the nerve membrane. *The Bulletin of Mathematical Biophysics* 17(4), 257–278.
- Gabbay, D. M., P. Thagard, J. Woods, P. S. Bandyopadhyay, and M. R. Forster (2011). *Philosophy of statistics*, Volume 7. Elsevier.
- Große Ruse, M., L. R. Søndergaard, S. Ditlevsen, M. Damgaard, S. Fuglsang, J. T. Ottesen, and J. L. Madsen (2015). Absorption and initial metabolism of 75 se-l-selenomethionine: a kinetic model based on dynamic scintigraphic data. *British Journal of Nutrition* 114(10), 1718–1723.
- Guedj, J., R. Thiébaud, and D. Commenges (2007). Maximum likelihood estimation in dynamical models of HIV. *Biometrics* 63(4), 1198–1206.
- Guy, R., C. Larédo, and E. Vergu (2015). Approximation of epidemic models by diffusion processes and their statistical inference. *Journal of Mathematical Biology* 70(3), 621–646.
- Hoadley, B. (1971). Asymptotic properties of maximum likelihood estimators for the independent not identically distributed case. *The Annals of Mathematical Statistics*, 1977–1991.
- Hodgkin, A. L. and A. F. Huxley (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology* 117(4), 500.
- Ibragimov, I. A. and R. Z. Has’minskii (2013). *Statistical Estimation: Asymptotic Theory*, Volume 16. Springer Science.

- Jensen, A. C., S. Ditlevsen, M. Kessler, and O. Papaspiliopoulos (2012). Markov chain monte carlo approach to parameter estimation in the fitzhugh-nagumo model. *Physical Review E* 86(4), 041114.
- Jiang, D., J. Yu, C. Ji, and N. Shi (2011). Asymptotic behavior of global positive solution to a stochastic SIR model. *Mathematical and Computer Modelling* 54(1), 221–232.
- Keeling, M. J. and P. Rohani (2008). *Modeling infectious diseases in humans and animals*. Princeton University Press.
- Kermack, W. O. and A. G. McKendrick (1927). A contribution to the mathematical theory of epidemics. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, Volume 115, pp. 700–721. The Royal Society.
- Kessler, M., A. Lindner, and M. Sørensen (2012). *Statistical methods for stochastic differential equations*. CRC Press.
- Klim, S., S. B. Mortensen, N. R. Kristensen, R. V. Overgaard, and H. Madsen (2009). Population stochastic modelling (PSM)—an R package for mixed-effects models based on stochastic differential equations. *Computer Methods and Programs in Biomedicine* 94(3), 279–289.
- Kondepudi, D. and I. Prigogine (2014). *Modern thermodynamics: from heat engines to dissipative structures*. John Wiley & Sons.
- Lavielle, M. (2014). *Mixed effects models for the population approach: models, tasks, methods and tools*. CRC Press.
- Le Cam, L. (2012). *Asymptotic methods in statistical decision theory*. Springer Science, New York.
- Leander, J., J. Almquist, C. Ahlström, J. Gabrielsson, and M. Jirstrand (2015). Mixed effects modeling using stochastic differential equations: illustrated by pharmacokinetic data of nicotinic acid in obese zucker rats. *The AAPS Journal* 17(3), 586–596.
- Leander, J., T. Lundh, and M. Jirstrand (2014). Stochastic differential equations as a tool to regularize the parameter estimation problem for continuous time dynamical systems given discrete time measurements. *Mathematical Biosciences* 251, 54–62.
- Lehmann, E. L. and J. P. Romano (2006). *Testing statistical hypotheses*. Springer Science & Business Media.
- Lindstrom, M. J. and D. M. Bates (1990). Nonlinear mixed effects models for repeated measures data. *Biometrics*, 673–687.
- Lo, A. W. (1988). Maximum likelihood estimation of generalized Itô processes with discretely sampled data. *Econometric Theory* 4(2), 231–247.
- Lorenz, E. N. (1963). Deterministic nonperiodic flow. *Journal of the atmospheric sciences* 20(2), 130–141.
- Møller, J. B., R. V. Overgaard, H. Madsen, T. Hansen, O. Pedersen, and S. H. Ingwersen (2010). Predictive performance for population models using stochastic differential equations applied on data from an oral glucose tolerance test. *Journal of Pharmacokinetics and Pharmacodynamics* 37(1), 85–98.
- Mortensen, S. B., S. Klim, B. Dammann, N. R. Kristensen, H. Madsen, and R. V. Overgaard (2007). A matlab framework for estimation of NLME models using stochastic differential equations. *Journal of Pharmacokinetics and Pharmacodynamics* 34(5), 623–642.

- Murray, J. D. (2002). *Mathematical Biology I: An Introduction*, Volume 17 of *Interdisciplinary Applied Mathematics*. Springer, New York, NY, USA,.
- Nagumo, J., S. Arimoto, and S. Yoshizawa (1962). An active pulse transmission line simulating nerve axon. *Proceedings of the IRE* 50(10), 2061–2070.
- Overgaard, R. V., N. Jonsson, C. W. Tornøe, and H. Madsen (2005). Non-linear mixed-effects models with stochastic differential equations: implementation of an estimation algorithm. *Journal of pharmacokinetics and pharmacodynamics* 32(1), 85–107.
- Pedersen, A. R. (1995). A new approach to maximum likelihood estimation for stochastic differential equations based on discrete observations. *Scandinavian journal of statistics*, 55–71.
- Phillips, P. C. and J. Yu (2009). Maximum likelihood and gaussian estimation of continuous time models in finance. In *Handbook of financial time series*, pp. 497–530. Springer, New York.
- Picchini, U., A. De Gaetano, and S. Ditlevsen (2010). Stochastic differential mixed-effects models. *Scandinavian Journal of Statistics* 37(1), 67–90.
- Picchini, U. and S. Ditlevsen (2011). Practical estimation of high dimensional stochastic differential mixed-effects models. *Computational Statistics & Data Analysis* 55(3), 1426–1444.
- Pinheiro, J. and D. Bates (2006). *Mixed-effects models in S and S-PLUS*. Springer Science, New York.
- Ribba, B., N. H. Holford, P. Magni, I. Trocóniz, I. Gueorguieva, P. Girard, C. Sarr, M. Elishmereni, C. Kloft, and L. E. Friberg (2014). A review of mixed-effects models of tumor growth and effects of anticancer drug treatment used in population analysis. *CPT: Pharmacometrics & Systems Pharmacology* 3(5), 1–10.
- Tornøe, C. W., H. Agersø, E. N. Jonsson, H. Madsen, and H. A. Nielsen (2004). Non-linear mixed-effects pharmacokinetic/pharmacodynamic modelling in NLME using differential equations. *Computer Methods and Programs in Biomedicine* 76(1), 31–40.
- Tornøe, C. W., R. V. Overgaard, H. Agersø, H. A. Nielsen, H. Madsen, and E. N. Jonsson (2005). Stochastic differential equations in NONMEM: implementation, application, and comparison with ordinary differential equations. *Pharmaceutical Research* 22(8), 1247–1258.
- Van der Vaart, A. W. (2000). *Asymptotic statistics*, Volume 3. Cambridge University Press.
- Wang, J. (2007). EM algorithms for nonlinear mixed effects models. *Computational Statistics & Data Analysis* 51(6), 3244–3256.
- Wolfinger, R. (1993). Laplace’s approximation for nonlinear mixed models. *Biometrika* 80(4), 791–795.