

CONSISTENT MAXIMUM LIKELIHOOD ESTIMATION USING SUBSETS WITH APPLICATIONS TO MULTIVARIATE MIXED MODELS

BY KARL OSKAR EKVALL* AND GALIN L. JONES**

*School of Statistics, University of Minnesota, *ekvall@umn.edu; **galin@umn.edu*

We present new results for consistency of maximum likelihood estimators with a focus on multivariate mixed models. Our theory builds on the idea of using subsets of the full data to establish consistency of estimators based on the full data. It requires neither that the data consist of independent observations, nor that the observations can be modeled as a stationary stochastic process. Compared to existing asymptotic theory using the idea of subsets, we substantially weaken the assumptions, bringing them closer to what suffices in classical settings. We apply our theory in two multivariate mixed models for which it was unknown whether maximum likelihood estimators are consistent. The models we consider have nonstochastic predictors and multivariate responses which are possibly mixed-type (some discrete and some continuous).

1. Introduction. Mixed models are frequently used in applications and have been the subject of numerous articles and books [7, 15, 20]. Yet, it was unknown until recently whether maximum likelihood estimators (MLEs) are consistent even in some simple generalized linear mixed models (GLMMs) [17]. What complicates proving consistency in some mixed models is the dependence among response variables induced by certain random effect designs. Of course, not all types of dependence between responses are problematic—there is a vast literature on maximum likelihood estimation with dependent observations [1, 6, 13, 14, 25, 27, 30]. But, as we will discuss in more detail below, for some commonly used random effect designs such as those with crossed random effects, existing conditions for consistency of MLEs are hard to verify [17]. In a few GLMMs with crossed random effects, consistency has been proved using a novel argument that relates the likelihood for the full data to that of a subset consisting of independent and identically distributed (i.i.d.) random variables, “the subset argument” [16].

Fundamentally, however, the issue is not unique to GLMMs or even mixed models; any other parametric model appropriate for the same settings may present similar difficulties. Accordingly, it was recognized in the first work on consistency using subsets that the idea has the potential to be extended to more general models [16]. We address this by establishing weaker conditions, based in part on the use of subsets, that are sufficient for consistency of MLEs, without assuming a particular model. They help explain formally what makes the subset argument work, why it is useful in some settings where more classical ones are not, and when it can fail. We illustrate the usefulness of our conditions by proving consistency of MLEs in two multivariate GLMMs (MGLMMs) to which existing theory has not been applied successfully.

To fix ideas, let Θ denote a parameter set, f_{θ}^n a joint density for the random vector $Y = (Y_1, \dots, Y_n)$, and θ^0 the “true” parameter. Let also $L_n(\theta; Y) = f_{\theta}^n(Y)/f_{\theta^0}^n(Y)$ and

Received October 2018; revised February 2019.

MSC2010 subject classifications. Primary 62F12; secondary 62J12.

Key words and phrases. Maximum likelihood, consistency, MGLMM, GLMM, crossed random effects, subset argument.

$\Lambda_n(\theta; Y) = \log L_n(\theta; Y)$. If Θ is a finite set, then since $L_n(\theta^0; Y) = 1$, a necessary and sufficient condition for consistency of MLEs is that, as $n \rightarrow \infty$,

$$(1) \quad \mathbb{P}(L_n(\theta; Y) \geq 1) \rightarrow 0 \quad \text{for all } \theta \neq \theta^0.$$

When Θ is not a finite set, (1) needs to be amended by a uniformity argument to be sufficient, but the main ideas are the same. There are many ways to establish (1). With i.i.d. observations and regularity conditions, (1) or stronger results follow from the law of large numbers applied to $n^{-1}\Lambda_n(\theta; Y)$ [5, 8, 10, 28]. If Y is a more general stochastic process, $\Lambda_n(\theta; Y)$ may still, suitably scaled, satisfy an ergodic theorem, leading again to (1) under regularity conditions. In the literature on maximum likelihood estimation with dependent observations, it is often assumed that some such limit law holds, either for $\Lambda_n(\theta; Y)$ or its derivatives [6, 13, 14], or that the moments of $\Lambda_n(\theta; Y)$ converge in an appropriate way [1, 25]. Unfortunately, in many practically relevant settings, it is not clear that any such convergence holds and proving that it does is arguably the main obstacle to establishing consistency of MLEs. Let us illustrate using an MGLMM, commonly considered both in statistics and applied sciences [3, 4, 11, 19, 29].

Let $X = [x_1, \dots, x_n]^T \in \mathbb{R}^{n \times p}$ be a matrix of nonstochastic predictors, $Z = [z_1, \dots, z_n]^T \in \mathbb{R}^{n \times r}$ a nonstochastic design matrix and $U \in \mathbb{R}^r$ a multivariate normal vector of random effects, with mean zero and covariance matrix Σ . For the MGLMM, $\Theta \subseteq \mathbb{R}^d$, for some $d \geq 1$, $\beta = \beta(\theta)$, and $\Sigma = \Sigma(\theta)$. The responses Y_1, \dots, Y_n are conditionally independent given U , with conditional exponential family densities

$$f_{\theta,i}(y_i | u) = k_i(y_i, \tau_i) \exp\left(\frac{y_i[x_i^T \beta + z_i^T u] - c_i(x_i^T \beta + z_i^T u)}{\tau_i}\right),$$

where, for $i = 1, \dots, n$, c_i is the conditional cumulant function, τ_i a dispersion parameter and $k_i(y_i, \tau_i)$ ensures $f_{\theta,i}(y_i | u)$ integrates to one. Conditional independence implies $f_{\theta}^n(y | u) = \prod_{i=1}^n f_{\theta,i}(y_i | u)$. Several of the responses could be from the same subject, hence the “multivariate,” and they can be of mixed type, some continuous and some discrete, for example.

The dependence among the linear predictors is easily characterized since $X\beta + ZU \sim \mathcal{N}(X\beta, Z\Sigma Z^T)$. The relevant density for maximum likelihood estimation, however, is the marginal density,

$$(2) \quad f_{\theta}^n(y) = \int_{\mathbb{R}^r} f_{\theta}^n(y | u) \phi_{\theta}^r(u) du,$$

where ϕ_{θ}^r denotes the r -dimensional multivariate normal density with mean zero and covariance matrix $\Sigma = \Sigma(\theta)$. The density $f_{\theta}^n(y)$ typically does not admit a closed form expression. Moreover, the dependence among responses it implies is in general less transparent than that among the linear predictors. What we can say in general is that two responses are dependent only if their corresponding linear predictors are. That is, response component i and j are independent if $z_i^T \Sigma z_j = 0$.

It is convenient if $Z\Sigma Z^T$ is, upon possible reordering of the responses, block diagonal since in that case the full vector of responses can be partitioned into independent subvectors. If these are of fixed length as n grows, then one is back in the classical setting where the full data consists only of an increasing number of independent vectors. This setting is common to many articles on asymptotic theory in mixed models [12, 22, 23, 26]. Unfortunately, in applications the number of independent response vectors—the number of diagonal blocks in $Z\Sigma Z^T$ —is often small. For example, Sung and Geyer [26] note that in the famous salamander data [18] there are 3 independent vectors, each of length 120. Thus, in their notation there are $n = 3$ independent observations, but in our notation there are $n = 3 \times 120 = 360$ possibly

dependent observations. It seems more reasonable, then, to consider large sample properties that do not assume the response vector Y consists only of a large number of independent subvectors. The type of limiting process we consider has, in the context of mixed models, previously only been investigated carefully in special cases that do not allow for predictors or mixed-type responses [16, 21]. To be sure, Jiang's [16] general theory does allow for predictors, but the specific applications do not. Due to the inclusion of nonstochastic predictors, uniform convergence results needed for our theory, which in simpler models can be established using classical laws of large numbers, are in one of our applications verified using empirical process theory.

The intuition behind the usefulness of the subset argument can be understood by considering the following simple LMM with crossed random effects. Suppose $Y_{i,j} = \theta + U_i^{(1)} + U_j^{(2)} + E_{i,j}$, where $U_i^{(1)}$, $U_j^{(2)}$ and $E_{i,j}$ are all i.i.d. standard normal, $i = 1, \dots, N$, $j = 1, \dots, N$. It is easy to check that the $Y_{i,j}$'s cannot be partitioned into independent subsets. However, there are many subsets that, even though there is dependence among them, consist of independent random variables. For example, the two subsets $(Y_{1,1}, Y_{2,2}, \dots, Y_{N,N})$ and $(Y_{1,2}, Y_{2,3}, \dots, Y_{N-1,N})$ are dependent, but taken separately they both consist of i.i.d. random variables. The MLE of θ based on either subset, that is, a subset sample mean, is consistent as $N \rightarrow \infty$. Intuitively, then, the MLE based on all of the N^2 variables should be, too. Of course, the subset argument is not needed to prove that in this simple example, but the intuition is the same for models where a direct proof is harder. How to formalize this intuition in more general models, without actually having to require the subset components to be either independent or identically distributed, is the topic of Section 2.

After developing a general theory, we will return to mixed models in Section 3 and prove consistency of MLEs in two MGLMMs. The first is a longitudinal linear mixed model with autoregressive temporal dependence and crossed random effects. In this model, the integral in (2) has a closed form solution which makes it easier to demonstrate some of the intricacies of the subset argument. For the second MGLMM we consider, which includes both binary and continuous responses, $f_\theta^n(y)$ does not admit a closed form expression. The subset argument is especially useful in this setting since the considered subsets have likelihoods that are more amenable to analysis than that of the full data.

The rest of the paper is organized as follows. We develop theory using subsets in Section 2. In Section 3, we apply the theory from Section 2 to two MGLMMs. Section 4 contains a brief discussion of our results. Many technical details are deferred to the [Appendix](#) and the [Supplementary Material](#) [9].

2. Consistency using subsets of the full data. Recall that $Y = (Y_1, \dots, Y_n)$ denotes a collection of random variables and let $W = (W_1, \dots, W_m)$ be a collection of random variables that form a subset of those in Y , that is, $\{W_1, \dots, W_m\} \subseteq \{Y_1, \dots, Y_n\}$. We will henceforth call W a subcollection of Y to avoid confusion with other subsets introduced later. The main results in this section give conditions for when subcollections can be used to prove consistency of maximizers of $L_n(\theta; Y)$. Unless otherwise noted, all convergence statements are as n tends to infinity and the number of elements in a subcollection, $m = m(n)$, tends to infinity as a function of n .

All discussed random variables are defined on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$, with the elements of Ω denoted ω . The parameter set Θ is assumed to be a subset of a metric space $(\mathcal{T}, d_{\mathcal{T}})$. We write, for any $t \in \mathcal{T}$ and $\delta > 0$, $B_\delta(t) = \{t' \in \mathcal{T} : d_{\mathcal{T}}(t, t') < \delta\}$. For any $A \subseteq \mathcal{T}$, \bar{A} denotes its closure and ∂A its boundary. We assume the true parameter θ^0 is the same for all n but the joint density $f_\theta^n(y)$ of Y , against a dominating, σ -finite product measure $\nu = \nu_n$, can depend on n in an arbitrary manner. In particular, our setting allows for

a triangular array of responses, $Y_{n,1}, \dots, Y_{n,n}$, though for convenience we do not make this explicit in the notation.

By θ^0 being the true parameter, we mean that $P(Y \in A) = \int_A f_{\theta^0}^n(y) \nu(dy)$ for any measurable A in the range space of Y . That is, expectations and probabilities with respect to P are the same as those taken with respect to distributions indexed by θ^0 . Densities for the subcollection and its components are denoted by g in place of f ; for example, $L_m(\theta; W) = g_{\theta}^m(W)/g_{\theta^0}^m(W)$.

We will establish the following sufficient condition for consistency of maximizers of $L_n(\theta; Y)$:

$$(3) \quad P\left(\sup_{\theta \in \Theta \cap B_{\varepsilon}(\theta^0)^c} L_n(\theta; Y) \geq 1\right) \rightarrow 0 \quad \forall \varepsilon > 0.$$

That is, the probability that there exists a maximizer of the likelihood outside an arbitrarily small ball around the true parameter tends to zero. We now discuss the use of subcollections and the assumptions used to achieve (3), which eventually leads to the main results in Theorems 2.3 and 2.4 presented at the end of the section.

The appeal of using subcollections to prove (3), instead of directly working with the full data likelihood $L_n(\theta; Y)$, can be explained using the following lemma.

LEMMA 2.1. *For every $c \in (0, \infty)$, $\theta \in \Theta$, and subcollection W , P -almost surely,*

$$P(L_n(\theta; Y) \geq c \mid W) \leq c^{-1} L_m(\theta; W).$$

Versions of Lemma 2.1 are well known [16, 17], but the Supplementary Material [9] contains a proof for completeness. From the lemma it follows that if $L_m(\theta; W) \rightarrow 0$, then $E[P(L_n(\theta; Y) \geq 1 \mid W)] = P(L_n(\theta; Y) \geq 1) \rightarrow 0$ by dominated convergence. That is, up to a uniformity argument, (3) can be established by showing that the likelihood of the subcollection converges to zero in probability, outside of a neighborhood of θ^0 . Uniform versions of that convergence will play a crucial role in our results.

DEFINITION 2.1. We say that a subset $A \subseteq \Theta$ is identified by a subcollection W if $\sup_{\theta \in A} L_m(\theta; W) \xrightarrow{P} 0$. If $\sup_{\theta \in A} L_m(\theta; W) = O_P(a_n)$ for some sequence of constants $\{a_n\}$, $n = 1, 2, \dots$, we call a_n an identification rate.

To understand this definition better, consider the case where the subcollection W consists of m i.i.d. random variables with common marginal density $g_{\theta,1}$. Suppose also that there is no $\theta \in A$ for which $g_{\theta,1} = g_{\theta^0,1}$ ν -almost everywhere. That is, θ^0 is an identified parameter in the classical sense if we restrict attention to the parameter set $A \cup \{\theta^0\}$. Then, under regularity conditions [10], Theorems 16 and 17, one has $\sup_{\theta \in A} E[\Lambda_m(\theta; W)] < 0$ and, by a uniform strong law of large numbers,

$$\lim_{m \rightarrow \infty} m^{-1} \sup_{\theta \in A} |\Lambda_m(\theta; W) - E[\Lambda_m(\theta; W)]| = 0.$$

Using this, it is straightforward to show that A is identified by W with an identification rate that is exponentially fast in m . That is, with i.i.d. components and regularity conditions, the classical definition of an identified parameter implies identification in the sense of Definition 2.1. However, we want to allow for subcollections that do not consist of i.i.d. components, and in that case the classical definition is not as useful. For example, we have independent but not identically distributed components in one of our MGLMMs. In this and more general cases, a parameter could be identified in the classical sense for all sample sizes n , but loosely speaking, the difference between the distributions for W indexed by some $\theta \in A$

and that indexed by θ^0 could vanish asymptotically, preventing W from identifying A in our sense. Finally, notice also that A being identified by W is essentially equivalent to MLEs based on W with the restricted parameter set $A \cup \{\theta^0\}$ being consistent.

We can now be more precise about how to use subcollections to establish (3). The strategy is to first find a subcollection W that identifies $B_\varepsilon(\theta^0)^c \cap \Theta$ for every $\varepsilon > 0$, and then use Lemma 2.1 to get the convergence for the full likelihood in (3). For this strategy to be useful, showing that W identifies $B_\varepsilon(\theta^0)^c \cap \Theta$ has to be easier than showing that Y does since the latter would directly imply (3). That is, one has to be able to pick out a subcollection with more convenient properties than the full data. Our applications in Section 3 illustrate how this can be done.

It is useful to allow for several subcollections $W^{(i)}$, consisting of m_i components, and subsets $A_i, i = 1, \dots, s$. By doing so, different subcollections can be used to identify different subsets of the parameter set. For example, if the parameter set is a product space, as is common in applications, then different subcollections can be used to, loosely speaking, identify different elements of the parameter vector. Assumption 1 makes precise what we need to identify $\Theta \cap B_\varepsilon(\theta^0)^c$ using several subcollections.

ASSUMPTION 1. For every small enough $\varepsilon > 0$, there are subsets $A_i = A_i(\varepsilon) \subseteq \Theta$ and corresponding subcollections $W^{(i)}, i = 1, \dots, s$, such that $\bigcup_{i=1}^s A_i \supseteq \Theta \cap B_\varepsilon(\theta^0)^c$ and each A_i is identified by $W^{(i)}$ with some identification rate $a_{n,i}, n = 1, 2, \dots, i = 1, \dots, s$.

This assumption is somewhat similar to assumptions A2 and A3 made by Jiang [16], which are also assumptions about parameter identification using several subcollections. However, those assumptions are stated in terms of $E(\Lambda_{m_i}(\theta; W^{(i)}))$ and $\text{Var}(\Lambda_{m_i}(\theta; W^{(i)}))$, $i = 1, \dots, s$. The fact that we do not have to assume anything about the variances of the log-likelihood ratios is an important improvement. For example, if subcollection i consists of i.i.d. components, the convergence of $m_i^{-1} \Lambda_{m_i}(\theta; W^{(i)})$ is immediate from the law of large numbers, but calculating its variance may be difficult.

For finite parameter sets, Assumption 1 is enough to give consistency of MLEs via Lemma 2.1. For more general cases, we also need to control the regularity of the log-likelihood for the full data. The following two assumptions are made to ensure that the uniformity of the convergence detailed in Assumption 1 and Definition 2.1 carries over to $\Lambda_n(\theta; Y)$, in the sense of (3).

ASSUMPTION 2. For every $i \in \{1, \dots, s\}$ and $n \in \{1, 2, \dots\}$, $\Lambda_n(\theta; Y)$ is \mathbf{P} -almost surely Lipschitz continuous in θ on the A_i defined in Assumption 1; that is, there exists a random variable $K_{n,i}$ not depending on θ such that, \mathbf{P} -almost surely and for every $\theta, \theta' \in A_i$,

$$|\Lambda_n(\theta; Y) - \Lambda_n(\theta'; Y)| \leq K_{n,i} d_{\mathcal{T}}(\theta, \theta').$$

ASSUMPTION 3. Each A_i from Assumption 1 can be covered by $M_{n,i}$ balls of radius $\delta_{n,i}$ such that

$$K_{n,i} \delta_{n,i} \xrightarrow{\mathbf{P}} 0 \quad \text{and} \quad M_{n,i} a_{n,i} \rightarrow 0,$$

where $a_{n,i}$ and $K_{n,i}, i = 1, \dots, s, n = 1, 2, \dots$, are the same as in Assumptions 1 and 2, respectively.

There is an interplay between Assumption 1–3 where the rates in Assumption 1 need to be sufficiently fast in comparison to the growth of the Lipschitz constants in Assumption 2; Assumption 3 specifies how the rates should align. Why these rates work will be clear from

the proof of Lemma 2.2, but the intuition is as follows. To get uniformity in θ , we cover A_1 (say) with balls small enough that the likelihood is approximately constant on them, so that one can work pointwise in θ in each ball instead of uniformly. If the likelihood changes much on A_1 in the sense that $K_{n,1}$ is large, then the radius $\delta_{n,1}$ needs to be small; this is what the first rate condition says. The second rate condition illustrates that there is a price for picking small balls, namely that if many balls are needed to cover A_1 , then fast identification rates are needed.

The assumptions give us the convergence in (3), and consequently, the following lemma.

LEMMA 2.2. *If Assumptions 1–3 hold, then the probability that there exists a global maximizer of $\Lambda_n(\theta; Y)$ in $B_\varepsilon(\theta^0)^c \cap \Theta$ tends to zero as $n \rightarrow \infty$, for every $\varepsilon > 0$.*

PROOF. We give an outline here and a detailed proof in the Appendix. Without loss of generality, we may assume $s = 1$, so there is one subcollection W that identifies $A = \Theta \cap B_\varepsilon(\theta^0)^c$, for arbitrary, small $\varepsilon > 0$, with rate a_n . It suffices to prove that $\mathbb{P}(\sup_{\theta \in A} L_n(\theta; Y) \geq 1) \rightarrow 0$. For $j = 1, \dots, M_n$ let θ^j be a point in the intersection of A and the j th ball in the cover of A given by Assumption 3. Some algebra and Assumption 2 gives

$$\mathbb{P}\left(\sup_{\theta \in A} L_n(\theta; Y) \geq 1\right) \leq \mathbb{P}\left(\max_{j \leq M_n} L_n(\theta^j; Y) \geq 1/2\right) + \mathbb{P}(e^{K_n \delta_n} \geq 2).$$

The second term is $o(1)$ by Assumption 3. It remains to deal with the first. By conditioning on the subcollection and using Lemma 2.1, one gets

$$\mathbb{P}\left(\max_{j \leq M_n} L_n(\theta^j; Y) \geq 1/2 \mid W\right) \leq 2M_n \sup_{\theta \in A} L_m(\theta; W).$$

The right-hand side is $o_{\mathbb{P}}(1)$ by Assumption 3, so the expectation of the left-hand side is $o(1)$ by dominated convergence, which completes the proof. \square

We will use Lemma 2.2 to establish both a Wald-type consistency, meaning consistency of sequences of global maximizers of $L_n(\theta; Y)$ and a Cramér-type consistency, meaning consistency of a sequence of roots to the likelihood equations $\nabla \Lambda_n(\theta; Y) = 0$. It follows almost immediately from the lemma that if $L_n(\theta; Y)$ has a global maximizer $\hat{\theta}_n$, \mathbb{P} -almost surely for every n , then $\hat{\theta}_n \xrightarrow{\mathbb{P}} \theta^0$. In particular, if Θ is compact one gets Wald-type consistency with an additional continuity assumption. Since Assumption 2 implies $L_n(\theta; Y)$ is continuous at every point except possibly θ^0 , assuming continuity also at the unknown θ^0 should be insignificant in any application of interest.

THEOREM 2.3. *If Θ is compact, $L_n(\theta; Y)$ is \mathbb{P} -almost surely continuous on Θ for every n , and Assumptions 1–3 hold, then a maximizer $\hat{\theta}_n$ of $L_n(\theta; Y)$ exists \mathbb{P} -almost surely for every n , and $\hat{\theta}_n \xrightarrow{\mathbb{P}} \theta^0$ for any sequence of such maximizers.*

PROOF. Since continuous functions attain their suprema on compact sets, $L_n(\theta; Y)$ has a maximizer on Θ , \mathbb{P} -almost surely. By Lemma 2.2, all maximizers are in $B_\varepsilon(\theta^0)$ with probability tending to one, for all small enough $\varepsilon > 0$. \square

Though compactness is a common assumption [14, 31], it is sometimes too restrictive or even unnecessary. If $L_n(\theta; Y)$, or more commonly $\Lambda_n(\theta; Y)$, is strictly concave in θ on a convex Θ , then it is enough to verify the assumptions on a neighborhood of θ^0 (cf. Theorem 2.4) to get consistency of the unique global maximizer. However, a global maximizer need not exist even as $n \rightarrow \infty$, or perhaps the assumptions cannot be verified for other reasons. With a few additional assumptions, Lemma 2.2 can then be used to get the weaker Cramér-type consistency, which also only requires verifying assumptions for neighborhoods of θ^0 .

THEOREM 2.4. *If $\Theta \subseteq \mathbb{R}^d$ for some $d \geq 1$, $L_n(\theta; Y)$ is almost surely differentiable in θ on a neighborhood of an interior θ^0 for every n , and Assumptions 1–3 hold with Θ replaced by $\bar{B}_\varepsilon(\theta^0)$ for all small enough $\varepsilon > 0$, then, with probability tending to one as $n \rightarrow \infty$, there exists a local maximizer of $L_n(\theta; Y)$, and hence a root to the likelihood equation $\nabla \Lambda_n(\theta; Y) = 0$, in $B_\varepsilon(\theta^0)$, for all small enough $\varepsilon > 0$.*

PROOF. Since θ^0 is interior, we may assume $\varepsilon > 0$ is small enough that all points of $\bar{B}_\varepsilon(\theta^0)$ are interior. Almost sure differentiability of $L_n(\theta; Y)$ implies almost sure continuity. Thus, $L_n(\theta; Y)$ attains a local maximum on the compact $\bar{B}_\varepsilon(\theta^0)$, \mathbb{P} -almost surely. By Lemma 2.2, with probability tending to one, there are no such maximizers in $\bar{B}_\varepsilon(\theta^0) \setminus B_\varepsilon(\theta^0) = \partial B_\varepsilon(\theta^0)$. Thus, with probability tending to one, there exists a local maximizer in $B_\varepsilon(\theta^0)$. Since $L_n(\theta; Y)$, and hence $\Lambda_n(\theta; Y)$ is \mathbb{P} -almost surely differentiable, any such maximizer must be a root to the likelihood equation $\nabla \Lambda_n(\theta; Y) = 0$. \square

In the next section, we apply Theorem 2.4 to two special cases of the MGLMM described in Section 1. We also discuss in more detail how to think about the subcollections and subsets in specific models.

3. Application to multivariate mixed models.

3.1. Longitudinal linear mixed model. The first model we consider is an extension of a variance components model that has been studied previously [21]. In addition to dependence between subjects induced by crossed random effects, the model incorporates autoregressive temporal dependence between measurements from the same subject. To make the discussion clearer, we assume easy-to-specify fixed and random effect structures. This allows us to focus on the core issues, that is, on how to select subcollections and subsets that can be used to verify the conditions of our theory. Our model includes a baseline mean and a treatment effect. A general fixed effect design matrix could be treated the same way as in our second example, discussed in Section 3.2. Before establishing consistency, we discuss the model definition and how to select appropriate subcollections.

Suppose for subjects (i, j) , $i = 1, \dots, N$ and $j = 1, \dots, N$, and time points $t = 1, \dots, T$, we observe the response $Y_{i,j,t}$, where for convenience we assume both N and T are even. Let the stacked vector of responses be

$$Y = [Y_{1,1,1}, \dots, Y_{1,1,T}, Y_{1,2,1}, \dots, Y_{N,N,T}]^T \in \mathbb{R}^n, \quad n = TN^2.$$

Recall from the Introduction that the MGLMM is specified by the conditional distribution $f_\theta^n(y | u)$ and the distribution of the random effects, $\phi_\theta^r(u)$. For a linear mixed model, we let $f_\theta^n(y | u)$ be the density of a multivariate normal distribution with mean $X\beta + Zu$ and covariance matrix $\theta_3 I_n$, $\theta_3 > 0$, where the two components of $\beta = [\theta_1, \theta_2]^T \in \mathbb{R}^2$ are a baseline mean and a treatment effect, respectively, and I_k denotes the $k \times k$ identity matrix. Note, in the notation of the Introduction, the dispersion parameter in the conditional distribution is $\tau_i = \theta_3$, for all i . We treat θ_3 as a parameter to be estimated and not as known, which is otherwise common in the literature.

Let h_n be a vector of zeros and ones where the i th element is one if it corresponds to an observation in time $t \leq T/2$ and zero otherwise and let 1_n denote an n -vector of ones. We take $X = [1_n, h_n] \in \mathbb{R}^{TN^2 \times 2}$, which corresponds to a treatment being applied in the first half of the experiment. Unless T is fixed, which we do not assume, this setup implies the predictors change with n . Indeed, as T grows, a particular observation can go from being made in the latter half of the experiment to the earlier half. Thus, the responses form a triangular array.

Partition U into three independent subvectors, $U^{(1)} \sim \mathcal{N}(0, \theta_4 I_N)$, $U^{(2)} \sim \mathcal{N}(0, \theta_5 I_N)$ and $U^{(3)} \sim \mathcal{N}(0, \theta_6 I_{N^2} \otimes \Psi)$, where $\Psi = (\Psi_{i,j}) = (\theta_7^{|i-j|})$ is a first-order autoregressive correlation matrix, $\theta_i > 0$, $i = 4, 5, 6$ and $\theta_7 \in (-1, 1)$. We will use $U^{(1)}$ and $U^{(2)}$ as crossed random effects, inducing dependence between subjects, and $U^{(3)}$ to get temporal dependence within subjects. To that end, let $Z_1 = I_N \otimes 1_N \otimes 1_T$, $Z_2 = 1_N \otimes I_N \otimes 1_T$ and $Z = [Z_1, Z_2, I_{TN^2}]$. Then, with $J_k = 1_k 1_k^T$, the covariance matrix of the linear predictors $X\beta + ZU$ is

$$Z\Sigma Z^T = \theta_4 I_N \otimes J_{NT} + \theta_5 J_N \otimes I_N \otimes J_T + \theta_6 I_{N^2} \otimes \Psi.$$

More transparently, for the elements of $E(Y | U) = X\beta + ZU$, it holds that

$$\text{cov}[E(Y_{i,j,t} | U), E(Y_{i',j',t'} | U)] = \begin{cases} \theta_4 + \theta_5 + \theta_6 \theta_7^{|t-t'|} & i = i', j = j', \\ \theta_4 & i = i', j \neq j', \\ \theta_5 & i \neq i', j = j', \\ 0 & \text{otherwise.} \end{cases}$$

The marginal density $f_\theta^n(y)$ admits a closed form expression in this example. Specifically, the marginal distribution for Y is multivariate normal with mean $m(\theta) = X\beta(\theta)$ and covariance matrix $C(\theta) = \theta_3 I_{TN^2} + Z\Sigma(\theta)Z^T$. Note that the structure of $C(\theta)$ is similar to that of the covariance matrix of the linear predictors just discussed. In particular, there are many zeros in the covariance matrix $C(\theta)$, that is, there are many independent observations, but Y cannot be partitioned into independent vectors.

3.1.1. *Subcollection selection.* The model definitions imply that $\Theta = \mathbb{R} \times \mathbb{R} \times (0, \infty) \times (0, \infty) \times (0, \infty) \times (0, \infty) \times (-1, 1)$, a subset of \mathbb{R}^7 , which we equip with the metric induced by the Euclidean norm $\| \cdot \|$. We write $\theta = (\theta_1, \dots, \theta_7)$.

Subcollections are selected for the purpose of verifying Assumption 1. The main idea guiding selection is suggested by the fact that identification follows, under regularity conditions, if the subcollection’s log-likelihood satisfies a law of large numbers. We will use $s = 2$ such subcollections and require that they together identify θ in the classical sense. By this, we mean that, letting ν_θ^i denote the distribution of subcollection i implied by parameter θ ,

$$\{\theta \in \Theta : \nu_\theta^1 = \nu_{\theta_0}^1\} \cap \{\theta \in \Theta : \nu_\theta^2 = \nu_{\theta_0}^2\} = \{\theta^0\}.$$

With these properties in mind, we take $W^{(1)}$ to consist of the vectors

$$W_i^{(1)} = (Y_{2i-1,2i-1,1}, Y_{2i,2i,T}) \in \mathbb{R}^2, \quad i = 1, \dots, N/2.$$

Because these vectors do not share any random effects, they are independent. In fact, they are i.i.d. multivariate normal with common mean $m_1(\theta) = [\theta_1 + \theta_2, \theta_1]^T$ and common covariance matrix $C_1(\theta) = I_2(\theta_3 + \theta_4 + \theta_5 + \theta_6)$. Clearly, θ_1 and θ_2 are identified in the classical sense by this subcollection, but not $\theta_3, \dots, \theta_7$. Note that even though the predictors, and hence the distributions, do not change with N for this subcollection, it is strictly speaking a triangular array unless T is fixed.

To identify the remaining parameters, take $W^{(2)}$ to consist of the vectors

$$W_i^{(2)} = (Y_{2i-1,2i-1,1}, Y_{2i-1,2i-1,2}, Y_{2i-1,2i-1,3}, Y_{2i-1,2i,1}, Y_{2i,2i-1,1}),$$

$i = 1, \dots, N/2$. These are also i.i.d. multivariate normal, with common mean $m_2(\theta) = (\theta_1 + \theta_2)1_5$ and common covariance matrix

$$C_2(\theta) = \begin{bmatrix} \sum_{i=3}^6 \theta_i & \theta_4 + \theta_5 + \theta_6\theta_7 & \theta_4 + \theta_5 + \theta_6\theta_7^2 & \theta_4 & \theta_5 \\ \cdot & \sum_{i=3}^6 \theta_i & \theta_4 + \theta_5 + \theta_6\theta_7 & \theta_4 & \theta_5 \\ \cdot & \cdot & \sum_{i=3}^6 \theta_i & \theta_4 & \theta_5 \\ \cdot & \cdot & \cdot & \sum_{i=3}^6 \theta_i & 0 \\ \cdot & \cdot & \cdot & \cdot & \sum_{i=3}^6 \theta_i \end{bmatrix}.$$

It is straightforward to check that $C_2(\theta) = C_2(\theta')$ implies $\theta_i = \theta'_i, i = 3, \dots, 7$.

In summary, the two subcollections together identify θ in the classical sense. Moreover, since both subcollections consist of i.i.d. multivariate normal vectors, their log-likelihoods satisfy a law of large numbers as $N \rightarrow \infty$. With this we are equipped to verify that Assumptions 1–3 hold locally, leading to the main result of the section in Theorem 3.4.

3.1.2. Consistency. The purpose of this section is to verify the conditions of Theorem 2.4. The interesting part of that is to check that Assumptions 1–3 hold with Θ replaced by $\bar{B}_\varepsilon(\theta^0)$, for all small enough $\varepsilon > 0$. For this purpose, we will first prove two lemmas that roughly correspond to Assumptions 1 and 2. The limiting process we consider is that N tends to infinity while T can be fixed or tend to infinity with N , at rates discussed below. Thus, the statements $n \rightarrow \infty$ and $N \rightarrow \infty$ are equivalent. We will need the following result which is proved in the Supplementary Material [9].

PROPOSITION 3.1. *If Θ is compact, $L_{m_i}(\theta; w^{(i)})$ is continuous in θ on Θ for every $w^{(i)}$ in the support of $W^{(i)}, i = 1, \dots, s$, and $\bigcap_{i=1}^s \{\theta \in \Theta : v_\theta^i = v_{\theta^0}^i\} = \{\theta^0\}$, then for any $\varepsilon > 0$ there are compact sets $\tilde{A}_1, \dots, \tilde{A}_s$ such that $\{\theta \in \Theta : v_\theta^i = v_{\theta^0}^i\} \cap \tilde{A}_i = \emptyset, i = 1, \dots, s$, and $\bigcup_{i=1}^s \tilde{A}_i = \Theta \cap B_\varepsilon(\theta^0)^c$.*

Note, when applying the proposition in the present application, $m_i = N, s = 2$, and Θ is replaced by $\bar{B}_\varepsilon(\theta^0)$. As we will see in the proof of the following lemma, the proposition is useful because the \tilde{A}_i 's it gives are compact. Lemma 3.2 formalizes verification of Assumption 1.

LEMMA 3.2. *If θ^0 is an interior point of Θ , then for all small enough $\varepsilon > 0$ there exist subsets A_1 and A_2 such that $A_1 \cup A_2 = \partial B_\varepsilon(\theta^0)$:*

1. $N^{-1} \sup_{\theta \in A_i} \mathbf{E}[\Lambda_{N/2}(\theta; W^{(i)})] = \sup_{\theta \in A_i} \mathbf{E}[\Lambda_1(\theta; W_1^{(i)})]/2 < 0$,
2. \mathbf{P} -almost surely, $N^{-1} \sup_{\theta \in A_i} |\Lambda_{N/2}(\theta; W^{(i)}) - \mathbf{E}[\Lambda_{N/2}(\theta; W^{(i)})]| \rightarrow 0$, and, consequently;
3. A_i is identified by $W^{(i)}$ with an identification rate $a_{n,i} = o(e^{-\epsilon N(n)})$ for some $\epsilon > 0, i = 1, 2$.

PROOF. We give an outline here and a detailed proof in the Supplementary Material [9]. It is easy to check that the requirements of Proposition 3.1 are satisfied with Θ replaced by $\bar{B}_\varepsilon(\theta^0)$. By taking the A_i 's to be the \tilde{A}_i 's given by Proposition 3.1, proving points 1–2 is similar to proving that MLEs based on subcollection i are consistent if the parameter set is restricted to the compact set $A_i \cup \{\theta^0\}$, $i = 1, 2$. Since the subcollection components are i.i.d., this is straightforward using classical ideas [10], Theorems 16 and 17. The only difference from the referenced work is that one subcollection is a triangular array and so we use a different strong law. Point 3 follows from points 1 and 2. \square

Note that, in this lemma and elsewhere, ϵ is a small number that is defined in context whereas ε always denotes the radius of the neighborhood of θ^0 we are considering. It remains to verify the assumptions concerned with the regularity of the log-likelihood of the full data. When the log-likelihood is differentiable, Lipschitz continuity follows from the mean value theorem if the gradient is bounded. The following lemma uses that to verify Assumption 2. The resulting Lipschitz constant, that is, the bound of the gradient, is the same for both A_1 and A_2 . The lemma also gives a probabilistic bound on the order of this Lipschitz constant as $n \rightarrow \infty$ that will be useful when verifying Assumption 3.

LEMMA 3.3. *If θ^0 is an interior point of Θ , then for every n and small enough $\varepsilon > 0$ there exists a random variable K_n such that, \mathbb{P} -almost surely,*

$$\sup_{\theta \in \bar{B}_\varepsilon(\theta^0)} \|\nabla \Lambda_n(\theta; Y)\| \leq K_n = o_{\mathbb{P}}(n^b),$$

for some $b > 0$.

Proving Lemma 3.3 (see the Supplementary Material [9]) is largely an exercise in bounding the eigenvalues of the covariance matrix $C(\theta)$ and its inverse on interior points of Θ . We are ready for the main result of the section.

THEOREM 3.4. *If θ^0 is an interior point of Θ and $T = O(N^k)$ for some $k \geq 0$ as $N \rightarrow \infty$, then, \mathbb{P} -almost surely, there exists a sequence $\hat{\theta}_n$ of roots to the likelihood equations $\nabla \Lambda_n(\theta; Y) = 0$ such that $\hat{\theta}_n \xrightarrow{\mathbb{P}} \theta^0$.*

PROOF. We verify the conditions of Theorem 2.4. Fix an arbitrary $\varepsilon > 0$. Since θ^0 is interior, we may assume ε is small enough that all points in $\bar{B}_\varepsilon(\theta^0)$ are interior points of Θ . As is proven in the Supplementary Material [9], $\ell_n(\theta; Y) = \log(f_\theta^n(Y))$ is \mathbb{P} -almost surely differentiable on $\bar{B}_\varepsilon(\theta^0)$, so $\Lambda_n(\theta; Y) = \ell_n(\theta; Y) - \ell_n(\theta^0; Y)$ is too. By Lemma 3.2, Assumption 1 holds with what is there denoted Θ replaced by $\bar{B}_\varepsilon(\theta^0)$. The identification rate is exponentially fast in $N/2$, $a_n = o(e^{-N\epsilon})$ for some $\epsilon > 0$. Lemma 3.3 shows that $\Lambda_n(\theta; Y)$ is K_n -Lipschitz on both A_1 and A_2 , and that $K_n = o_{\mathbb{P}}(n^b)$ for some $b > 0$. This verifies Assumption 2. It remains only to verify that the rate conditions in Assumption 3 hold. The δ -covering number of the sphere $\partial B_\varepsilon(\theta^0)$ is $O([\varepsilon/\delta]^{d-1})$ as $\delta \rightarrow 0$ [2], Lemma 1. Thus, since $A_i \subseteq \partial B_\varepsilon(\theta^0)$, by picking $\delta_{i,n} = n^{-b}$ we can have $M_{n,i} = O(n^{[d-1]b})$ as $n \rightarrow \infty$, $i = 1, 2$. Our choice of $\delta_{n,i}$ ensures $K_{n,i}\delta_{n,i} = K_n\delta_n = o_{\mathbb{P}}(1)$, which is the first rate condition. Since the identification rate is exponential in $N/2$ for both subcollections, we have that $M_{n,i}a_{n,i} = O(N^{2b[d-1]}T^{b[d-1]}e^{-\epsilon N})$ for some $\epsilon > 0$, which is $o(1)$ as $N \rightarrow \infty$ since T is of (at most) polynomial order in N . \square

We expect the proof technique used here to work in many other models. Essentially, all that is needed is that the subcollections' log-likelihoods satisfy uniform strong laws, that the

gradient of the full data log-likelihood is of polynomial order, and that the number of observations in any subcollection grows faster than logarithmically in the total sample size. Here, to show that the gradient of the log-likelihood is of polynomial order (Lemma 3.3) we worked with the closed form expression $2 \log f_{\theta}^n(Y) = -\log \det(C(\theta)) - (Y - m(\theta))^T C^{-1}(\theta)(Y - m(\theta)) - n \log(2\pi)$, using that $\|Y\|$ is of polynomial order and that the eigenvalues of $C(\theta)$ are appropriately bounded by polynomials in n on $B_{\varepsilon}(\theta^0)$. This illustrates that, in order to determine if the gradient in a given model is of polynomial order or not, one in general has to consider both the stochastic properties of the data and the particular parameterization. Uniform strong laws for the subcollections' log-likelihoods, leading to Lemma 3.2, hold here because the subcollections consist of i.i.d. random variables. This is clearly not necessary; in the next example, we consider subcollections with independent but not identically distributed variables, and, similarly, strong laws for stationary stochastic processes may apply if one has a model with subcollections consisting of dependent but identically distributed variables.

It is possible that the assumption that $T = O(N^k)$, $k \geq 0$, could be relaxed by picking other subcollections that also make use of the variation in the time dimension. It is not trivial, however, since the dependence between any two responses sharing a random effect does not vanish as time between the observations increases. Indeed, it is crucial that $N \rightarrow \infty$ in this model: if $T \rightarrow \infty$ but N is fixed, then the data consist of a fixed number (N^2) of vectors of T equicorrelated random variables. In that case, it is not possible to find a subcollection that consists of an increasing number of independent variables. Accordingly, one can show that even if one was to simplify our model so that the N^2 vectors were independent and one was estimating only a mean parameter, the MLE would not be consistent. In the next section, we examine how predictors and mixed-type responses affect the argument.

3.2. Logit-normal MGLMM. The model we consider in this section is an extension in several ways of the logistic GLMMs for which the technique based on subcollections was first developed [16]. The random effect structures are similar, that is, crossed, but we have multivariate, mixed-type responses and predictors. The main ideas for verifying the assumptions of the theory from Section 2 are the same as in our LMM example. However, due to the inclusion of predictors, we use results from empirical process theory in place of the more classical strong laws used for the LMM. Showing existence of appropriate subsets of the parameter space that the subcollections identify also requires more work than with i.i.d. components. As before, we discuss the model definition and subcollection selection before establishing consistency.

Suppose for subjects (i, j) , $i = 1, \dots, N$, $j = 1, \dots, N$, there are two responses, $Y_{i,j,1}$ which is continuous and $Y_{i,j,2}$ which is binary. The vector of all responses is

$$Y = [Y_{1,1,1}, Y_{1,1,2}, Y_{1,2,1}, \dots, Y_{N,N,2}]^T \in \mathbb{R}^n, \quad n = 2N^2.$$

For each subject, we observe a vector of nonstochastic predictors $x_{i,j} \in \mathbb{R}^p$, the same for both responses. Similarly, $z_{i,j} \in \mathbb{R}^r$ is the same for both responses. Let $\eta_{i,j,k} = x_{i,j}^T \beta_k + z_{i,j}^T u$ be the linear predictor, $i = 1, \dots, N$, $j = 1, \dots, N$, $k = 1, 2$, where $\beta_1 = [\theta_1, \dots, \theta_p]^T$, $\beta_2 = [\theta_{p+1}, \dots, \theta_{2p}]^T$. We assume that $\|x_{i,j}\| \leq 1$ for all i, j . In practice, this only rules out the possibility that $\|x_{i,j}\| = \infty$ since our setting allows for the standardization of predictors. The conditional density of the responses given the random effects that we consider is, up to scaling by $(2\pi)^{-n/2}$,

$$f_{\theta}^n(y | u) \propto \exp \left[\sum_{i,j} -(y_{i,j,1} - \eta_{i,j,1})^2 / 2 + y_{i,j,2} \eta_{i,j,2} - \log(1 + e^{\eta_{i,j,2}}) \right].$$

Given the random effects, $Y_{i,j,1}$ is normal with mean $\eta_{i,j,1}$ and variance 1, and $Y_{i,j,2}$ is Bernoulli with success probability $1/(1 + e^{-\eta_{i,j,2}})$ —a logistic GLMM. The choice of $\tau_i = 1$

for all i is made for identifiability reasons for the Bernoulli responses, and for convenience for the normal responses. Setting the τ_i 's to some other known constants does not fundamentally change the results.

Suppose $U^{(1)} \sim \mathcal{N}(0, \theta_d I_N)$ and $U^{(2)} \sim \mathcal{N}(0, \theta_d I_N)$, independently, with corresponding design matrices $Z_1 = I_N \otimes 1_N \otimes 1_2$ and $Z_2 = 1_N \otimes I_n \otimes 1_2$. Taking $U = [U^{(1)\top}, U^{(2)\top}]^\top$ and $Z = [Z_1, Z_2]$, the linear predictors are $\eta_{i,j,k} = x_{i,j}^\top \beta_k + u_i^{(1)} + u_j^{(2)}$. Thus, responses from the same subject share two random effects, responses from different subjects with one of the first two indexes in common share one random effect, and other responses share no random effects and are hence independent. The covariance matrix for the linear predictors is easily computed in the same way as in the LMM. The covariance matrix for responses, however, is less transparent. It is for simplicity that we assume in this section that all random effects have the same variance. It is not necessary for our theory to be operational but this simplification shortens proofs considerably and allows us to focus on the main ideas.

3.2.1. *Subcollection selection.* With p predictors the $(2p + 1)$ -dimensional parameter set is $\Theta = \mathbb{R}^p \times \mathbb{R}^p \times (0, \infty)$, a subset of \mathbb{R}^d , again equipped with the usual Euclidean metric. The intuition behind the selection of subcollections is that the normal responses should identify the coefficient β_1 and the variance parameter θ_d . Similarly, the Bernoulli responses should identify the coefficient vector β_2 . With that in mind, we take for $i = 1, 2$,

$$W^{(i)} = (Y_{1,1,i}, Y_{2,2,i}, \dots, Y_{N,N,i}).$$

Both of these subcollections consist of independent but not identically distributed random variables—independence follows from the fact that no components in the same subcollection share random effects. Notice that these subcollections are in practice often triangular arrays since the predictors may need to be scaled by $1/\max_{i \leq N, j \leq N} \|x_{i,j}\|$ to satisfy $\|x_{i,j}\| \leq 1$. All responses in the first subcollection have marginal normal distributions and all responses in the second have marginal Bernoulli distributions.

Identification is more complicated than in our previous example. One issue is that there can be many θ_d and β_2 that give the same marginal success probability for the components in the second subcollection. A second issue is that, since the predictors can change with n , classical identification for a fixed n does not necessarily lead to identification in the sense of Definition 2.1. Additionally, the approach used in the LMM to find appropriate subsets A_1 and A_2 by means of Proposition 3.1 only works in general when the subcollection components are i.i.d. Thus, we take a slightly different route to establishing consistency compared to the LMM.

3.2.2. *Consistency.* In this section, we verify the conditions of Theorem 2.4. The limiting process is that $N \rightarrow \infty$, which is equivalent to $n \rightarrow \infty$ since $n = 2N^2$. We will first prove two lemmas that roughly correspond to Assumptions 1 and 2.

Let $\lambda_{\min}(\cdot)$ denote the minimum eigenvalue of its matrix argument.

LEMMA 3.5. *If θ^0 is an interior point of Θ and*

$$\liminf_{N \rightarrow \infty} \lambda_{\min} \left(N^{-1} \sum_{i=1}^N x_{i,i} x_{i,i}^\top \right) > 0,$$

then for all small enough $\varepsilon > 0$ there exist A_1 and A_2 such that $A_1 \cup A_2 = \partial B_\varepsilon(\theta^0)$:

1. $\limsup_{N \rightarrow \infty} N^{-1} \sup_{\theta \in A_i} \mathbb{E}[\Lambda_N(\theta; W^{(i)})] < 0$,
2. $\sup_{\theta \in A_i} N^{-1} |\Lambda_N(\theta; W^{(i)}) - \mathbb{E}[\Lambda_N(\theta; W^{(i)})]| \xrightarrow{P} 0$, and consequently;

3. A_i is identified by $W^{(i)}$ with an identification rate $a_{n,i} = o(e^{-\epsilon N})$ for some $\epsilon > 0$, $i = 1, 2$.

PROOF. A detailed proof is the [Appendix](#); we here give the proof idea. Let $A_2 = \partial B_\epsilon(\theta^0) \cap \{\theta : |\theta_d - \theta_d^0| \leq \zeta\} \cap \{\|\beta_2 - \beta_2^0\| \geq \epsilon/2\}$, for some small $\zeta > 0$. Let A_1 be the closure of $\partial B_\epsilon(\theta^0) \cap A_2^c$. The idea is that if ζ is small enough, so that $\theta_d \approx \theta_d^0$ and $\|\beta_2 - \beta_2^0\| \geq \epsilon/2$ on A_2 , then the distributions of $W^{(2)}$ implied by $\theta \in A_2$ and θ^0 are different if $X = [x_{1,1}, x_{2,2}, \dots, x_{N,N}]^T$ has full column rank. That is, $W^{(2)}$ should be able to distinguish every $\theta \in A_2$ from θ^0 . Moreover, one can show that on A_1 it holds either that $|\theta_d - \theta_d^0| \geq \min(\zeta, \epsilon/4)$ or that $\|\beta_1 - \beta_1^0\| \geq \epsilon/4$. In either case, $W^{(1)}$ should be able to distinguish $\theta \in A_1$ from θ^0 . Formalizing this idea leads to point 1. Point 2 follows from checking the conditions of a uniform law of large numbers [24], Theorem 8.2, and point 3 from points 1 and 2. \square

The explicit construction of the subsets A_1 and A_2 , as opposed to using Proposition 3.1, warrants an additional comment. Recall, the proposition gives compact \tilde{A}_1 and \tilde{A}_2 such that $\tilde{A}_1 \cup \tilde{A}_2 = \partial B_\epsilon(\theta^0)$ and $v_\theta^i \neq v_{\theta^0}^i$, $\theta \in \tilde{A}_i$, $i = 1, 2$. If one takes $A_i = \tilde{A}_i$, then point 1 in Lemma 3.2 follows. Moreover, when the subcollection components are i.i.d., this in turn leads to point 1 in Lemma 3.5, which is what is really needed. However, when the distributions of the subcollection components are not identical, this last implication is not true in general.

Having selected appropriate subcollections and subsets it remains only to check that the log-likelihood for the full data satisfies the regularity conditions in Assumptions 2–3. The following lemma verifies Assumption 2 and establishes a rate needed for the verification of Assumption 3.

LEMMA 3.6. *If θ^0 is an interior point of Θ , then for every n and small enough $\epsilon > 0$ there exists a random variable K_n such that, \mathbf{P} -almost surely,*

$$\sup_{\theta \in \tilde{B}_\epsilon(\theta^0)} \|\nabla \Lambda_n(\theta; Y)\| \leq K_n = o_{\mathbf{P}}(n^b),$$

for some $b > 0$.

Upon inspecting the proof (Supplementary Material [9]) one sees that b can be taken to be $1 + \epsilon$, for any $\epsilon > 0$. This is a better (slower) rate than that obtained in the linear mixed model (see the proof of Lemma 3.3). We are now ready to state the main result of the section.

THEOREM 3.7. *If θ^0 is an interior point of Θ and*

$$\liminf_{N \rightarrow \infty} \lambda_{\min} \left(N^{-1} \sum_{i=1}^N x_{i,i} x_{i,i}^T \right) > 0,$$

then \mathbf{P} -almost surely, there exists a sequence $\hat{\theta}_n$ of roots to the likelihood equations $\nabla \Lambda_n(\theta; Y) = 0$ such that $\hat{\theta}_n \xrightarrow{\mathbf{P}} \theta^0$.

PROOF. The proof is similar to that of Theorem 3.4 so we skip some details. We may assume all points in $\tilde{B}_\epsilon(\theta^0)$ are interior points of Θ . As is proven in the Supplementary Material [9], $\Lambda_n(\theta; Y)$ is differentiable on $\tilde{B}_\epsilon(\theta^0)$. By Lemma 3.5, the identification rate is exponentially fast in N and Lemma 3.6 shows that $\Lambda_n(\theta; Y)$ is K_n -Lipschitz on both A_1 and A_2 , and that $K_n = o_{\mathbf{P}}(n^b)$ for some $b > 0$. This verifies Assumption 2. By picking $\delta_{i,n} = n^{-b}$, we can have $M_{n,i} = O(n^{[d-1]b})$ as $n \rightarrow \infty$, $i = 1, 2$. Thus, $K_n \delta_n = o_{\mathbf{P}}(1)$ and $M_{n,i} a_{n,i} = O(N^{2b[d-1]} e^{-\epsilon N})$ for some $\epsilon > 0$, which is $o(1)$ as $N \rightarrow \infty$ since $n = 2N^2$. \square

4. Discussion. Our theory develops the current state-of-the-art asymptotic theory based on subcollections to cover more general cases. The assumptions we make highlight what makes the use of subcollections work. In particular, the interplay between the identification rates of subcollections and the regularity of the likelihood function for the full data is made precise. We note that when the subcollections consist of $m \in \{1, 2, \dots\}$ independent random variables, as in our examples, then if $n = o(m^b)$ for some $b > 0$ and $\nabla \Lambda_n(\theta; Y) = o_{\mathbb{P}}(n^{b'})$ for some $b' > 0$, uniformly on a compact Θ , the rate conditions are satisfied. This is so because, under regularity conditions, the identification rate in a subcollection with m independent random variables is exponential in $m = n^{1/b}$. Since this argument works for arbitrarily large b and b' our theory is operational in a wide range of models. Loosely speaking, if the score function is of less than exponential order in the sample size and there are subcollections of independent random variables that grow faster than logarithmically in the sample size, the MLE is consistent. The conditions should be verifiable in many models since they often require only standard asymptotic tools. For example, in the LMM example nothing more than a uniform law of large numbers and strict positivity of the K–L divergence between distributions corresponding to distinct, identified parameters is needed. Though not pursued here, by inspecting the assumptions of our theory one also sees that it has the potential to be extended to allow the dimension of the parameter set, d , grow with n . The rates required in our assumptions could be satisfied also if d grows, at least if at a slow enough rate. The wide applicability of empirical process theory, which we use in the second application, also suggests that it may be possible to verify our conditions in yet more complicated models.

Consistency of MLEs has not previously been established in either of the two models to which we apply the general theory. In particular, previous work on asymptotic theory for MLEs in mixed models often either assumes independent replications of a response vector, that there are no predictors or no mixed-type responses. We have tried to keep the models here as simple as possible while still illustrating key ideas. Crossed random effects, temporal dependence and predictors are included because they are challenging theoretically and are commonly used in practice. We have refrained from including things that do not require any new methods but make ideas less transparent. For example, it would be straightforward to include random effects that are not crossed, possibly at the expense of using more subcollections or subcollections consisting of independent vectors of larger dimension than what is now necessary. Similarly, adding several crossed random effects does not make things much harder, only less transparent.

Avenues for future research includes the rate of convergence of the MLEs as well as their asymptotic distribution. Intuitively, one expects MLEs based on the full data to converge at least as fast as the slowest of the subcollection MLEs, that is, the estimators one gets from using only a subset of the full data. There is some evidence of this, namely that, under regularity conditions, the Fisher information in the full data is always larger than that in any subcollection [16]. On the other hand, it is easy to show that, for the simple LMM example in the [Introduction](#), the full data MLE converges at the same rate as that based on a subcollection of $N = \sqrt{n}$ i.i.d. observations; that is, at the rate $n^{1/4}$. Given the similarities of the random effect structures, that convergence rate may in future work be a reasonable working hypothesis for MLEs in the MGLMM considered here.

APPENDIX: PROOFS

PROOF OF LEMMA 2.2. Fix some arbitrary $\varepsilon > 0$. If $\sup_{\theta \in A_i} L_n(\theta; Y) < 1$ for $i = 1, \dots, s$, then since $L_n(\theta^0; Y) = 1$, there are no global maximizers in $\bigcup_{i=1}^s A_i \supseteq \Theta \cap B_\varepsilon(\theta^0)^c$.

Thus, it suffices to prove

$$\mathbb{P}\left(\bigcup_{i=1}^s \left\{ \sup_{\theta \in A_i} L_n(\theta; Y) \geq 1 \right\}\right) \leq \sum_{i=1}^s \mathbb{P}\left(\sup_{\theta \in A_i} L_n(\theta; Y) \geq 1\right) \rightarrow 0.$$

Since s is fixed it is enough that $\mathbb{P}(\sup_{\theta \in A_i} L_n(\theta; Y) \geq 1) \rightarrow 0$ for every $i = 1, \dots, s$. Without loss of generality, consider $i = 1$. Pick a cover of A_1 as given by Assumption 3 and, for every ball in the cover, pick a θ^j in the intersection of that ball with A_1 . If there are some balls that do not intersect A_1 , they may be discarded from the cover, so we assume without loss of generality that all balls do intersect A_1 . We then get $M_{n,1}$ points such that every point in A_1 is within $\delta_{n,1}$ of at least one of them. For any $\theta \in A_1$, let $\theta^j(\theta)$ denote the θ^j closest to it (pick an arbitrary one if there are many). Using the Lipschitz continuity given by Assumption 2 and that $x \mapsto e^x$ is increasing we have

$$\begin{aligned} \mathbb{P}\left(\sup_{\theta \in A_1} L_n(\theta; Y) \geq 1\right) &= \mathbb{P}\left(\sup_{\theta \in A_1} \Lambda_n(\theta; Y) \geq 0\right) \\ &= \mathbb{P}\left(\sup_{\theta \in A_1} \ell_n(\theta; Y) \geq \ell_n(\theta^0; Y)\right) \end{aligned}$$

which is upper bounded by

$$\mathbb{P}\left(\sup_{\theta \in A_1} [\ell_n(\theta^j(\theta); Y) + K_{n,1}d_{\mathcal{T}}(\theta, \theta^j(\theta))] \geq \ell_n(\theta^0; Y)\right).$$

Because there are only $M_{n,1}$ points θ^j , and $d_{\mathcal{T}}(\theta^j(\theta), \theta) \leq \delta_{n,1}$ since $\theta^j(\theta)$ is the one closest to θ , we get that the last inline equation is upper bounded by

$$\begin{aligned} &\mathbb{P}\left(\max_{j \leq M_{n,1}} f_{\theta^j}(Y) e^{K_{n,1}\delta_{n,1}} \geq f_{\theta^0}(Y)\right) \\ &\leq \mathbb{P}\left(2 \max_{j \leq M_{n,1}} f_{\theta^j}(Y) \geq f_{\theta^0}(Y)\right) + \mathbb{P}(e^{K_{n,1}\delta_{n,1}} \geq 2) \\ &= \mathbb{P}\left(2 \max_{j \leq M_{n,1}} f_{\theta^j}(Y) \geq f_{\theta^0}(Y)\right) + o(1), \end{aligned}$$

where the last line uses Assumption 3. The remaining term,

$$\mathbb{P}\left(2 \max_{j \leq M_{n,1}} f_{\theta^j}(Y) \geq f_{\theta^0}(Y)\right) = \mathbb{P}\left(\max_{j \leq M_{n,1}} L_n(\theta^j; Y) \geq 1/2\right),$$

we will deal with using Lemma 2.1 and dominated convergence. After conditioning on $W^{(1)}$, we have

$$\begin{aligned} \mathbb{P}\left(\max_{j \leq M_{n,1}} L_n(\theta^j; Y) \geq 1/2 \mid W^{(1)}\right) &\leq \sum_{i=1}^{M_{n,1}} 2L_{m_1}(\theta^j; W^{(1)}) \\ &\leq 2M_{n,1} \sup_{\theta \in A_1} L_{m_1}(\theta, W^{(1)}), \end{aligned}$$

\mathbb{P} -almost surely, where the first inequality is by subadditivity and Lemma 2.1, and the second uses that $L_n(\theta^j; W^{(1)}) \leq \sup_{\theta \in A_1} L_{m_1}(\theta; W^{(1)})$ by definition. The expression in the last line vanishes as $n \rightarrow \infty$ by Assumption 3. Thus,

$$\mathbb{P}\left(\max_{j \leq M_{n,1}} L_n(\theta^j; Y) \geq 1/2\right) \rightarrow 0$$

by dominated convergence. The dominating function can be the constant 1. This completes the proof. \square

Let $C(\delta, G, \|\cdot\|)$ denote the δ -covering number of the set G under the distance associated with the norm $\|\cdot\|$. We will use the following result due to Pollard [24], Theorem 8.2, here stated in terms of covering numbers instead of packing numbers.

LEMMA A.1. *Let $h_1(\omega, \theta), h_2(\omega, \theta), \dots, \theta \in A \subseteq \Theta$, be independent processes with integrable envelopes $H_1(\omega), H_2(\omega), \dots$, meaning $|h_i(\omega, \theta)| \leq H_i(\omega)$, for all i and $\theta \in A$. Let $H = (H_1, \dots, H_N)$ and*

$$\mathcal{H}_{N,\omega} = \{[h_1(\omega, \theta), \dots, h_N(\omega, \theta)] \in \mathbb{R}^N : \theta \in A\}.$$

If for every $\epsilon > 0$, there exists a $K > 0$ such that:

1. $N^{-1} \sum_{i=1}^N \mathbf{E}[H_i I(H_i > K)] < \epsilon$ for all N , and
2. $\log C(\epsilon \|H\|_1, \mathcal{H}_{N,\omega}, \|\cdot\|_1) = o_{\mathbb{P}}(N)$ as $N \rightarrow \infty$,

then

$$\sup_{\theta \in A} N^{-1} \left| \sum_{i=1}^N h_i(\omega, \theta) - \mathbf{E}(h_i(\omega, \theta)) \right| \xrightarrow{\mathbb{P}} 0.$$

PROOF OF LEMMA 3.5. Let us first prove that, given $\epsilon > 0$, there exists a $\zeta > 0$, and hence $A_i = A_i(\epsilon, \zeta), i = 1, 2$, such that point 1 in the lemma holds. The definition of $A_i(\epsilon, \zeta)$ is as in the main text. Let $c(t) = \log(1 + e^t)$ denote the cumulant function in the conditional distribution of $Y_{i,i,2}$ given the random effects and define

$$p_i(\beta_2, \theta_d) = \mathbf{E} \left[c' \left(x_{i,i}^\top \beta_2 + \sqrt{\theta_d / \theta_d^0} (U_i^{(1)} + U_j^{(2)}) \right) \right].$$

Recall, \mathbf{E} denotes expectation with respect to the distributions indexed by θ^0 , so $p_i(\beta_2, \theta_d)$ is the success probability of $Y_{i,i,2}$ when β_2 and θ_d are the true parameters.

Note that because the components in $W^{(2)}$ are independent, we can write $\mathbf{E}[\Lambda_N(\theta; W^{(2)})]$ as a sum of N terms, each summand being the negative K–L divergence between two Bernoulli variables with parameters $p_i(\beta_2, \theta_d)$ and $p_i(\beta_2^0, \theta_d^0)$. Thus (see the Supplementary Material [9]),

$$N^{-1} \mathbf{E}[\Lambda_N(\theta; W^{(2)})] \leq -2N^{-1} \sum_{i=1}^N [p_i(\beta_2, \theta_d) - p_i(\beta_2^0, \theta_d^0)]^2$$

which one can show is upper bounded by

$$(4) \quad -2 \left[N^{-1} \sum_{i=1}^N |p_i(\beta_2, \theta_d) - p_i(\beta_2, \theta_d^0)| - N^{-1} \sum_{i=1}^N |p_i(\beta_2^0, \theta_d^0) - p_i(\beta_2, \theta_d^0)| \right]^2.$$

Let us work separately with the averages in the last line. We will show that the second can be made arbitrarily small on A_2 by selecting ζ small enough, and that the first is bounded away from zero on the same A_2 , leading to an asymptotic upper bound on $\sup_{\theta \in A_2} N^{-1} \mathbf{E}[\Lambda_N(\theta; W^{(2)})]$ away from zero. We start with the first average.

Let H be a compact subset of \mathbb{R} such that $x_{i,i}^\top \beta_2 \in H$ for all i and $\theta \in \bar{B}_\epsilon(\theta^0)$. Such H exists because the predictors are bounded and β_2 is bounded on $\bar{B}_\epsilon(\theta^0)$. Then, defining $\tilde{p}_i(\gamma, \theta_d)$ as $p_i(\beta_2, \theta_d)$ but with $x_{i,i}^\top \beta_2$ replaced by γ , we get

$$\sup_{\theta \in A_2} |p_i(\beta_2, \theta_d) - p_i(\beta_2, \theta_d^0)| \leq \sup_{\theta \in A_2} \sup_{\gamma \in H} |\tilde{p}_i(\gamma, \theta_d) - \tilde{p}_i(\gamma, \theta_d^0)|.$$

Since the random variable in the expectation defining \tilde{p}_i is bounded by 1 (it is the mean of a Bernoulli random variable), \tilde{p}_i is continuous by dominated convergence. Thus, since H is compact, $\sup_{\gamma \in H} |\tilde{p}_i(\gamma, \theta_d) - \tilde{p}_i(\gamma, \theta_d^0)|$ is continuous in θ_d . That is, we can make $\sup_{\gamma \in H} |\tilde{p}_i(\gamma, \theta_d) - \tilde{p}_i(\gamma, \theta_d^0)|$ arbitrarily small on $A_2 = A_2(\zeta, \varepsilon)$ by picking ζ small enough, which is what we wanted to show. We next work with the second average in (4).

By the mean value theorem, for some $\tilde{\beta}_{2,i}$ between β_2 and β_2^0 , $|p_i(\beta_2^0, \theta_d^0) - p_i(\beta_2, \theta_d^0)| = |\mathbb{E}(c''(x_{i,i}^\top \tilde{\beta}_{2,i} + U_i^{(2)} + U_j^{(2)}))x_{i,i}^\top(\beta_2 - \beta_2^0)|$. Here, differentiation under the expectation is permissible since c'' is the variance of a Bernoulli random variable, hence bounded by $1/4$, and $|x_{i,i}^\top(\beta_2 - \beta_2^0)| \leq \|x_{i,i}\| \|\beta_2 - \beta_2^0\| \leq \varepsilon$ on $\bar{B}_\varepsilon(\theta^0)$. By the same bound on c'' , we get that $\mathbb{E}(c''(\gamma + U_i^{(1)} + U_j^{(2)}))$ is continuous in γ . Thus, $\inf_{\gamma \in H} \mathbb{E}(c''(\gamma + U_i^{(1)} + U_j^{(2)})) \geq c_1 > 0$. That c_1 must be positive follows from that c'' is strictly positive on all of \mathbb{R} . We have thus proven that $|p_i(\beta_2^0, \theta_d^0) - p_i(\beta_2, \theta_d^0)| \geq c_1 |x_{i,i}^\top(\beta_2 - \beta_2^0)|$, uniformly on $\bar{B}_\varepsilon(\theta^0)$. Using this and that $|x_{i,i}^\top(\beta_2 - \beta_2^0)| \leq \|x_{i,i}\| \|\beta_2 - \beta_2^0\| \leq \varepsilon \leq 1$ so that squaring it makes it smaller,

$$\begin{aligned} & N^{-1} \sum_{i=1}^N |p_i(\beta_2^0, \theta_d^0) - p_i(\beta_2, \theta_d^0)| \\ & \geq c_1 N^{-1} \sum_{i=1}^N |x_{i,i}^\top(\beta_2 - \beta_2^0)| \\ & \geq c_1 N^{-1} (\beta_2 - \beta_2^0)^\top \left(\sum_{i=1}^N x_{i,i} x_{i,i}^\top \right) (\beta_2 - \beta_2^0) \\ & \geq c_1 \|\beta_2 - \beta_2^0\|^2 N^{-1} \lambda_{\min} \left(\sum_{i=1}^N x_{i,i} x_{i,i}^\top \right) \end{aligned}$$

which lower limit as $N \rightarrow \infty$ is bounded below by some strictly positive constant, say c_2 , since $\liminf_{N \rightarrow \infty} N^{-1} \lambda_{\min}(\sum_{i=1}^N x_{i,i} x_{i,i}^\top) \geq c_3 > 0$, for some c_3 , and $\|\beta_2 - \beta_2^0\| \geq \varepsilon/2 > 0$ on A_2 . To summarize, we may pick ζ so small that the second average in (4) is less than $c_2/2$, say, and hence get $\sup_{\theta \in A_2} N^{-1} \mathbb{E}[\Lambda_N(\theta; W^{(2)})] \leq -2(c_2 - c_2/2)^2 < 0$, for all but at most finitely many N . This proves point 1 as it pertains to A_2 .

Consider next

$$A_1 = \partial B_\varepsilon(\theta^0) \cap (\{\theta : |\theta_d - \theta_d^0| \geq \zeta\} \cup \{\theta : \|\beta_2 - \beta_2^0\| \leq \varepsilon/2\})$$

and $W^{(1)}$. Due to independence, $\mathbb{E}[\Lambda_N(\theta; W^{(1)})]$ can be written as a sum of N terms in the form

$$(5) \quad -\frac{1}{2} \left[\log \left(\frac{1 + 2\theta_d}{1 + 2\theta_d^0} \right) + \frac{1 + 2\theta_d^0 + [x_i^\top(\beta_2 - \beta_2^0)]^2}{1 + 2\theta_d} - 1 \right],$$

which is the negative K–L divergence between two univariate normal distributions. Let us consider the possible values this can take for $\theta \in A_1$. If $|\theta_d - \theta_d^0| \geq \zeta$, then (5) is upper bounded by what is obtained when $\beta_1 = \beta_1^0$. This in turn is a continuous function in θ_d and hence attains its supremum on the compact set $\{\theta_d : \zeta \leq |\theta_d - \theta_d^0| \leq \varepsilon\}$, and hence on A_1 . This supremum is strictly positive because the divergence can be zero only if $\theta_d = \theta_d^0$. If instead $\|\beta_2 - \beta_2^0\| \leq \varepsilon/2$, then either $|\theta_d - \theta_d^0| \geq \varepsilon/4$ or $\|\beta_1 - \beta_1^0\| \geq \varepsilon/4$, for otherwise it cannot be that $\|\theta - \theta^0\| = \varepsilon$. If $|\theta_d - \theta_d^0| \geq \varepsilon/4$ the divergence in (5) has a lower bound away from zero by the same argument as for the cases $|\theta_d - \theta_d^0| \geq \zeta$. It remains to deal with the case $\|\beta_1 - \beta_1^0\| \geq \varepsilon/4$.

Writing

$$[x_{i,i}^\top(\beta_1^0 - \beta_1)]^2 = (\beta_1^0 - \beta_1)^\top x_i x_i^\top (\beta_1^0 - \beta_1)$$

we see that $-2N^{-1}\Lambda_N(\theta; W^{(1)})$ is equal to

$$\log\left(\frac{1 + 2\theta_d}{1 + 2\theta_d^0}\right) + \frac{1 + 2\theta_d^0 + N^{-1} \sum_{i=1}^N (\beta_1^0 - \beta_1)^\top x_i x_i^\top (\beta_1^0 - \beta_1)}{1 + 2\theta_d} - 1,$$

which has a lower limit that is greater than

$$\log\left(\frac{1 + 2\theta_d}{1 + 2\theta_d^0}\right) + \frac{1 + 2\theta_d^0 + c_3(\varepsilon/4)^2}{1 + 2\theta_d} - 1.$$

This expression is in turn maximized in θ_d at $\theta_d = \theta_d^0 + c_3(\varepsilon/16)^2$; this follows from a straightforward optimization in $1 + 2\theta_d$. The corresponding maximum evaluates to $\log(1 + 2\theta_d^0 + c_3(\varepsilon/4)^2) - \log(1 + 2\theta_d^0) > 0$. This completes the proof of point 1.

The proof of point 2 consists of checking the conditions of Lemma A.1. We first work with A_1 and $W^{(1)}$. Let $h_i(\omega, \theta) = \log[f_\theta(Y_{i,i,1}(\omega))/f_{\theta^0}(Y_{i,i,1}(\omega))]$ be the log-likelihood ratio for the i th observation in the first subcollection, $i = 1, \dots, N$. We equip $\mathcal{H}_{N,\omega}$ with the L_1 norm $\|\cdot\|_1$, and Θ is equipped with the L_2 norm as before. To facilitate checking the two conditions, we will first derive envelopes with the following properties: $\sup_{-\infty < i < \infty} \mathbb{E}H_i^k < \infty$ for every $k \geq 0$, $\sup_{-\infty < i < \infty} \mathbb{P}(H_i \geq K) \rightarrow 0$ as $K \rightarrow 0$, and each $h_i(\omega, \theta)$ is H_i -Lipschitz in θ on $\bar{B}_\varepsilon(\theta^0)$, and hence on A_1 , for every ω . We start with the Lipschitz property.

Let us use the slight abuse of notation that $y_{i,i,1} = Y_{i,i,1}(\omega)$. Since the distribution of $W^{(1)}$ does not depend on β_2 , we have $\nabla_{\beta_2} h_i(\omega, \theta) = 0$, and for some $c_1, c_2, c_3, c_4, c_5 > 0$ (depending on ε), and every $\theta \in \bar{B}_\varepsilon(\theta^0)$,

$$\begin{aligned} \|\nabla_{\beta_1} h_i(\omega, \theta)\| &= \|(y_{i,i,1} - x_{i,i}^\top \beta_1)x_{i,i}/(1 + 2\theta_d)\| \leq c_1|y_{i,i,1}| + c_2, \\ |\nabla_{\theta_d} h_i(\omega, \theta)| &= \frac{1}{2} \left| \frac{1}{1 + 2\theta_d} - (y_{i,i,1} - x_{i,i}^\top \beta_1)^2/(1 + 2\theta_d)^2 \right| \\ &\leq c_3 + c_4(|y_{i,i,1}| + c_5)^2. \end{aligned}$$

Let H_i be the sum of the bounds, that is,

$$H_i(\omega) = c_1|y_{i,i,1}| + c_2 + c_3 + c_4(|y_{i,i,1}| + c_5)^2.$$

By the mean value theorem, $|h_i(\omega, \theta) - h_i(\theta', \omega)| = |(\theta - \theta')^\top \nabla h_i(\omega, \tilde{\theta})| \leq \|\theta - \theta'\| H_i$ for some $\tilde{\theta}$ between θ and θ' . That is, h_i is H_i -Lipschitz on $\bar{B}_\varepsilon(\theta^0)$. That H_i is an envelope for h_i follows from noting that $h_i(\omega, \theta^0) = 0$ so by taking $\theta' = \theta^0$ in the previous calculation, $|h_i(\omega, \theta)| \leq H_i \|\theta - \theta^0\| \leq H_i$ on $\bar{B}_\varepsilon(\theta^0)$. That $\sup_i \mathbb{E}(H_i^k) < \infty$ for every $k > 0$ and $\sup_i \mathbb{P}(H_i > K) \rightarrow 0$ as $K \rightarrow \infty$ follow from that $Y_{i,i,1}$ is normally distributed with variance $1 + 2\theta_d^0$, not depending on i , and mean satisfying $-\|\beta_1^0\| \leq x_{i,i}^\top \beta_1^0 \leq \|\beta_1^0\|$. We are now ready to check the conditions of Lemma A.1.

By the Cauchy–Schwarz inequality and the properties just derived, we have for every fixed N that

$$N^{-1} \sum_{i=1}^N \mathbb{E}[H_i I(H_i > K)] \leq \sup_i \mathbb{E}[H_i^2] \sup_i \mathbb{P}(H_i \geq K) \rightarrow 0, \quad K \rightarrow \infty,$$

which verifies the first condition.

For the second condition, note that the derived Lipschitz property gives, for arbitrary $h = (h_1(\omega, \theta), \dots, h_N(\omega, \theta))$ and $h' = (h_1(\omega, \theta'), \dots, h_N(\omega, \theta'))$ in $\mathcal{H}_{N,\omega}$:

$$\begin{aligned} \|h - h'\|_1 &= \sum_{i=1}^N |h_i(\omega, \theta) - h_i(\omega, \theta')| \\ &= \|\theta - \theta'\| \|H\|_1. \end{aligned}$$

Thus, if we cover $\partial B_\epsilon(\theta^0)$ with ϵ -balls with centers $\theta^j, j = 1, \dots, M$, then the corresponding L_1 balls in \mathbb{R}^N of radius $\epsilon \|H\|_1$ with centers

$$h^j = (h_1(\omega, \theta^j), \dots, h_N(\omega, \theta^j))$$

cover $\mathcal{H}_{N,\omega}$. This is so because for every $\theta \in \partial B_\epsilon(\theta^0)$ there is a θ^j such that $\|\theta - \theta^j\| \leq \epsilon$, and hence by the Lipschitz property $\|h(\omega, \theta) - h(\omega, \theta^j)\|_1 \leq \|H\|_1 \epsilon$. Thus, $C(\epsilon \|H\|_1, \mathcal{H}_{N,\omega}, \|\cdot\|_1) \leq C(\epsilon, \partial B_\epsilon(\theta^0), \|\cdot\|)$. Since the covering number $C(\epsilon, \partial B_\epsilon(\theta^0), \|\cdot\|)$ is constant in N , the second condition of Lemma A.1 is verified for A_1 and $W^{(1)}$.

The arguments for A_2 and $W^{(2)}$ are similar, redefining $h_i(\omega, \theta)$ with $Y_{i,i,1}$ replaced by $Y_{1,1,2}$, taking A_2 in place of A_1 , and so on. We need only prove the existence of envelopes H_1, \dots, H_N with the desired properties. Using that $|y_{i,j,2} - c'(\eta_{i,2,1})| \leq 1$ and that $f_\theta(y_{i,i,2} | u) f_\theta(u) / f_\theta(y_{i,i,2}) = f_\theta(u | y_{i,i,2})$, one gets

$$\begin{aligned} \|\nabla_{\beta_2} h_i(\omega, \theta)\| &= \left\| \frac{1}{f_\theta(y_{i,i,2})} \int f_\theta(y_{i,i,2} | u) f_\theta(u) [y_{i,i,2} - c'(\eta_{i,j,2})] x_{i,i} \, du \right\| \\ &\leq \|x_{i,i}\| \leq 1. \end{aligned}$$

Using that $U_i^{(1)}$ and $U_j^{(2)}$ are the only random effects entering the linear predictor $\eta_{i,j,2}$, and that $f_\theta(y_{i,j,2} | u) \leq 1$,

$$\begin{aligned} |\nabla_{\theta_d} h_i(\omega, \theta)| &\leq \frac{1}{2\theta_d f_\theta(y_{i,i,2})} \int f_\theta(u_i^{(1)}, u_j^{(2)}) \left(\frac{(u_i^{(1)})^2 + (u_j^{(2)})^2}{\theta_d} \right) \, du + \frac{1}{\theta_d} \\ &= \frac{1}{\theta_d f_\theta(y_{i,j,2})} + \frac{1}{\theta_d}. \end{aligned}$$

Due to continuity and compactness, the quantity in the last line attains its supremum on $\bar{B}_\epsilon(\theta^0)$. This maximum is finite for both $y_{i,i,2} = 1$ and $y_{i,i,2} = 0$ since the marginal success probability cannot be one or zero on interior points of Θ . Thus, on $\bar{B}_\epsilon(\theta^0)$, $\|\nabla h_i(\omega, \theta)\|$ is bounded by a constant, say H , the largest needed for the two cases $y_{i,i,2} = 0$ and $y_{i,i,2} = 1$. By setting $H_i = H, i = 1, \dots, N$, we have envelopes with the right properties and this completes the proof of point 2.

Finally, we prove point 3. Consider without loss of generality the first subset and subcollection. For economical notation, we omit dependence on the subcollection and write $L_N(\theta) = L_N(\theta; W^{(1)})$ and $\Lambda_N(\theta) = \Lambda_N(\theta; W^{(1)})$. Point 1 gives that $\sup_{\theta \in A_1} \mathbb{E}[\Lambda_N(\theta)] < -3\epsilon$ for some $\epsilon > 0$ and all large enough N . Assuming that N is large enough that this holds, we get

$$\begin{aligned} \mathbb{P}\left(e^{\epsilon N} \sup_{\theta \in A_1} L_N(\theta) > e^{-\epsilon N}\right) &\leq \mathbb{P}\left(N^{-1} \sup_{\theta \in A_1} \Lambda_N(\theta) > \epsilon + \sup_{\theta \in A_1} \mathbb{E}[\Lambda_N(\theta)]\right) \\ &\leq \mathbb{P}\left(N^{-1} \sup_{\theta \in A_1} |\Lambda_N(\theta) - \mathbb{E}[\Lambda_N(\theta)]| > \epsilon\right), \end{aligned}$$

which vanishes as $N \rightarrow \infty$ by point 2. \square

Acknowledgments. We thank an Associate Editor and two reviewers for suggestions that improved the article, and are grateful to Charles Doss and Charlie Geyer for their comments on an earlier version.

The first author was supported by the American–Scandinavian Foundation.

SUPPLEMENTARY MATERIAL

Supplement to “Consistent maximum likelihood estimation using subsets with applications to multivariate mixed models” (DOI: [10.1214/19-AOS1830SUPP](https://doi.org/10.1214/19-AOS1830SUPP); .pdf). The supplemental material contains proofs and additional details.

REFERENCES

- [1] BAR-SHALOM, Y. (1971). On the asymptotic properties of the maximum-likelihood estimate obtained from dependent observations. *J. Roy. Statist. Soc. Ser. B* **33** 72–87. [MR0287630](https://doi.org/10.2307/2335333)
- [2] BRONŠTEĪN, E. M. and IVANOV, L. D. (1975). The approximation of convex sets by polyhedra. *Sibirsk. Mat. Zh.* **16** 1110–1112, 1132. [MR0400054](https://doi.org/10.1007/BF01223833)
- [3] CHEN, H.-C. and WEHRLY, T. E. (2015). Assessing correlation of clustered mixed outcomes from a multivariate generalized linear mixed model. *Stat. Med.* **34** 704–720. [MR3301605](https://doi.org/10.1002/sim.6374) <https://doi.org/10.1002/sim.6374>
- [4] COULL, B. A. and AGRESTI, A. (2000). Random effects modeling of multiple binomial responses using the multivariate binomial logit-normal distribution. *Biometrics* **56** 73–80.
- [5] CRAMÉR, H. (1946). *Mathematical Methods of Statistics*. Princeton Mathematical Series **9**. Princeton Univ. Press, Princeton, NJ. [MR0016588](https://doi.org/10.2307/2335333)
- [6] CROWDER, M. J. (1976). Maximum likelihood estimation for dependent observations. *J. Roy. Statist. Soc. Ser. B* **38** 45–53. [MR0403035](https://doi.org/10.2307/2335333)
- [7] DEMIDENKO, E. (2013). *Mixed Models: Theory and Applications with R*, 2nd ed. Wiley Series in Probability and Statistics. Wiley, Hoboken, NJ. [MR3235905](https://doi.org/10.2307/2335333)
- [8] DOOB, J. L. (1934). Probability and statistics. *Trans. Amer. Math. Soc.* **36** 759–775. [MR1501765](https://doi.org/10.2307/2335333) <https://doi.org/10.2307/1989822>
- [9] EKVALL, K. O. and JONES, G. L. (2020). Supplement to “Consistent maximum likelihood estimation using subsets with applications to multivariate mixed models.” <https://doi.org/10.1214/19-AOS1830SUPP>.
- [10] FERGUSON, T. S. (1996). *A Course in Large Sample Theory*. Texts in Statistical Science Series. CRC Press, London. [MR1699953](https://doi.org/10.1007/978-1-4899-4549-5) <https://doi.org/10.1007/978-1-4899-4549-5>
- [11] GUEORGIEVA, R. V. and AGRESTI, A. (2001). A correlated probit model for joint modeling of clustered binary and continuous responses. *J. Amer. Statist. Assoc.* **96** 1102–1112. [MR1947258](https://doi.org/10.1198/016214501753208762) <https://doi.org/10.1198/016214501753208762>
- [12] GÜVEN, B. (1995). Asymptotic properties of maximum likelihood estimation in the mixed analysis of variance model. *Statist. Papers* **36** 175–182. [MR1340623](https://doi.org/10.1007/BF02926030) <https://doi.org/10.1007/BF02926030>
- [13] HALL, P. and HEYDE, C. C. (1980). *Martingale Limit Theory and Its Application: Probability and Mathematical Statistics*. Academic Press, New York. [MR0624435](https://doi.org/10.1007/978-1-4899-4549-5)
- [14] HEIJMANS, R. D. H. and MAGNUS, J. R. (1986). Consistent maximum-likelihood estimation with dependent observations. The general (nonnormal) case and the normal case. *J. Econometrics* **32** 253–285. [MR0864928](https://doi.org/10.1016/0304-4076(86)90040-0) [https://doi.org/10.1016/0304-4076\(86\)90040-0](https://doi.org/10.1016/0304-4076(86)90040-0)
- [15] JIANG, J. (2007). *Linear and Generalized Linear Mixed Models and Their Applications*. Springer Series in Statistics. Springer, New York. [MR2308058](https://doi.org/10.1007/978-1-4899-4549-5)
- [16] JIANG, J. (2013). The subset argument and consistency of MLE in GLMM: Answer to an open problem and beyond. *Ann. Statist.* **41** 177–195. [MR3059414](https://doi.org/10.1214/13-AOS1084) <https://doi.org/10.1214/13-AOS1084>
- [17] JIANG, J. (2017). *Asymptotic Analysis of Mixed Effects Models: Theory, Applications, and Open Problems*. Monographs on Statistics and Applied Probability **155**. CRC Press, Boca Raton, FL. [MR3676375](https://doi.org/10.1007/978-1-4899-4549-5)
- [18] MCCULLAGH, P. and NELDER, J. A. (1989). *Generalized Linear Models*. Monographs on Statistics and Applied Probability. CRC Press, London. (Second edition of [MR0727836](https://doi.org/10.1007/978-1-4899-3242-6)). [MR3223057](https://doi.org/10.1007/978-1-4899-3242-6) <https://doi.org/10.1007/978-1-4899-3242-6>
- [19] MCCULLOCH, C. (2008). Joint modelling of mixed outcome types using latent variables. *Stat. Methods Med. Res.* **17** 53–73. [MR2420190](https://doi.org/10.1177/0962280207081240) <https://doi.org/10.1177/0962280207081240>
- [20] MCCULLOCH, C. E., SEARLE, S. R. and NEUHAUS, J. M. (2008). *Generalized, Linear, and Mixed Models*, 2nd ed. Wiley Series in Probability and Statistics. Wiley, Hoboken, NJ. [MR2431553](https://doi.org/10.1007/978-1-4899-4549-5)
- [21] MILLER, J. J. (1977). Asymptotic properties of maximum likelihood estimates in the mixed model of the analysis of variance. *Ann. Statist.* **5** 746–762. [MR0448661](https://doi.org/10.2307/2335333)

- [22] MIN, M. (2007). *Asymptotic Normality in Generalized Linear Mixed Models*. ProQuest LLC, Ann Arbor, MI. Thesis (Ph.D.)—Univ. Maryland, College Park, MD. [MR2711359](#)
- [23] NIE, L. (2006). Strong consistency of the maximum likelihood estimator in generalized linear and nonlinear mixed-effects models. *Metrika* **63** 123–143. [MR2242536](#) <https://doi.org/10.1007/s00184-005-0001-3>
- [24] POLLARD, D. (1990). *Empirical Processes: Theory and Applications. NSF-CBMS Regional Conference Series in Probability and Statistics* **2**. IMS, Hayward, CA; Amer. Statist. Assoc., Alexandria, VA. [MR1089429](#)
- [25] SILVEY, S. D. (1961). A note on maximum-likelihood in the case of dependent random variables. *J. Roy. Statist. Soc. Ser. B* **23** 444–452. [MR0138158](#)
- [26] SUNG, Y. J. and GEYER, C. J. (2007). Monte Carlo likelihood inference for missing data models. *Ann. Statist.* **35** 990–1011. [MR2341695](#) <https://doi.org/10.1214/009053606000001389>
- [27] WALD, A. (1948). Asymptotic properties of the maximum likelihood estimate of an unknown parameter of a discrete stochastic process. *Ann. Math. Stat.* **19** 40–46. [MR0024114](#) <https://doi.org/10.1214/aoms/1177730288>
- [28] WALD, A. (1949). Note on the consistency of the maximum likelihood estimate. *Ann. Math. Stat.* **20** 595–601. [MR0032169](#) <https://doi.org/10.1214/aoms/1177729952>
- [29] WARTON, D. I., BLANCHET, F. G., O’HARA, R. B., OVASKAINEN, O., TASKINEN, S., WALKER, S. C. and HUI, F. K. (2015). So many variables: Joint modeling in community ecology. *Trends Ecol. Evol.* **30** 766–779.
- [30] WEISS, L. (1971). Asymptotic properties of maximum likelihood estimators in some nonstandard cases. *J. Amer. Statist. Assoc.* **66** 345–350. [MR0309242](#)
- [31] WHITE, H. (1982). Maximum likelihood estimation of misspecified models. *Econometrica* **50** 1–25. [MR0640163](#) <https://doi.org/10.2307/1912526>