

CONFIDENCE REGIONS NEAR SINGULAR INFORMATION AND BOUNDARY POINTS WITH APPLICATIONS TO MIXED MODELS

BY KARL OSKAR EKVALL^a AND MATTEO BOTTAI^b

*Division of Biostatistics, Institute of Environmental Medicine, Karolinska Institutet, ^akarl.oskar.ekvall@ki.se,
^bmatteo.bottai@ki.se*

We propose confidence regions with asymptotically correct uniform coverage probability of parameters whose Fisher information matrix can be singular at important points of the parameter set. Our work is motivated by the need for reliable inference on scale parameters close or equal to zero in mixed models, which is obtained as a special case. The confidence regions are constructed by inverting a continuous extension of the score test statistic standardized by expected information, which we show exists at points of singular information under regularity conditions. Similar results have previously only been obtained for scalar parameters, under conditions stronger than ours, and applications to mixed models have not been considered. In simulations our confidence regions have near-nominal coverage with as few as $n = 20$ observations, regardless of how close to the boundary the true parameter is. It is a corollary of our main results that the proposed test statistic has an asymptotic chi-square distribution with degrees of freedom equal to the number of tested parameters, even if they are on the boundary of the parameter set.

1. Introduction. In mixed models, the importance of a random effect is often assessed by inference on a variance or scale parameter. A parameter near zero typically indicates a weak effect and many tests for whether a variance is equal to zero have been proposed (Stram and Lee (1994, 1995), Baey, Cournède and Kuhn (2019), Chen, Xiao and Staicu (2019), Crainiceanu and Ruppert (2004), Drikvandi et al. (2013), Fitzmaurice, Lipsitz and Ibrahim (2007), Giampaoli and Singer (2009), Greven et al. (2008), Hall and Præstgaard (2001), Lin (1997), Qu, Guennel and Marshall (2013), Saville and Herring (2009), Sinha (2009), Stern and Welsh (2000), Verbeke and Molenberghs (2003), Wiencierz, Greven and Küchenhoff (2011), Wood (2013), Zhu and Zhang (2006)). In addition to being of practical interest, this case is of theoretical interest because the parameter is a boundary point of the parameter set and, consequently, asymptotic distributions of common test statistics are nonstandard. For example, the asymptotic distribution of the likelihood ratio test statistic for a variance equal to zero is a nontrivial mixture of chi-square distributions (Geyer (1994), Self and Liang (1987), Stram and Lee (1994)), whereas for a strictly positive variance it is a chi-square distribution with one degree of freedom. More generally, the asymptotic distributions of common test statistics under a sequence of parameters tending to a boundary point as the sample size increases, can be different depending on the rate of that convergence (Bottai (2003), Rotnitzky et al. (2000)). While this need not be an issue when testing a point null hypothesis, it complicates more ambitious inference: coverage probabilities of confidence regions obtained by inverting such test statistics often depend substantially on how close to the boundary the true parameter is, leading to unreliable inference. The coverage of boundary points is addressed by existing methods, but the coverage of points near the boundary is not. We address both using a connection between boundary points and points where the Fisher information matrix is singular, which we call critical points. More specifically, we show many boundary points

Received March 2021; revised January 2022.

MSC2020 subject classifications. Primary 62F25, 62F05; secondary 62J05.

Key words and phrases. Singular information, boundary points, confidence regions, mixed models.

of interest are also critical points and use this to construct confidence regions that (i) have asymptotically correct uniform coverage probability, (ii) have empirical coverage close to nominal in simulations, and (iii) are straightforward to implement for many mixed models, including the ubiquitous linear mixed model. We know of no other confidence regions with properties (i) and (ii) in the settings we consider. Because not all critical points are boundary points, the proposed regions are useful in many settings where methods for inference on the boundary do not apply.

To be more precise about the connections between boundary points, critical points and mixed models, suppose a parameter $\theta \in \mathbb{R}$ scales a random effect with mean zero and unit variance in a mixed model, implying θ^2 is a variance. For example, θ can be the coefficient of a random effect in a generalized linear mixed model. If the random effect has a distribution asymmetric around zero, inference on both the sign and magnitude of θ may be possible, in which case $\theta = 0$ is not a boundary point. In other settings, the sign is unidentifiable and inference on $\theta \geq 0$ and θ^2 essentially equivalent; $\theta = 0$ is a boundary point. Either way, we will show that in quite general mixed models $\theta = 0$ is a critical point. More generally, when θ is a vector of parameters whose j th element θ_j is a scale parameter, θ is often a critical point if $\theta_j = 0$. Whether in a mixed model or not, inference near critical points is known to be difficult: the likelihood ratio test statistic and the maximum likelihood estimator behave quite differently than under classical conditions (Rotnitzky et al. (2000)) and confidence regions obtained by inverting common test statistics such as the Wald, likelihood ratio and score standardized by observed information have incorrect coverage probabilities (Bottai (2003)). By contrast, we show that, under regularity conditions, (i) the score test statistic standardized by expected Fisher information has a continuous extension at critical points and, when inverted, (ii) that test statistic gives a confidence region with asymptotically correct uniform coverage probability on compact sets. That is, the confidence region $\mathcal{R}_n(\alpha)$ based on n observations with nominal level $(1 - \alpha) \in (0, 1)$ satisfies, for any compact subset C of the parameter set,

$$(1) \quad \lim_{n \rightarrow \infty} \inf_{\theta \in C} P_{\theta} \{ \theta \in \mathcal{R}_n(\alpha) \} = 1 - \alpha,$$

where the subscript θ on P indicates the data on which \mathcal{R}_n is based have the distribution indexed by θ . Importantly, C can include neighborhoods of boundary and critical points. It is an immediate corollary that the test rejecting a null hypothesis $\theta = \theta_0$ when $\theta_0 \notin \mathcal{R}_n(\alpha)$ has asymptotic size α for any θ_0 . These results apply to but are not restricted to mixed models. Moreover, in contrast to many methods for testing variance parameters in mixed models, ours in general does not require the implementation of simulation algorithms or computing the maximum likelihood estimator, which can be complicated in nonlinear mixed models.

The connection between singular information and boundary settings has been noticed previously (Chesher (1984), Cox and Hinkley (2000), Lee and Chesher (1986)), but results similar to ours have only been obtained for settings with a single scalar parameter (Bottai (2003)). We recover those results as special cases, and under weaker conditions. Asymptotic properties of maximum likelihood estimators and likelihood ratio test statistics have been established for the special case where the rank of the Fisher information matrix is one less than full (Rotnitzky et al. (2000)), but confidence regions were not considered. Notably, our theory does not require the Fisher information to have a particular rank and, indeed, we will see that in mixed models the rank is often full minus the number of scale parameters equal to zero.

We end this section with a simple example that illustrates how critical points often appear in mixed models. After the example, we give additional background and develop theory in Section 2. In Section 3, we discuss the application to mixed models and verify the conditions of the theory from Section 2 in two such models. Section 4 presents simulation results, Section 5 contains a data example and Section 6 concludes.




FIG. 1. *Log-likelihoods for two independent samples of $n = 100$ independent observations each, generated with $\theta = 1/\sqrt{10}$ and $r = 1$.*

EXAMPLE 1. Suppose, for $i = 1, \dots, n$ and $j = 1, \dots, r$,

$$Y_{i,j} = \theta W_i + E_{i,j},$$

where $\theta \in [0, \infty)$ and all W_i and $E_{i,j}$ are independent standard normal random variables. For example, r can be the number of observations in a cluster, n the number of clusters and the random effect W_i used to model heterogeneity between clusters or dependence between observations in the same cluster. The $Y_i = [Y_{i1}, \dots, Y_{ir}]^\top, i = 1, \dots, n$, are independent and multivariate normally distributed with mean zero and common covariance matrix $\Sigma(\theta) = \theta^2 1_r 1_r^\top + I_r$, where 1_r is an r -vector of ones and I_r the $r \times r$ identity matrix. With some algebra (Supplementary Material, Ekvall and Bottai (2022)), one can show the log-likelihood for one observation $y_i \in \mathbb{R}^r$ is

$$\log f_\theta(y_i) = -\frac{1}{2} \log(1 + \theta^2 r) - \frac{1}{2} \{y_i^\top y_i - (y_i^\top 1_r)^2 \theta^2 / (1 + r\theta^2)\}.$$

Differentiating with respect to θ gives the score for one observation:

$$s(\theta; y_i) = -\frac{r\theta}{1 + r\theta^2} + (y_i^\top 1_r)^2 \frac{\theta}{(1 + r\theta^2)^2}.$$

At $\theta = 0$, this score is zero for any $y_i \in \mathbb{R}^r$, and hence the Fisher information is zero; that is, $\theta = 0$ is a critical point. There are no other critical points because the second term of $s(\theta; Y_i), Y_i \sim f_\theta$, has positive variance when $\theta \neq 0$.

Figure 1 shows two (pseudo) randomly generated realizations of the log-likelihood in this example. For one data set, the critical point is a global maximizer and for the other a local minimizer. One can show that if the true θ is small, both types of outcomes have probability approximately 1/2. In particular, the score always vanishes at $\theta = 0$ and the maximum likelihood estimator for θ is zero with probability approximately 1/2. The maximum likelihood estimator’s mass at zero gives some intuition for why confidence regions that directly or indirectly use asymptotic normality of that estimator can have poor coverage properties near the critical point (see Rotnitzky et al., 2000 and Bottai, 2003, for details). In this example, the critical point is at the boundary since we assumed $\theta \geq 0$ for identifiability, but $\theta = 0$ would still be a critical point if the W_i had an asymmetric distribution and the sign of θ were identifiable.

2. Inference near critical points.

2.1. *Definitions and assumptions.* Suppose, independently for $i = 1, \dots, n, Y_i \in \mathbb{R}^{r_i}, r_i \geq 1$, has density f_θ^i against a dominating measure γ_i , for θ in a parameter set $\Theta \subseteq \mathbb{R}^d$. For simplicity, we often write $f_\theta(y_i)$ in place of $f_\theta^i(y_i)$, where y_i is an arbitrary realization of Y_i . Let $y^n = (y_1, \dots, y_n)$ be a realization of $Y^n = (Y_1, \dots, Y_n)$ and ∇ denote the derivative operator with respect to θ . Let also $\ell^i(\theta; y_i) = \log f_\theta(y_i), \ell_n(\theta; y^n) = \sum_{i=1}^n \ell^i(\theta; y_i), s^i(\theta; y_i) =$

$\nabla \ell^i(\theta; y_i)$, $s_n(\theta; y^n) = \nabla \ell_n(\theta; y^n)$, $\mathcal{I}^i(\theta) = \text{cov}_\theta\{s^i(\theta; Y_i)\}$ and $\mathcal{I}_n(\theta) = \text{cov}_\theta\{s_n(\theta; Y^n)\}$, where the subscript on cov indicates $Y^n = (Y_1, \dots, Y_n)$ has the distribution indexed by θ ; we may omit such subscripts when it has been explicitly stated which distribution the random variables have. Focus will be on inference near points where the Fisher information \mathcal{I}_n is singular.

DEFINITION 1. We say a point $\theta \in \Theta$ is critical if $\mathcal{I}_n(\theta)$ is singular and noncritical otherwise.

Denote an arbitrary set of orthonormal eigenvectors of $\mathcal{I}_n(\theta)$ by $\{v_{\theta 1}^n, \dots, v_{\theta d}^n\}$. Then Definition 1 says, equivalently, that θ is a critical point if $s_n(\theta; Y^n)^\top v_{\theta j}^n$ is constant almost surely \mathbb{P}_θ for at least one j . In our motivating examples, θ is the parameter vector in a mixed model but there are other potential applications for the theory in this section. For example, Azzalini (2014), Section 3.1.3, consider a three-parameter multivariate skew-normal distribution with singular Fisher information.

As mentioned in the Introduction, inverting common test statistics typically does not give confidence regions satisfying (1) for subsets of the parameter set including critical points (Bottai (2003)). To address this, we will consider the score test statistic standardized by expected Fisher information:

$$(2) \quad T_n(\theta; y^n) = s_n(\theta; y^n)^\top \mathcal{I}_n(\theta)^{-1} s_n(\theta; y^n).$$

The right-hand side of (2) is undefined at critical θ since $\mathcal{I}_n(\theta)$ is not invertible there. When θ is a scalar parameter, the following calculation can be formalized (proof of Theorem 2.1) to define a continuous extension at critical θ :

$$\begin{aligned} \lim_{\delta \rightarrow 0} T_n(\theta + \delta; y^n) &= \lim_{\delta \rightarrow 0} \left[\{s_n(\theta + \delta; y^n)/\delta\}^2 \text{var}_{\theta+\delta}\{s_n(\theta + \delta; Y^n)/\delta\}^{-1} \right] \\ &= \{\nabla s_n(\theta; y^n)\}^2 \text{var}_\theta\{\nabla s_n(\theta; Y^n)\}^{-1}, \end{aligned}$$

where the first line results from dividing and multiplying $T_n(\theta + \delta; y^n)$ by δ^2 , assuming $\theta + \delta$ is a noncritical point, and the second line uses the definition of derivative of $\theta \mapsto s(\theta; \cdot)$ at θ together with regularity conditions, ensuring among other things that $s_n(\theta; y^n) = 0$ at critical θ . That is, at critical points the continuous extension of the test statistic for one parameter is based on the second derivative of the log-likelihood. Notably, if expected information is replaced by observed, which is not in general zero at critical points, the limit is typically zero.

Now, our interest is twofold, namely conditions that ensure (i) T_n has a continuous extension to critical points when $d \geq 1$ and (ii) confidence regions obtained by inverting that extension satisfy (1). Both (i) and (ii) are substantially more complicated when $d > 1$ than when $d = 1$ since the eigenvectors and the rank of the Fisher information become important.

We will use the following assumptions.

ASSUMPTION 1. For every $i = 1, 2, \dots$, the distribution $f_\theta(y_i)\gamma_i(dy_i)$ has the same null sets for every $\theta \in \Theta$.

In light of Assumption 1, we will often write “almost every y_i ” without specifying the measure, implicitly referring to $f_\theta(y_i)\gamma_i(dy_i)$ for any $\theta \in \Theta$ or the corresponding product measure when such statements are about y^n . Notably, the null sets can be different for different i . Indeed, the Y_i need not even take values in the same space.

ASSUMPTION 2. For every $\theta' \in \Theta$, there exists an open ball $B = B(\theta') \subseteq \mathbb{R}^d$ centered at θ' on which, for every $i = 1, 2, \dots$ and almost every y_i , partial derivatives of $\theta \mapsto \ell^i(\theta; y_i)$

of an order $k = k(\theta') \geq 2$ exist and are jointly continuous in (θ, y_i) . Moreover, there exists a $\delta = \delta(\theta') > 0$ such that every partial derivative of order at most k satisfies

$$(3) \quad \sup_{i=1,2,\dots} \sup_{\theta \in B \cap \Theta, \tilde{\theta} \in B} \int \left| \frac{\partial^l}{\partial \theta_1^{l_1} \dots \partial \theta_d^{l_d}} \ell^i(\theta; y_i) \Big|_{\theta=\tilde{\theta}} \right|^{2+\delta} f_\theta(y_i) \gamma_i(dy_i) < \infty,$$

where l_1, \dots, l_d are nonnegative integers summing to $l \leq k$.

The ball B in Assumption 2 can include points not in Θ . Then the assumption should be understood as saying there exist extensions of the partial derivatives in (3) to $B \times \mathcal{Y}_i$, for some set \mathcal{Y}_i of full measure, satisfying the outlined conditions. We have sacrificed some generality for clarity in that, as will be clear later, the moment condition (3) can be weakened to apply to only certain partial derivatives of the order k . One consequence of Assumption 2 is that $\mathbb{E}_\theta \{s_n(\theta; Y^n)\} = 0$ (Lemma 0.3 in the Supplementary Material, Ekvall and Bottai (2022)). Hence, $\mathcal{I}_n(\theta)v_{\theta_j}^n = 0$ is equivalent to $s_n(\theta; y^n)^\top v_{\theta_j}^n = 0$ for almost every y^n .

For symmetric matrices A and B , we write $A \succeq B$ or $B \preceq A$ if $A - B$ is positive semi-definite.

ASSUMPTION 3. There exist continuous $c_1, c_2 : \Theta \rightarrow (0, \infty)$ such that, for every $i = 1, 2, \dots$ and $\theta \in \Theta$,

$$c_1(\theta)\mathcal{I}^1(\theta) \preceq \mathcal{I}^i(\theta) \preceq c_2(\theta)\mathcal{I}^1(\theta).$$

Assumption 3 holds with $c_1 = c_2 \equiv 1$ if the Y_i are identically distributed since in that case $\mathcal{I}^i = \mathcal{I}^1$. More generally, it controls the null space of \mathcal{I}_n : if $\mathcal{I}_n(\theta)v = 0$ for some n and $v \in \mathbb{R}^d$, then $\mathcal{I}^1(\theta)v = 0$ and, by Assumption 3, $\mathcal{I}^i(\theta)v = 0$ for every i . Thus, the critical points in Definition 1 do not depend on n . At noncritical points, Assumption 3 says, loosely speaking, that information in individual observations does not grow without bound or tend to zero. The choice of \mathcal{I}^1 is arbitrary in the sense that Assumption 3 holds as stated if and only if it holds with \mathcal{I}^1 replaced by any other \mathcal{I}^i .

The next assumption requires some more notation. Let $\nabla_j^l, j = 1, \dots, d, l = 0, 1, \dots$, denote the l th order derivative operator with respect to θ_j , with $\nabla_j = \nabla_j^1$. For example, $\nabla_j^l \ell_n(\tilde{\theta}; y^n) = \partial^l \ell_n(\theta; y^n) / \partial \theta_j^l |_{\theta=\tilde{\theta}}$. For every $\theta \in \Theta$ and $n = 1, 2, \dots$, let $k_j = k_j(\theta, n) \in \{1, 2, \dots\}$ ($j = 1, \dots, d$) be integers to be specified shortly and define, for $i = 1, \dots, n$,

$$\tilde{s}_n^i(\theta; y_i) = [\{\nabla_1^{k_1-1} s^i(\theta; y_i)\}^\top v_{\theta_1}^n, \dots, \{\nabla_d^{k_d-1} s^i(\theta; y_i)\}^\top v_{\theta_d}^n]^\top \in \mathbb{R}^d.$$

Let also

$$\tilde{s}_n(\theta; y^n) = \sum_{i=1}^n \tilde{s}_n^i(\theta; y_i); \quad \tilde{\mathcal{I}}_n(\theta) = \text{cov}_\theta \{\tilde{s}_n(\theta; Y^n)\}.$$

We call $\tilde{s}_n(\theta; y^n)$ a modified score because it can be understood as a two-step modification of $s_n(\theta; y^n)$: First, $s_n(\theta; y^n)$ is rotated to $[v_{\theta_1}^n, \dots, v_{\theta_d}^n]^\top s_n(\theta; y^n)$, in which the elements are uncorrelated linear combinations of first-order derivatives of $\ell_n(\theta; y^n)$. Second, if $k_j > 1$ for some j , the linear combination in the j th element is replaced by a linear combination of higher-order partial derivatives. The idea to be formalized is to replace elements in $[v_{\theta_1}^n, \dots, v_{\theta_d}^n]^\top s_n(\theta; y^n)$ that are zero for almost every y^n at a critical point by linear combinations of higher-order derivatives that are not. Note also $\tilde{s}_n(\theta; y_i)$ and $\tilde{\mathcal{I}}_n(\theta)$ depend on the $v_{\theta_j}^n$. In particular, the set $\{v_{\theta_1}^n, \dots, v_{\theta_d}^n\}$ is not uniquely determined in general and different choices lead to different $\tilde{s}_n(\theta; y^n)$. Our next assumption says three technical conditions must hold for at least one choice.

ASSUMPTION 4. For every $\theta \in \Theta$ and $n = 1, 2, \dots$, there exist integers $k_j = k_j(\theta, n) \in \{1, \dots, d\}$ ($j = 1, \dots, d$) and a set of orthonormal eigenvectors $\{v_{\theta_1}^n, \dots, v_{\theta_d}^n\}$ of $\mathcal{I}_n(\theta)$ such that:

- (i) $\tilde{\mathcal{I}}_n(\theta)$ is positive definite;
- (ii) $\{\nabla_j^l s_n(\tilde{\theta}; y^n)\}^\top v_{\theta_j}^n = 0$, $l = 0, \dots, k_j - 2$, for every $\tilde{\theta}$ with $\tilde{\theta}_j = \theta_j$ and almost every y^n ;
- (iii) $k_j(\theta, n)$ is upper bounded by the $k = k(\theta)$ in Assumption 2 for every j, θ and n .

Condition (i) is essentially a more general, multidimensional version of the requirement by Bottai (2003) that the second derivative of the log-likelihood has positive variance at critical points. Indeed, at noncritical θ (i) holds with $k_j = 1$ for all j since $\tilde{s}_n(\theta; Y^n) = [v_{\theta_1}^n, \dots, v_{\theta_d}^n]^\top s_n(\theta; Y^n)$ has diagonal covariance matrix with the eigenvalues of $\mathcal{I}_n(\theta)$ on the diagonal. Conversely, at critical θ at least one of those eigenvalues is zero, and hence it must be that $k_j > 1$ for at least one j in order for (i) to hold. For example, we will see mixed models where (i) holds at a critical θ with $v_{\theta_1}^n = e_1 = [1, 0, \dots, 0]^\top \in \mathbb{R}^d$, $k_1 = 2$ and $k_j = 1$, $j = 2, \dots, d$. Then the first element of $\tilde{s}_n(\theta, y^n)$ is $\nabla_1^2 \ell(\theta; y^n)$ while the remaining elements are linear combinations of the elements of $s_n(\theta; y^n)$. While our theory only requires existence, implementing our method in practice can, at least in some models, require specification of a set of k_j and numerical computation of a set of v_{θ_j} satisfying (i); we discuss this further in Section 3.

Condition (ii) is vacuously satisfied at noncritical points since it only applies to $k_j \geq 2$. Consequently, Assumption 4 is weaker than assuming $\mathcal{I}_n(\theta)$ is positive definite for every θ as is common in classical theory. To understand condition (ii) more generally, suppose θ is a critical point, $k_1 = 2$ and $v_{\theta_1}^n$ one of the eigenvectors of $\mathcal{I}_n(\theta)$ corresponding to the eigenvalue zero. As noted following Assumption 2, this implies $s_n(\theta; y^n)^\top v_{\theta_1}^n = 0$ for almost every y^n ; condition (ii) says the same must hold if the score is evaluated at any other $\tilde{\theta}$ with first element $\tilde{\theta}_1 = \theta_1$. This motivates the following definition.

DEFINITION 2. We say θ_j is a critical element with corresponding critical (eigen)vector v if $\mathcal{I}_n(\tilde{\theta})v = 0$ at every $\tilde{\theta}$ with $\tilde{\theta}_j = \theta_j$.

In general, an eigenvector of the Fisher information with vanishing eigenvalue need not be a critical vector. For example, if v_1 and v_2 are critical vectors with corresponding critical elements θ_1 and θ_2 , then any linear combination of v_1 and v_2 is also an eigenvector with vanishing eigenvalue at every $\tilde{\theta}$ where both $\tilde{\theta}_1 = \theta_1$ and $\tilde{\theta}_2 = \theta_2$; but there is not a corresponding critical element.

The possibility that $k_j > 2$ in condition (ii) means that, if condition (i) is not satisfied with $k_j \in \{1, 2\}$ for all j , then one may pass to higher order derivatives as long as the corresponding ones of lower order are zero for almost every y^n .

Finally, we note Assumption 4 implies the rank of $\mathcal{I}_n(\theta)$ is d minus the number of critical elements and that the assumption is sensitive to parameterization. For example, the model considered by Azzalini (2014), Section 3.1, does not satisfy Assumption 4 in the first parameterization discussed by the authors, but can be made to by a simple reparameterization. Because the test statistic in (2) is invariant under differentiable reparameterizations with full rank Jacobian, it suffices to verify the conditions in one parameterization for the results to apply more generally.

ASSUMPTION 5. The noncritical points are dense in Θ ; that is, for any $\theta \in \Theta$, there exists a sequence $\{\theta_m\} \in \Theta$ of noncritical points tending to θ .

In many settings, the sets of critical elements are discrete subsets of \mathbb{R} . Then, if θ is a critical point with only critical element θ_j , it is often possible to verify Assumption 5 with, for example, $\theta_m = \theta + m^{-1}e_j$.

2.2. *Continuous extension.* Our purpose in this section is to prove the following theorem.

THEOREM 2.1. *If Assumptions 1–2 and 4–5 hold, then for any $n \geq 1$ there is a set \mathcal{Y}^n of full measure such that $T_n(\cdot; \cdot)$ has a continuous extension on $\Theta \times \mathcal{Y}^n$, and that extension is*

$$T_n(\theta; y^n) = \tilde{s}_n(\theta; y^n)\tilde{\mathcal{I}}_n(\theta)^{-1}\tilde{s}_n(\theta; y^n).$$

Here and in what follows, we use the same notation for the test statistic in (2) and its continuous extension given by Theorem 2.1. The expressions in Theorem 2.1 and (2) agree at noncritical θ since $\tilde{s}_n(\theta; y^n) = [v_{\theta_1}^n, \dots, v_{\theta_d}^n]^\top s_n(\theta; y^n)$ there, and the premultiplication of any invertible matrix is canceled when standardizing by the covariance matrix. Observe the continuity in Theorem 2.1 is jointly in the parameter and data. Before giving a proof, we state and discuss some intermediate results used in that proof. The first is a lemma which, loosely speaking, says that if the noncritical points are dense, then it is enough to establish continuity along sequences of noncritical points for it to hold more generally. Proofs of formally stated results are in the Supplementary Material (Ekvall and Bottai (2022)) if not given here.

LEMMA 2.2. *Suppose Assumption 5 holds and that, for every $\theta \in \Theta$ and y^n in some $\mathcal{Y}^n \subseteq \mathbb{R}^{r_1} \times \dots \times \mathbb{R}^{r_n}$, with $n \geq 1$ fixed, $\lim_{m \rightarrow \infty} T_n(\theta_m; y_m^n)$ exists and is the same for all sequences $\{\theta_m\} \in \Theta$ of noncritical points tending to θ and sequences $\{y_m^n\} \in \mathcal{Y}^n$ tending to y^n ; then $T_n(\cdot; \cdot)$ has a continuous extension on $\Theta \times \mathcal{Y}^n$.*

In order to use Lemma 2.2 to prove Theorem 2.1, one must show that T_n converges along sequences of noncritical points tending to critical points. Evaluating the score at such sequences gives a sequence of score vectors whose covariance matrices are nonsingular but tend to a singular limit. The following lemma says those score vectors can be scaled to tend to a limit with positive definite covariance matrix.

LEMMA 2.3. *Suppose Assumptions 1, 2 and 4 hold and let $\{\theta_m\} \in \Theta$ be a sequence of noncritical points tending to some $\theta \in \Theta$. Then there exist sequences of nonzero constants $\{a_{mj}\}$, $j = 1, \dots, d$, and a set \mathcal{Y}^n of full measure such that, for any $\{y_m^n\} \in \mathcal{Y}^n$ tending to a $y^n \in \mathcal{Y}^n$ as $m \rightarrow \infty$,*

$$a_{mj}s_n(\theta_m; y_m^n)^\top v_{\theta_j}^n \rightarrow \tilde{s}_{nj}(\theta; y^n),$$

where \tilde{s}_{nj} is the j th element of \tilde{s}_n in Assumption 4.

PROOF. Let \mathcal{Y}^n be the intersection of the sets of full measure in Assumptions 2 and 4. Since the limit point θ and n are fixed, denote $v_j = v_{\theta_j}^n$ for simplicity. For an arbitrary j and all large enough m , condition (ii) of Assumption 2 lets us apply Taylor’s theorem with Lagrange-form remainder to the map $\theta_{mj} \mapsto v_j^\top s_n(\theta_m; y_m^n)$ to get that $v_j^\top s_n(\theta_m; y_m^n)$ is equal to

$$\sum_{l=0}^{k_j-2} \frac{(\theta_{mj} - \theta_j)^l v_j^\top \{\nabla_j^l s_n(\theta_m^{(j)}; y_m^n)\}}{l!} + \frac{(\theta_{mj} - \theta_j)^{k_j-1} v_j^\top \{\nabla_j^{k_j-1} s_n(\tilde{\theta}_m^{(j)}; y_m^n)\}}{(k_j - 1)!},$$

where $\theta_m^{(j)}$ and $\tilde{\theta}_m^{(j)}$ are θ_m with θ_{mj} replaced by, respectively, θ_j and a point between θ_{mj} and θ_j ; and k_j is selected in accordance with Assumption 4 at θ . By condition (ii) of that

assumption, the first $k_j - 1$ terms in the last display vanish. Now set $a_{mj} = (k_j - 1)! / (\theta_{mj} - \theta_j)^{k_j - 1}$, which is always defined since θ_m is noncritical, and hence $\theta_{mj} \neq \theta_j$ by condition (ii) of Assumption 4. Then $a_{mj} v_j^\top s_n(\theta_m; y_m^n) = v_j^\top \{\nabla_j^{k_j - 1} s_n(\tilde{\theta}_m^{(j)}; y_m^n)\}$, which has the desired limit by continuity of partial derivatives given by Assumption 2. \square

The importance of controlling the behavior of the eigenvectors of the Fisher information near critical points is highlighted by the proof of Lemma 2.3: the first $k_j - 1$ terms in the Taylor expansions need not vanish if the eigenvectors depend on θ in a way violating condition (ii) of Assumption 4. We are ready to prove Theorem 2.1.

PROOF OF THEOREM 2.1. Fix $n, \theta \in \Theta$ and $y^n \in \mathcal{Y}^n$, where \mathcal{Y}^n is the intersection of the sets of full measure given by Assumptions 2 and 4. Pick $V_n = [v_{\theta_1}^n, \dots, v_{\theta_d}^n] \in \mathbb{R}^{d \times d}$ satisfying Assumption 4. By Lemma 2.2, it suffices to show

$$\lim_{m \rightarrow \infty} s_n(\theta_m; y_m^n)^\top \mathcal{I}(\theta_m)^{-1} s_n(\theta_m; y_m^n) = \tilde{s}(\theta; y^n)^\top \tilde{\mathcal{I}}_n(\theta)^{-1} \tilde{s}(\theta; y^n)$$

for an arbitrary sequence of noncritical points $\{\theta_m\}$ tending to θ and $\{y_m^n\}$ tending to y^n . We note it is enough to establish the limit for any one V_n satisfying Assumption 4 since the sequence does not depend on the choice, and hence neither can the limit. Let $A_m = \text{diag}(a_{m1}, \dots, a_{md})$ be defined by the sequences of constants given by Lemma 2.3. Then $A_m V_n^\top$ is invertible and consequently

$$T_n(\theta_m; y_m^n) = \{A_m V_n^\top s_n(\theta_m; y_m^n)\}^\top \{A_m V_n^\top \mathcal{I}_n(\theta_m) A_m V_n\}^{-1} A_m V_n^\top s_n(\theta_m; y_m^n).$$

Lemma 2.3 says, for the \tilde{s}_n in Assumption 4, $A_m V_n^\top s_n(\theta_m, y_m^n) \rightarrow \tilde{s}_n(\theta, y^n)$. Thus, we are done if we can show $A_m V_n^\top \mathcal{I}(\theta_m) V_n A_m = \text{cov}\{A_m V_n^\top s(\theta_m, Y_m^n)\} \rightarrow \tilde{\mathcal{I}}_n(\theta)$, where $Y_m^n = (Y_{m1}, \dots, Y_{mn})$ has the distribution indexed by θ_m . To that end, note $\theta_m \rightarrow \theta$ implies $f_{\theta_m}^i \rightarrow f_\theta^i$ pointwise in y_i for every i by continuity implied by Assumption 2. Hence, the joint density for Y_m^n tends pointwise to that of a Y^n with distribution indexed by θ . Thus, $Y_m^n \rightarrow Y^n$ in total variation by Scheffe’s theorem (Billingsley (1995), Theorem 16.12), and hence also in distribution. To show the desired convergence of covariance matrices, we may thus assume, by Skorokhod’s representation theorem (Billingsley (1999), Theorem 6.7), that $Y_m^n \rightarrow Y^n$ almost surely. But then $A_m V_n^\top s_n(\theta_m; Y_m^n) \rightarrow \tilde{s}(\theta, Y^n)$ almost surely by Lemma 2.3. Now convergence of the covariance matrices follows if the elements of the sequence $\{A_m V_n^\top s_n(\theta_m; Y_m^n) s_n(\theta_m; Y_m^n)^\top V_n A_m\}$ are uniformly integrable (Billingsley (1995), Theorem 25.12). To show they are, note the j th element of $A_m V_n^\top s_n(\theta_m, Y_m^n)$ is $\{\nabla_j^{k_j - 1} s_n(\tilde{\theta}_m^{(j)}; Y_m^n)\}^\top v_{\theta_j}$, with $\tilde{\theta}_m^{(j)}$ selected as in Lemma 2.3. Thus, the (j, l) th element of $A_m V_n^\top s_n(\theta_m; Y_m^n) s_n(\theta_m; Y_m^n)^\top V A_m$ is

$$\{\nabla_j^{k_j - 1} s_n(\tilde{\theta}_m^{(j)}; Y_m^n)\}^\top v_{\theta_j} \{\nabla_l^{k_l - 1} s_n(\tilde{\theta}_m^{(l)}; Y_m^n)\}^\top v_{\theta_l},$$

which has uniformly bounded $(1 + \delta/2)$ th moment by the Cauchy–Schwarz inequality and Assumption 2. From this the desired uniform integrability follows (Billingsley (1995), 25.13) and that completes the proof. \square

2.3. *Asymptotic uniform coverage probability.* We now turn to confidence regions obtained by inverting the continuous extension T_n . Specifically, for $\alpha \in (0, 1)$, define

$$(4) \quad \mathcal{R}_n(\alpha) = \{\theta \in \Theta : T_n(\theta; Y^n) \leq q_{d, 1-\alpha}\},$$

where $q_{d, 1-\alpha}$ is the $(1 - \alpha)$ th quantile of the chi-square distribution with d degrees of freedom. We have the following main result of the section.

THEOREM 2.4. *Under Assumptions 1–5, the confidence region $\mathcal{R}_n(\alpha)$ in (4) has asymptotically correct uniform coverage probability on compact sets; that is, it satisfies (1).*

We need some intermediate results before proving Theorem 2.4. Our strategy will be to prove that, for any compact $C \subseteq \Theta$ and $\alpha \in (0, 1)$,

$$(5) \quad \lim_{n \rightarrow \infty} \sup_{\theta \in C} |\mathbb{P}_\theta \{\theta \in \mathcal{R}_n(\alpha)\} - (1 - \alpha)| = 0,$$

which implies (1). The following two lemmas let us focus on convergence along sequence of noncritical points as in the previous section, but now taking the stochastic properties of the data into account.

LEMMA 2.5. *Equation (5) holds for every compact $C \subseteq \Theta$ if and only if, for every convergent sequence $\{\theta_n\} \in \Theta$ as $n \rightarrow \infty$,*

$$(6) \quad T_n(\theta_n; Y_n^n) \rightsquigarrow \chi_d^2,$$

where $Y_n^n = (Y_{n1}, \dots, Y_{nn})$ has the distribution indexed by θ_n .

The proof of Lemma 2.5 essentially amounts to showing continuous convergence is equivalent to uniform convergence on compact sets, the function of interest being $\theta \mapsto \mathbb{P}_\theta \{\theta \in \mathcal{R}_n(\alpha)\}$. Lemma 2.5 suggests, roughly speaking, that to get reliable confidence regions the distribution of the test statistic should be the same regardless of how close to the critical point the true parameter is.

The next lemma says it suffices to consider sequences of noncritical points. Specifically, we need not treat the asymptotic distribution at critical points separately.

LEMMA 2.6. *If Assumptions 1–5 hold and (6) holds for every convergent sequence $\{\theta_n\} \in \Theta$ of noncritical points, then (6) holds for any convergent sequence in Θ .*

PROOF. Let F_n denote the cumulative distribution function of $T_n(\theta_n; Y_n^n)$ and let F denote that of χ_d^2 . Assumption 5 says that, for every fixed n , we can pick a sequence of noncritical points $\{\theta_n^m\}$ tending to θ_n as $m \rightarrow \infty$ with n fixed. Assumptions 1–5 ensure Theorem 2.1 holds and this implies, essentially by Slutsky and continuous mapping theorems (Lemma 0.2 in the Supplementary Material, Ekvall and Bottai (2022)),

$$T_n(\theta_n^m; Y_n^m) \rightsquigarrow T_n(\theta_n; Y_n^n), \quad m \rightarrow \infty,$$

where $Y_n^m = (Y_{m1}, \dots, Y_{mn})$ has the distribution indexed by θ_n^m . Thus, the corresponding cumulative distribution functions $\{F_n^m\}$ tend to F_n at every point of continuity of F_n as $m \rightarrow \infty$ with n fixed. Let D_n be the set of discontinuities of F_n and $D = \bigcup_n D_n$. Since any cumulative distribution function has at most countably many discontinuities, D is countable as a countable union of countable sets. Now for any $t \in \mathbb{R} \setminus D$, we can pick, for every n , an $m = m(n)$ large enough that $\|\theta_n^{m(n)} - \theta_n\| \leq 1/n$ and $|F_n^{m(n)}(t) - F_n(t)| \leq 1/n$. We then have by the triangle inequality,

$$|F_n(t) - F(t)| \leq 1/n + |F_n^{m(n)}(t) - F(t)|,$$

which tends to zero as $n \rightarrow \infty$ by the assumption that (6) holds along sequences of noncritical points since F is continuous; in particular, t is a point of continuity of F . The proof is completed by observing that, since D is countable, $\mathbb{R} \setminus D$ is dense in \mathbb{R} , and hence the convergence in fact holds at every $t \in \mathbb{R}$ (Fristedt and Gray (1997), Proposition 2, Chapter 14). □

We are ready to prove Theorem 2.4

PROOF OF THEOREM 2.4. By Lemma 2.6, it suffices to consider an arbitrary sequence $\{\theta_n\}$ of noncritical points tending to a $\theta \in \Theta$. Let $U_n = A_n V_n^\top s_n(\theta_n; Y_n^n)$, where $V_n = [v_{\theta_1}^n, \dots, v_{\theta_d}^n] \in \mathbb{R}^{d \times d}$ satisfies Assumption 4 at θ for every n , $A_n = \text{diag}(a_{n1}, \dots, a_{nd})$ is a scaling matrix defined by the $\{a_{nj}\}$ given by Lemma 2.3 (with the index $m = n$), and $Y_n^n = (Y_{n1}, \dots, Y_{nn})$ has the distribution indexed by θ_n . By reordering the elements of θ if necessary, we may partition $V_n = [V_{n1}, V_{n2}]$ and have that the columns of V_{n1} are the critical vectors at θ . As noted following Assumption 3, these actually do not depend on n , but V_{n2} may. With these definitions, $T_n(\theta_n; Y_n^n) = U_n^\top \text{cov}(U_n)^{-1} U_n$, and hence the continuous mapping theorem (Billingsley (1999), Theorem 2.7) says it suffices to show $\text{cov}(U_n)^{-1/2} U_n \rightsquigarrow \mathcal{N}(0, I_d)$.

Let $U_{nt} = t^\top U_n$ for an arbitrary $t \in \mathbb{R}^d$ with $\|t\| = 1$. We start by verifying Lyapunov’s conditions (Billingsley (1995), Theorem 27.3) for $U_{nt}/\sigma_{nt} \rightsquigarrow \mathcal{N}(0, 1)$, where $\sigma_{nt}^2 = \text{var}(U_{nt})$. First, $\mathbb{E}(U_{nt}) = 0$ since $\mathbb{E}\{s_n(\theta_n; Y_n^n)\} = 0$. Thus, it suffices to show $\sigma_{nt}^2 \geq \epsilon n$ for some $\epsilon > 0$ and $\mathbb{E}(|U_{nti}|^{2+\delta}) \leq M$ for some $\delta > 0$ and $M < \infty$, where $U_{nti} = t^\top A_n V_n^\top s^i(\theta_n; Y_{ni})$ is the i th summand of U_{nt} . For the latter we have, since $\|t\| = \|v_{nj}\| = 1$, using the triangle and generalized mean inequalities, with $\tilde{\theta}^{(j)}$ as in the proof of Lemma 2.3,

$$|U_{nti}|^{2+\delta} \leq \left(\sum_{j=1}^d |v_{nj}^\top \nabla_j^{k_j-1} s^i(\tilde{\theta}^{(j)}; Y_{ni})| \right)^{2+\delta} \leq 4^{1+\delta} \sum_{j=1}^d \sum_{l=1}^d |\nabla_j^{k_j-1} \nabla_l \ell^i(\tilde{\theta}^{(j)}; Y_{ni})|^{2+\delta},$$

whose expectation is less than some $M < \infty$ for all large enough n by Assumptions 2 and 4.

Next, by Assumption 3,

$$\sigma_{nt}^2 \geq nc_1(\theta_n) t^\top A_n V_n^\top \mathcal{I}^1(\theta_n) V_n A_n t \geq nc_1(\theta_n) \underline{\epsilon}\{A_n V_n^\top \mathcal{I}^1(\theta_n) V_n A_n\},$$

where $\underline{\epsilon}(\cdot)$ is the smallest eigenvalue. Since c_1 is continuous and positive at θ , the right-hand side is greater than ϵn for some $\epsilon > 0$ if $\liminf_{n \rightarrow \infty} \underline{\epsilon}\{A_n V_n^\top \mathcal{I}^1(\theta_n) V_n A_n\} > 0$. We will prove this by contradiction. To that end, suppose $\liminf_{n \rightarrow \infty} \underline{\epsilon}\{A_n V_n^\top \mathcal{I}^1(\theta_n) V_n A_n\} = 0$ and extract a subsequence tending to zero. Then, by compactness of the set of semiorthogonal matrices and the Bolzano–Weierstrass property (Folland (2007), Theorem 0.25), there is a further subsequence along which V_{n2} tends to a semiorthogonal V_2 . Along this subsequence, by arguments almost identical to those in the proof of Theorem 2.1, $A_n V_n^\top s_1(\theta_n; Y_{n1})$ tends in distribution to a random vector with positive definite covariance matrix, and $A_n V_n^\top \mathcal{I}^1(\theta_n) V_n A_n$ tends to that covariance matrix. Thus, along the subsequence, by Weyl’s inequalities (Bhatia (1997), Section III.2), $\underline{\epsilon}\{A_n V_n^\top \mathcal{I}^1(\theta_n) V_n A_n\}$ tends to a strictly positive number, which is the desired contradiction.

We have proven $U_{nt}/\sigma_{nt} \rightsquigarrow \mathcal{N}(0, 1)$, from which $\text{cov}(U_n)^{-1/2} U_n \rightsquigarrow \mathcal{N}(0, I_d)$ follows (Biscio, Poinas and Waagepetersen (2018), Lemma 2.1) if

$$0 < \liminf_{n \rightarrow \infty} \underline{\epsilon}\{n^{-1} \text{cov}(U_n)\} \leq \limsup_{n \rightarrow \infty} \bar{\epsilon}\{n^{-1} \text{cov}(U_n)\} < \infty,$$

where $\bar{\epsilon}(\cdot)$ is the maximum eigenvalue. For the first inequality, Assumption 3 gives $n^{-1} \text{cov}(U_n) \geq c_1(\theta_n) A_n V_n^\top \mathcal{I}^1(\theta_n) V_n A_n$, and we have already shown the right-hand side has smallest eigenvalue asymptotically bounded away from zero. A similar argument, using that Assumption 3 implies $\bar{\epsilon}\{n^{-1} \text{cov}(U_n)\} \leq c_2(\theta_n) \bar{\epsilon}\{A_n V_n^\top \mathcal{I}^1(\theta_n) V_n A_n\}$, establishes the upper bound and this completes the proof. \square

Theorems 2.1 and 2.4 have the following corollary, which recovers a result of Bottai (2003) but with several conditions weakened. The proof is a straightforward verification of Assumptions 2–5, and hence omitted.

COROLLARY 2.7. *If $d = 1$; the Y_i are identically distributed; Assumption 1 holds; for every $\theta' \in \Theta$ there exists an open ball $B = B(\theta') \subseteq \mathbb{R}$ centered at θ' such that (i) $\nabla^2 \ell^1(\theta; y_1)$ exists on B for almost every y_1 and is continuous in (θ, y_1) , and (ii) it holds that*

$$\sup_{\theta \in B \cap \Theta, \tilde{\theta} \in B} \int |\nabla^2 \ell^1(\tilde{\theta}; y_1)|^{2+\delta} f_{\theta}(y_1) \gamma_1(dy_1) < \infty$$

for some $\delta = \delta(\theta') > 0$; and $\text{var}_{\theta} \{\nabla^2 \ell^1(\theta; Y_1)\} > 0$ for every critical θ ; then the conclusions of Theorems 2.1 and 2.4 hold.

We end this section by illustrating the introduced ideas in the example from the Introduction. More complicated mixed models are considered in the next section.

EXAMPLE 1 (continued). Recall that the $Y_i \in \mathbb{R}^r, i = 1, \dots, n$, are independent and multivariate normally distributed with mean 0 and covariance matrix $\Sigma(\theta) = \theta^2 1_r 1_r^\top + I_r$, and that the score for one observation is

$$s^i(\theta; y_i) = -\frac{r\theta}{1+r\theta^2} + (y_i^\top 1_r)^2 \frac{\theta}{(1+r\theta^2)^2} = \frac{\theta}{1+r\theta^2} \left(-r + \frac{(y_i^\top 1_r)^2}{1+r\theta^2} \right).$$

Since Y_i has positive variance for all θ , the score is equal to zero for almost every y_i if and only if $\theta = 0$, verifying Assumption 5. Assumption 1 holds since $f_{\theta}(y_i) > 0$ for all $y_i \in \mathbb{R}^r$, with $\gamma_i \equiv \gamma$ being Lebesgue measure on \mathbb{R}^r . To verify Assumption 4, observe that at $\theta = 0$ the second derivative of the log-likelihood is

$$\lim_{\theta \rightarrow 0} \frac{s^i(\theta; y_i) - s^i(0; y_i)}{\theta} = \lim_{\theta \rightarrow 0} \left[\frac{1}{1+r\theta^2} \left(-r + \frac{(y_i^\top 1_r)^2}{1+r\theta^2} \right) \right] = -r + (y_i^\top 1_r)^2,$$

which has variance bounded away from zero when evaluated at $y_i = Y_i$. Thus, Assumption 4 holds at $\theta = 0$ with $k_1 = 2$ and at $\theta \neq 0$ with $k_1 = 1$. It is straightforward to verify Assumptions 2 and 3, and hence conclude Theorems 2.1 and 2.4 apply. For additional insight, we also provide a more direct argument for why the score test standardized by expected information works in this example while other common test statistics do not. At $\theta \neq 0$ (see the Supplementary Material, Ekvall and Bottai (2022) for details),

$$T_n(\theta; Y^n) = \frac{1}{2rn} \left\{ -rn + \sum_{i=1}^n (Y_i^\top 1_r)^2 / (1+r\theta^2) \right\}^2 \sim \frac{(-rn + r\chi_n^2)^2}{2rn}.$$

It is immediate from the middle expression that $T_n(\cdot; \cdot)$ has a continuous extension on $[0, \infty) \times \mathbb{R}^{nr}$. Essentially, standardizing by expected information cancels the leading factor $\theta/(1+r\theta^2)$ in the expression for $s^i(\theta; y_i)$, which was the reason for the singularity at $\theta = 0$. Notably, this cancellation does not happen if one instead standardizes by observed information. The last expression shows that, in this example, the distribution of the proposed test statistic is in fact independent of the parameter, and hence it is almost immediate that (6), and hence Theorem 2.4 hold; writing χ_n^2 as the sum of n independent χ_1^2 and an appeal to the classical central limit theorem is all that is needed. To emphasize the fact that other common test statistics do not enjoy the same asymptotic properties, we show in Theorem C.1 of the Supplementary Material (Ekvall and Bottai (2022)) that the asymptotic distribution of the score test statistic standardized by observed information evaluated at the true θ_n , is different depending on how $\{\theta_n\}$ tends to 0.

3. Inference near critical points in mixed models.

3.1. *Scale parameters at zero.* Suppose momentarily the Y_i are identically distributed; that is, they are independent copies of some $Y \in \mathbb{R}^r$. Then $\mathcal{I}_n(\theta) = n\mathcal{I}_1(\theta)$ and we write $\mathcal{I}(\theta) = \mathcal{I}_1(\theta)$. Similarly, we drop the observation index on $y \in \mathbb{R}^r$, $s(\theta; y) = \nabla \log f_\theta(y)$ and γ . Suppose further Y has conditional density $f_\theta(y | w)$ against γ given a vector of random effects $W \in \mathbb{R}^q$ whose elements are independent with mean zero and unit variance. Denote the distribution of W by ν . We partition the parameter vector as $\theta = (\lambda, \psi) \in \mathbb{R}^{d_1} \times \mathbb{R}^{d_2}$ and assume, for some parameter matrix $\Lambda = \Lambda(\lambda) \in \mathbb{R}^{q \times q}$ and known $h : \mathbb{R}^r \times \mathbb{R}^{d_2} \times \mathbb{R}^q \rightarrow [0, \infty)$,

$$f_\theta(y | w) = h(y, \psi, \Lambda(\lambda)w).$$

We call λ a scale parameter since it determines the matrix $\Lambda(\lambda)$ scaling W . The vector ψ includes any other parameters in the distribution of $Y | W$. With these assumptions, the marginal distribution of Y has density against γ given by

$$(7) \quad f_\theta(y) = \int f_\theta(y | w)\nu(dw) = \int h(y, \psi, \Lambda w)\nu(dw).$$

For example, a generalized linear mixed model with linear predictor $X\psi + ZU$, where $U = \Lambda W \sim \mathcal{N}(0, \Lambda\Lambda^\top)$, and $X \in \mathbb{R}^{r \times d_2}$ and $Z \in \mathbb{R}^{r \times q}$ are design matrices, satisfies (7) for an appropriate choice of h .

In practice it is often possible to define ψ, λ, Λ and h so that any critical points are so because the leading $d_1 \times d_1$ block of $\mathcal{I}(\theta)$ is singular. That is, all critical points are due to linear combinations of the scores for the scale parameters being equal to zero almost surely. To investigate when the latter can happen, let $\nabla_{(3)}h(y, \psi, \Lambda w)$ denote the gradient of $h(y, \psi, \cdot)$ evaluated at Λw , the subscript (3) here indicating gradient with respect to the third argument vector. Then, assuming the derivatives exist and can be moved inside the integral, the j th element of $s(\theta; y)$, $j = 1, \dots, d_1$ is

$$(8) \quad \frac{1}{f_\theta(y)} \int \nabla_j f_\theta(y | w)\nu(dw) = \frac{1}{f_\theta(y)} \int \{\nabla_{(3)}h(y, \psi, \Lambda w)\}^\top \{\nabla_j \Lambda(\lambda)\} w \nu(dw),$$

where ∇_j acts elementwise on matrices. The following result gives a sufficient condition for linear combinations of these scores to be equal to zero for almost every y .

PROPOSITION 3.1. *If (8) holds at $\theta \in \Theta$ and for a $v = [v_1, \dots, v_{d_1}, 0, \dots, 0]^\top \in \mathbb{R}^d$, it holds that*

$$\sum_{j=1}^{d_1} v_j \{\nabla_j \Lambda(\lambda)\} W \quad \text{and} \quad \Lambda(\lambda)W$$

are independent under θ ; then $\mathcal{I}(\theta)v = 0$.

PROOF. It suffices to show that $v^\top s(\theta; y) = 0$ for almost every y . We have

$$v^\top s(\theta; y) = \frac{1}{f_\theta(y)} \int \{\nabla_{(3)}h(y, \psi, \Lambda w)\}^\top \sum_{j=1}^{d_1} v_j \{\nabla_j \Lambda(\lambda)\} w \nu(dw).$$

By the assumptions, the integral is the expectation of the inner product of two independent random vectors. Thus, since (measurable) functions of independent random variables are independent, the integral in the last display is

$$\int \{\nabla_3 h(y, \psi, \Lambda w)\}^\top \nu(dw) \int \sum_{j=1}^{d_1} v_j \{\nabla_j \Lambda(\lambda)\} w \nu(dw),$$

which is equal to zero since $\int w \nu(dw) = 0$, and that completes the proof. \square

Proposition 3.1 can help identify critical points. In fact, even though the condition that $\sum_{j=1}^{d_1} v_j \{\nabla_j \Lambda(\lambda)\} W$ and ΛW are independent random variables is not necessary, we will see that in practice it often identifies all critical points.

In what follows, to facilitate further analysis, we assume Λ is diagonal. Specifically, we assume every diagonal element of Λ is one of the λ_j . Then the λ_j are scale parameters in the usual sense and

$$(9) \quad \Lambda(\lambda) = \text{diag}(\lambda_{(1)}, \dots, \lambda_{(q)}),$$

where $\lambda_{(l)}$ means the λ_j ($j = 1, \dots, d_1$) in the l th diagonal element of $\Lambda(\lambda)$ ($l = 1, \dots, q$). For example, if $q = 3$ one possibility is that $\Lambda(\lambda) = \text{diag}(\lambda_1, \lambda_2, \lambda_1)$ so that the first and third element of W have the same scale parameter; then $d_1 = 2$, $\lambda_{(1)} = \lambda_1$, $\lambda_{(2)} = \lambda_2$ and $\lambda_{(3)} = \lambda_1$. Assuming (9) is common in practice and is less restrictive than it may first seem: it allows for the possibility that w affects the conditional density $f_\theta(y | w)$ through $H\Lambda w$ for some $H \in \mathbb{R}^{q \times q}$, which may depend on ψ . Then $U = H\Lambda W$ is a vector of dependent random effects whose scales are determined by λ and whose dependence is determined by H . In particular, by letting H be an orthogonal matrix the λ_j^2 are the eigenvalues of $\text{cov}_\theta(U) = H\Lambda(\lambda)^2 H^\top$.

With (9), Proposition 3.1 has the following corollary which says standard basis vectors are critical vectors for scale parameters at zero.

COROLLARY 3.2. *If (8) holds at $\theta \in \Theta$, $\Lambda(\lambda)$ satisfies (9), and $\theta_j = \lambda_j = 0$; then $\mathcal{I}(\theta)e_j = 0$, where e_j is the j th standard basis vector in \mathbb{R}^d .*

PROOF. With (9), $\nabla_j \Lambda(\lambda)$, $j = 1, \dots, d_1$, is a diagonal matrix whose l th diagonal element ($l = 1, \dots, q$) is 1 if the l th diagonal element of $\Lambda(\lambda)$ is λ_j , and zero otherwise. Thus, at θ such that $\lambda_j = 0$, $\{\nabla_j \Lambda(\lambda)\} W$ is a function of the elements of W scaled by λ_j , and $\Lambda(\lambda) W$ is a function of the elements of W not scaled by λ_j . Thus, Proposition 3.1 is satisfied with $v = e_j$. \square

We illustrate the wide applicability of Corollary 3.2 using another example.

EXAMPLE 2 (Generalized linear mixed model). Let $X \in \mathbb{R}^{r \times d_2}$ and $Z \in \mathbb{R}^{r \times q}$ be design matrices and suppose

$$f_\theta(y | w) = h(y, \psi, \Lambda w) = \exp\{y^\top(X\psi + Z\Lambda w) - c(X\psi + Z\Lambda w)\},$$

where $c : \mathbb{R}^r \rightarrow \mathbb{R}$ is the sum of the cumulant functions for the r responses (see, e.g., McCulloch, Searle and Neuhaus (2008), for definitions). For example, the j th element of $\nabla c(X\psi + Z\Lambda w)$, the gradient of c evaluated at the linear predictor, is the conditional mean of the j th response given w . Assume (9) and, for $j = 1, \dots, d_1$, let $[j]$ denote the set of $k \in \{1, \dots, q\}$ such that $\Lambda_{kk} = \lambda_j$. Then

$$\nabla_j Z\Lambda(\lambda)w = \sum_{k \in [j]} Z^k w_k,$$

where Z^k is the k th column of Z . Hence, assuming differentiation under the integral is permissible,

$$s_j(\theta; y) = \frac{1}{f_\theta(y)} \int f_\theta(y | w) \{y - \nabla c(X\psi + Z\Lambda w)\}^\top \sum_{k \in [j]} Z^k w_k \nu(dw).$$

Corollary 3.2 suggests $s_j(\theta; y)$ is zero for almost every y at θ such that $\lambda_j = 0$, which can here be verified directly: the sum in the integral is a function of the w_k scaled by λ_j , while the

remaining part of the integrand is a function of the w_k not scaled by λ_j . Hence, the integral is

$$\int f_\psi(y | \Lambda w)[y - \nabla c(X\psi + Z\Lambda w)]^\top v(dw) \int \sum_{k \in [j]} Z^k w_k v(dw),$$

which is equal to zero for every y since the second integral is the expectation of a linear combination of the W_k , which have mean zero.

It is clear from Example 2 that critical points occur in many mixed models and that standard basis vectors are often critical vectors. When all critical vectors are standard basis vectors, the calculations required to verify Assumption 4 are simplified. The following result illustrates this point and will be useful in examples.

PROPOSITION 3.3. *If Assumptions 1–2 hold, (i) Y_1, \dots, Y_n are independent copies of $Y \in \mathbb{R}^r$, (ii) for every $\theta \in \Theta$ the null space of $\mathcal{I}(\theta)$ is spanned by $\{e_j, j : \lambda_j = 0\}$ and (iii) the random vector*

$$\bar{s}(\theta; Y) = [\nabla_1^{k_1} \log f_\theta(Y), \dots, \nabla_d^{k_d} \log f_\theta(Y)]^\top, \quad Y \sim f_\theta,$$

where $k_j = k_j(\theta) = 2$ if $\theta_j = \lambda_j = 0$ and $k_j = 1$ otherwise, has positive definite covariance matrix with finite entries at every $\theta \in \Theta$; then Assumption 4 holds.

PROOF OF PROPOSITION 3.3. Let $\{v_{\theta_1}, \dots, v_{\theta_d}\}$ be orthonormal eigenvectors of $\mathcal{I}(\theta)$ and $V = [v_{\theta_1}, \dots, v_{\theta_d}]$. Suppose the rank of $\mathcal{I}(\theta)$ is $d - d_*$, $0 \leq d_* \leq d$. By condition (i), upon reordering θ if necessary, we may assume $v_{\theta_j} = e_j$, $j = 1, \dots, d_*$. In the definition of $\tilde{s}_n(\theta; y^n)$, set $k_1 = \dots = k_{d_*} = 2$ and $k_{d_*+1} = \dots = k_d = 1$. Then, since v_{θ_j} has zeros in the first d_* entries for $j \geq d_* + 1$ by orthogonality, $\tilde{s}_n(\theta; Y^n) = V^\top \sum_{i=1}^n \bar{s}(\theta; Y_i)$. Thus, $\tilde{\mathcal{I}}_n(\theta) = nV^\top \text{cov}_\theta\{\bar{s}(\theta; Y)\}V$, which is positive definite if and only if $\text{cov}_\theta\{\bar{s}(\theta; Y)\}$ is; this shows condition (i) of Assumption 4 holds.

Condition (ii) of Assumption 4 holds because $s_n(\tilde{\theta}; y^n)^\top e_j = 0$ for almost every y^n is equivalent to $\mathcal{I}_n(\tilde{\theta}) = n\mathcal{I}(\tilde{\theta})e_j = 0$, and this holds at any $\tilde{\theta}$ with $\tilde{\theta}_j = \theta_j = 0$ by the assumption that $\{e_j, j : \lambda_j = 0\}$ spans the null space of $\mathcal{I}(\theta)$.

Condition (iii) of Assumption 4 holds because we have assumed $k_j \leq 2$, which completes the proof. \square

In the following two sections, we verify the conditions of Theorems 2.1 and 2.4 in two mixed models. The first is an exponential mixed model with independent and identically distributed observations of a vector of correlated, positive responses. It has one scale parameter, one fixed effect parameter and a nonnormal random effect. This example illustrates our theory in nonlinear mixed models with nonstandard random effect distributions. The second model is a quite general version of the linear mixed model with normally distributed random effects. It has general design matrices X_i and Z_i , possibly different for different $i = 1, \dots, n$, and hence nonidentically distributed observations; and several fixed and random effect parameters.

3.2. Exponential mixed model with uniform random effect. Suppose the $Y_i, i = 1, \dots, n$, are independent copies of a $Y \in \mathbb{R}^2$ which has conditionally independent elements given $W \in \mathbb{R}$ with conditional densities

$$(10) \quad f_\theta(y_j | w) = (\psi + \lambda w) \exp\{-y_j(\psi + \lambda w)\} \mathbb{I}(y_j \geq 0); \quad \theta = (\lambda, \psi) \in \Theta \subseteq \mathbb{R}^2,$$

where y_j is the j th element of $y \in \mathbb{R}^2$; we omit observation indexes for the remainder of the section and work only with the generic $Y \in \mathbb{R}^2$. The density $f_\theta(y_j | w)$ is that of an

exponential random variable with mean $1/(\psi + \lambda w)$. We have selected $r = 2$ responses to simplify calculations, but $r \geq 2$ presents no fundamental difficulties. The specification clearly requires $\psi + \lambda W$ be positive almost surely. There are many potentially useful specifications satisfying this, but to be concrete suppose W is uniform on $(-\sqrt{3}, \sqrt{3})$, so that $\nu(dw) = f(w) dw$ with $f(w) = 12^{-1/2} \mathbb{I}(|w| \leq \sqrt{3})$; and the parameter set is $\Theta = \{(\lambda, \psi) \in [0, \infty) \times \mathbb{R} : \psi > \sqrt{3}\lambda\}$. The log-likelihood for one observation is, ignoring additive constants,

$$\log f_\theta(y) = \log \int f_\theta(y | w) \nu(dw) = \log \int (\psi + \lambda w)^2 \exp\{-(\psi + \lambda w)y_\bullet\} \nu(dw),$$

where $y_\bullet = y_1 + y_2$. The score for λ is

$$s_\lambda(\theta; y) = -\frac{1}{f_\theta(y)} \int f_\theta(y | w) \{y_\bullet - 2/(\psi + \lambda w)\} w \nu(dw).$$

LEMMA 3.4. *In the exponential mixed model (10), it holds for any $\theta = (0, \psi) \in \Theta$ that $\text{var}_\theta\{s_\lambda(\theta, Y)\} = 0$ and $\text{var}_\theta\{\nabla_1^2 \ell(\theta; Y)\} > 0$.*

PROOF. Setting $\theta = (0, \psi)$ in the expression for $s_\lambda(\theta; y)$ and moving terms that do not depend on w outside the integral gives $s_\lambda(0, \psi; y) = -f_\theta(y)^{-1} \psi \exp(-\psi y_\bullet) \{y_\bullet - 2/\psi\} \int w \nu(dw) = 0$. Moreover, when $s_\lambda(\theta; y) = 0$,

$$\nabla_1^2 \log f_\theta(Y) = -\frac{1}{f_\theta(Y)} \int f_\theta(Y | w) [\{Y_\bullet - 2/(\psi + \lambda w)\}^2 - 2/(\psi + \lambda w)^2] w^2 \nu(dw),$$

where $\log f_\theta(Y)$ means f_θ evaluated at $Y \sim f_\theta$. When $\lambda = 0$, this simplifies to $(Y_\bullet - 2/\psi)^2 - 2/\psi^2$, which has positive variance under $\theta = (0, \psi)$. \square

Lemma 3.4 shows $\theta = (0, \psi)$ is a critical point for any ψ , agreeing with Corollary 3.2. It also shows the corresponding second derivative of the log-likelihood has positive variance if $\lambda = 0$, suggesting it may be possible to verify Assumption 4 with $k_1 = 2$ at θ with $\lambda = 0$ and $k_1 = 1$ elsewhere. The following lemma will be helpful to that end.

LEMMA 3.5. *In the exponential mixed model (10), the information matrix $\mathcal{I}(\theta)$ has rank one if $\lambda = 0$ and rank two otherwise.*

PROOF. The score for ψ is $s_\psi(\theta; y) = -f_\theta(y)^{-1} \int f_\theta(y | w) \{y_\bullet - 2/(\psi + \lambda w)\} \nu(dw)$. Thus, when $\lambda = 0$, $s_\psi(\theta; Y) = Y_\bullet - 2/\psi$, which has positive variance. In conjunction with Lemma 3.4, this shows the rank of $\mathcal{I}(\theta)$ is one when $\lambda = 0$. Suppose $\lambda > 0$ and make the change of variables $t = (\psi + \lambda w)y_\bullet$ in the integral in the definition of the log-likelihood. Letting $G(t) = e^{-t}(t^2 + 2t + 2)$, which is an antiderivative of $g(t) = -t^2 e^{-t}$, gives

$$\log f_\theta(y) = -\log(\lambda) + \log\{G(\psi - \sqrt{3}y_\bullet\lambda) - G(\psi + \sqrt{3}y_\bullet\lambda)\}.$$

Differentiating with respect to λ and ψ , taking an arbitrary linear combination given by $v = [v_1, v_2]^T \in \mathbb{R}^2$, and observing the (Lebesgue) density for Y_\bullet is positive on $(0, \infty)$ shows $\mathcal{I}(\theta)v = 0$ only if

$$t \mapsto \frac{v_1 \lambda^{-1} t \{g(\psi - t) + g(\psi + t)\} - v_2 \{g(\psi - t) - g(\psi + t)\}}{G(\psi - t) - G(\psi + t)}$$

is constant on $(0, \infty)$, possibly except on a Lebesgue null set. Verifying this map is indeed nonconstant on a set of positive Lebesgue measure is routine so we omit the details. \square

We are ready to verify Assumptions 1–5, giving the main result of the section.

THEOREM 3.6. *The conclusions of Theorems 2.1 and 2.4 hold in the exponential mixed model (10).*

PROOF. Assumption 1 holds with γ being Lebesgue measure since $f_\theta(y) > 0$ for all $y \in (0, \infty)^2$ and $\theta \in \Theta$.

To verify the moment condition in Assumption 2, it suffices by Lemma 3.5 to find locally uniform bounds of $\int |s_\lambda(\tilde{\theta}; y)|^3 f_\theta(y) dy$ and $\int |s_\psi(\tilde{\theta}; y)|^3 f_\theta(y) dy$ around an arbitrary $\theta' \in \Theta$, and of $\int |\nabla_1^2 \log f_{\tilde{\theta}}(y)|^3 f_\theta(y) dy$ around θ' with $\lambda' = 0$. For the former, observe that for any $k \geq 1$, by Jensen’s inequality and using $f_\theta(y | w) f(w) / f_\theta(y) = f_\theta(w | y)$,

$$\begin{aligned} \int |s_\lambda(\tilde{\theta}; y)|^k f_\theta(y) dy &= \int \left| \int \{y_\bullet - 2/(\tilde{\psi} + \tilde{\lambda}w)\} w \frac{f_{\tilde{\theta}}(y | w) f(w)}{f_{\tilde{\theta}}(y)} dw \right|^k f_\theta(y) dy \\ &\leq \int \int |\{y_\bullet - 2/(\tilde{\psi} + \tilde{\lambda}w)\} w|^k f_{\tilde{\theta}}(w | y) dw f_\theta(y) dy. \end{aligned}$$

Now for any $\theta' \in \Theta$, since $\psi > \lambda\sqrt{3} \geq \lambda|w|$, we can find an $M < \infty$ and a ball B centered at θ' small enough that $1/(\psi + \lambda w) \leq M$ for all $(\psi, \lambda) \in B$ and $w \in (-\sqrt{3}, \sqrt{3})$. Then for $\tilde{\theta} \in B$, by Jensen’s inequality (twice), the inner integral less than $3^{k/2} 2^{k-1} (|y_\bullet|^k + |2M|^k) \leq 3^{k/2} 2^{k-1} \{2^{k-1} (|y_1|^k + |y_2|^k) + |2M|^k\}$. Thus, the last line in the last display is bounded uniformly for $(\tilde{\theta}, \theta) \in B \times B$ if $\mathbb{E}_\theta(Y_j^k)$ is, $j = 1, 2$. But for $\theta \in B$, $\mathbb{E}_\theta(Y_j^k) = \mathbb{E}_\theta\{\mathbb{E}_\theta(Y_j^k | W)\} = k! \mathbb{E}\{1/(\psi + \lambda W)^k\} \leq k! M^k$. The other moment bounds can be handled similar to show Assumption 2 holds; we omit the details.

Assumption 3 holds because the Y_i are identically distributed.

We use Proposition 3.3 to verify Assumption 4. Condition (i) of that proposition is by assumption and condition (ii) is established in Lemma 3.5. To verify condition (iii), it suffices by Lemma 3.5 to show that, when $\lambda = 0$, $v_1 \nabla_1^2 \log f_\theta(y) + v_2 s_\psi(\theta; y)$ is constant for almost every y only if $v_1 = v_2 = 0$; this clearly holds since the linear combination is $v_1 (y_\bullet - 2/\psi)^2 - 2v_1/\psi^2 + v_2 (y_\bullet - 2\psi)$.

Assumption 5 holds because the only critical points are those with $\lambda = 0$, and this completes the proof. \square

3.3. Linear mixed models. Consider a linear mixed model which assumes, for some $\psi \in \mathbb{R}^{d_2}$ and Λ satisfying (9), independently for $i = 1, \dots, n$,

$$(11) \quad Y_i | X_i, W_i \sim \mathcal{N}(X_i \psi + Z_i \Lambda W_i, \sigma^2 I_{r_i}), \quad W_i | X_i \sim \mathcal{N}(0, I_{r_i}),$$

where $Z_i \in \mathbb{R}^{r_i \times q}$ is a design matrix and $X_i \in \mathbb{R}^{r_i \times d_2}$ a matrix of predictors. The predictors can be nonstochastic or, more generally, have a distribution not depending on θ , possibly different for different i . Assume also for simplicity that $\sigma^2 > 0$ is known. When $\sigma^2 = 1$ and observations are identically distributed, this model is a special case of the generalized linear mixed model in Example 2.

We will first obtain a reliable confidence region for $\theta = (\lambda, \psi) \in \Theta = [0, \infty)^{d_1} \times \mathbb{R}^{d_2}$ by verifying the conditions of Theorem 2.4. Then we show how that result can be modified to give a reliable confidence region for λ only, with ψ a nuisance parameter. The model implies the distribution of $Y_i | X_i$ is multivariate normal with mean $X_i \psi$ and covariance matrix

$$\Sigma_i = \Sigma_i(\lambda) = \sigma^2 I_{r_i} + \sum_{j=1}^{d_1} \lambda_j^2 H_j^i,$$

where $H_j^i = \sum_{k \in [j]} Z_i^k (Z_i^k)^\top \in \mathbb{R}^{r_i \times r_i}$ is the sum of outer products of columns of Z_i corresponding to random effects scaled by λ_j , $j = 1, \dots, d_1$. The log-likelihood is, ignoring

additive terms not depending on θ ,

$$\ell_n(\theta; y^n, X^n) = \sum_{i=1}^n \left\{ -\frac{1}{2} \log |\Sigma_i(\lambda)| - \frac{1}{2} (y_i - X_i \psi)^\top \Sigma_i(\lambda)^{-1} (y_i - X_i \psi) \right\},$$

where $X^n = (X_1, \dots, X_n)$. Differentiating with respect to λ_j gives, for $j = 1, \dots, d_1$,

$$s_j^i(\theta; y_i, X_i) = \lambda_j \operatorname{tr} \{ \Sigma_i^{-1} H_j^i - \Sigma_i^{-1} (y_i - X_i \psi) (y_i - X_i \psi)^\top \Sigma_i^{-1} H_j^i \},$$

which is equal to zero for every i when $\lambda_j = 0$. Thus, scale parameters at zero are critical points, agreeing with Corollary 3.2. We also have the following partial converse, essentially saying that only scale parameters can be critical elements.

LEMMA 3.7. *The Fisher information $\mathcal{I}^i(\theta)$ for one observation in the linear mixed model (11) is block-diagonal and, if $\underline{e}\{\mathbb{E}(X_i^\top X_i)\} > 0$, is singular if and only if its leading block $\mathcal{I}_\lambda^i(\theta) = \operatorname{cov}_\theta \{s_\lambda^i(\theta; Y_i, X_i)\}$ is, where $\underline{e}(\cdot)$ is the smallest eigenvalue.*

The lemma holds as stated with nonstochastic predictors, but notation can be simplified since $\mathbb{E}(X_i^\top X_i) = X_i^\top X_i$ in that case. The following lemma gives a lower bound on the (conditional) variance of any linear combination of the score for the scale parameters. Together with the previous lemma, it can be used to identify all critical points.

LEMMA 3.8. *In the linear mixed model (11), for any $v \in \mathbb{R}^{d_1}$,*

$$\operatorname{var}_\theta \{v^\top s_\lambda^i(\theta; Y_i, X_i) \mid X_i\} \geq 2\bar{e}(\Sigma_i)^{-2} \underline{e}(Z_i^\top Z_i)^2 \max_j (\lambda_j v_j)^2$$

and

$$\operatorname{var}_\theta \{v^\top s_\lambda^i(\theta; Y_i, X_i) \mid X_i\} \leq 2r_i \bar{e}(\Sigma_i)^{-2} \bar{e}(Z_i^\top Z_i)^2 \max_j (\lambda_j v_j)^2,$$

where $\bar{e}(\cdot)$ is the largest eigenvalue.

Together with Lemma 3.7, Lemma 3.8 ensures that, as long as $Z_i^\top Z_i$ and $\mathbb{E}(X_i^\top X_i)$ are invertible, $\{e_j, j : \lambda_j = 0\}$ spans the null space of $\mathcal{I}^i(\theta)$. As observed following Proposition 3.3, this makes checking the assumptions in Section 2 easier. Similarly, an implication is that, when implementing the proposed method, we can take $k_j(\theta) = 2$ when $\lambda_j = 0$ and $k_j(\theta) = 1$ otherwise. We are ready for the first main result of the section.

Let $\|\cdot\|$ be the spectral norm when the argument is a matrix.

THEOREM 3.9. *If $\limsup_{i \rightarrow \infty} r_i < \infty$ and there exists an $M \in (0, \infty)$ such that, for every $i = 1, 2, \dots$, $M^{-1} \leq \underline{e}(Z_i^\top Z_i) \leq \bar{e}(Z_i^\top Z_i) \leq M$, $M^{-1} \leq \underline{e}\{\mathbb{E}(X_i^\top X_i)\}$ and $\mathbb{E}(\|X_i\|^{4+\delta}) \leq M$ for some $\delta > 0$; then Assumptions 1–5, and hence the conclusions of Theorems 2.1 and 2.4, hold in the linear mixed model (11).*

PROOF OF THEOREM 3.9. We prove the theorem with nonstochastic X_i ; the calculations for stochastic X_i are similar. Then, since $\mathbb{E}(\|X_i\|^{4+\delta}) = \|X_i\|^{4+\delta}$ we may assume $\bar{e}(X_i^\top X_i) \leq M$, upon increasing M as needed. Assumption 1 is satisfied since $f_\theta(y_i)$ is strictly positive on \mathbb{R}^{r_i} for every $\theta \in \Theta$, with γ_i being Lebesgue measure. We verify Assumption 2 with $\delta > 0$ and $k = 2$ at critical θ' . Define $\xi^i(\theta; y_i, X_i) \in \mathbb{R}^d$ by

$$\xi_j^i(\theta; y_i, X_i) = \operatorname{tr} \{ \Sigma_i^{-1} H_j^i - \Sigma_i^{-1} (y_i - X_i \psi) (y_i - X_i \psi)^\top \Sigma_i^{-1} H_j^i \}, \quad j \leq d_1,$$

and the remaining d_2 elements equal to $s_{\psi}^i(\theta; y_i, X_i)$. When $\lambda_j > 0$, $\xi_j^i(\theta; y_i, X_i) = s_j^i(\theta; Y_i, X_i)/\lambda_j$. Because $\xi_j^i(\theta; y_i, X)$ is continuous in θ since $\sigma^2 > 0$, when $\lambda_j = 0$ it holds that $\xi_j^i(\theta; y_i, X_i) = \nabla_j^2 \ell^i(\theta; y_i, X_i)$. More generally, for $j, l \leq d_1$,

$$\begin{aligned} \nabla_{jl}^2 \ell^i(\theta; y_i, X_i) &= \mathbb{I}(j = l)\xi_j^i(\theta; y_i, X_i) - \lambda_j \operatorname{tr}\{\Sigma_i^{-1} H_k^i \Sigma_i^{-1} H_j^i\} \\ &\quad + \lambda_j \operatorname{tr}\{\Sigma_i^{-1} H_k^i \Sigma_i^{-1} (y_i - X_i \psi)(y_i - X_i \psi)^\top \Sigma_i^{-1} H_j^i\} \\ &\quad + \lambda_j \operatorname{tr}\{\Sigma^{-1} (y_i - X_i \psi)(y_i - X_i \psi)^\top \Sigma_i^{-1} H_k^i \Sigma_i^{-1} H_j^i\}. \end{aligned}$$

Since $\sigma^2 > 0$, we can pick for any $\theta' \in \Theta$ a small enough open ball $B \subseteq \mathbb{R}^d$ centered at θ' and large enough $M < \infty$ to have, on B , that $\sigma^{-2} \leq M$, $\|\Sigma_i^{-1}\| \leq M$ and $\lambda_j \leq M$ for all j . Thus, using submultiplicativity of the spectral norm and that the trace is upper bounded by r_i times the spectral norm, on B ,

$$\begin{aligned} |\nabla_{jl}^2 \ell^i(\theta; y_i, X_i)| &\leq r_i \{M \|H_j^i\| + M^2 \|H_j^i\| \|y_i - X_i \psi\|^2\} + r_i M^3 \|H_i^i\| \|H_j^i\| \\ &\quad + 2r_i M^4 \|y_i - X_i \psi\|^2 \|H_j^i\| \|H_i^i\|. \end{aligned}$$

Now, $\|H_j^i\| \leq \|Z_i^\top Z_i\| \leq M$ and $\|y_i - X_i \psi\| \leq \|y_i\| + \|X_i \psi\| \leq \|y_i\| + M \|\psi\|$ for all (i, j) . Thus, since $\|\psi\|$ is bounded on B , the expectation in Assumption 2 to be bounded is, upon increasing M as needed, less than

$$M \{1 + \sup_{\theta \in \Theta \cap B} \mathbb{E}_\theta (\|Y_i\|^{4+2\delta})\},$$

which is bounded uniformly in i since Y_i is a normal vector whose mean and covariance matrix are bounded uniformly in i and $\theta \in B \cap \Theta$. The calculations for $k = 1$ and $j > d_1$ or $l > d_1$ are very similar, and hence omitted.

To verify Assumption 3, it suffices to find c_1, c_2 such that, for any $v \in \mathbb{R}^d$,

$$c_1(\theta) v^\top \mathcal{T}^1(\theta) v \leq v^\top \mathcal{I}^i(\theta) v \leq c_2(\theta) v^\top \mathcal{T}^1(\theta) v.$$

Let $v = [v_1^\top, v_2^\top]^\top$ with $v_1 \in \mathbb{R}^{d_1}$. By Lemmas 3.7 and 3.8, using that r_i and the eigenvalues of $X_i^\top X_i, Z_i^\top Z_i$, and Σ_i^{-1} are bounded, there exist $M \in (0, \infty)$ and $c_3, c_4 : \Theta \rightarrow (0, \infty)$ not depending on i such that

$$M^{-1} \left\{ \|v_2\| + c_3(\theta) \max_{j \leq d_1} (\lambda_j v_j)^2 \right\} \leq v^\top \mathcal{I}^i(\theta) v \leq M \left\{ \|v_2\| + c_4(\theta) \max_{j \leq d_1} (\lambda_j v_j)^2 \right\}.$$

Thus, the desired inequalities hold with $c_1(\theta) = M^{-2} \{1 + c_4(\theta)\}^{-1} \min\{1, c_3(\theta)\}$ and $c_2(\theta) = M^2 \min\{1, c_3(\theta)\}^{-1} \{1 + c_4(\theta)\}$.

To verify Assumption 4, we argue as in the proof of Proposition 3.3, with minor modifications to take the observation index into account. Thus, adding an observation index to the vector \bar{s} defined in Proposition 3.3, and making the dependence on the predictors explicit, we get that $\xi^i(\theta; Y_i, X_i)$ is equal to $\bar{s}^i(\theta; Y_i, X_i)$ with each element corresponding to a $\lambda_j > 0$ scaled by $1/\lambda_j$. Thus, the covariance matrix of $\xi^i(\theta; Y_i, X_i)$ is positive definite if and only if that of $\bar{s}^i(\theta; Y_i, X_i)$ is. With that, the verification of Assumption 4 follows as in the proof of Proposition 3.3.

Assumption 5 holds because, for $j = 1, \dots, d_1, \lambda_j = 0$ is the only critical element in $[0, \infty)$, and that completes the proof. \square

Lastly, in this section we consider confidence regions for λ only. To that end, let

$$T_n^\lambda(\lambda; \psi, Y^n, X^n) = \left\{ \sum_{i=1}^n s_\lambda^i(\theta; Y_i, X_i)^\top \right\} \mathcal{I}_n^\lambda(\theta)^{-1} \left\{ \sum_{i=1}^n s_\lambda^i(\theta; Y_i, X_i) \right\},$$

where $\mathcal{I}_n^\lambda(\theta) = \text{cov}_\theta\{s_n^\lambda(\theta; Y^n, X^n)\}$ is the leading $d_1 \times d_1$ block of $\mathcal{I}_n(\theta)$. For such regions to be practically useful, or feasible, ψ has to be known or estimated. Our next result says ψ can be estimated by any square root n -consistent estimator without affecting the asymptotic coverage probability; this result assumes identically distributed observations for simplicity. To be more specific, let

$$\mathcal{R}_n^\lambda(\alpha) = \{\lambda : T_n^\lambda(\lambda; \psi, Y^n, X^n) \leq q_{d_1, 1-\alpha}\}$$

and let $\hat{\mathcal{R}}_n^\lambda(\alpha)$ be that confidence region with an estimator $\hat{\psi}$ in place of ψ in T_n .

THEOREM 3.10. *If the $(Y_i, X_i), i = 1, 2, \dots,$ are identically distributed and the conditions of Theorem 3.9 hold, the confidence region $\mathcal{R}_n^\lambda(\alpha)$ satisfies, for any compact $C \subseteq \Theta$ and $\alpha \in (0, 1)$,*

$$\lim_{n \rightarrow \infty} \inf_{\theta \in C} \mathbb{P}_\theta\{\lambda \in \mathcal{R}_n^\lambda(\alpha)\} = 1 - \alpha;$$

and if in addition $\sqrt{n}\|\hat{\psi}_n - \psi_n\| = O_P(1)$ under any convergent $\{\theta_n = (\lambda_n, \psi_n)\} \in \Theta$, then the same holds for $\hat{\mathcal{R}}_n^\lambda(\alpha)$.

In practice, the estimator $\hat{\psi}_n$ can be, for example, the least squares estimator

$$\hat{\psi}_n = \left(\sum_{i=1}^n X_i^\top X_i\right)^{-1} \sum_{i=1}^n X_i^\top Y_i.$$

This estimator is square root- n consistent under convergent $\{\theta_n\}$ under weak conditions. More specifically, it is multivariate normally distributed given (X_1, \dots, X_n) with mean ψ_n and covariance matrix

$$\left(\sum_{i=1}^n X_i^\top X_i\right)^{-1} \left(\sum_{i=1}^n X_i^\top \Sigma_n X_i\right) \left(\sum_{i=1}^n X_i^\top X_i\right)^{-1},$$

which tends to a matrix of zeros under weak conditions.

3.4. Practical considerations. In many mixed models, the likelihood does not admit an analytical expression, which complicates implementing likelihood-based methods in practice. In particular, evaluating the log-likelihood, or its derivatives and their expectations, may require the numerical evaluation or approximation of integrals. A detailed treatment of numerical integration is outside the scope of the present article, but we briefly discuss two possible practical issues: (i) the critical points and their corresponding critical vectors are known, but the proposed test statistic does not admit a convenient expression at those points; and (ii) it is difficult to establish which points are critical by analytical means. For clarity, we discuss identically distributed observations and assume $k_j \leq 2$ suffices to satisfy Assumption 4, which in our experience is often the case in practice; other settings can be handled similarly.

Issue (i) can occur when, for example, Corollary 3.2 says a scale parameter $\lambda_j = 0$ is a critical point with critical vector e_j in a generalized linear mixed model. To compute our test statistic at such points, it is useful to note that, with the \bar{s} defined in Proposition 3.3 and under regularity conditions,

$$T_n(\theta; Y^n) = n^{-1} \left\{ \sum_{i=1}^n \bar{s}(\theta; Y_i) \right\} \text{cov}_\theta\{\bar{s}(\theta; Y_1)\}^{-1} \left\{ \sum_{i=1}^n \bar{s}(\theta; Y_i) \right\}.$$

Thus, the user needs to compute first and second order partial derivatives of the log-likelihood and their covariance matrix. A routine calculation shows, under regularity conditions, the j th element of $\bar{s}(\theta, y_i)$ is

$$(12) \quad \frac{1}{f_\theta(y_i)} \int \nabla_j^{k_j} f_\theta(y_i | w) v(dw),$$

where $k_j = 2$ for $\theta_j = \lambda_j = 0$ and $k_j = 1$ otherwise. The derivative inside the integral can typically be computed analytically for any k_j , and hence there are no fundamental computational differences between $k_j = 1$ and $k_j = 2$. We illustrate these calculations in detail in an example in the Supplementary Material (Ekvall and Bottai (2022)). Software packages for mixed models typically approximate integrals like that in (12) with $k_j = 1$ using Laplace approximations, adaptive Gaussian quadrature or Monte Carlo (e.g., Bates et al. (2015), Knudson et al. (2021)). The integrals with $k_j = 2$ and those required to compute the covariance matrix can be handled by similar methods. When the critical vectors are known but are not standard basis vectors, the computations are similar but with different linear combinations of the derivatives.

By contrast, issue (ii) is not merely computational. If one cannot establish the existence or lack of critical points, then one does not know whether the Fisher information is invertible, and hence whether the theory motivating many other methods applies. This highlights a practical advantage of the proposed procedure: if the inversion of the Fisher information fails, the user is effectively warned they are attempting inference at a previously unknown critical point. Having identified a critical point, the critical vectors can be calculated numerically by spectral decomposition of the Fisher information, and then one is essentially back to issue (i).

4. Numerical experiments. We examine finite sample properties of the proposed method in a linear mixed model. The Supplementary Material (Ekvall and Bottai (2022)) contains similar simulations in a generalized linear mixed model for binary responses with asymmetrically distributed random effects, and the results there are qualitatively similar.

Suppose that for $i = 1, \dots, n$ and $j = 1, \dots, r$,

$$(13) \quad Y_{ij} = \psi_1 + \psi_2 X_{ij} + U_{1i} + U_{2i} X_{ij} + E_{ij},$$

where $[U_{1i}, U_{2i}]^\top \sim \mathcal{N}\{0, \text{diag}(\lambda_1^2, \lambda_2^2)\}$, independently for $i = 1, \dots, n$ and independent of $E = [E_{1,1}, \dots, E_{n,r}]^\top \sim \mathcal{N}(0, \sigma^2 I_{nr})$. To conform with the previous section, we assume $\sigma^2 = 1$ is known. Thus, the parameter set is $\Theta = \{\theta = (\lambda, \psi) \in [0, \infty)^2 \times \mathbb{R}^2\}$. We study confidence regions for λ with ψ estimated, that is, a nuisance parameter. The Supplementary Material (Ekvall and Bottai (2022)) contains simulations where ψ is known, and simulations where σ is unknown and confidence regions are created for (λ, σ) with ψ estimated. We compare the proposed confidence region to those obtained by inverting the likelihood ratio test statistic and the Wald test statistic standardized by expected information evaluated at the (unconstrained) maximum likelihood estimates. Specifically,

$$T_n^L(\lambda; y^n) = 2\{\ell_n(\hat{\theta}) - \ell_n([\lambda, \tilde{\psi}]^\top, y^n)\},$$

where $\hat{\theta} \in \arg \max_{\theta \in \Theta} \ell_n(\theta; y^n)$ and $\tilde{\psi} = \tilde{\psi}(\lambda) \in \arg \max_{\psi \in \mathbb{R}^2} \ell_n([\lambda, \psi]^\top, y^n)$; and

$$T_n^W(\lambda; y^n) = (\hat{\lambda} - \lambda)^\top \mathcal{I}_n^\lambda(\hat{\theta})(\hat{\lambda} - \lambda),$$

where $\mathcal{I}_n^\lambda(\theta)$ is the 2×2 block of $\mathcal{I}_n(\theta)$ corresponding to λ .

Our theory suggests the proposed confidence region may have good finite sample coverage near critical points because the proposed test statistic has the same asymptotic distribution under any sequence of parameters tending to a critical point as the sample size increases.

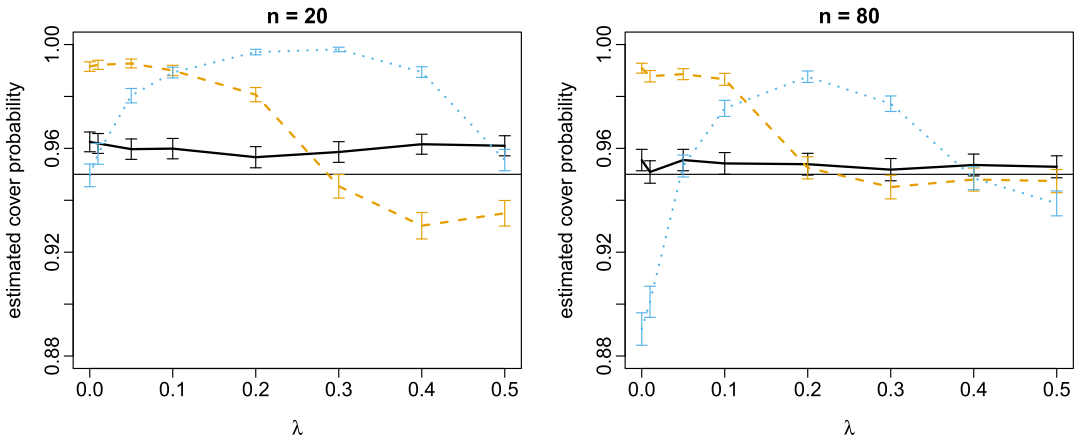


FIG. 2. Monte Carlo estimates of coverage probabilities of confidence regions from inverting the modified score (solid), likelihood ratio (dashed) and Wald (dotted) test statistics. The straight horizontal line indicates the nominal 0.95 coverage probability and vertical bars denote ± 2 times Monte Carlo standard errors.

Conversely, we expect the coverage of the other confidence regions may depend on how close to a critical point the true parameter is. To examine this, we consider true λ at different distances from the origin:

$$\lambda_1 = \lambda_2 \in \{10^{-6}, 0.01, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5\}.$$

Because these λ correspond to interior points of the parameter set, a level $1 - \alpha$ confidence region covers λ if the test statistic at λ is smaller than the $(1 - \alpha)$ th quantile of the chi-square distribution with two degrees of freedom. The use of that distribution is motivated by classical asymptotic theory for $T_n^L(\lambda; Y^n)$ and $T_n^W(\lambda; Y^n)$, and by our theory for the proposed test statistic. Different reference distributions should be used for the Wald and likelihood ratio statistics at boundary points (Baey, Cournède and Kuhn (2019), Geyer (1994), Self and Liang (1987)), while the proposed method uses the same reference distribution at every point of the parameter set.

We performed a Monte Carlo experiment with 10,000 replications. We set $\psi = 1_2$ and, for every (λ, ψ) considered, generated stochastic predictors as independent draws from a uniform distribution on $[-1, 2]$. Responses were then generated according to (13). The sample sizes were $n \in \{20, 80\}$ and $r = 10$. Code for reproducing the results is available at <https://github.com/koekvall/conf-crit-suppl>.

Figure 2 summarizes the results. Notably, the proposed confidence region has near-nominal estimated coverage probability in all considered settings. By contrast, the estimated coverage probabilities for the likelihood ratio and Wald confidence regions are substantially different from nominal for many settings. Moreover, their coverage is sometimes lower and sometimes higher than nominal. That is, for those methods, the quality of the chi-square distribution as a reference distribution depends on how close to the critical point the true parameter is. Simulations in the Supplementary Material (Ekvall and Bottai (2022)) indicate the proposed method give near-nominal coverage probabilities also of critical points, including ones at the boundary.

Figure 3 examines the agreement between sample quantiles of the three considered test statistics and the theoretical quantiles of a chi-square distribution with two degrees of freedom. The first plot, corresponding to small but nonzero scale parameters, shows the agreement is poor near critical points for the likelihood ratio and Wald test statistics. For the larger scale parameters, agreement between sample and theoretical quantiles is decent for all three test statistics.

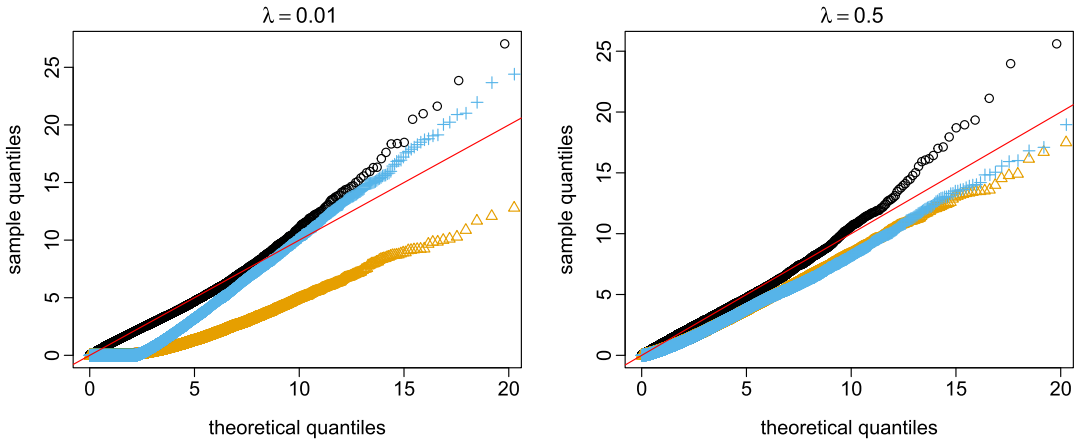


FIG. 3. *Quantile-quantile plots for modified score (circles), likelihood ratio (triangles) and Wald (plus signs) test statistics. The theoretical quantiles are from the chi-square distribution with 2 degrees of freedom and the sample quantiles from 10,000 Monte Carlo replications with $(n, r) = (80, 10)$.*

For additional insight into test statistics’ behavior near critical points, Figure 4 shows estimated rejection probabilities (size and power) for tests of the null hypothesis that $\lambda_1 = \lambda_2 = 10^{-6}$. The data generating settings are the same as those used for Figure 2. The power curves are not directly comparable because, as was also shown in Figure 2, the different tests have different sizes. Nevertheless, the power curves behave similarly as the true λ moves away from the null hypothesis value. This indicates the differences in coverage observed in Figure 2 is not in general due to how large the different confidence regions are.

We considered several configurations in addition to those reported, including both larger and smaller values of n and λ , and the results were remarkably consistent. To compute the proposed test statistic, we used the `lmmstest` R package written by the first author. To fit the model and compute the likelihood ratio and Wald test statistics, we used the `lme4` R package (Bates et al. (2015)). The Supplementary Material (Ekvall and Bottai (2022)) includes times for computing the test statistics in the simulations on which Figure 2 is based. The proposed test statistic was about 3–4 times faster to evaluate than the likelihood ratio and Wald statistics

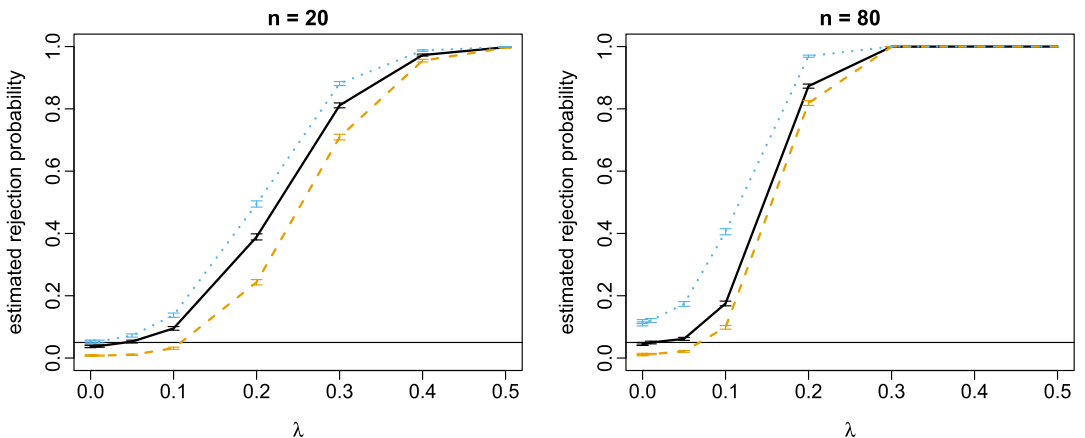


FIG. 4. *Monte Carlo estimates of rejection probabilities for testing the null hypothesis $\lambda = 10^{-6}1_2$ using the modified score (solid), likelihood ratio (dashed) and Wald (dotted) test. The straight horizontal line indicates the size of the tests (0.05) and vertical bars denote ± 2 times Monte Carlo standard errors.*

TABLE 1

Maximum likelihood estimates and 95% confidence intervals (CI) for scale parameters in the linear mixed model (14)

Parameter	Estimate	Mod. Score CI	Lik. Rat. CI	Wald CI
λ_1 (intercept)	0	(0, 0.0482)	(0, 0.0604)	(0, 0.126)
λ_2 (age)	0.0201	(0.0188, 0.0219)	(0.0183, 0.0222)	(0.0181, 0.0222)
σ (error)	0.156	(0.152, 0.162)	(0.151, 0.162)	(0.151, 0.162)

on average, but we expect both relative and absolute computing times to vary substantially between settings and implementations.

We also note there is a Stata routine for calculation of the proposed confidence region for a single variance parameter in linear mixed models with a random intercept (Bottai and Orsini (2004)), a special case which can also be treated using exact finite sample methods (Crainiceanu and Ruppert (2004)).

5. Data example. We illustrate using a data set presented by Fitzmaurice, Laird and Ware (2011), which contains a subset of the pulmonary function data collected in the Six Cities Study (Dockery et al. (1983)). The data include a pulmonary measure called the forced expiratory volume in the first second (FEV1), height (ht) and age obtained from a randomly selected subset of the female participants living in Topeka, Kansas. The sample includes $n = 300$ girls and young women, with a minimum of one and a maximum of twelve observations over time. Age and height are believed to be associated with the ability to take in and force out air. To model these data, consider the linear mixed model

$$(14) \quad \text{FEV1}_{it} = \psi_1 + \psi_2 \text{age}_{it} + \psi_3 \text{ht}_{it} + \psi_4 \text{age}_{i1} + \psi_5 \text{ht}_{i1} + U_{i1} + U_{i2} \text{age}_{it} + E_{it},$$

where $U_{i1} \sim \mathcal{N}(0, \lambda_1^2)$, $U_{i2} \sim \mathcal{N}(0, \lambda_2^2)$, and $E_{it} \sim \mathcal{N}(0, \sigma^2)$ are mutually independent for all individuals indexed by i and time points indexed by t . This is a model considered by Fitzmaurice, Laird and Ware (2011), modified slightly to fit our setting. The parameter set is $\Theta = \{\theta = (\psi, \lambda, \sigma) \in \mathbb{R}^5 \times [0, \infty)^2 \times (0, \infty)\}$.

The maximum likelihood estimate of ψ is $\hat{\psi} = (-2.2, 0.078, 2.80, -0.040, -0.19)$, but we focus on the scale parameters whose estimates are in Table 1. Notably, the maximum likelihood estimate of λ_1 is zero, indicating common confidence regions may be unreliable. We present a confidence region for λ (Figure 5) and three componentwise confidence regions (Table 1). When creating a confidence region for a subvector or component of θ , say λ_1 , the other components are effectively nuisance parameters. Then, at noncritical points we standardize the score for λ_1 by the Schur complement (efficient information) $\tilde{\mathcal{I}}_n / \tilde{\mathcal{I}}_n^{-\lambda_1}$, where $\tilde{\mathcal{I}}_n^{-\lambda_1}$ is $\tilde{\mathcal{I}}_n$ with rows and columns corresponding to λ_1 removed and evaluate at estimates of the nuisance parameters (see, e.g., Fewster and Jupp (2013) or Bickel et al. (1998), Chapter 2, for motivation). Similarly, at critical points we standardize the modified score for λ_1 by $\tilde{\tilde{\mathcal{I}}}_n / \tilde{\tilde{\mathcal{I}}}_n^{-\lambda_1}$; we comment further on nuisance parameters in Section 6.

For context, we also present likelihood ratio intervals based on the profile likelihood computed using the `confint` function in `lme4`, and Wald intervals based on maximum likelihood estimates from `lme4` and the expected Fisher information evaluated at estimates. These use the chi-square distribution with one degree of freedom as reference. To decide whether to include the boundary points $\lambda_1 = 0$ and $\lambda_2 = 0$ in the componentwise confidence intervals, we note $\hat{\lambda}_1 = 0$ so 0 is in confidence intervals based on the profile likelihood or Wald statistic for any reference distribution. Conversely, the test statistics for λ_2 are so large at 0 that 0 should not be in either region for any relevant reference distribution. To validate this, we




FIG. 5. Componentwise test statistics for λ_1 (first plot) and λ_2 (second plot), and joint 80–99% confidence regions for λ (third plot), based on the chi-square distribution with two degrees of freedom. The dashed lines in the first two plots mark the 0.95th quantile of a chi-square distribution with one degree of freedom. The dashed lines in the third plot mark maximum likelihood estimates.

used the `varTestnlme` R package (Baey and Kuhn (2019)) to test, separately, $\lambda_1^2 = 0$ and $\lambda_2^2 = 0$, and got the p -values 0.5 and 4×10^{-78} , respectively.

The proposed interval for λ_1 is substantially smaller than that based on the likelihood ratio (Table 1). This is consistent with our simulations where the latter had greater than nominal empirical coverage of small scale parameters. The Wald interval for λ_1 is even wider than the likelihood ratio-based interval. The proposed interval for λ_2 is smaller than the other two, and its left endpoint is further from zero. Thus, there are indications the proposed procedure leads to not only reliable but more precise inference. The intervals for σ only differ in the third significance digit. This is consistent with both theory and simulations since the estimate of σ is further from zero than those of λ_1 and λ_2 , and hence the different test statistics are expected to behave similarly.

Figure 5 shows plots of the componentwise test statistics based on the proposed method for a range of λ_1 and λ_2 . The values of λ_1 and λ_2 such that the graph of the corresponding test statistic is below the critical value 3.84 (the 0.95th quantile of the chi-square distribution with one degree of freedom), give the confidence regions in Table 1. The graphs indicate the test statistics are convex in λ_1 and λ_2 , respectively. The proposed regions for λ (Figure 5, third plot), which use the chi-square distribution with two degrees of freedom as reference distribution, can be used to assess which values of λ are supported by the data. We may, for example, reject the joint null hypothesis that $\lambda_1 = \lambda_2 = 0$ at conventional levels of significance. To create these graphs and the corresponding confidence regions in Table 1, we evaluated the test statistics at a grid of 50 values each for λ_1 and λ_2 and included in the confidence regions those points where the test statistics were less than the desired quantile of the reference distributions. Thus, for the first and second plot we evaluated componentwise test statistics 50 times each, and for the third plot we evaluated the test statistic for λ at $50 \times 50 = 2500$ points. The coarseness of the grid can be adjusted depending on desired accuracy and computing times.

6. Final remarks. Linking the boundary problem to the singular information problem allows a deeper understanding of the behavior of the likelihood function in shrinking neighborhoods of the boundary of a parameter set. Perhaps more importantly, it permits the construction of confidence regions that have asymptotically correct uniform coverage probability.

The advantages of using the proposed modified score test in constructing confidence regions are many-fold: the proposed procedure does not in general require a consistent point estimator, which can be troublesome when the parameter is at or near the boundary; it does not in general rely on simulation algorithms, which typically need to be programmed for the

specific problem at hand; it can be applied to a broad variety of models, including the linear mixed model; it allows inference on scale parameters when the random effects follow an asymmetric distribution, which gives insight about the sign and the magnitude of the skewness. In addition, to the best of our knowledge, the asymptotic behavior under a sequence of parameters of the Wald and likelihood ratio test statistics for a scale parameter, has not been described for settings in which the Fisher information has any rank less than full.

Our work suggests several avenues for future research: First, more theory is needed on settings with nuisance parameters not orthogonal to the parameters of interest. Existing theory suggests replacing the Fisher information by the efficient Fisher information as we did in Section 5, but this has not been formalized for inference near critical points. Based on simulations (Supplementary Material, Ekvall and Bottai (2022)) and intuition, we conjecture our results may be adapted to settings where the block of the Fisher information corresponding to the nuisance parameters is not nearly singular, but that there may be additional challenges otherwise. Second, in some mixed models, for example, with crossed random effects, there are few independent observations even as the total number of observations grows. Then a different asymptotic theory may be of interest. There are results on the consistency of maximum likelihood estimators in such settings (Ekvall and Jones (2020), Jiang (2013)), but the properties of test statistics are largely unknown. Third, efficient software implementations of the proposed method for popular mixed models are needed: creating the proposed confidence region in practice often requires inverting the test statistic numerically, which can be computationally expensive. This is in contrast to the Wald statistic, which can be inverted analytically, but similar to the likelihood ratio statistic, which in general requires numerical inversion. A natural starting point when implementing a numerical procedure would be to consider a grid of parameter values centered at some reasonable estimates. For example, even in nonlinear mixed models where maximum likelihood estimation can be computationally expensive, fast approximate maximum likelihood estimates and Wald confidence regions are often available through penalized quasi-likelihood or Laplace approximation of the likelihood. Some further remarks on computing are in the Supplementary Material (Ekvall and Bottai (2022)).

Acknowledgments. The authors thank two referees and an Associate Editor for comments that helped improve the manuscript substantially.

SUPPLEMENTARY MATERIAL

Supplement to “Confidence regions near singular information and boundary points with applications to mixed models” (DOI: [10.1214/22-AOS2177SUPP](https://doi.org/10.1214/22-AOS2177SUPP); .pdf). The supplemental material contains proofs and additional details and results.

REFERENCES

- AZZALINI, A. (2014). *The Skew-Normal and Related Families. Institute of Mathematical Statistics (IMS) Monographs* **3**. Cambridge Univ. Press, Cambridge. [MR3468021](https://doi.org/10.1017/9781107325647)
- BAEY, C., COURNÈDE, P.-H. and KUHN, E. (2019). Asymptotic distribution of likelihood ratio test statistics for variance components in nonlinear mixed effects models. *Comput. Statist. Data Anal.* **135** 107–122. [MR3915362 https://doi.org/10.1016/j.csda.2019.01.014](https://doi.org/10.1016/j.csda.2019.01.014)
- BAEY, C. and KUHN, E. (2019). varTestnlme: Variance components testing in mixed-effect models.
- BATES, D., MÄCHLER, M., BOLKER, B. and WALKER, S. (2015). Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* **67**.
- BHATIA, R. (1997). *Matrix Analysis. Graduate Texts in Mathematics* **169**. Springer, New York. [MR1477662 https://doi.org/10.1007/978-1-4612-0653-8](https://doi.org/10.1007/978-1-4612-0653-8)
- BICKEL, P. J., KLAASSEN, C. A. J., RITOV, Y. and WELLNER, J. A. (1998). *Efficient and Adaptive Estimation for Semiparametric Models*. Springer, New York. [MR1623559](https://doi.org/10.1007/978-1-4612-0653-8)

- BILLINGSLEY, P. (1995). *Probability and Measure*, 3rd ed. *Wiley Series in Probability and Mathematical Statistics*. Wiley, New York. MR0534323
- BILLINGSLEY, P. (1999). *Convergence of Probability Measures*, 2nd ed. *Wiley Series in Probability and Statistics: Probability and Statistics*. Wiley, New York. MR1700749 <https://doi.org/10.1002/9780470316962>
- BISCIO, C. A. N., POINAS, A. and WAAGEPETERSEN, R. (2018). A note on gaps in proofs of central limit theorems. *Statist. Probab. Lett.* **135** 7–10. MR3758254 <https://doi.org/10.1016/j.spl.2017.11.009>
- BOTTAI, M. (2003). Confidence regions when the Fisher information is zero. *Biometrika* **90** 73–84. MR1966551 <https://doi.org/10.1093/biomet/90.1.73>
- BOTTAI, M. and ORSINI, N. (2004). Confidence intervals for the variance component of random-effects linear models. *The Stata Journal: Promoting Communications on Statistics and Stata* **4** 429–435.
- CHEN, S. T., XIAO, L. and STAIUCU, A.-M. (2019). An approximate restricted likelihood ratio test for variance components in generalized linear mixed models. Available at arXiv:1906.03320 [stat].
- CHESHER, A. (1984). Testing for neglected heterogeneity. *Econometrica* **52** 865–872.
- COX, D. R. and HINKLEY, D. V. (2000). *Theoretical Statistics*. Chapman & Hall/CRC, Boca Raton.
- CRAINICEANU, C. M. and RUPPERT, D. (2004). Likelihood ratio tests in linear mixed models with one variance component. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **66** 165–185. MR2035765 <https://doi.org/10.1111/j.1467-9868.2004.00438.x>
- DOCKERY, D. W., BERKEY, C. S., WARE, J. H., SPEIZER, F. E. and FERRIS JR., B. G. (1983). Distribution of forced vital capacity and forced expiratory volume in one second in children 6 to 11 years of age. *Amer. Rev. Respir. Dis.* **128** 405–412.
- DRIKVANDI, R., VERBEKE, G., KHODADADI, A. and PARTOVI NIA, V. (2013). Testing multiple variance components in linear mixed-effects models. *Biostatistics* **14** 144–159.
- EKVALL, K. O. and BOTTAI, M. (2022). Supplement to “Confidence regions near singular information and boundary points with applications to mixed models.” <https://doi.org/10.1214/22-AOS2177SUPP>
- EKVALL, K. O. and JONES, G. L. (2020). Consistent maximum likelihood estimation using subsets with applications to multivariate mixed models. *Ann. Statist.* **48** 932–952. MR4102682 <https://doi.org/10.1214/19-AOS1830>
- FEWSTER, R. M. and JUPP, P. E. (2013). Information on parameters of interest decreases under transformations. *J. Multivariate Anal.* **120** 34–39. MR3072716 <https://doi.org/10.1016/j.jmva.2013.05.010>
- FITZMAURICE, G. M., LAIRD, N. M. and WARE, J. H. (2011). *Applied Longitudinal Analysis*, 2nd ed. *Wiley Series in Probability and Statistics*. Wiley, Hoboken, NJ. MR2830137
- FITZMAURICE, G. M., LIPSITZ, S. R. and IBRAHIM, J. G. (2007). A note on permutation tests for variance components in multilevel generalized linear mixed models. *Biometrics* **63** 942–946. MR2395815 <https://doi.org/10.1111/j.1541-0420.2007.00775.x>
- FOLLAND, G. B. (2007). *Real Analysis: Modern Techniques and Their Applications*, 2nd ed. *Pure and Applied Mathematics (New York)*. Wiley, New York. MR1681462
- FRISTEDT, B. and GRAY, L. (1997). *A Modern Approach to Probability Theory. Probability and Its Applications*. Birkhäuser, Inc., Boston, MA. MR1422917 <https://doi.org/10.1007/978-1-4899-2837-5>
- GEYER, C. J. (1994). On the asymptotics of constrained M -estimation. *Ann. Statist.* **22** 1993–2010. MR1329179 <https://doi.org/10.1214/aos/1176325768>
- GIAMPAOLI, V. and SINGER, J. M. (2009). Likelihood ratio tests for variance components in linear mixed models. *J. Statist. Plann. Inference* **139** 1435–1448. MR2485137 <https://doi.org/10.1016/j.jspi.2008.06.016>
- GREVEN, S., CRAINICEANU, C. M., KÜCHENHOFF, H. and PETERS, A. (2008). Restricted likelihood ratio testing for zero variance components in linear mixed models. *J. Comput. Graph. Statist.* **17** 870–891. MR2649072 <https://doi.org/10.1198/106186008X386599>
- HALL, D. B. and PRÆSTGAARD, J. T. (2001). Order-restricted score tests for homogeneity in generalised linear and nonlinear mixed models. *Biometrika* **88** 739–751. MR1859406 <https://doi.org/10.1093/biomet/88.3.739>
- 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>
- KNUDSON, C., BENSON, S., GEYER, C. and JONES, G. (2021). Likelihood-based inference for generalized linear mixed models: Inference with the R package `glm`. *Stat* **10** Paper No. e339. MR4218938 <https://doi.org/10.1002/sta4.339>
- LEE, L. F. and CHESHER, A. (1986). Specification testing when score test statistics are identically zero. *J. Econometrics* **31** 121–149. MR0846287 [https://doi.org/10.1016/0304-4076\(86\)90045-X](https://doi.org/10.1016/0304-4076(86)90045-X)
- LIN, X. (1997). Variance component testing in generalised linear models with random effects. *Biometrika* **84** 309–326. MR1467049 <https://doi.org/10.1093/biomet/84.2.309>
- 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

- QU, L., GUENNEL, T. and MARSHALL, S. L. (2013). Linear score tests for variance components in linear mixed models and applications to genetic association studies. *Biometrics* **69** 883–892. MR3146784 <https://doi.org/10.1111/biom.12095>
- ROTNITZKY, A., COX, D. R., BOTTAI, M. and ROBINS, J. (2000). Likelihood-based inference with singular information matrix. *Bernoulli* **6** 243–284. MR1748721 <https://doi.org/10.2307/3318576>
- SAVILLE, B. R. and HERRING, A. H. (2009). Testing random effects in the linear mixed model using approximate Bayes factors. *Biometrics* **65** 369–376. MR2751460 <https://doi.org/10.1111/j.1541-0420.2008.01107.x>
- SELF, S. G. and LIANG, K.-Y. (1987). Asymptotic properties of maximum likelihood estimators and likelihood ratio tests under nonstandard conditions. *J. Amer. Statist. Assoc.* **82** 605–610. MR0898365
- SINHA, S. K. (2009). Bootstrap tests for variance components in generalized linear mixed models. *Canad. J. Statist.* **37** 219–234. MR2531828 <https://doi.org/10.1002/cjs.10012>
- STERN, S. E. and WELSH, A. H. (2000). Likelihood inference for small variance components. *Canad. J. Statist.* **28** 517–532. MR1793108 <https://doi.org/10.2307/3315962>
- STRAM, D. O. and LEE, J. W. (1994). Variance components testing in the longitudinal mixed effects model. *Biometrics* **50** 1171–1177.
- STRAM, D. O. and LEE, J. W. (1995). Corrections: Variance component testing in the longitudinal mixed effects model. *Biometrics* **51** 1196.
- VERBEKE, G. and MOLENBERGHS, G. (2003). The use of score tests for inference on variance components. *Biometrics* **59** 254–262. MR1987392 <https://doi.org/10.1111/1541-0420.00032>
- WIENCIERZ, A., GREVEN, S. and KÜCHENHOFF, H. (2011). Restricted likelihood ratio testing in linear mixed models with general error covariance structure. *Electron. J. Stat.* **5** 1718–1734. MR2870148 <https://doi.org/10.1214/11-EJS654>
- WOOD, S. N. (2013). A simple test for random effects in regression models. *Biometrika* **100** 1005–1010. MR3142347 <https://doi.org/10.1093/biomet/ast038>
- ZHU, H. and ZHANG, H. (2006). Generalized score test of homogeneity for mixed effects models. *Ann. Statist.* **34** 1545–1569. MR2278367 <https://doi.org/10.1214/009053606000000380>