

TRACTABLE BAYESIAN VARIABLE SELECTION: BEYOND NORMALITY

DAVID ROSSELL AND FRANCISCO J. RUBIO

ABSTRACT. Bayesian variable selection for continuous outcomes often assumes normality, and so do its theoretical studies. There are sound reasons behind this assumption, particularly for large p : ease of interpretation, analytical and computational convenience. More flexible frameworks exist, including semi- or non-parametric models, often at the cost of losing some computational or theoretical tractability. We propose a simple extension of the Normal model that allows for skewness and thicker-than-normal tails but preserves its tractability. We show that a classical strategy to induce asymmetric Normal and Laplace errors via two-piece distributions leads to easy interpretation and a log-concave likelihood that greatly facilitates optimization and integration. We also characterize asymptotically its maximum likelihood estimator and Bayes factor rates under model misspecification. Our work focuses on the likelihood and can thus be combined with any likelihood penalty or prior, but here we adopt non-local priors, a family that induces extra sparsity and which we characterize under misspecification for the first time. Under suitable conditions Bayes factor rates are of the same order as those that would be obtained under the correct model, but we point out a potential loss of sensitivity to detect truly active covariates. Our examples show how a novel approach to infer the error distribution leads to substantial gains in sensitivity, thus warranting the effort to go beyond normality, whereas for near-normal data one can get substantial speedups relative to assuming unnecessarily flexible models. The methodology is available as part of R package `mombf`.

KEYWORDS: Variable selection, two-piece errors, Bayes factors, model misspecification, robust regression

1. INTRODUCTION

The rise of high-dimensional problems has generated a renewed interest in simple models. Beyond the obvious issue that modest sample sizes limit the number of parameters that can be learned accurately, simple models remain a central choice due to their analytical and computational tractability, ease of interpretation, and the fact that they often work well in practice. There is, however, a pressing need to seek extensions which, while retaining the aforementioned advantages, incorporate additional flexibility and can be studied without unrealistically assuming that the posed model is correct. Ideally such extensions should detect when the added flexibility is not needed so that one can fall back onto simpler models. Here we tackle canonical variable selection in linear regression from a Bayesian standpoint, although some results may also prove useful for penalized likelihood methods. Our work is based on two-piece distributions, an easily interpretable

family that has a long history and which we fully characterize in the linear model case (synthesizing and extending current results) under model misspecification. Our main contributions are showing that two-piece errors (specifically when applied to the Normal and Laplace families) lead to highly tractable inference, proposing simple computational algorithms, and characterizing the resulting variable selection under model misspecification, including when this likelihood is combined with non-local priors (NLPs). We show that in the presence of asymmetries or heavy tails the normal model incurs a significant loss of power, and propose a formal strategy to detect such departures from normality. When these departures are negligible our model collapses onto Normal errors, leading to an increased computational efficiency relative to more flexible models, otherwise the additional runtime is justified on the basis of increased sensitivity.

To fix ideas, we consider the linear regression model

$$(1) \quad y = X\theta + \epsilon,$$

where $y = (y_1, \dots, y_n)^T$ is the observed outcome for n individuals, X is an $n \times p$ matrix with potential predictors, $\theta = (\theta_1, \dots, \theta_p)^T \in \mathbb{R}^p$ are regression coefficients and $\epsilon = (\epsilon_1, \dots, \epsilon_n)^T$ are independent and identically distributed errors. The goal is to determine the non-zero coefficients in θ under an arbitrary data-generating distribution for the ϵ_i 's, building a simple framework that remains applicable to large p . To this end let $\gamma_j = I(\theta_j \neq 0)$ for $j = 1, \dots, p$ be the variable inclusion indicators, $p_\gamma = \sum_{j=1}^p \gamma_j$ the number of active variables and, to consider that residuals may be asymmetric and/or have thicker-than-normal tails $\gamma_{p+1} = 1$ denotes the presence of asymmetry ($\gamma_{p+1} = 0$ otherwise) and $\gamma_{p+2} = 1$ that of thick tails ($\gamma_{p+2} = 0$ for Normal tails). Thus $\gamma = (\gamma_1, \dots, \gamma_{p+2})$ denotes the assumed model. X_γ and θ_γ are the corresponding submatrix of X and subvector of θ , respectively. We denote the i^{th} row in X and X_γ by $x_i^T \in \mathbb{R}^p$ and $x_{\gamma i}^T \in \mathbb{R}^{p_\gamma}$.

There are a number of proposals to relax the normality assumption. Within the frequentist literature Wang et al. (2007) proposed median regression with LASSO penalties (LASSO-LAD) and Wang and Li (2009) with rank-based SCAD penalties. Arslan (2012) extended LASSO median regression by weighting observations and Fan et al. (2014) considered adaptive LASSO quantile regression. These approaches are formally connected to assuming either Laplace or (two-piece) asymmetric Laplace errors. There are also model-free methods based on M-estimation, *e.g.* combining Huber's loss with an adaptive LASSO penalty Lambert-Lacroix (2011), sparse trimmed-means LASSO Alfons et al. (2013), and non-negative garrote extensions to induce robustness to outliers Gijbels and Vrinssen (2015). Theoretical characterizations also exist, *e.g.* Mendelson (2014) proved the consistency and asymptotic normality of high-dimensional M-estimators and Loh (2015) extended the results to generalized M-estimators with non-convex loss functions. Within the Bayesian framework, Gottardo and Raftery (2007) and Wang et al. (2016) consider variable selection after transforming y_i and/or x_i , the former allowing for t errors and the latter inducing NLPs on θ via the transformation's Jacobians. While certainly interesting, the transformed conditional mean $E(y_i | x_i)$ is no longer linear in x_i and parameter interpretation and prior elicitation is less straightforward, our main interest here is in linear predictors with simple error distributions. Along these lines, Yu et al. (2013)

proposed Gibbs sampling for model choice in Bayesian quantile regression using a latent scale augmentation, and Yan and Kottas (2015) extended Azzalini’s skew Normal to Laplace errors within Bayesian quantile regression, which leads to easily-implementable MCMC, and induced sparsity via LASSO penalties. Related to our work Rubio and Genton (2016) and Rubio and Yu (2016; in press) employ skew-symmetric and two-piece errors in linear regression, respectively, albeit the set of covariates is fixed and they focus on prediction and censored responses. Yet another possible avenue is to pose highly flexible errors, *e.g.* Chung and Dunson (2009) set a non-parametric model to simultaneously learn the effect of x_i on the mean and on the shape of the residual distribution, and Kundu and Dunson (2014) proposed variable selection with non-parametric symmetric residuals, for which notably Chae et al. (2016) proved model selection consistency and concentration rates under model misspecification and p growing with n . Most Bayesian work in this area uses Markov Chain Monte Carlo (MCMC) for parameter estimation and computation of marginal likelihoods and does not collapse onto the Normal model when warranted by the data, hampering its computational scalability as p or n grow, further the theoretical study is often either absent or M-closed.

In contrast, we show that simpler parametric error models equipped with efficient analytical approximations to the integrated likelihood (thus bypassing the need to use MCMC) also achieve selection consistency under model misspecification, and embed these models within a framework that when appropriate may collapse onto normality to give further speed-ups. We also describe how model misspecification can produce marked drops in the sensitivity to detect truly active variables, in particular under the presence of asymmetry or heavy tails. This effect of model misspecification on Bayesian variable selection complements the examples in Grünwald and van Ommen (2014), who illustrated how under the presence of inliers model misspecification may favour the addition of spurious variables, and is in agreement with the example in Figure 1 of Kundu and Dunson (2014). We view this combination of NLPs to discard spurious coefficients with flexible likelihoods to increase power as a promising strategy. The manuscript is structured as follows. Section 2 reviews two-piece distributions and establishes the concavity of the log-likelihood in the asymmetric Normal and Laplace cases. Section 3 proposes a prior formulation based on NLPs Johnson and Rossell (2010) that enforces variable selection sparsity and discards degrees of asymmetry that are irrelevant in practice. Section 4 tackles maximum likelihood and posterior mode estimation, specifically giving asymptotic distributions and optimization algorithms that capitalize on the likelihood tractability. Section 5 outlines a framework to select both variables and the residual distribution, proposes convenient approximations to the integrated likelihood and characterizes asymptotically the associated Bayes factors. Finally, the methodology is illustrated in simulated and experimental data in Section 6 and concluding remarks are offered in Section 7. The supplementary material contains all proofs and further results.

2. LOG-LIKELIHOOD

We recall the definition of a two-piece distribution for model (1) and predictors X_γ .

Definition 1. A random variable $y_i \in \mathbb{R}$ following a two-piece distribution with location $x_{\gamma_i}^T \theta_\gamma$, scale $\sqrt{\vartheta} \in \mathbb{R}^+$ and asymmetry α has density function $s(y_i; x_{\gamma_i}^T \theta_\gamma, \vartheta, \alpha) =$

$$(2) \quad \frac{2}{\sqrt{\vartheta}[a(\alpha) + b(\alpha)]} \left[f\left(\frac{y_i - x_{\gamma_i}^T \theta_\gamma}{\sqrt{\vartheta}a(\alpha)}\right) I(y_i < x_{\gamma_i}^T \theta_\gamma) + f\left(\frac{y_i - x_{\gamma_i}^T \theta_\gamma}{\sqrt{\vartheta}b(\alpha)}\right) I(y_i \geq x_{\gamma_i}^T \theta_\gamma) \right],$$

where $f(\cdot)$ is a symmetric unimodal density with mode at 0 and support on \mathbb{R} , and $a(\alpha), b(\alpha) \in \mathbb{R}^+$.

Two-piece distributions induce asymmetry by (continuously) merging two symmetric densities that have the same mode $x_{\gamma_i}^T \theta_\gamma$ but different scale parameters $\sqrt{\vartheta}a(\alpha)$, $\sqrt{\vartheta}b(\alpha)$ on each side of the mode. Some popular parameterizations for these are the inverse scale factors $\{a(\alpha), b(\alpha)\} = \{\alpha, 1/\alpha\}$ for $\alpha \in \mathbb{R}^+$ Fernández and Steel (1998) or the epsilon-skew parameterization $\{a(\alpha), b(\alpha)\} = \{1-\alpha, 1+\alpha\}$ for $\alpha \in [-1, 1]$ Mudholkar and Hutson (2000). We adopt the latter as it leads to orthogonality in the expected log-likelihood hessian between α and ϑ , also it allows easy interpretation as the total variation distance between $s(y_i; x_{\gamma_i}^T \theta_\gamma, \vartheta, \alpha)$ and its symmetric counterpart $s(y_i; x_{\gamma_i}^T \theta_\gamma, \vartheta, 0)$ is $|\alpha|/2$ Dette et al. (2016). Further, a classical skewness coefficient proposed by Arnold-Groeneveld defined as $AG = 1 - 2F(x_{\gamma_i}^T \theta_\gamma) \in [-1, 1]$ for a univariate random variable with mode at $x_{\gamma_i}^T \theta_\gamma$ and cumulative distribution function $F(\cdot)$, is equal to $AG = -\alpha$ Rubio and Steel (2014).

Two-piece distributions are appealing for regression given that the mode of $s(\cdot)$ is $x_{\gamma_i}^T \theta_\gamma$, its mean (when defined) depends on x_{γ_i} only through $x_{\gamma_i}^T \theta_\gamma$ and its variance is proportional to ϑ (see below for specific expressions), facilitating interpretation and prior elicitation. Despite these properties and their being a classical strategy with a fascinating history, proposed at least as early as 1897 and rediscovered multiple times Wallis (2014), their popularity has been limited due to practical concerns, *e.g.* log-likelihood maximization may be hampered by discontinuous gradients or Hessians. For this reason we focus on two-piece Normal and Laplace errors, for which we prove log-concavity and thus analytical and computational tractability, giving a practical mechanism to capture asymmetry and heavier-than-normal tails. Specifically, the two-piece Normal is obtained by letting $f(z) = N(z; 0, 1)$ in (2) be the standard normal density, and gives $E(y_i | x_{\gamma_i}) = x_{\gamma_i}^T \theta_\gamma - \alpha \sqrt{8\vartheta/\pi}$, $\text{Var}(y_i | x_{\gamma_i}) = \vartheta[(3 - 8/\pi)\alpha^2 + 1]$ and a median that is also linear in x_{γ_i} Mudholkar and Hutson (2000). The corresponding likelihood has the simple expression $\log L_1(\theta_\gamma, \vartheta, \alpha) =$

$$(3) \quad -\frac{n}{2} \log(2\pi) - \frac{n}{2} \log(\vartheta) - \frac{1}{2\vartheta} \left(\sum_{i \in A(\theta_\gamma)} \frac{(y_i - x_{\gamma_i}^T \theta_\gamma)^2}{(1 + \alpha)^2} + \sum_{i \notin A(\theta_\gamma)} \frac{(y_i - x_{\gamma_i}^T \theta_\gamma)^2}{(1 - \alpha)^2} \right) = \\ = -\frac{n}{2} \log(2\pi) - \frac{n}{2} \log(\vartheta) - \frac{1}{2\vartheta} (y - X_\gamma \theta_\gamma)^T W^2 (y - X_\gamma \theta_\gamma).$$

where $A(\theta_\gamma) = \{i : y_i < x_{\gamma_i}^T \theta_\gamma\}$ are the observations with negative residuals, $W = \text{diag}(w)$, $w_i = |1 + \alpha|^{-1}$ if $i \in A(\theta_\gamma)$ and $w_i = |1 - \alpha|^{-1}$ if $i \notin A(\theta_\gamma)$. For later convenience we denote by \bar{w} the signed weight vector with $\bar{w}_i = w_i$ if $i \in A(\theta_\gamma)$ and $\bar{w}_i = -w_i$

if $i \notin A(\theta_\gamma)$, by $w^k = (w_1^k, \dots, w_n^k)$ the element-wise k^{th} power of a vector, $\bar{w}^k = (\text{sign}(\bar{w}_1)|\bar{w}_1|^k, \dots, \text{sign}(\bar{w}_n)|\bar{w}_n|^k)^T$ and $\bar{W}^k = \text{diag}(\bar{w}^k)$. Note that (3) is linked to asymmetric least square regression and reduces to the Normal likelihood for $\alpha = 0$.

The two-piece Laplace is obtained by setting $f(z) = 0.5 \exp(-|z|)$ in (2). This distribution is more commonly referred to as asymmetric Laplace, we denote it $y_i \sim \text{AL}(x_{\gamma_i}^T \theta_\gamma, \vartheta, \alpha)$ and note that $E(y_i | x_{\gamma_i}, \theta_\gamma, \vartheta, \alpha) = x_{\gamma_i}^T \theta_\gamma - 2\alpha\sqrt{\vartheta}$ and $\text{Var}(y_i | x_{\gamma_i}) = 2\vartheta(1 + \alpha^2)$ Arellano-Valle et al. (2005). For coherency from here onwards we also refer to the two-piece Normal as asymmetric Normal and denote $y_i \sim \text{AN}(x_{\gamma_i}^T \theta_\gamma, \vartheta, \alpha)$. The asymmetric Laplace log-likelihood is $\log L_2(\theta_\gamma, \vartheta, \alpha) =$

$$(4) \quad -n \log(2) - \frac{n}{2} \log(\vartheta) - \frac{1}{\sqrt{\vartheta}} \left(\sum_{i \in A(\theta_\gamma)} \frac{|y_i - x_{\gamma_i}^T \theta_\gamma|}{1 + \alpha} + \sum_{i \notin A(\theta_\gamma)} \frac{|y_i - x_{\gamma_i}^T \theta_\gamma|}{1 - \alpha} \right).$$

The symmetric Laplace case is obtained for $\alpha = 0$, in which case optimization of (4) with respect to θ_γ is equivalent to median regression, whereas for fixed $\alpha \neq 0$ it leads to quantile regression. We remark that in our framework α is a parameter to be learnt from the data, which is intrinsically different from and may lead to different selected covariates than classical quantile regression.

Our first results regarding the tractability of (3)-(4) are given in Propositions 1-2 (Proposition 1(i) was already shown by Mudholkar and Hutson (2000)).

Proposition 1. *The asymmetric normal log-likelihood in (3) satisfies:*

(i) *Its gradient is continuous and is given by*

$$g_1(\theta_\gamma, \vartheta, \alpha) = \begin{pmatrix} \frac{1}{\vartheta} X_\gamma^T W^2 (y - X_\gamma \theta_\gamma) \\ -\frac{n}{2\vartheta} + \frac{1}{2\vartheta^2} (y - X_\gamma \theta_\gamma)^T W^2 (y - X_\gamma \theta_\gamma) \\ \frac{1}{\vartheta} (y - X_\gamma \theta_\gamma)^T \bar{W}^3 (y - X_\gamma \theta_\gamma) \end{pmatrix}.$$

(ii) *Its Hessian with respect to θ_γ is continuous everywhere except on the zero Lebesgue measure set $\{\theta_\gamma \in \mathbb{R}^p : x_{\gamma_i}^T \theta_\gamma = y_i \text{ for some } i = 1, \dots, n\}$, and is $H_1(\theta_\gamma, \vartheta, \alpha) = \vartheta^{-1} \times$*

$$\begin{pmatrix} -X_\gamma^T W^2 X_\gamma & \frac{1}{\vartheta} X_\gamma^T W^2 (X_\gamma \theta_\gamma - y) & -2X_\gamma^T \bar{W}^3 (y - X_\gamma \theta_\gamma) \\ \frac{1}{\vartheta} (X_\gamma \theta_\gamma - y)^T W^2 X_\gamma & \frac{n}{2\vartheta} - \frac{(y - X_\gamma \theta_\gamma)^T \bar{W}^2 (y - X_\gamma \theta_\gamma)}{\vartheta^2} & -\frac{1}{\vartheta} (y - X_\gamma \theta_\gamma)^T \bar{W}^3 (y - X_\gamma \theta_\gamma) \\ -2(y - X_\gamma \theta_\gamma)^T \bar{W}^3 X_\gamma & -\frac{1}{\vartheta} (y - X_\gamma \theta_\gamma)^T \bar{W}^3 (y - X_\gamma \theta_\gamma) & -3(y - X_\gamma \theta_\gamma)^T W^4 (y - X_\gamma \theta_\gamma) \end{pmatrix},$$

(iii) *If $\text{rank}(X_\gamma) = p_\gamma$ then $H_1(\theta_\gamma, \vartheta, \alpha)$ is strictly negative definite with respect to (θ_γ, α) and (3) has a unique maximum $(\hat{\theta}_\gamma, \hat{\vartheta}, \hat{\alpha})$. Alternatively, if $\text{rank}(X_\gamma) < p_\gamma$ then $H_1(\theta_\gamma, \vartheta, \alpha)$ is negative semidefinite.*

The implication is that, analogously to Normal errors, when X_γ has full rank (3) is continuous and concave almost everywhere in (θ_γ, α) . This fact combined with $\log L_1$ having a continuous gradient guarantees overall concavity and hence a unique maximum (see the proof for a formal argument). Further, inspection of (1) reveals that $\log L_1$ is locally quadratic as a function of θ_γ within regions of constant $A(\theta_\gamma)$ and that its

maximizer with respect to (θ_γ, α) does not depend on ϑ , two observations that facilitate optimization.

Proposition 2 shows that, although $\log L_2$ is piecewise-linear in θ_γ and thus has a singular hessian, one can prove concavity and uniqueness of a maximum in terms of (θ_γ, α) as in Proposition 1, extending the well-known result of concavity with respect to only θ_γ Koenker (2005). In Sections 4-5 we describe how the result facilitates computation, in particular leading to simple optimization and analytical approximations to integrated likelihoods, and asymptotic characterizations.

Proposition 2. *The asymmetric Laplace log-likelihood in (4) satisfies:*

(i) *It is continuously differentiable with gradient*

$$g_2(\theta_\gamma, \vartheta, \alpha) = \vartheta^{-\frac{1}{2}} \times \begin{pmatrix} -X_\gamma^T \bar{w} \\ -\frac{n}{2\vartheta^{\frac{1}{2}}} + \frac{1}{2\vartheta} w^T |y - X_\gamma \theta_\gamma| \\ |y - X_\gamma \theta_\gamma|^T \bar{w}^2 \end{pmatrix},$$

except on the zero Lebesgue measure set $\{\theta_\gamma \in \mathbb{R}^p : x_{\gamma_i}^T \theta_\gamma = y_i \text{ for some } i = 1, \dots, n\}$, where the gradient is undefined.

(ii) *Its Hessian with respect to θ_γ is continuous everywhere except on the zero Lebesgue measure set $\{\theta_\gamma \in \mathbb{R}^p : x_{\gamma_i}^T \theta_\gamma = y_i \text{ for some } i = 1, \dots, n\}$, and is $H_2(\theta_\gamma, \vartheta, \alpha) = \vartheta^{-1/2} \times$*

$$\begin{pmatrix} 0 & \frac{1}{2\vartheta} X_\gamma^T \bar{w} & X_\gamma^T w^2 \\ \frac{1}{2\vartheta} \bar{w}^T X_\gamma & \frac{n}{2\vartheta^{\frac{3}{2}}} - \frac{3}{4\vartheta^2} w^T |y - X_\gamma \theta_\gamma| & -\frac{1}{2\vartheta} |y - X_\gamma \theta_\gamma|^T \bar{w}^2 \\ (X_\gamma^T w^2)^T & -\frac{1}{2\vartheta} |y - X_\gamma \theta_\gamma|^T \bar{w}^2 & -2|y - X_\gamma \theta_\gamma|^T \bar{w}^3 \end{pmatrix}.$$

(iii) *If $\text{rank}(X_\gamma) = p_\gamma$ then (4) is strictly concave in (θ_γ, α) and has a unique maximum $(\hat{\theta}_\gamma, \hat{\vartheta}, \hat{\alpha})$. Alternatively, if $\text{rank}(X_\gamma) < p_\gamma$ then it is non-strictly concave in (θ_γ, α) .*

3. PRIOR FORMULATION

We complete the Bayesian model via priors on the model indicators γ and the model-specific parameters (θ_γ, α) . For $p(\gamma)$ by default we adopt the standard Beta-Binomial(a_γ, b_γ) prior Scott and Berger (2010) where $a_\gamma, b_\gamma > 0$ are known constants (by default $a_\gamma = b_\gamma = 1$), although our implementation also incorporates uniform and Binomial priors. The four posed residual distributions (Normal, asymmetric Normal, Laplace and asymmetric Laplace) are assigned equal prior probability independently from the variable inclusions. Therefore

$$(5) \quad p(\gamma) = \frac{1}{4} \frac{B(a_\gamma + \sum_{j=1}^p \gamma_j, b_\gamma + p - \sum_{j=1}^p \gamma_j)}{B(a_\gamma, b_\gamma)},$$

where $B(\cdot)$ is the Beta function. Any model with $p_\gamma > n$ is assigned $p(\gamma) = 0$, as it would result in data interpolation.

Regarding $p(\theta_\gamma | \gamma)$, given that the mode, mean and median of y_i are linear in $x_{\gamma_i}^T \theta_\gamma$ the usual prior specification strategies under Normal errors remain sensible. The possibilities

are too numerous to list here, see *e.g.* Bayarri et al. (2012) or Mallick and Nengjun (2013) and references therein. We focus on the class of NLPs introduced by Johnson and Rossell (2010), as these lead to stronger sparsity than conventional (local) priors and (under suitable conditions) consistency of posterior model probabilities in high-dimensional Normal regression where $p = o(n)$ Johnson and Rossell (2012) or $p = o(e^n)$ Shin et al. (2015). The basic intuition is that under model γ all elements in θ_γ are assumed to be non-zero, thus $p(\theta_\gamma | \gamma)$ should vanish as any element in θ_γ approaches 0. We focus on two specific choices Johnson and Rossell (2012); Rossell et al. (2013)

$$(6) \quad p_M(\theta_\gamma | \vartheta, \gamma) = \prod_{\gamma_j=1} \frac{\theta_j^2}{k g_\theta \vartheta} N(\theta_j; 0, g_\theta k \vartheta),$$

$$(7) \quad p_E(\theta_\gamma | \vartheta, \gamma) = \prod_{\gamma_j=1} \exp \left\{ \sqrt{2} - \frac{g_\theta k \vartheta}{\theta_j^2} \right\} N(\theta_j; 0, g_\theta k \vartheta),$$

called product MOM and eMOM priors (respectively), where g_θ is a known prior dispersion. For Normal or asymmetric Normal errors $k = 1$, and for the Laplace or asymmetric Laplace $k = 2$ as then $\text{Var}(\epsilon_i)$ is proportional to 2ϑ . Along the same lines for the scale parameter we set a standard inverse gamma $p(\vartheta | \gamma) = \text{IG}(\vartheta; a_\vartheta/2, k b_\vartheta/2)$ (in our examples $a_\vartheta = b_\vartheta = 0.01$). MOM vanishes at a quadratic speed around the origin and accelerates polynomial Bayes factor sparsity rates, whereas eMOM vanishes exponentially and leads to quasi-exponential rates Johnson and Rossell (2010); Rossell and Telesca (2015), a result we extend here for our new class of models and under model misspecification (Section 5). In our examples we follow the default recommendation in Johnson and Rossell (2010) and set $g_\theta = 0.348, 0.119$ for MOM and eMOM (respectively), under the rationale that they assign 0.01 prior probability to $|\theta_i/\sqrt{\vartheta}| < 0.2$, *i.e.* effect sizes often deemed practically irrelevant. Naturally, whenever prior information is available we recommend using it to set g_θ . The supplementary material describes a third prior class called iMOM that provides a thick-tailed counterpart to the eMOM. Although the iMOM is implemented in our software we do not consider it further here given that its performance was very similar to the eMOM but it has the unappealing property of leading to non-convex optimization (akin to other thick-tailed priors, *e.g.* Cauchy), and when considering $p(\alpha)$ (see below) it leads to a density that diverges on the boundary ($\alpha = -1$ or $\alpha = 1$).

To set $p(\alpha | \gamma_{p+1} = 1)$ ($\alpha = 0$ under $\gamma_{p+1} = 0$) we reparameterize $\tilde{\alpha} = \text{atanh}(\alpha) \in \mathbb{R}$ as in Rubio and Steel (2014). These authors proposed $0.5(1 + \alpha) \sim \text{Beta}(2, 2)$, which places the prior mode at $\alpha = 0$ and thus defines a local prior. Our goal here is to detect situations where the degree of asymmetry is practically relevant and to otherwise allow the posterior to collapse on the symmetric model. To achieve this we consider $p_M(\tilde{\alpha} | \gamma_{p+1} = 1) = \tilde{\alpha}^2 \phi(\tilde{\alpha}/\sqrt{g_\alpha})/\sqrt{g_\alpha}$, and $p_E(\tilde{\alpha} | \gamma_{p+1} = 1) = e^{\sqrt{2}-g_\alpha/\tilde{\alpha}^2} N(\tilde{\alpha}; 0, g_\alpha)$ where $g_\alpha \in \mathbb{R}^+$ is a fixed prior dispersion parameter. To set g_α by default we consider that Arnold-Groeneveld asymmetry coefficients $|\alpha| < 0.2$ are often practically irrelevant, thus we set g_α such that $P(|\alpha| \geq 0.2) = 0.99$. Also note that $\alpha = 2$ gives a total variation distance of $|\alpha|/2 = 0.1$, *i.e.* the largest difference $|P(\epsilon_i \in A | \alpha = 0) - P(\epsilon_i \in A | \alpha)|$ for any set $A \subseteq [-1, 1]$ is 0.1, which we view as typically irrelevant. Since $\text{atanh}(0.2) = 0.203$,

a direct calculation gives that $P(|\tilde{\alpha}| \geq 0.203) = 0.99$ when $g_\alpha = 0.357, 0.122$ under MOM and eMOM. To assess sensitivity in our examples we also considered g_α such that $P(|\alpha| \geq 0.1) = 0.99$ (total variation distance=0.05), giving $g_\alpha = 0.087, 0.030$. Figure 1 depicts $p(\alpha)$ under these settings. Our results showed that variable selection is typically robust to choices of g_α within this range.




FIGURE 1. Default priors for α .

4. PARAMETER ESTIMATION

We obtain some results for parameter estimation under a given γ that will also prove useful for variable selection (Section 5). Section 4.1 gives the limiting distribution of $(\hat{\theta}_\gamma, \hat{\vartheta}_\gamma, \hat{\alpha}_\gamma) = \arg \max_{\theta_\gamma, \vartheta, \alpha} \log L_k(\theta_\gamma, \vartheta, \alpha)$ as $n \rightarrow \infty$ for asymmetric Normal ($k = 1$) and Laplace ($k = 2$) when data are generated from the linear model (1) but the error distribution may be misspecified. Briefly, as is typically the case in such settings we obtain parameter estimation consistency and asymptotic normality, albeit there is a loss of efficiency and an underestimation of uncertainty. Section 4.2 presents novel optimization algorithms for maximum likelihood and posterior mode estimation we designed to improve the computational scalability of current related methods.

4.1. Asymptotic distributions. We lay out technical conditions for our results to hold.

- A1.** The parameter space $\Gamma \subset \mathbb{R}^p \times \mathbb{R}_+ \times (-1, 1)$ is compact and convex.
- A2.** Data are truly generated as $y_i = x_i^T \theta^* + \epsilon_i$ for some $\theta^* \in \mathbb{R}^p$, fixed $p_{\gamma^*} = \sum_{j=1}^p \mathbf{I}(\theta_j^* \neq 0)$ and ϵ_i are *i.i.d.* and independent of x_i . Let the data-generating $y_i | x_i \stackrel{i.i.d.}{\sim} S_0(\cdot | x_i)$ with density $s_0(y_i | x_i) > 0$ for all y_i .
- A3.** For all γ there is some n_0 such that $X_\gamma^T X_\gamma$ is strictly positive definite almost surely for all $n > n_0$.

A4. Denote by $x_i \stackrel{i.i.d.}{\sim} \Psi(\cdot)$ the generating process of the covariates (which can be either stochastic or deterministic).

$$\int |y_1|^j dS_0(y_1|x_1) d\Psi(x_1) < \infty,$$

$$\int \|x_1\|^j d\Psi(x_1) < \infty,$$

where $j = 1, 2,$ or $4,$ and we specify the order j of interest in each of the results below, and $\|\cdot\|$ denotes the Euclidean distance $\|z\| = (\sum z_i^2)^{\frac{1}{2}}$.

A5. For $\eta \in \Gamma$

$$\int \frac{\partial}{\partial \eta_j} \left[\int m_\eta(y_1, x_1) dS_0(y_1|x_1) \right] d\Psi(x_1) = \frac{\partial}{\partial \eta_j} \int \int m_\eta(y_1, x_1) dS_0(y_1|x_1) d\Psi(x_1),$$

$$\int \frac{\partial^2}{\partial \eta_i \partial \eta_j} \left[\int m_\eta(y_1, x_1) dS_0(y_1|x_1) \right] d\Psi(x_1) = \frac{\partial^2}{\partial \eta_i \partial \eta_j} \int \int m_\eta(y_1, x_1) dS_0(y_1|x_1) d\Psi(x_1).$$

These conditions are in line with those in classical robust regression, *e.g.* see Huber (1973) or Koenker and Bassett (1982). Condition A1 is made out of technical convenience, naturally one may take an arbitrarily large Γ . Condition A2 states that data truly arise from a linear model, where the key assumption is that the residuals are independent. Condition A3 holds whenever the rows of X are regarded as a deterministic sequence satisfying the condition, or for instance when x_i are independent and identically distributed from an underlying distribution of fixed dimension with positive-definite $\text{Cov}(x_1)$, as then $X^T X$ converges almost surely to a positive-definite matrix by the strong law of large numbers. We focus on fixed p , extensions to p growing with n are possible along the lines in Mendelson (2014), but its detailed treatment is beyond the scope of this paper. Condition A4 requires existence of moments up to certain order. Condition A5 requires being able to exchange integration and differentiation, and is needed only to prove asymptotic normality.

Our results summarize and extend classical studies focusing on θ_γ in least squares, median and quantile regression to consider the whole parameter vector $(\theta_\gamma, \vartheta, \alpha)$. Briefly, Eicker (1964) and Srivastava (1971) showed that the least squares estimator ($k = 1, \alpha = 0$) satisfies $\sqrt{n}V^T(\hat{\theta}_\gamma - \theta_0) \xrightarrow{D} N(0, \text{Var}(\epsilon_1)I)$, where θ_0 minimizes Kullback-Leibler divergence to the data-generating truth and $VV^T = X_\gamma^T X_\gamma/n$, assuming that $\text{Var}(\epsilon_1) < \infty$ and minimum conditions on $X_\gamma^T X_\gamma$. To our knowledge the asymmetric Normal has been much less studied, *e.g.* Kimber (1985), Mudholkar and Hutson (2000) and Arellano-Valle et al. (2005) considered the case with no covariates and no checks of the conditions required by large sample theory are shown, which are non-trivial given that $H_1(\theta_\gamma, \vartheta, \alpha)$ is discontinuous. Regarding Laplace errors ($k = 2, \alpha = 0$) Pollard (1991) and Knight (1999) showed $2f_0\sqrt{n}V^T(\hat{\theta}_\gamma - \theta_0) \xrightarrow{D} N(0, I)$ where $f_0 = p(\epsilon_0)$ and ϵ_0 is the median of $s_0(\epsilon_i)$ under mild conditions on $X_\gamma^T X_\gamma$ and $f_0 > 0$. Koenker (1994) generalized the result to the asymmetric Laplace obtaining $2f_0\sqrt{n/(1-\alpha^2)}V^T(\hat{\theta}_\gamma - \theta_0) \xrightarrow{D} N(0, I)$, where $f_0 = p(\epsilon)$ evaluated at the τ^{th} quantile $\epsilon = S_0^{-1}(\tau)$, where in our parameterization $\tau = (1 + \alpha)/2$. Proposition

3 establishes the consistency of the maximum likelihood estimator $\widehat{\eta}_\gamma = (\widehat{\theta}_\gamma, \widehat{\vartheta}_\gamma, \widehat{\alpha}_\gamma)$ to the Kullback-Leibler optimal parameter values when assuming asymmetric Normal and asymmetric Laplace errors, whereas Proposition 4 gives asymptotic normality.

Proposition 3. *Assume Conditions A1–A4 with $p < n$, where $j = 2$ in A4 when $k = 1$ and $j = 1$ when $k = 2$. Then, the function $M_k(\theta_\gamma, \vartheta, \alpha) = \mathbb{E}[\log L_k(y_1 | x_1^T \theta_\gamma, \vartheta, \alpha)]$ has a unique maximizer $(\theta_\gamma^*, \vartheta_\gamma^*, \alpha_\gamma^*) = \operatorname{argmax}_\Gamma M_k(\theta_\gamma, \vartheta, \alpha)$. Moreover, the maximum likelihood estimator $(\widehat{\theta}_\gamma, \widehat{\vartheta}_\gamma, \widehat{\alpha}_\gamma) \xrightarrow{P} (\theta_\gamma^*, \vartheta_\gamma^*, \alpha_\gamma^*)$ as $n \rightarrow \infty$.*

Proposition 4. *Assume Conditions A1–A5, with $j = 4$ in A4 when $k = 1$ and $j = 2$ when $k = 2$. Denote $\eta = (\theta_\gamma, \vartheta, \alpha)$, $m_\eta(y_1, x_1) = \log s_k(y_1 | x_1^T \theta_\gamma, \vartheta, \alpha)$, $Pm_\eta = \mathbb{E}[m_\eta(y_1, x_1)]$, and $\eta_\gamma^* = (\theta_\gamma^*, \vartheta_\gamma^*, \alpha_\gamma^*) = \operatorname{argmax}_\Gamma Pm_\eta$. Then, the sequence $\sqrt{n}(\widehat{\eta}_\gamma - \eta_\gamma^*)$ is asymptotically normal with mean 0 and covariance matrix $V_{\eta_\gamma^*}^{-1} \mathbb{E}[\dot{m}_{\eta_\gamma^*} \dot{m}_{\eta_\gamma^*}^T] V_{\eta_\gamma^*}^{-1}$, where $\dot{m}_{\eta_\gamma^*}$ is the gradient of $m_\eta(\cdot)$, with respect to η , evaluated at η_γ^* and $V_{\eta_\gamma^*}$ is the second derivative matrix of Pm_η evaluated at η_γ^* .*

The sandwich covariance $V_{\eta_\gamma^*}^{-1} \mathbb{E}[\dot{m}_{\eta_\gamma^*} \dot{m}_{\eta_\gamma^*}^T] V_{\eta_\gamma^*}^{-1}$ is typically an inflated version of that obtained when the true model is assumed ($V_{\eta_\gamma^*}^{-1}$), implying the well-known consequence of model misspecification that parameter estimation suffers a loss of efficiency and uncertainty is underestimated. To gain insight Corollary 5 gives specific asymptotic variances under various model misspecification cases. For instance, when truly $\epsilon_i \sim N(0, \vartheta)$ wrongly assuming Laplace errors increases the variance by a factor $\pi/2$, and a similar phenomenon is observed when ignoring the presence of residual asymmetry. We defer discussion of the implications for variable selection to Section 5 and the examples in Section 6.

Corollary 5. *The asymptotic distribution of $\widehat{\theta}_\gamma$ obtained by maximizing either the Normal, ANormal, Laplace or ALaplace likelihood is $V(\widehat{\theta}_\gamma - \theta_\gamma^*) \xrightarrow{D} N(0, vI)$ for some $v > 0$. The asymptotic variances v , when ϵ_i truly arise i.i.d. under four specific distributions, are given below.*

	Maximized log-likelihood			
True model	Normal	ANormal	Laplace	ALaplace
$N(0, \vartheta)$	ϑ	ϑ	$\frac{\pi}{2}\vartheta$	$\frac{\pi}{2}\vartheta$
$AN(0, \vartheta, \alpha)$	$\vartheta(1 + 0.454\alpha^2)$	$\vartheta(1 - \alpha^2)$ (★)	$\frac{\pi}{2}\vartheta k_\alpha$	$\frac{\pi}{2}\vartheta(1 - \alpha_\gamma^{*2})$
$L(0, \vartheta)$	2ϑ	2ϑ	ϑ	ϑ
$AL(0, \vartheta, \alpha)$	$2\vartheta(1 + \alpha^2)$	$2\vartheta w_{\alpha, \alpha_\gamma^*}$ (★)	$\vartheta(1 + \alpha)^2$	$\vartheta(1 - \alpha^2)$

where $k_\alpha = \exp \left\{ \left[\Phi^{-1} \left(\frac{1}{2(1+|\alpha|)} \right) \right]^2 \right\} \geq 1$, $w_{\alpha, \alpha_\gamma^*} = \frac{(1 + \alpha)^2 - 2\alpha(1 + \alpha_\gamma^*)}{(1 - \alpha^2)^2} \in [0, 1]$, and α_γ^* is as in Proposition 4. Cases marked (★) were derived assuming that covariates have zero mean.

4.2. Optimization. We outline simple, efficient algorithms to obtain $(\widehat{\theta}_\gamma, \widehat{\vartheta}_\gamma, \widehat{\alpha}_\gamma) = \operatorname{argmax}_{\theta_\gamma, \vartheta, \alpha} \log L_k(\theta_\gamma, \vartheta, \alpha)$, where $k \in \{1, 2\}$ are the asymmetric Normal and Laplace log-likelihoods (3)–(4). We also consider the corresponding posterior modes $(\tilde{\theta}_\gamma, \tilde{\vartheta}_\gamma, \tilde{\alpha}_\gamma) =$

$\arg \max_{\theta_\gamma, \vartheta, \alpha} \log L_k(\theta_\gamma, \vartheta, \alpha) + \log p(\theta_\gamma, \vartheta, \alpha | \gamma)$, where $p(\theta_\gamma, \vartheta, \alpha | \gamma)$ is the prior density (Section 3). The algorithms are useful to obtain parameter estimates or Laplace approximations to the integrated likelihood. Mudholkar and Hutson (2000) and Arellano-Valle et al. (2005) gave an algorithm to obtain $\hat{\theta}_\gamma$ for $\log L_1$ in the case with no covariates ($p_\gamma = 1$). To tackle point discontinuities in the derivatives their algorithm requires solving n separate optimization problems, which does not scale up with increasing n , or alternatively using method of moments estimators. Maximum likelihood estimation of θ_γ under the asymmetric Laplace and fixed α is connected to quantile regression (see below). Regarding Bayesian frameworks, most rely on MCMC for parameter estimation but this is too computationally costly when we wish to consider a potentially large number of models. Instead, we propose a generic framework for jointly obtaining $(\hat{\theta}_\gamma, \hat{\vartheta}_\gamma, \hat{\alpha}_\gamma)$ or $(\tilde{\theta}_\gamma, \tilde{\vartheta}_\gamma, \tilde{\alpha}_\gamma)$ applicable to both the asymmetric Normal and Laplace. The key result we exploit is concavity of the log-likelihood given by Propositions 1-2, which allows iteratively optimizing first θ_γ and then (ϑ, α) . Optimization with respect to (ϑ, α) has closed form, whereas setting θ_γ can be seen as weighted least squares for the asymmetric Normal and as quantile regression for the asymmetric Laplace. The latter task of maximizing $\log L_2$ with respect to θ_γ is a classical problem that can be framed as linear programming, for which simplex and interior-point methods are available. However these are not applicable to the posterior mode as the target is no longer piecewise linear and even efficient implementations have computational complexity greater than cubic in p and supra-linear in n Koenker (2005).

We outline two simple algorithms that have lower complexity and can be readily adapted to obtain the posterior mode. Briefly, in Algorithm 4.2, Step 2 follows from setting first derivatives to zero and directly extends Mudholkar and Hutson (2000) (Proposition 4.4) and Arellano-Valle et al. (2005) (Section 4.2). Step 3 is essentially a Levenberg-Marquardt algorithm Levenberg (1944); Marquardt (1963) exploiting gradient continuity. g_θ and H_θ denote the gradient and hessian with respect to θ_γ as in Propositions 1-2, where for $\log L_2()$ we use the asymptotic hessian $X^T X / (\vartheta(1 - \alpha^2))$. Its updates are in between those of a Newton-Raphson and gradient descent algorithms and can be interpreted as restricting the Newton-Raphson step to a trust region where the quadratic approximation is accurate Sorensen (1982). For large regularization parameter λ the update δ converges to the gradient algorithm, which by continuity is guaranteed to increase the target, whereas for small λ it converges to the Newton-Raphson algorithm, achieving quadratic convergence as $\theta_\gamma^{(\ell)}$ approaches the optimum.

Given a good initial guess $\hat{\theta}_\gamma^{(0)}$, the fact that $\log L_k$ are locally well approximated by a quadratic function in θ_γ ($\log L_1$ is exactly locally quadratic) results in Algorithm 4.2 usually converging after a few iterations. As usual with second-order optimization each iteration requires a matrix inversion that is costly when p is large. As an alternative, Algorithm 4.2 uses coordinate descent to optimize each θ_{γ_j} sequentially, which only requires univariate updates, where updating the set $A(\theta_\gamma)$ for each θ_{γ_j} implies that Step 3 has cost $O(np)$. In contrast Algorithm 4.2 determines $A(\theta_\gamma)$ once per iteration and performs matrix inversion, with total cost $O(n + p^3)$ per iteration. Hence although Algorithm 4.2

Optimization via Levenberg-Marquardt

- (1) Initialize $\hat{\theta}_\gamma^{(0)} = (X^T X)^{-1} X^T y$, $\lambda = 0$. Set $t = 1$
- (2) Let $s_1 = \sum_{i \in A(\hat{\theta}_\gamma^{(t-1)})} |y_i - x_{\gamma i}^T \hat{\theta}_\gamma^{(t-1)}|^{3-k}$, $s_2 = \sum_{i \notin A(\hat{\theta}_\gamma^{(t-1)})} |y_i - x_{\gamma i}^T \hat{\theta}_\gamma^{(t-1)}|^{3-k}$. Update

$$\hat{\alpha}^{(t)} = \frac{s_1^{\frac{k}{2+k}} - s_2^{\frac{k}{2+k}}}{s_1^{\frac{k}{2+k}} + s_2^{\frac{k}{2+k}}}; \hat{\vartheta}^{(t)} = \frac{1}{4n^k} \left(s_1^{\frac{k}{2+k}} + s_2^{\frac{k}{2+k}} \right)^{2+k}$$

- (3) Propose $m = \theta_\gamma^{(t-1)} + \delta$, where

$$\delta = -(H_\theta + \lambda \text{diag}(H_\theta))^{-1} g_\theta$$

and g_θ, H_θ are the subsets of $g_k(\hat{\theta}_\gamma^{(t-1)}, \hat{\vartheta}^{(t)}, \hat{\alpha}^{(t)})$ and $H_k(\hat{\theta}_\gamma^{(t-1)}, \hat{\vartheta}^{(t)}, \hat{\alpha}^{(t)})$ corresponding to θ_γ . If $\log L_k(m, \vartheta^{(t)}, \alpha^{(t)}) > \log L_k(\theta_\gamma^{(t-1)}, \vartheta^{(t)}, \alpha^{(t)})$ set $\theta_\gamma^{(t)} = m$ and $\lambda = \lambda/2$, else update $\lambda = 1 + \lambda$ and repeat Step 3.

usually requires fewer iterations than Algorithm 4.2, for large p the latter typically preferable. A related study of computational cost is offered in Breheny and Huang (2011) in the context of penalized likelihood optimization, who found that coordinate descent is often preferable to multivariate updates. These results show that, contrary to historical beliefs, two-piece distributions lead to convenient optimization. R package `mombf` Rossell et al. (2016) incorporates both algorithms but our examples are based on Algorithm 4.2, the results were essentially identical to those of Algorithm 4.2 but the running time was substantially shorter.

We adapted both algorithms to find the posterior mode by simply redefining g_k and H_k to be the gradient and Hessian of $\log L_k(\theta_\gamma, \vartheta, \alpha) + \log p(\theta_\gamma, \vartheta, \alpha | \gamma)$. The corresponding expressions for $\log p(\theta_\gamma, \vartheta, \alpha | \gamma)$ are in Supplementary Section 10.2. We remark that due to the penalty around the origin NLPs such as $p_M()$ and $p_E()$ in (6)-(7) are not log-concave, however this is not an issue as they are symmetric and log-concave in each quadrant (fixed $\text{sign}(\theta_\gamma, \alpha)$). Thus $\log p(\theta_\gamma, \vartheta, \alpha | y, \gamma)$ is concave in each quadrant, its unique global mode lies in the same quadrant as the maximum likelihood estimator and we may initialize the algorithm at $(\tilde{\theta}_\gamma^{(0)}, \tilde{\vartheta}^{(0)}, \tilde{\alpha}^{(0)}) = (\hat{\theta}_\gamma, \hat{\vartheta}_\gamma, \hat{\alpha}_\gamma)$. Convergence is typically achieved after a few iterations.

Optimization via coordinate descent

- (1) Set an arbitrary $c > 1$ and initialize $\theta_\gamma^{(0)}$, $\lambda = 0$ as in Algorithm 4.2.
 - (2) Update $(\hat{\vartheta}^{(t)}, \hat{\alpha}^{(t)})$ as in Algorithm 4.2.
 - (3) For $j = 1, \dots, p_\gamma$, let $m = \theta_{\gamma j}^{(t-1)} - \frac{g_j}{h_{jj}(1+\lambda)}$, where g_j is the j^{th} element in $g_1(\theta_\gamma)$ and h_{jj} the (j, j) element in $H_1(\theta_\gamma)$ at $\theta_\gamma = (\theta_{\gamma 1}^{(t)}, \dots, \theta_{\gamma j-1}^{(t)}, \theta_{\gamma j}^{(t-1)}, \dots, \theta_{\gamma p_\gamma}^{(t-1)})$. If L_k evaluated at $\theta_{\gamma j}^{(t)} = m$ increases set $\theta_{\gamma j}^{(t)} = m$, $\lambda = \lambda/c$, else iteratively update $\lambda = c + \lambda$ and m until L_k increases.
-

5. MODEL SELECTION

Under a standard Bayesian framework $p(\gamma | y) = p(y | \gamma)p(\gamma)/p(y)$, with integrated likelihood

$$\begin{aligned}
 p(y | \gamma) &= \int L_1(\theta_\gamma, \vartheta, 0)p(\theta_\gamma, \vartheta)d\theta_\gamma d\vartheta, \text{ if } \gamma_{p_\gamma+1} = 0, \gamma_{p_\gamma+2} = 0, \\
 p(y | \gamma) &= \int L_1(\theta_\gamma, \vartheta, \alpha)p(\theta_\gamma, \vartheta, \alpha)d\theta_\gamma d\vartheta d\alpha, \text{ if } \gamma_{p_\gamma+1} = 1, \gamma_{p_\gamma+2} = 0, \\
 p(y | \gamma) &= \int L_2(\theta_\gamma, \vartheta, 0)p(\theta_\gamma, \vartheta)d\theta_\gamma d\vartheta, \text{ if } \gamma_{p_\gamma+1} = 0, \gamma_{p_\gamma+2} = 1, \\
 (8) \quad p(y | \gamma) &= \int L_2(\theta_\gamma, \vartheta, \alpha)p(\theta_\gamma, \vartheta, \alpha)d\theta_\gamma d\vartheta d\alpha, \text{ if } \gamma_{p_\gamma+1} = 1, \gamma_{p_\gamma+2} = 1.
 \end{aligned}$$

Section 5.1 discusses how to compute $p(y | \gamma)$ and Section 5.2 the asymptotic properties of the associated Bayes factors and a stochastic model search algorithm that can be used when p is too large to allow exhaustive enumeration of the 2^{p+2} models.

5.1. Integrated likelihood. Computing (8) in the case $\gamma_{p+1} = \gamma_{p+2} = 0$ corresponds to Normal linear regression, for existing methods are typically available, *e.g.* Johnson and Rossell (2012) gave closed-form expressions for the MOM and Laplace approximations for the eMOM and iMOM. The three remaining cases require numerical evaluation, for which we propose Laplace and Monte Carlo approximations. The former are appealing due to log-likelihood concavity and asymptotic normality (Section 4), indeed in our examples they delivered very similar inference and were orders of magnitude faster than Monte Carlo. Hence by default we recommend Laplace approximations over Monte Carlo, except in small p situations where the latter is still practical. To ensure that the parameter support is on the real numbers Laplace approximations are based on the reparameterization $\eta = (\theta_\gamma, \log(\vartheta), \text{atanh}(\alpha))$ and given by

$$(9) \quad \hat{p}(y | \gamma) = \exp\{\log L_k(\tilde{\eta}) + \log p(\tilde{\eta})\} \frac{(2\pi)^{\sum_{j=1}^{p+2} \gamma_j/2}}{|H_k(\tilde{\eta})|^{1/2}},$$

where $k = 1, 2$ for $\gamma_{p+2} = 0, 1$ respectively, $\tilde{\eta}$ and \tilde{H}_k are the posterior mode and hessian of $\log L_k(\eta) + \log p(\eta)$. The specific expressions are given in Supplementary Section 10. Expression (9) simply requires the posterior mode (Algorithms 4.2-4.2) and evaluating the hessian. The latter is straightforward for $k = 1$, but for $k = 2$ it is singular in θ_γ , requiring some care. The reasoning behind (9) is to approximate the log-integrand in (8) by a smooth function that has strictly positive definite hessian in θ_γ , which is facilitated in our setting by $\log L_2$ concavity and asymptotic normality. We found that a simple yet effective strategy is to replace H_2 by the asymptotic expected hessian \bar{H}_2 obtained under independent asymmetric Laplace errors.

Although we did not find the following concern to be a practical issue in our examples, we remark that in principle \bar{H}_2 may underestimate the underlying uncertainty in θ_γ and thus inflate $|\bar{H}_2|$, *e.g.* under truly non-Laplacian independent and identically distributed errors one needs to add a multiplicative constant (Section 4.1) whereas independent

but heteroskedastic errors require a matrix-reweighting adjustment Kocherginsky et al. (2005). Typical strategies to improve the estimated curvature rely either on direct estimation under the assumption of independent and identically distributed errors, indirect estimation via inversion of score tests, although these only provide univariate confidence intervals and their cost does not scale well with p , or sampling-based methods such as bootstrap or Monte Carlo. As a practical alternative here we consider that the goal is really to approximate the actual curvature of $\log L_2$, which can be easily done with a few point evaluations of $\log L_2$ in a neighbourhood of $\tilde{\eta}_\gamma$. Briefly, we consider the adjustment $D\overline{H}_2D$ where D is a diagonal matrix such that its element d_{ii} gives the best approximation of $\log L_2$ as a quadratic function of θ_i in the least squares sense. $D\overline{H}_2D$ matches the actual curvature in $\log L_2$ and is thus less dependent on asymptotic theory than other strategies, and has the advantage that D can be computed quickly. See Supplementary Section 10.3 for further details and Supplementary Figure 5 for an example. Given that the unadjusted \overline{H}_2 performed well in our examples and the associated results were practically indistinguishable to those based on Monte Carlo, unless otherwise stated our results are based on \overline{H}_2 .

As our Monte Carlo alternative we implemented an importance sampling estimator based on multivariate T draws with heavy tails and proposal covariance matching the asymptotic posterior covariance. Specifically, let $\eta^{(b)} \sim T_3(\tilde{\eta}, \tilde{H}_k^{-1}/3)$ for $b = 1, \dots, B$ where B is a large integer, then

$$(10) \quad \hat{p}_I(y | \gamma) = B^{-1} \sum_{b=1}^B L_k(\eta^{(b)}) p(\eta^{(b)}) / T_3(\eta^{(b)}; \tilde{\eta}, \tilde{H}_k^{-1}/3).$$

We remark that NLPs are multimodal in (θ_γ, α) , thus some care is needed when using Laplace approximations. To give an honest characterization of the properties of our preferred computational method, in Section 5 we obtain asymptotic rates for Bayes factors based on $\hat{p}(y | \gamma)$ in (9). Rossell and Telesca (2015) studied the discrepancies between $p(y | \gamma)$ and $\hat{p}(y | \gamma)$ for MOM, iMOM and eMOM priors and Normal errors. Briefly, given that secondary modes vanish asymptotically for truly active covariates but not for spurious covariates, $\hat{p}(y | \gamma)$ imposes a stronger penalty on spurious variables relative to $p(y | \gamma)$, however for such models $p(y | \gamma)$ decreases fast enough that both approximations typically lead to very similar inference.

5.2. Bayes factor rates. Let $\gamma^* = (\text{I}(\theta_1^* \neq 0), \dots, \text{I}(\theta_p^* \neq 0), \text{I}(\alpha^* \neq 0), \text{I}(k^* = 2))$ be the optimal model, that is $(\theta^*, \vartheta^*, \alpha^*, k^*) = \arg \max_{\Gamma, k} M_k(\theta, \vartheta, \alpha)$ maximize the expected log-likelihood across $k = 1, 2$, and the expectation is with respect to the data-generating density in Condition A1. We indicate by $\gamma^* \subset \gamma$ that γ^* is a submodel of γ , *i.e.* $\gamma_j^* \leq \gamma_j$ for $j = 1, \dots, p+1$, and by $\gamma^* \not\subset \gamma$ that $\gamma_j^* > \gamma_j$ for some j . If the data were truly generated from the assumed error distribution it is well-known that the Bayes factor in favour of γ decreases exponentially with n when $\gamma^* \not\subset \gamma$ (γ is missing important variables). Conversely when γ adds spurious variables to γ^* the Bayes factor is only $O_p(n^{-(p_\gamma - p_{\gamma^*})/2})$ under local priors, an imbalance that is ameliorated under NLPs, which achieve faster polynomial or quasi-exponential rates depending on their chosen parametric

form Johnson and Rossell (2010, 2012). Proposition 6 gives an extension under model misspecification, the first result of this kind for NLPs. The rates apply directly to the Laplace approximations in (9), which has the advantage of characterizing our preferred computational framework and reducing the amount of technical details in proving the result. A critical condition for the result to hold is that the prior density be strictly positive at the optimum, $p(\theta_{\gamma^*}, \alpha_{\gamma^*}^* | \gamma^*) > 0$, which is trivially satisfied by pMOM and peMOM priors. It also holds for local priors, which for simplicity we define as $p(\theta_\gamma, \vartheta, \alpha | \gamma) > 0$ for all $(\theta_\gamma, \vartheta, \alpha) \in \Gamma_\gamma$ and we assume to be continuous.

Proposition 6. *Suppose that Conditions A1-A3 hold, fixed p_γ, p_{γ^*} and $n \rightarrow \infty$. If $\gamma^* \not\subset \gamma$ then $\log(\hat{p}(y | \gamma)/\hat{p}(y | \gamma^*)) = O_p(n)$ for local, pMOM and peMOM priors. Conversely, if $\gamma^* \subset \gamma$ the following holds.*

- (i) *If $p(\theta_\gamma, \alpha | \gamma)$ is a local prior then $\hat{p}(y | \gamma)/\hat{p}(y | \gamma^*) = O_p(n^{-(p_\gamma - p_{\gamma^*})/2})$.*
- (ii) *If $p(\theta_\gamma, \alpha | \gamma)$ is the pMOM prior then $\hat{p}(y | \gamma)/\hat{p}(y | \gamma^*) = O_p(n^{-3(p_\gamma - p_{\gamma^*})/2})$.*
- (iii) *If $p(\theta_\gamma, \alpha | \gamma)$ is the peMOM prior then $\hat{p}(y | \gamma)/\hat{p}(y | \gamma^*) = O_p(e^{-\sqrt{n}})$.*

Proposition 6 implies strong model selection consistency with Bayes factor rates that are the same as when the correct model is assumed. We emphasize that this does not imply that there is no cost due to assuming an incorrect model: the coefficients governing the exponential or polynomial rates are affected. For instance, as shown in our examples model misspecification often causes a decrease in the power to detect truly active variables. Fully characterizing this issue theoretically is complicated as the behaviour of Bayes factors depends on the unknown data-generating truth, however it is possible to provide some intuition as to when this loss of power occurs.

Consider an arbitrary variable configuration $(\gamma_1, \dots, \gamma_p) \not\subset (\gamma_1^*, \dots, \gamma_p^*)$ that is missing some truly active variables. Suppose that as in Condition A1 truly $\epsilon_i \sim s_0(\epsilon_i) = s(\epsilon_i | \xi_0)$ for some error density family $s(\epsilon_i | \xi)$, $\xi \in \Xi$ and fixed $\xi_0 \in \Xi$. Denote by $L_0(\theta_\gamma, \xi)$ the likelihood under the correct $\epsilon_i \sim s(\epsilon_i | \xi)$ and $p_0(y | \gamma) = \int L_0(y | \theta_\gamma, \xi) p(\theta_\gamma, \xi) d\theta_\gamma d\xi$ the associated integrated likelihood under some prior $p(\theta_\gamma, \xi) > 0$. The interest is in comparing the correct Bayes factor $p_0(y | \gamma^*)/p_0(y | \gamma)$ to the misspecified $\hat{p}(y | \gamma^*)/\hat{p}(y | \gamma)$. Under fairly general conditions

$$\log(p_0(y | \gamma^*)/p_0(y | \gamma)) \approx nD_0(p_0(y | \theta_\gamma^*, \xi_\gamma^*, \gamma))$$

plus lower order terms, where $D_0(p_0(y | \theta_\gamma^*, \xi_\gamma^*, \gamma))$ is the Kullback-Leibler divergence between the data-generating $p_0(y | \theta_\gamma^*, \xi_0, \gamma^*)$ and the KL-optimal $p_0(y | \theta_\gamma^*, \xi_\gamma^*, \gamma)$ under γ . In the proof of Proposition 6 we saw a similar result for $\log(\hat{p}(y | \gamma^*)/\hat{p}(y | \gamma))$ when $\gamma \not\subset \gamma^*$ both under local and non-local priors, which after trivial algebra gives

(11)

$$\log\left(\frac{p_0(y | \gamma^*)/p_0(y | \gamma)}{\hat{p}(y | \gamma^*)/\hat{p}(y | \gamma)}\right) \approx n(D_0(p_0(y | \theta_\gamma^*, \xi_\gamma^*, \gamma)) + D_0(p(y | \eta_{\gamma^*}^*, \gamma^*)) - D_0(p(y | \eta_\gamma^*, \gamma))).$$

The sign of the right hand side in (11) determines whether the misspecified Bayes factor has lower or greater asymptotic power than the correct Bayes factor. A precise study of (11) deserves separate treatment elsewhere, but the expression can be loosely interpreted as a type of triangle inequality. If the divergence due to simultaneously using the wrong

error distribution and γ instead of γ^* , $D_0(p(y | \eta_{\gamma}^*, \gamma))$, is smaller than the sum of the divergences due to only using the wrong error distribution plus that of only using γ instead of γ^* . Then, misspecifying the error distribution results in slower (but still exponential) Bayes factor rates to detect truly active variables. There is no guarantee that (11) is positive in general for any combination of assumed model and data-generating truth, however all our examples exhibited such a loss of power, suggesting that this is often the case.

5.3. Model exploration. Algorithm 5.3 describes a novel Gibbs sampling that can be used when p_{γ} is too large for exhaustive enumeration of all $2^{p_{\gamma}+2}$ models. Although conceptually simple, Algorithm 5.3 extends a method that delivered good results for high-dimensional variable selection under Normal errors Johnson and Rossell (2012), and is designed to spend most iterations in the Normal model whenever it is a good enough approximation. That is, as illustrated in our examples the computational effort adapts automatically to the nature of the data, so that the cost associated to abandoning the Normal model is only incurred when this is required to improve inference. Our implementation also allows the user to fix $(\gamma_{p+1}, \gamma_{p+2})$, so that one can condition on Normal, asymmetric Normal, Laplace or asymmetric Laplace errors whenever this is desired.

The number of iterations T should ideally be large enough for the chain to converge, see for instance Johnson (2013) for a discussion of formal convergence diagnostics based on coupling methods. In practice it usually suffices to monitor some posterior quantities of interest, for instance in the setting of variable selection with NLPs Rossell and Telesca (2015) found useful to set T large enough so that sampling-based estimates of $p(\gamma_j = 1 | y)$ are close enough to estimates based on renormalizing posterior probabilities across the models visited so far.

Gibbs model space search.

(1) Let $\gamma_{p+1}^{(0)} = \gamma_{p+2}^{(0)} = 0$ and set $\gamma_1^{(0)}, \dots, \gamma_p^{(0)}$ using the greedy forward-backward initialization algorithm in Johnson and Rossell (2012). Set $t = 1$.

(2) For $j = 1, \dots, p$ update $\gamma_j^{(t)} = 1$ with probability

$$\frac{p(\gamma_1^{(t)}, \dots, \gamma_{j-1}^{(t)}, 1, \gamma_{j+1}^{(t-1)}, \dots, \gamma_p^{(t-1)} | y)}{\sum_{\gamma_j=0}^1 p(\gamma_1^{(t)}, \dots, \gamma_{j-1}^{(t)}, \gamma_j, \gamma_{j+1}^{(t-1)}, \dots, \gamma_p^{(t-1)} | y)}$$

(3) Update $(\gamma_{p+1}^{(t)}, \gamma_{p+2}^{(t)}) = (l, m)$ with probability

$$\frac{p(\gamma_1^{(t)}, \dots, \gamma_p^{(t)}, l, m | y)}{\sum_{\gamma_{p+1}=0}^1 \sum_{\gamma_{p+2}=0}^1 p(\gamma_1^{(t)}, \dots, \gamma_p^{(t)}, \gamma_{p+1}, \gamma_{p+2} | y)}.$$

If $t \leq T$ set $t = t + 1$ and go back to Step 2, otherwise stop.

6. RESULTS

We studied via simulations the practical implications of model misspecification on variable selection, both on small and large p (Sections 6.1-6.2), as well as the ability of our framework to detect asymmetries ($\gamma_{p+1} = 1$) and heavier-than-normal tails ($\gamma_{p+2} = 1$). Computations were carried out using function `modelSelection` in R package `mombf` 1.8.0 Rossell et al. (2016), using default prior settings (Section 3) and Laplace approximations to $p(y | \gamma)$ unless otherwise stated. Although our goal is to build a Bayesian framework to cope with simple departures from normality, for comparison we included penalized likelihood methods with available R implementation: standard LASSO penalties on least squares regression (LASSO-LS, Tibshirani (1996)), LASSO penalties on least absolute deviation (LASSO-LAD, Wang and Li (2009)) and SCAD penalties on least squares Fan and Li (2001). For LASSO-LS, LASSO-LAD and SCAD we set the penalization parameter with 10-fold cross-validation using functions `mylars`, `slim` and `ncvreg` in R packages `parcor` 0.2.6, `flare` 1.5.0 and `ncvreg` 3.4.0 (respectively) with default parameters. For LASSO-LAD we increased the default grid of penalization parameter values in `slim` from 5 to 10 points. All R code is provided as supplementary material.

6.1. Low-dimensional simulation. We started by simulating 200 data sets from a linear model with normal residuals, each with $n = 100$, $p = 6$, $\theta = (0, 0.5, 1, 1.5, 0, 0)$ ($\theta_1 = 0$ corresponds to the intercept), $\vartheta = 2$. Covariate values were generated from a multivariate Normal centered at 0, with unit variances and all pairwise correlations $\rho_{ij} = 0.5$. We compared the results under assumed Normal, asymmetric Normal, Laplace and asymmetric Laplace errors, and also when inferring the residual distribution with our framework (Section 5). Throughout we used MOM priors with default $g_\theta = 0.348$, $g_\alpha = 0.357$ and uniform model probabilities $p(\gamma) \propto 1$. Given that p is small we enumerated and computed $p(\gamma | y)$ for all models. Figure 2 (top left) shows the marginal posterior probabilities $p(\gamma_j = 1 | y)$. These were almost identical under assumed Normal and asymmetric Normal errors. Both models were preferable to Laplace or asymmetric Laplace errors, mainly in giving higher $p(\gamma_j = 1 | y)$ for truly active variables.

We repeated the simulation study, this time generating $\epsilon_i \sim \text{AN}(0, 2, -0.5)$, $\epsilon_i \sim L(0, 2)$ and finally $\epsilon_i \sim \text{AL}(0, 2, -0.5)$. Here we observed more marked differences than under $\epsilon_i \sim N(0, 2)$, specifically failing to account for thick tails caused a substantial drop in $p(\gamma_j = 1 | y)$ for truly active predictors. As an example, when truly $\epsilon_i \sim \text{AL}(0, 4, -0.5)$ the mean $p(\theta_3 \neq 0 | y)$ increased from 0.63 under assumed normal errors to 0.89 under asymmetric Laplace errors. These results suggest that wrongly assuming normal errors may have more pronounced consequences on inference than using more robust error distributions. Interestingly, including asymmetry in the model had no noticeable adverse effects on inference even when residuals were truly symmetric, and improved power when residuals were truly asymmetric. Hence the reasoning for adopting symmetric models seems mostly computational.

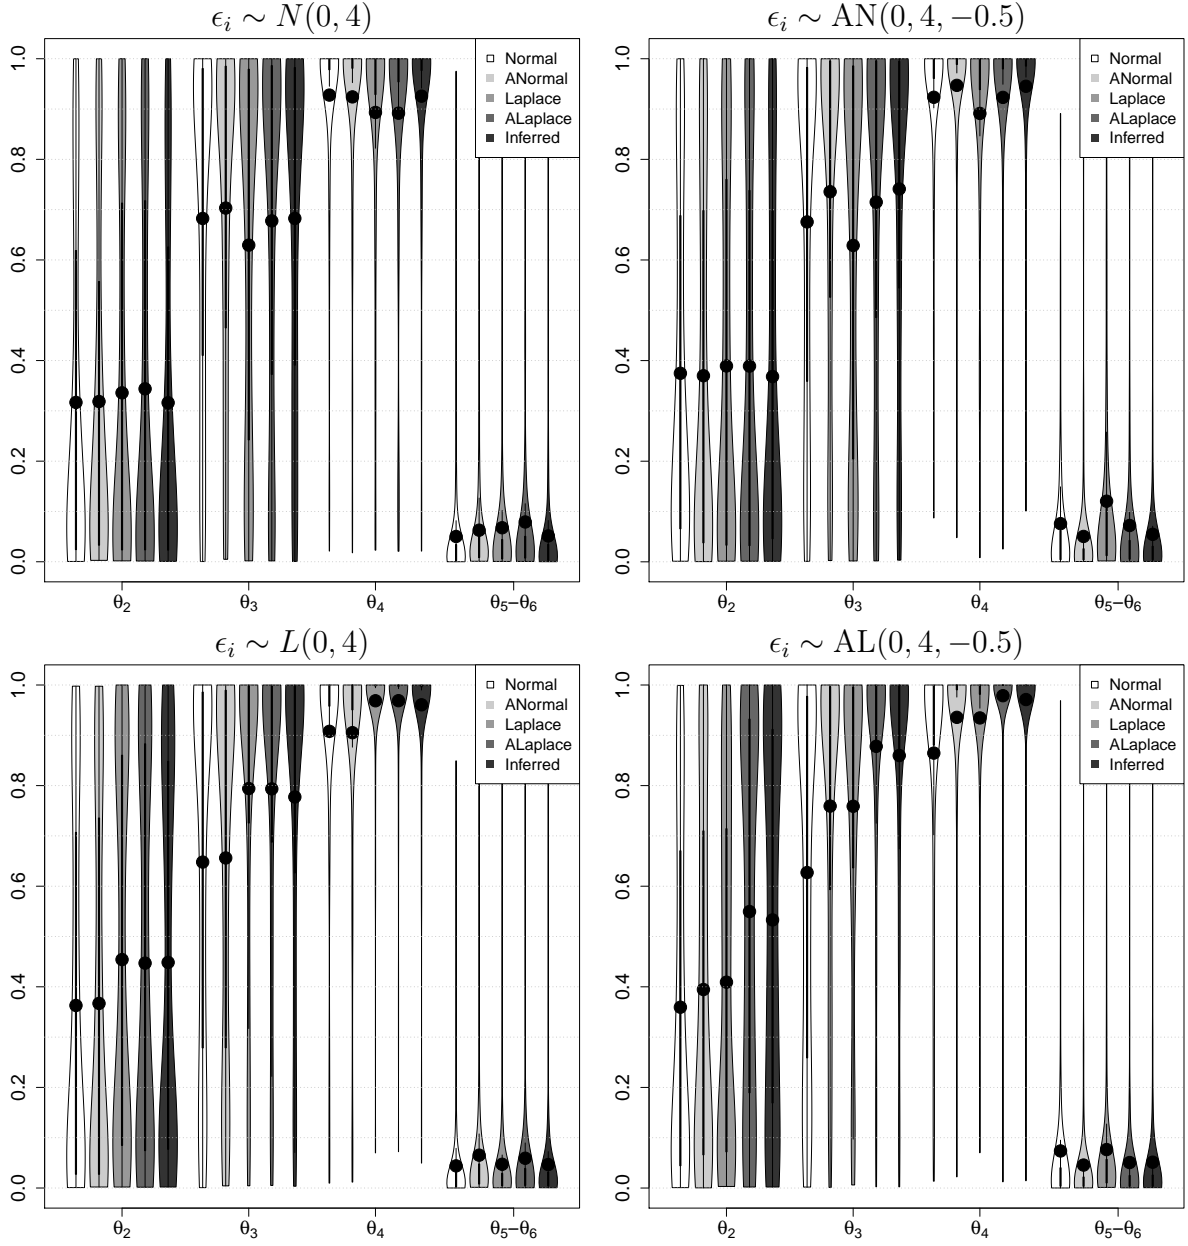


FIGURE 2. $P(\theta_i \neq 0 \mid y)$ for $p = 6$, $\theta = (0, 0.5, 1, 1.5, 0, 0)$, $n = 100$, $\rho_{ij} = 0.5$. Black circles show the mean.

Our framework based on inferring $(\gamma_{p+1}, \gamma_{p+2})$ showed a highly competitive behaviour, usually fairly close to assuming the true distribution (Figure 2). The posterior probability assigned to the true error distribution was always > 0.8 (Supplementary Table 3), indicating that the desired departures from normality were effectively detected.

	Simulation truth			
	$N(0, 4)$	$AN(0, 4, -0.5)$	$L(0, 4)$	$AL(0, 4, -0.5)$
Normal	1s	1s	1s	1s
ANormal	34s	25s	27s	25s
Laplace	45s	66s	45s	80s
ALaplace	71s	87s	68s	61s
Inferred	4s	29s	33s	50s

TABLE 1. CPU time on 2.6GHz Intel Xeon processor, 32Gb RAM, Linux 3.6.16. $p = 100$, $\vartheta = 2$, $\theta = (0, 0.5, 0.75, 1, 0, \dots, 0)$, $n = 100$, $\rho_{ij} = 0.5$.

We repeated all the analyses above first using Monte Carlo estimates of $p(y | \gamma)$ based on $B = 10,000$ importance samples, and then again using our alternative default $g_\alpha = 0.087$. Supplementary Table 3 shows that inference on the error distribution remained remarkably stable, albeit as expected reducing $g_\alpha = 0.357$ to 0.087 increases slightly $p(\alpha \neq 0 | y)$ in all settings. Supplementary Figures 6-7 show $p(\gamma_j = 1 | y)$. These are virtually indistinguishable from those in Figure 2, indicating that the results are robust to these implementation details.

6.2. High dimensional simulation. We repeated the simulation study in Section 6.1 by adding 95 spurious predictors with $\theta = (0, 0.5, 1, 1.5, 0, \dots, 0)$ for a total of $p = 100$ covariates, and subsequently 400 more spurious predictors for a total $p = 500$. Given that the model space is too large for a full enumeration, we run the Gibbs algorithm in Section 5.3 with $T = 10,000$ iterations. To initialize the chain we used the greedy Gibbs algorithm from Johnson and Rossell (2012), which starts at the null model and keeps adding or removing individual covariates until a local mode is found, under the assumption of Normal errors. We set $p(\gamma)$ to the default Beta-Binomial(1,1) and left all other settings as in Section 6.1.

We conducted one first set of simulations under $\vartheta = 1$. Figure 3 shows the proportion of simulations in which the posterior mode $\hat{\gamma} = \arg \max_\gamma p(\gamma | y)$ was equal to the simulation truth $\gamma_0 = (0, 1, 1, 1, 0, \dots, 0)$. The main finding was that assuming the wrong error distribution had a marked detrimental impact on Bayesian variable selection, particularly in the presence of asymmetries or thicker-than-normal tails. Supplementary Table 5 gives the exact figures, as well as the number of false and true positives. All Bayesian formulations compared favourably to LASSO-LS, LASSO-LAD and LASSO-SCAD, mainly due to the latter incurring an inflated number of false positives. This is in agreement with earlier findings Johnson and Rossell (2012); Rossell and Telesca (2015) when comparing NLPs to penalized likelihoods, and likely partially related to the fact that cross-validation focuses on predictive ability and thus tends to favour the inclusion of a few spurious covariates. Interestingly, in our study LASSO-LAD showed little advantages over LASSO-LS, even under truly Laplace errors. Analogously to the $p = 6$ case in Figure 2, when $p = 101, 501$ the marginal inclusion probabilities for truly active variables suffered a drop when ignoring the presence of asymmetry or heavy tails (Supplementary Figures 8-9).

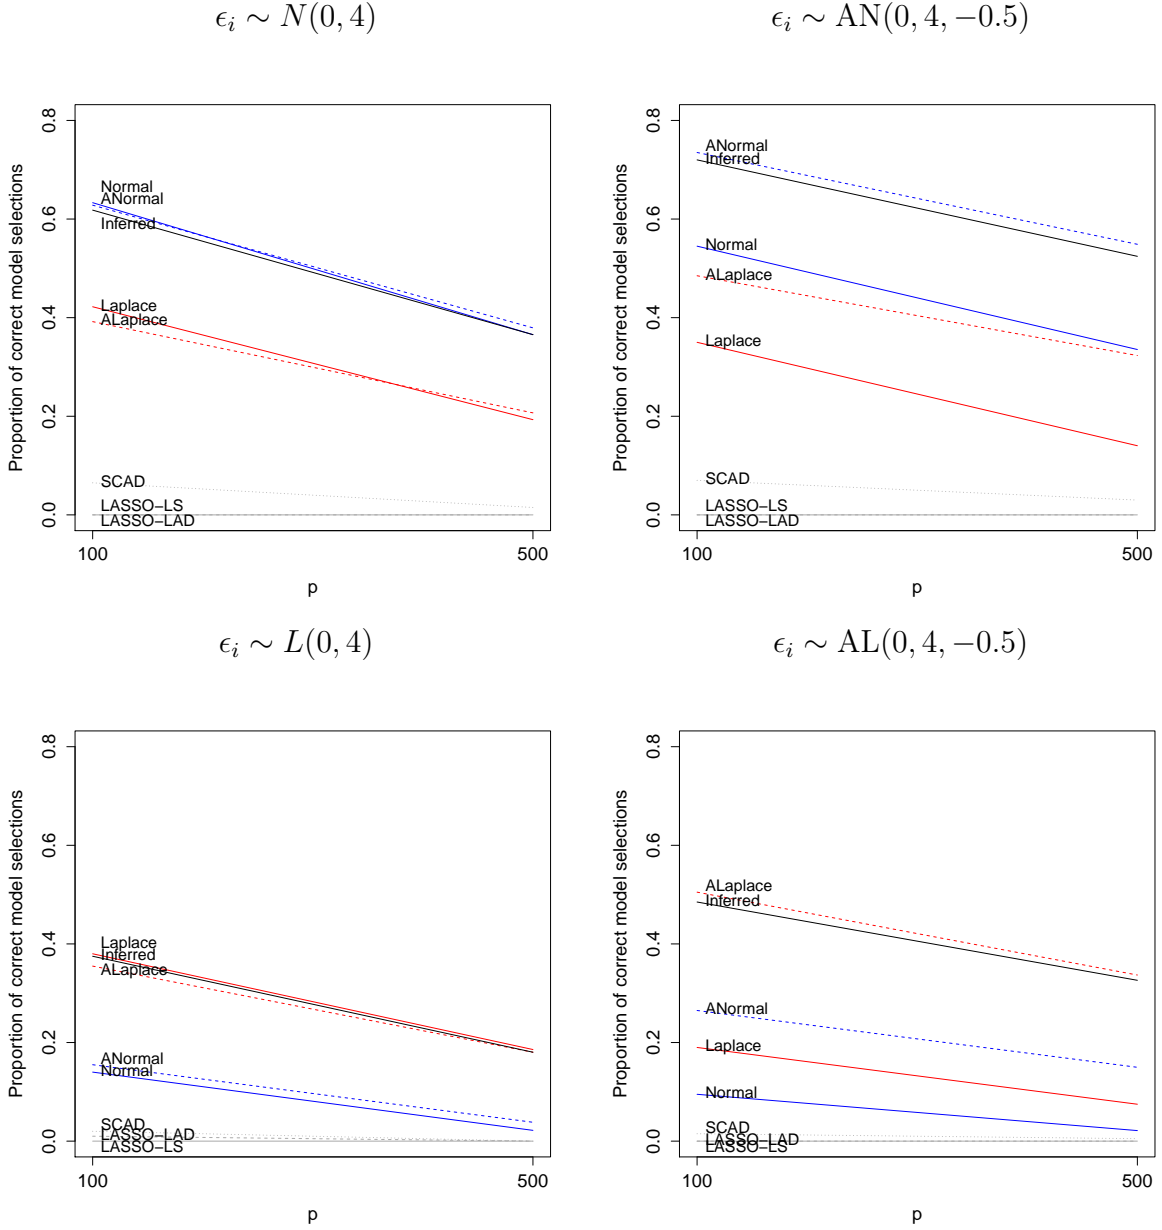


FIGURE 3. Proportion of correct model selections $p(\hat{\gamma} = \gamma_0 \mid y)$. $p = 500$, $\vartheta = 1$, $\theta = (0, 0.5, 1, 1.5, 0, \dots, 0)$, $n = 100$, $\rho_{ij} = 0.5$.

Our framework to infer the error distribution delivered highly competitive inference. Given the practical importance of computational speed, Table 1 indicates CPU times for $p = 100$. As expected, the availability of closed-form expressions for the Normal model resulted in faster computation, although as we just saw this comes at the cost of




FIGURE 4. QQ normal plot for TGF B (left) DLD (right) data.

a remarkable drop in variable selection when errors are non-normal. It is also interesting that computations were faster for the asymmetric Normal than for the Laplace or asymmetric Laplace, likely due to Taylor expansions of $\log L_1$ used by our optimization algorithms being more precise than for the non-differentiable $\log L_2$. Our framework to infer the error distribution had similar runtimes to the Normal model when this was an adequate representation for the data, and adaptively increased the computational effort under departures from normality to be able to deliver more precise inference.

To emulate a situation with lower signal-to-noise ratio we repeated the simulation study under $\vartheta = 2$. The results are shown in Supplementary Table 6 and Supplementary Figures 10-11. Briefly, the performance of all methods suffered in this more challenging setting due to a drop in the power to detect truly active predictors, however the relative performance across methods was largely analogous to those for $\vartheta = 1$.

6.3. TGF B data. We illustrate our methodology with the human microarray gene expression data in colon cancer patients from Calon et al. (2012). Briefly, following upon Rossell and Telesca (2015) we aim to detect which amongst $p = 10,172$ candidate genes have an effect on the expression levels of TGF B, a gene known to play an important role in colon cancer progression. These data contain moderately correlated covariates with absolute Pearson correlations ranging in $(0, 0.956)$ and 99% of them being in the interval $(0, 0.375)$. Both response and predictors were standardized to zero mean and unit variance. The dataset along with further information are provided in Rossell and Telesca (2015).

We start by considering inference under the normal model, i.e. conditional on $\gamma_{p+1} = \gamma_{p+2} = 0$. We ran 1,000 Gibbs iterations (*i.e.* $10^3 \times 10,172$ model updates), which was

Gene symbol	$p(\gamma y)$	
	Normal	Inferred
ARL4C,AOC3,URB2,FAM89B,PCGF2,CCDC102B	0.299	0.304
ARL4C,CNRIP1,AOC3,PCGF2	0.165	0.167
ARL4C,CNRIP1,PCGF2	0.161	0.163
ARL4C,CNRIP1,AOC3,PCGF2,RPS6KB2	0.045	0.046
ARL4C,AOC3,PCGF2,CCDC102B	0.028	0.028
ARL4C,AOC3,FAM89B,PCGF2,CCDC102B	0.025	0.025

TABLE 2. TGFB data. Highest probability models under normal and inferred error distribution.

deemed sufficient for practical convergence (see supplementary material in Rossell and Telesca (2015)) and required 1 minute and 35 seconds. Table 2 shows the highest posterior probability models. The top model included the 6 genes ARL4C, AOC3, URB2, FAM89B, PCGF2, CCDC102B and had an estimated $p(\gamma | y) = 0.299$. Alternatively, selecting variables with marginal $p(\gamma_j = 1 | y) > 0.5$ Barbieri and Berger (2004) returned 5 out of these 6 genes ($p(\gamma_j = 1 | y) = 0.482$ for FAM89B). Briefly, according to genecards.org FAM89B is a TGFB regulator, ARL4C and PCGF2 have been related to various cancer types and AOC3 is used to alleviate cancer symptoms, reinforcing the plausibility that these genes may be indeed related to TGFB. URB2 and CCDC102B have no known relation to cancer, although the latter is connected to ARL4D in the STRING interaction networks.

We next considered the possibility that the normal model might not be adequate for these data. As an exploratory check, a quantile-quantile plot based on the residuals under the top model revealed no strong departure from normality (Figure 4). Although this is somewhat reassuring one cannot discard a lack of normality under a different set of predictors, as the top model was selected under assumed normality. To conduct a more formal analysis we run Algorithm 5.3 ($T = 1,000$ iterations) now including $\gamma_{p+1}, \gamma_{p+2}$, which required 1 minute and 34 seconds. The posterior probabilities for Normal, asymmetric Normal, Laplace and asymmetric Laplace errors were 0.998, 0.0002, 0.0018 and 1.3×10^{-27} respectively. The six top models and their posterior probabilities closely matched those under the assumed normal model (Table 2) and the correlation between marginal inclusion probabilities under normal and inferred residuals was 0.96. These results support that when errors are approximately Normal our framework to infer $(\gamma_{p+1}, \gamma_{p+2})$ in Algorithm 5.3 indeed remains computationally competitive relative to conditioning upon Normal errors.

6.4. DLD data. We consider another genomics study described in Yuan et al. (2016). In contrast with that in Section 6.3 here RNA-sequencing was used to measure gene expression, a newer and more precise technology than microarrays. The study included 100 colorectal, 36 prostate and 6 pancreatic cancer and 50 healthy control patients, for a total of $n = 192$ patients. Briefly, the authors used a standard measure of expression called RPM. RPM considers the number of reads mapped to a given gene relative to the

gene length and may exhibit heavy tails or asymmetries, even after log or other transformations. We focus on the 58 messenger RNA genes identified in the exRNA species diversity analysis provided by the authors in Supplementary Table S1. To illustrate our methodology, we consider predicting the expression of gene DLD based on the remaining 57 genes and the 3 binary variables indicating the patient type (colorectal, prostate, pancreatic). According to genecards.org, the protein encoded by DLD can perform mechanistically distinct functions, it can regulate the energy metabolism and has been found to be associated with dehydrogenase and leukocyte adhesion deficiencies.

We first applied our methodology conditioning on normal residuals ($\gamma_{p+1} = \gamma_{p+2} = 0$). We used 10,000 Gibbs iterations, which required 1 second of CPU time. The highest posterior probability model had $p(\gamma | y) = 0.52$ and contained 6 genes (C6orf226, ECH1, CSF2RA, FBXL19, RRP1B), however its residuals showed a clear departure from normality (Figure 4, right). We run again our Gibbs algorithm, this time inferring γ_{p+1} and γ_{p+2} , which required 3 minutes and 40 seconds. The analysis returned an overwhelming $p(\gamma_{p+1} = 1, \gamma_{p+2} = 0 | y) = 0.9998$ in favour of Laplace-distributed residuals. The posterior mode contained the same 6 predictors but there were some interesting differences in terms of its having a diminished $p(\gamma | y) = 0.378$, and the second best model highlighting a gene MTMR1 that was missed by the analysis based on Normal residuals. MTMR1 encodes a protein related to the myotubularin family containing consensus sequences for protein tyrase phosphatases, whereas the response gene DLD has a post-translational modification based on tyrosine phosphorylation, thus giving a plausible biological mechanism connecting MTMR1 and DLD. Supplementary Table 7 lists the six largest marginal variable inclusion probabilities under normal and inferred error distribution.

7. DISCUSSION

Most efforts in Bayesian variable selection focus either on the Normal model or on flexible alternatives that require MCMC. Our framework represents a middle-ground to add flexibility in a parsimonious manner that remains analytically and computationally tractable, facilitating applications where either p is large or n is too moderate to fit more flexible models accurately. Our results show that failing to account for certain deviations from normality, particularly asymmetries and heavy tails, is a non-ignorable issue that has important consequences. These empirical observations agree with asymptotic characterizations, namely a lower estimation precision and a detrimental effect on variable selection. Although Bayes factors achieve the same asymptotic rates than those under the correct model, the coefficients governing these rates often lead to decreased sensitivity. From this standpoint it appears promising to combine flexible likelihoods with non-local priors, the former increasing power and the latter penalizing the addition of spurious parameters.

Interesting venues for future work include extensions that allow polynomial tails, heteroskedasticity or dependent errors, while maintaining tractability. A promising strategy along our lines here is to pose models that can adaptively discard the extra flexibility whenever it is not needed, to strike a balance between the tractability offered by simple models and the ultimate goal of providing accurate inference. Another interesting venue

is to obtain more sophisticated optimization or model search strategies, our goal here was to illustrate that even relatively simple methods can offer a competitive behaviour. These computational issues become particularly meaningful in increasingly challenging settings, *e.g.* large graphical or spatio-temporal models. Overall, we hope to have provided a basic framework that others can build on to tackle these exciting applications.

ACKNOWLEDGMENTS

DR was partially funded by the NIH grant R01 CA158113-01 and the Ramón y Cajal Fellowship RYC-2015-18544.

SUPPLEMENTARY MATERIAL

8. IMOM PRIOR

The product iMOM prior density on θ_γ Johnson and Rossell (2012) is given by

$$(12) \quad p_I(\theta_\gamma \mid \vartheta, \gamma) = \prod_{\gamma_j=1} \frac{(g_\theta \vartheta)^{\frac{1}{2}}}{\sqrt{\pi} \theta_j^2} \exp \left\{ -\frac{g_\theta \vartheta}{\theta_j^2} \right\},$$

where by default $g_\theta = 0.133$ assigns $p(|\theta/\vartheta^{1/2}| > 0.2) = 0.99$. Regarding the asymmetry parameter $\tilde{\alpha} = \operatorname{atanh}(\alpha)$, the prior is $p_I(\tilde{\alpha} \mid \gamma_{p+1} = 1) = \tilde{\alpha}^{-2} \sqrt{g_\alpha/\pi} e^{-g_\alpha/\tilde{\alpha}^2}$ and the default prior dispersions are $g_\alpha = 0.033$ to obtain $P(|\tilde{\alpha}| \geq 0.1) = 0.99$ and $g_\alpha = 0.136$ for $P(|\alpha| \geq 0.2) = 0.99$.

9. PROOFS

For simplicity we drop γ from the notation in the proof of Propositions 1-4 and Corollary 5, given that all arguments are conditional on a given model γ .

9.1. Proof of Proposition 1. We start by stating a useful lemma stating that positive definite hessian plus continuous gradient guarantees concavity.

Lemma 7. *Let $f(\theta)$ be a function with continuous gradient $g(\theta)$ for all θ and negative definite hessian $H(\theta)$ almost everywhere with respect to the Lebesgue measure. then $f(\theta)$ is strictly concave. If $H(\theta)$ is negative semidefinite then $f(\theta)$ is concave.*

Proof. Let θ_1 and θ_2 be two arbitrary values and denote $\theta_w = (1-w)\theta_1 + w\theta_2$ where $w \in [0, 1]$. Define $h(w) = -f(\theta_w)$, to show that $f(\theta)$ is concave it suffices to see that $h(w)$ is convex for arbitrary (w, θ_1, θ_2) . Straightforward algebra shows that $\frac{\partial}{\partial w} h(w) = -g(\theta_w)(\theta_2 - \theta_1)$ and further derivation shows that

$$\frac{\partial^2}{\partial w^2} h(w) = -(\theta_2 - \theta_1)^T H(\theta_w)(\theta_2 - \theta_1) > 0,$$

since $H(\theta)$ is negative definite (≥ 0 for negative semidefinite).

The second derivative $\frac{\partial^2}{\partial w^2} h(w) > 0$ almost everywhere and the first derivative $\frac{\partial}{\partial w} h(w)$ is continuous, which implies that $\frac{\partial}{\partial w} h(w)$ is strictly increasing in w and hence $h(w)$ is strictly convex (non-strictly convex when $H(\theta)$ is negative semidefinite). \square

Proof of Proposition 1, Part (i)

The gradient $g_1(\theta, \vartheta, \alpha)$ follows from straightforward algebra, which is obviously continuous with respect to $\vartheta \in \mathbb{R}^+$ and $\alpha \in [-1, 1]$. To see continuity of $g_1(\theta, \vartheta, \alpha)$ with respect to θ consider increasing a single θ_j for some $j \in \{1, \dots, p\}$ and fix the remaining elements in θ , which we denote $\theta_{(-j)}$. Also denote $x_{i(-j)}$ the subvector of x_i obtained by removing x_{ij} . Clearly $\log L_1(\theta, \vartheta, \alpha)$ is quadratic in θ_j with coefficients that stay constant until θ_j increases beyond a value t such that an observation i^* is added to or removed from $A(\theta)$, i.e. $y_{i^*} < x_{i^*(-j)}^T \theta_{(-j)} + x_{i^*j} \theta_j$ for $\theta_j \leq t$ and $y_{i^*} > x_{i^*(-j)}^T \theta_{(-j)} + x_{i^*j} \theta_j$ for $\theta_j > t$.

Taking the limit of the contribution of i^* to $\log(L_1(\theta, \vartheta, \alpha))$ as either $\theta_j \rightarrow t^-$ or $\theta_j \rightarrow t^+$ we obtain

$$\lim_{\theta_j \rightarrow t^+} \frac{(y_{i^*} - x_{i^*} \theta_i)^2}{(1 + \alpha)^2} = \lim_{\theta_j \rightarrow t^-} \frac{(y_{i^*} - x_{i^*} \theta_i)^2}{(1 - \alpha)^2} = 0,$$

i.e. $\log(L_1(\theta, \vartheta, \alpha))$ is continuous. Similarly, taking the limits for the contribution to the first partial derivative with respect to θ_j gives

$$\lim_{\theta_j \rightarrow t^+} \frac{2(y_{i^*} - x_{i^*} \theta_i)}{(1 + \alpha)^2} = \lim_{\theta_j \rightarrow t^-} \frac{2(y_{i^*} - x_{i^*} \theta_i)}{(1 - \alpha)^2} = 0,$$

which proves that $g_1(\theta, \vartheta, \alpha)$ is continuous.

Proof of Proposition 1, Part (ii)

The form of $H_1(\theta, \vartheta, \alpha)$ follows from easy algebra.

Proof of Proposition 1, Part (iii)

We start by noting that the maximum of the asymmetric-normal log-likelihood with respect to (θ, α) does not depend on ϑ , hence we simply need to see that

$$(13) \quad H = \begin{pmatrix} X^T W^2 X & 2X^T \bar{W}^3 (y - X\theta) \\ 2(y - X\theta)^T \bar{W}^3 X & 3(y - X\theta)^T W^2 (y - X\theta) \end{pmatrix},$$

is positive definite for almost all (θ, α) . Once we show this, by Part (i) and Lemma 7 we have that there is a unique maximum.

To see that H is positive definite, we shall show that all its leading principal minors are positive. Note that $X^T W^2 X$ is the gram matrix corresponding to WX and is hence positive definite when $\text{rank}(WX) = p$, or equivalently when $\text{rank}(X) = p$ given that the effect of W is to simply re-scale the rows of X . If $\text{rank}(WX) < p$ then $X^T W^2 X$ is positive semidefinite. Therefore, we just need to check that $\det(H) > 0$. Now, the usual formula for determinant based on submatrices gives that $\det(H) = \det(X^T W^2 X) \det(B)$, where $B =$

$$(14) \quad \begin{aligned} & 3(y - X\theta)^T W^4 (y - X\theta) - 4(y - X\theta)^T \bar{W}^3 X (X^T W^2 X)^{-1} X^T \bar{W}^3 (y - X\theta) \\ & = 3(y - X\theta)^T W^2 \left(I - \frac{4}{3} \bar{W} X (X^T W^2 X)^{-1} X^T \bar{W} \right) W^2 (y - X\theta), \end{aligned}$$

is a scalar, I is the $n \times n$ identity matrix, as usual W is an $n \times n$ diagonal matrix with entries $1/(1 \pm \alpha)^2$ where the \pm depends on whether $i \in A(\theta)$ or $i \notin A(\theta)$, and similarly \bar{W} is diagonal with entries $\pm(1 \pm \alpha)$. All that is left is to see that $B > 0$. For ease of

notation let us define $Z = \overline{WX}$, given that $\overline{W\overline{W}} = \text{diag}(1/(1 \pm \alpha)^2) = W^2$ we can write

$$\begin{aligned} B &= 3(y - X\theta)^T W^2 \left(I - \frac{4}{3} Z(Z^T Z)^{-1} Z^T \right) W^2 (y - X\theta) = \\ &4(y - X\theta)^T W^2 (I - Z(Z^T Z)^{-1} Z^T) W^2 (y - X\theta) - (y - X\theta)^T W^2 W^2 (y - X\theta) > 0 \\ (15) \quad &\Leftrightarrow 4 \frac{(y - X\theta)^T W^2 (I - Z(Z^T Z)^{-1} Z^T) W^2 (y - X\theta)}{(y - X\theta)^T W^2 W^2 (y - X\theta)} - 1 > 0. \end{aligned}$$

To complete the proof, note that $a = W^2(y - X\theta) \in \mathbb{R}^n$ is simply a vector and that the hat matrix $Z(Z^T Z)^{-1} Z^T$ is symmetric and idempotent, which implies that it has $\text{rank}(Z)$ eigenvalues equal to 1 and $n - \text{rank}(Z)$ eigenvalues equal to 0. Thus $I - Z(Z^T Z)^{-1} Z^T$ has $n - \text{rank}(Z)$ eigenvalues equal to 1 and the remaining $\text{rank}(Z)$ eigenvalues equal to 0. Given that $n > \text{rank}(Z)$ by assumption, $I - Z(Z^T Z)^{-1} Z^T$ has at least one non-zero eigenvalue, which allows us to bound

$$\min_{a \in \mathbb{R}^n} \frac{a(I - Z(Z^T Z)^{-1} Z^T)a}{a^T a} \geq 1,$$

which from (15) gives that $B \geq 3$ and hence that H is positive definite.

9.2. Proof of Proposition 2. Parts (i) and (ii) follow from straightforward algebra. For Part (iii) we first show that $\log L_2(\theta, \vartheta, \alpha)$ is (non-strictly) concave in (θ, α) and then that when $\text{rank}(X) = p$ it is strictly concave. To see non-strict concavity note that $-|y_i - x_i^T \theta| / (\sqrt{\vartheta}(1 + \alpha)) = -\max\{y_i - x_i^T \theta, x_i^T \theta - y_i\} / (\sqrt{\vartheta}(1 + \alpha))$ is the maximum of two (non-strictly) concave functions in (θ, α) and hence also concave, from which it follows that $L_3(\theta, \vartheta, \alpha)$ is a sum of concave functions and thus concave.

For ease of notation let $\eta = (\theta, \vartheta, \alpha)$, we now show that $\log L_2(\eta)$ is strictly concave at any arbitrary $\eta_1 = (\theta_1, \vartheta, \alpha_1)$ as long as $\text{rank}(X) = p$. It is useful to note that $H_2(\theta, \vartheta, \alpha)$ is strictly negative definite in α , as the corresponding minor $-2|W^3(y - X\theta)| / \sqrt{\vartheta} < 0$. From the definition of concavity and continuity of the log-likelihood, if $\log L_2(\eta)$ were concave but non-strictly concave at $\eta = \eta_1$ then for some $\eta_2 = (\theta_2, \vartheta, \alpha_2) \neq \eta_1$ we would have that $\log L_2(a\eta_1 + (1 - a)\eta_2) = a \log L_2(\eta_1) + (1 - a) \log L_2(\eta_2)$ for all $a \in [0, 1]$, *i.e.* $\log L_2(\eta)$ would be locally linear (in fact, constant) along the direction defined by $\eta_2 - \eta_1$, and in particular $\log L_2(\eta_1) = \log L_2(\eta_2)$. From its form

$$\log L_2(\eta) = -\frac{n}{2} \log(\vartheta) - \frac{1}{\vartheta} \left(\frac{\sum_{i \in A(\theta)} |y_i - x_i^T \theta|}{1 + \alpha} + \frac{\sum_{i \notin A(\theta)} |y_i - x_i^T \theta|}{1 - \alpha} \right)$$

is locally linear in θ but clearly non-linear in α , implying that $\alpha_2 = \alpha_1$. More formally, it is easy to see that for fixed $\theta_1 \neq \theta_2$ the roots of $\log L_2(\eta_1) = \log L_2(\eta_2)$ in terms of α_2 are given by the roots of a quadratic polynomial that are not linear in θ_2 , thus the only possible linear solution is $\alpha_2 = \alpha_1$. The problem is hence reduced to showing that there is no θ_2 sufficiently close to θ_1 such that

$$(16) \quad |W(y - X\theta_1)| = |W(y - X\theta_2)|,$$

where $|\cdot|$ denotes the L_1 norm and as usual W is a diagonal matrix with (i, i) element $(1 + \alpha)^{-1}$ if $i \in A(\theta_1)$ and $(1 - \alpha)^{-1}$ if $i \notin A(\theta_1)$, where we note that $A(\theta_2) = A(\theta_1)$

for θ_2 sufficiently close to θ_1 and thus the same weighting matrix W can be used in left and right hand sides of (16). Expression (16) is the L_1 error function featuring in median regression with re-scaled $\tilde{y} = Wy$ and $\tilde{X} = WX$, which is concave as long as $p = \text{rank}(WX) = \text{rank}(X)$, as we wished to prove.

9.3. Proof of Proposition 3.

Two-piece normal errors ($k = 1$). The proof strategy is as follows: we first show that the average log-likelihood $M_n(\theta_\gamma, \vartheta, \alpha) = \frac{1}{n} \log L_1(\theta_\gamma, \vartheta, \alpha)$ converges to its expected value $M(\theta_\gamma, \vartheta, \alpha)$ uniformly across $(\theta_\gamma, \vartheta, \alpha) \in \Gamma$, and later show that $M(\theta_\gamma, \vartheta, \alpha)$ has a unique maximum $(\theta_\gamma^*, \vartheta_\gamma^*, \alpha_\gamma^*)$, which jointly satisfy the conditions in Theorem 5.7 from van der Vaart (1998) for consistency of $(\hat{\theta}_\gamma, \hat{\vartheta}_\gamma, \hat{\alpha}_\gamma) \xrightarrow{P} (\theta_\gamma^*, \vartheta_\gamma^*, \alpha_\gamma^*)$.

We remark that Condition A3 is met for instance by deterministic sequences $\{x_i\}$ satisfying the stated positive-definiteness condition and also by $x_i \stackrel{i.i.d.}{\sim} \Psi$ as long as $E(x_1 x_1^T) = \Sigma$ for some positive definite Σ , since then $n^{-1} X^T X \xrightarrow{a.s.} \Sigma$ by the strong law of large numbers, and given that eigenvalues are continuous functions of $X^T X$ by the continuous mapping theorem $X^T X$ is positive definite almost surely as $n \rightarrow \infty$. Finally, Γ is assumed to contain the maximizer $(\theta_\gamma^*, \vartheta_\gamma^*, \alpha_\gamma^*)$.

By the law of large numbers and the *i.i.d.* assumption, we have that $M_n(\theta_\gamma, \vartheta, \alpha) \xrightarrow{P} M(\theta_\gamma, \vartheta, \alpha)$, for each $(\theta_\gamma, \vartheta, \alpha) \in \Gamma$. Next, we prove that the limit M is finite for all $(\theta_\gamma, \vartheta, \alpha) \in \Gamma$.

$$\begin{aligned} |M(\theta_\gamma, \vartheta, \alpha)| &= \left| \mathbb{E} [\log s_1(y_1 | x_1^T \theta_\gamma, \vartheta, \alpha)] \right| \leq \mathbb{E} [|\log s_1(y_1 | x_1^T \theta_\gamma, \vartheta, \alpha)|] \\ &= \int \int |\log s_1(y_1 | x_1^T \theta_\gamma, \vartheta, \alpha)| dS_0(y_1 | x_1) d\Psi(x_1) \\ &= \int \int_{y_1 < x_1^T \theta_\gamma} \left| \log \frac{1}{\sqrt{\vartheta}} \phi \left(\frac{y_1 - x_1^T \theta_\gamma}{\sqrt{\vartheta}(1 + \alpha)} \right) \right| dS_0(y_1 | x_1) d\Psi(x_1) \\ &+ \int \int_{y_1 \geq x_1^T \theta_\gamma} \left| \log \frac{1}{\sqrt{\vartheta}} \phi \left(\frac{y_1 - x_1^T \theta_\gamma}{\sqrt{\vartheta}(1 - \alpha)} \right) \right| dS_0(y_1 | x_1) d\Psi(x_1). \end{aligned}$$

For the first term in the last inequality we obtain, by integrating over the whole space, assumption A4 with $j = 2$, and the triangle inequality, the following upper bound

$$\begin{aligned} &\int \int \left| \log \frac{1}{\sqrt{\vartheta}} \phi \left(\frac{y_1 - x_1^T \theta_\gamma}{\sqrt{\vartheta}(1 + \alpha)} \right) \right| dS_0(y_1 | x_1) d\Psi(x_1) \\ &\leq |\log \sqrt{2\pi\vartheta}| + \int \int \frac{(y_1 - x_1^T \theta_\gamma)^2}{2\vartheta(1 + \alpha)^2} dS_0(y_1 | x_1) d\Psi(x_1) < \infty. \end{aligned}$$

Analogously for the second term. Now, let $\vartheta = \vartheta^*$ be an arbitrary fixed value for the (squared) scale parameter. The aim now is to first show that the average log-likelihood

$M_n(\theta_\gamma, \vartheta^*, \alpha) = n^{-1} \log L_1(\theta_\gamma, \vartheta^*, \alpha)$ converges to its expected value $M(\theta_\gamma, \vartheta^*, \alpha)$ uniformly in (θ_γ, α) , which implies that $(\widehat{\theta}_\gamma, \widehat{\alpha}_\gamma) \xrightarrow{P} (\theta_\gamma^*, \alpha_\gamma^*)$, and to then exploit that $\widehat{\vartheta}_\gamma$ and ϑ_γ^* have simple expressions to show that $\widehat{\vartheta}_\gamma \xrightarrow{P} \vartheta_\gamma^*$. To see that $M_n(\theta_\gamma, \vartheta^*, \alpha)$ converges to $M(\theta_\gamma, \vartheta^*, \alpha)$ uniformly in (θ_γ, α) we use the result in Proposition 1 that for positive-definite $X^T X$ (which holds for $n > n_0$) we have that $M_n(\theta_\gamma, \vartheta^*, \alpha)$ is a sequence of concave functions in (θ_γ, α) , which by the convexity lemma in Pollard (1991) (see also Theorem 10.8 from Rockafellar (2015)) implies that

$$(17) \quad \sup_{(\theta_\gamma, \alpha) \in K} |M_n(\theta_\gamma, \vartheta^*, \alpha) - M(\theta_\gamma, \vartheta^*, \alpha)| \xrightarrow{P} 0,$$

for each compact set $K \subseteq \Gamma$, and also that $M(\theta_\gamma, \vartheta^*, \alpha)$ is finite and concave in (θ_γ, α) and thus has a unique maximum $(\theta_\gamma^*, \alpha_\gamma^*)$. That is, for a distance measure $d(\cdot)$ and every $\varepsilon > 0$ we have

$$(18) \quad \sup_{d((\theta_\gamma^*, \vartheta^*, \alpha_\gamma^*), (\theta, \vartheta^*, \alpha)) \geq \varepsilon} M(\theta_\gamma, \vartheta^*, \alpha) < M(\theta_\gamma^*, \vartheta^*, \alpha_\gamma^*).$$

The consistency of $(\widehat{\theta}_\gamma, \widehat{\alpha}_\gamma) \xrightarrow{P} (\theta_\gamma^*, \alpha_\gamma^*)$ follows directly from (17) and (18) together with Theorem 5.7 from van der Vaart (1998). To see that $\widehat{\vartheta}_\gamma \xrightarrow{P} \vartheta_\gamma^*$, note first that from

$$(19) \quad M(\theta_\gamma, \vartheta^*, \alpha) = -\log(\sqrt{2\pi\vartheta^*}) - \frac{1}{2\vartheta^*} \int \left[\frac{(y_1 - x_1^T \theta_\gamma)^2}{(1 + \alpha)^2} I(y_1 < x_1^T \theta_\gamma) + \frac{(y_1 - x_1^T \theta_\gamma)^2}{(1 - \alpha)^2} I(y_1 \geq x_1^T \theta_\gamma) \right] dS_0(y_1 | x_1) d\Psi(x_1),$$

we see that $(\theta_\gamma^*, \alpha_\gamma^*)$ does not depend on ϑ^* , thus $(\theta_\gamma^*, \alpha_\gamma^*)$ is a global maximum. From (19) $M(\theta_\gamma^*, \vartheta, \alpha_\gamma^*)$ trivially has the maximizer

$$\vartheta_\gamma^* = \int \left[\frac{(y_1 - x_1^T \theta_\gamma^*)^2}{(1 + \alpha_\gamma^*)^2} I(y_1 < x_1^T \theta_\gamma^*) + \frac{(y_1 - x_1^T \theta_\gamma^*)^2}{(1 - \alpha_\gamma^*)^2} I(y_1 \geq x_1^T \theta_\gamma^*) \right] dS_0(y_1 | x_1) d\Psi(x_1),$$

and from the likelihood equations we have that

$$(20) \quad \widehat{\vartheta}_\gamma = \frac{1}{n} \left(\sum_{i=1}^n \frac{(y_i - x_i^T \widehat{\theta}_\gamma)^2}{(1 + \widehat{\alpha}_\gamma)^2} I(y_i \leq x_i^T \widehat{\theta}_\gamma) + \frac{(y_i - x_i^T \widehat{\theta}_\gamma)^2}{(1 - \widehat{\alpha}_\gamma)^2} I(y_i > x_i^T \widehat{\theta}_\gamma) \right).$$

In order to simplify notation, let us define

$$\rho(y_i, x_i, \theta_\gamma, \alpha) = \frac{(y_i - x_i^T \theta_\gamma)^2}{(1 + \alpha)^2} I(y_i \leq x_i^T \theta_\gamma) + \frac{(y_i - x_i^T \theta_\gamma)^2}{(1 - \alpha)^2} I(y_i > x_i^T \theta_\gamma).$$

Then, by the triangle inequality

$$\left| \widehat{\vartheta}_\gamma - \vartheta_\gamma^* \right| \leq \left| \widehat{\vartheta}_\gamma - \frac{1}{n} \sum_{i=1}^n \rho(y_i, x_i, \theta_\gamma^*, \alpha_\gamma^*) \right| + \left| \frac{1}{n} \sum_{i=1}^n \rho(y_i, x_i, \theta_\gamma^*, \alpha_\gamma^*) - \vartheta_\gamma^* \right|.$$

For the second term it follows, by the law of large numbers, that

$$\left| \frac{1}{n} \sum_{i=1}^n \rho(y_i, x_i, \theta_\gamma^*, \alpha_\gamma^*) - \vartheta_\gamma^* \right| \xrightarrow{P} 0.$$

For the first term we have

$$\begin{aligned} \left| \widehat{\vartheta}_\gamma - \frac{1}{n} \sum_{i=1}^n \rho(y_i, x_i, \theta_\gamma^*, \alpha_\gamma^*) \right| &= \left| M_n(\widehat{\theta}_\gamma, 1/2, \widehat{\alpha}_\gamma) - M_n(\theta_\gamma^*, 1/2, \alpha_\gamma^*) \right| \\ &\leq \left| M_n(\widehat{\theta}_\gamma, 1/2, \widehat{\alpha}_\gamma) - M(\widehat{\theta}_\gamma, 1/2, \widehat{\alpha}_\gamma) \right| \\ &\quad + \left| M(\widehat{\theta}_\gamma, 1/2, \widehat{\alpha}_\gamma) - M(\theta_\gamma^*, 1/2, \alpha_\gamma^*) \right| \\ &\quad + \left| M(\theta_\gamma^*, 1/2, \alpha_\gamma^*) - M_n(\theta_\gamma^*, 1/2, \alpha_\gamma^*) \right| \\ &\leq 2 \sup_{(\theta_\gamma, \alpha) \in \Gamma} |M_n(\theta_\gamma, 1/2, \alpha) - M(\theta_\gamma, 1/2, \alpha)| \\ &\quad + \left| M(\widehat{\theta}_\gamma, 1/2, \widehat{\alpha}) - M(\theta_\gamma^*, 1/2, \alpha_\gamma^*) \right|. \end{aligned}$$

By using (17), the consistency of $(\widehat{\theta}_\gamma, \widehat{\alpha}_\gamma)$, and the continuous mapping theorem it follows that $\left| \widehat{\vartheta}_\gamma - \frac{1}{n} \sum_{i=1}^n \rho(y_i, x_i, \theta_\gamma^*, \alpha_\gamma^*) \right| \xrightarrow{P} 0$. Consequently, $\widehat{\vartheta} \xrightarrow{P} \vartheta_\gamma^*$, which completes the proof.

Two-piece Laplace errors ($k = 2$). The proof strategy is analogous to that with $k =$

1. Denote $M_n(\theta_\gamma, \vartheta, \alpha) = \frac{1}{n} \log L_2(\theta_\gamma, \vartheta, \alpha)$. By the law of large numbers, we have that $M_n(\theta_\gamma, \vartheta, \alpha) \xrightarrow{P} M(\theta_\gamma, \vartheta, \alpha)$, for each $(\theta_\gamma, \vartheta, \alpha) \in \Gamma$. Moreover,

$$\begin{aligned} |M(\theta_\gamma, \vartheta, \alpha)| &= \left| \mathbb{E} [\log s_2(y_1 | x_1^T \theta_\gamma, \vartheta, \alpha)] \right| \leq \mathbb{E} [|\log s_2(y_1 | x_1^T \theta_\gamma, \vartheta, \alpha)|] \\ &= \int |\log s_2(y_1 | x_1^T \theta_\gamma, \vartheta, \alpha)| dS_0(y_1 | x_1) d\Psi(x_1) \\ &= \int_{y < x^T \theta_\gamma} \left| \log \frac{1}{\sqrt{\vartheta}} f \left(\frac{y_1 - x_1^T \theta_\gamma}{\sqrt{\vartheta}(1 + \alpha)} \right) \right| dS_0(y_1 | x_1) d\Psi(x_1) \\ &\quad + \int_{y_1 \geq x_1^T \theta_\gamma} \left| \log \frac{1}{\sqrt{\vartheta}} f \left(\frac{y_1 - x_1^T \theta_\gamma}{\sqrt{\vartheta}(1 - \alpha)} \right) \right| dS_0(y_1 | x_1) d\Psi(x_1), \end{aligned}$$

where $f(z) = 0.5 \exp(-|z|)$. For the first term in the last inequality we have, by integrating over the whole space and the triangle inequality, the following upper bound

$$\begin{aligned} &\int \left| \log \frac{1}{\sqrt{\vartheta}} f \left(\frac{y_1 - x_1^T \theta_\gamma}{\sqrt{\vartheta}(1 + \alpha)} \right) \right| dS_0(y_1 | x_1) d\Psi(x_1) \\ &\leq |\log 2\sqrt{\vartheta}| + \int \frac{|y_1 - x_1^T \theta_\gamma|}{\sqrt{\vartheta}(1 + \alpha)} dS_0(y_1 | x_1) d\Psi(x_1) < \infty, \end{aligned}$$

where the finiteness follows by assumption A4 with $j = 1$. An analogous result is obtained for the second term. Now, let $\vartheta = \vartheta^*$ be an arbitrary fixed value for the (squared) scale parameter. From Proposition 2, it follows that for positive-definite $X^T X$ (which is guaranteed by assumption A2, for $n > n_0$) we have that $M_n(\theta_\gamma, \vartheta^*, \alpha)$ is concave in (θ, α) , which by the convexity lemma in Pollard (1991) implies that

$$(21) \quad \sup_{(\theta_\gamma, \alpha) \in K} |M_n(\theta_\gamma, \vartheta^*, \alpha) - M(\theta_\gamma, \vartheta^*, \alpha)| \xrightarrow{P} 0,$$

for any compact set $K \subseteq \Gamma$, and also that $M(\theta_\gamma, \vartheta^*, \alpha)$ is concave in (θ_γ, α) and thus has a unique maximum $(\theta_\gamma^*, \alpha_\gamma^*)$. That is, for a distance measure $d(\cdot)$ and every $\varepsilon > 0$ we have

$$(22) \quad \sup_{d((\theta_\gamma^*, \vartheta^*, \alpha_\gamma^*), (\theta_\gamma, \vartheta^*, \alpha)) \geq \varepsilon} M(\theta_\gamma, \vartheta^*, \alpha) < M(\theta_\gamma^*, \vartheta^*, \alpha_\gamma^*).$$

The consistency of $(\hat{\theta}_\gamma, \hat{\alpha}_\gamma) \xrightarrow{P} (\theta_\gamma^*, \alpha_\gamma^*)$ follows directly from (21) and (22) together with Theorem 5.7 from van der Vaart (1998). To see that $\hat{\vartheta}_\gamma \xrightarrow{P} \vartheta^*$, note first that from

$$(23) \quad M(\theta_\gamma, \vartheta^*, \alpha) = -\log(2\sqrt{\vartheta^*}) - \frac{1}{\sqrt{\vartheta^*}} \int \left[\frac{|y_1 - x_1^T \theta_\gamma|}{1 + \alpha} I(y_1 < x_1^T \theta_\gamma) + \frac{|y_1 - x_1^T \theta_\gamma|}{1 - \alpha} I(y_1 \geq x_1^T \theta_\gamma) \right] dS_0(y_1|x_1) \Psi(x_1),$$

we see that $(\theta_\gamma^*, \alpha_\gamma^*)$ does not depend on ϑ^* , thus $(\theta_\gamma^*, \alpha_\gamma^*)$ is a global maximum. From (19) $M(\theta_\gamma^*, \vartheta, \alpha_\gamma^*)$ trivially has the maximizer

$$\vartheta_\gamma^* = \left\{ \int \left[\frac{|y_1 - x_1^T \theta_\gamma^*|}{1 + \alpha_\gamma^*} I(y_1 < x_1^T \theta_\gamma^*) + \frac{|y_1 - x_1^T \theta_\gamma^*|}{1 - \alpha_\gamma^*} I(y_1 \geq x_1^T \theta_\gamma^*) \right] dS_0(y_1|x_1) d\Psi(x_1) \right\}^2,$$

and from the likelihood equations we have that

$$(24) \quad \hat{\vartheta}_\gamma = \left[\frac{1}{n} \left(\sum_{i=1}^n \frac{|y_i - x_i^T \hat{\theta}_\gamma|}{1 + \hat{\alpha}_\gamma} I(y_i \leq x_i^T \hat{\theta}_\gamma) + \frac{|y_i - x_i^T \hat{\theta}_\gamma|}{1 - \hat{\alpha}_\gamma} I(y_i > x_i^T \hat{\theta}_\gamma) \right) \right]^2.$$

Let us define

$$\rho(y_i, x_i, \theta_\gamma, \alpha) = \frac{|y_i - x_i^T \theta_\gamma|}{1 + \alpha} I(y_i \leq x_i^T \theta_\gamma) + \frac{|y_i - x_i^T \theta_\gamma|}{1 - \alpha} I(y_i > x_i^T \theta_\gamma).$$

Then, by the triangle inequality

$$\left| \sqrt{\hat{\vartheta}_\gamma} - \sqrt{\vartheta_\gamma^*} \right| \leq \left| \sqrt{\hat{\vartheta}_\gamma} - \frac{1}{n} \sum_{i=1}^n \rho(y_i, x_i, \theta_\gamma^*, \alpha_\gamma^*) \right| + \left| \frac{1}{n} \sum_{i=1}^n \rho(y_i, x_i, \theta_\gamma^*, \alpha_\gamma^*) - \sqrt{\vartheta_\gamma^*} \right|.$$

For the second term in the right-hand side of the last equation, it follows, by the law of large numbers and the continuous mapping theorem, that

$$\left| \frac{1}{n} \sum_{i=1}^n \rho(y_i, x_i, \theta_\gamma^*, \alpha_\gamma^*) - \sqrt{\vartheta_\gamma^*} \right| \xrightarrow{P} 0.$$

For the first term we have

$$\begin{aligned}
\left| \sqrt{\widehat{\vartheta}_\gamma} - \frac{1}{n} \sum_{i=1}^n \rho(y_i, x_i, \theta_\gamma^*, \alpha_\gamma^*) \right| &= \left| M_n(\widehat{\theta}_\gamma, 1, \widehat{\alpha}_\gamma) - M_n(\theta_\gamma^*, 1, \alpha_\gamma^*) \right| \\
&\leq \left| M_n(\widehat{\theta}_\gamma, 1, \widehat{\alpha}_\gamma) - M(\widehat{\theta}_\gamma, 1, \widehat{\alpha}_\gamma) \right| \\
&\quad + \left| M(\widehat{\theta}_\gamma, 1, \widehat{\alpha}_\gamma) - M(\theta_\gamma^*, 1, \alpha_\gamma^*) \right| \\
&\quad + \left| M(\theta_\gamma^*, 1, \alpha_\gamma^*) - M_n(\theta_\gamma^*, 1, \alpha_\gamma^*) \right| \\
&\leq 2 \sup_{(\theta_\gamma, \alpha) \in \Gamma} |M_n(\theta_\gamma, 1, \alpha) - M(\theta_\gamma, 1, \alpha)| \\
&\quad + \left| M(\widehat{\theta}_\gamma, 1, \widehat{\alpha}_\gamma) - M(\theta_\gamma^*, 1, \alpha_\gamma^*) \right|.
\end{aligned}$$

By using (21), the consistency of $(\widehat{\theta}_\gamma, \widehat{\alpha}_\gamma)$, and the continuous mapping theorem it follows that $\left| \sqrt{\widehat{\vartheta}_\gamma} - \frac{1}{n} \sum_{i=1}^n \rho(y_i, x_i, \theta_\gamma^*, \alpha_\gamma^*) \right| \xrightarrow{P} 0$. Consequently, $\widehat{\vartheta}_\gamma \xrightarrow{P} \vartheta_\gamma^*$, which completes the proof.

9.4. Proof of Proposition 4.

Two-piece normal errors ($k = 1$). The proof technique consists of showing first that $\dot{m}_\eta(y_1, x_1)$ is dominated by an L^2 function (square integrable), $K(y_1, x_1)$, for η in a neighborhood of η_γ^* . Then, we prove that the function Pm_η admits a second-order Taylor expansion at η_γ^* and that the matrix $V_{\eta_\gamma^*}$ is nonsingular. Finally, we appeal to the consistency result in Proposition 3 in order to apply Theorem 5.23 of van der Vaart (1998) to prove the asymptotic normality of $\widehat{\eta}_\gamma$.

We first note that under assumptions A1–A4, where A4 is assumed to be satisfied for $j = 4$ throughout, Proposition 3 implies the existence and uniqueness of η_γ^* . The gradient of $m_\eta(y_1, x_1)$, which is given by (i) in Proposition 1 (with $n = 1$), is bounded for all $\eta \in \Gamma$ and for each (y_1, x_1) , due to the compactness of Γ . Now, a direct application of the Minkowski inequality implies that $\|\dot{m}_\eta(y_1, x_1)\|$ is upper bounded by the sum of the absolute values of the entries of $\dot{m}_\eta(y_1, x_1)$. Let us now define $K(y_1, x_1) = \sup_{\eta \in \mathcal{B}_{\eta_\gamma^*}} \|\dot{m}_\eta(y_1, x_1)\|$, where $\mathcal{B}_{\eta_\gamma^*} \subset \Gamma$ is any neighborhood of η_γ^* , whose projection over θ coincides with $\mathcal{B}_{\theta_\gamma^*}$. Thus, from the expression of $\dot{m}_\eta(y_1, x_1)$ together with assumption A4, it follows that

$$\int K(y_1, x_1)^2 dS_0(y_1|x_1) d\Psi(x_1) < \infty,$$

Then, by using the mean value theorem and the Cauchy-Schwartz inequality, it follows that for $\eta_1, \eta_2 \in \mathcal{B}_{\eta_0}$, with probability 1,

$$\begin{aligned}
|m_{\eta_1}(y_1, x_1) - m_{\eta_2}(y_1, x_1)| &= |\dot{m}_{\eta_\star}(y_1, x_1)^T (\eta_1 - \eta_2)| \\
&\leq \|\dot{m}_{\eta_\star}(y_1, x_1)\| \cdot \|\eta_1 - \eta_2\| \\
&\leq K(y_1, x_1) \cdot \|\eta_1 - \eta_2\|,
\end{aligned}$$

where $\eta_\star = (1 - c)\eta_1 + c\eta_2$, for some $c \in (0, 1)$.

Now, for each x_1 :

$$\begin{aligned} Pm_{\eta|x_1} = \mathbb{E}[m_\eta|x_1] &= -\frac{1}{2} \log(2\pi) - \frac{1}{2} \log(\vartheta) \\ &- \frac{1}{2\vartheta(1+\alpha)^2} \int_{-\infty}^{x_1^T \theta_\gamma} (y_1 - x_1^T \theta_\gamma)^2 dS_0(y_1|x_1) \\ &- \frac{1}{2\vartheta(1-\alpha)^2} \int_{x_1^T \theta_\gamma}^{\infty} (y_1 - x_1^T \theta_\gamma)^2 dS_0(y_1|x_1). \end{aligned}$$

Thus, the gradient of $Pm_{\eta|x_1}$ is given by

$$\begin{aligned} \frac{\partial}{\partial \theta_\gamma} Pm_{\eta|x_1} &= -\frac{x_1}{\vartheta(1+\alpha)^2} I_1 + \frac{x_1}{\vartheta(1-\alpha)^2} I_2, \\ \frac{\partial}{\partial \vartheta} Pm_{\eta|x_1} &= -\frac{1}{2\vartheta} + \frac{I_3}{2\vartheta^2(1+\alpha)^2} + \frac{I_4}{2\vartheta^2(1-\alpha)^2}, \\ \frac{\partial}{\partial \alpha} Pm_{\eta|x_1} &= \frac{I_3}{\vartheta(1+\alpha)^3} - \frac{I_4}{\vartheta(1-\alpha)^3}, \end{aligned}$$

Then, the second derivative matrix is given by

$$\begin{aligned} \frac{\partial^2}{\partial \theta_\gamma^2} Pm_{\eta|x_1} &= -\frac{x_1 x_1^T [(1+\alpha)^2 - 4\alpha S_0(x_1^T \theta_\gamma|x_1)]}{\vartheta(1-\alpha^2)^2}, \\ \frac{\partial^2}{\partial \vartheta^2} Pm_{\eta|x_1} &= \frac{1}{2\vartheta^2} - \frac{I_3}{\vartheta^3(1+\alpha)^2} - \frac{I_4}{\vartheta^3(1-\alpha)^2}, \\ \frac{\partial^2}{\partial \alpha^2} Pm_{\eta|x_1} &= -\frac{3I_3}{\vartheta(1+\alpha)^4} - \frac{3I_4}{\vartheta(1-\alpha)^4}, \\ \frac{\partial^2}{\partial \vartheta \partial \theta_\gamma} Pm_{\eta|x_1} &= \frac{x_1}{\vartheta^2(1+\alpha)^2} I_1 - \frac{x_1}{\vartheta^2(1-\alpha)^2} I_2, \\ \frac{\partial^2}{\partial \alpha \partial \theta_\gamma} Pm_{\eta|x_1} &= \frac{2x_1}{\vartheta(1+\alpha)^3} I_1 + \frac{2x_1}{\vartheta(1-\alpha)^3} I_2, \\ \frac{\partial^2}{\partial \vartheta \partial \alpha} Pm_{\eta|x_1} &= -\frac{I_3}{\vartheta^2(1+\alpha)^3} + \frac{I_4}{\vartheta^2(1-\alpha)^3}, \end{aligned}$$

where $I_1 = \int_{-\infty}^{x_1^T \theta_\gamma} S_0(y_1|x_1) dy_1$, and $I_2 = \int_{x_1^T \theta_\gamma}^{\infty} [1 - S_0(y_1|x_1)] dy_1$, $I_3 = \int_{-\infty}^{x_1^T \theta_\gamma} (y_1 - x_1^T \theta_\gamma)^2 dS_0(y_1|x_1)$, and $I_4 = \int_{x_1^T \theta_\gamma}^{\infty} (y_1 - x_1^T \theta_\gamma)^2 dS_0(y_1|x_1)$. These entries are finite for all $\eta \in \Gamma$ by assumption A4. Note that $Pm_\eta = \mathbb{E}[Pm_{\eta|x_1}]$, where the expectation is taken over x_1 . Assumptions A1–A4 together with Proposition (3) imply that Pm_η is finite and

that this expectation is concave and has a unique maximum at η_γ^* . From assumption A5,

$$\begin{aligned} \left. \frac{\partial}{\partial \theta_\gamma} Pm_\eta \right|_{\eta=\eta_\gamma^*} &= \mathbb{E} \left[\left. \frac{\partial}{\partial \theta_\gamma} Pm_{\eta|x_1} \right] \right|_{\eta=\eta_\gamma^*} = 0, \\ \left. \frac{\partial}{\partial \alpha} Pm_\eta \right|_{\eta=\eta_\gamma^*} &= \mathbb{E} \left[\left. \frac{\partial}{\partial \alpha} Pm_{\eta|x_1} \right] \right|_{\eta=\eta_\gamma^*} = 0, \end{aligned}$$

which in turn imply that $\frac{\partial^2}{\partial \vartheta \partial \theta_\gamma} Pm_\eta = 0$ and $\frac{\partial^2}{\partial \vartheta \partial \alpha} Pm_\eta = 0$ at $\eta = \eta_\gamma^*$. Thus, the matrix of second derivatives evaluated at η_γ^* has the following structure:

$$V_\eta = \begin{pmatrix} \frac{\partial^2}{\partial \theta_\gamma^2} Pm_\eta & 0 & \frac{\partial^2}{\partial \vartheta \partial \alpha} Pm_\eta \\ 0 & \frac{\partial^2}{\partial \vartheta^2} Pm_\eta & 0 \\ \frac{\partial^2}{\partial \vartheta \partial \alpha} Pm_\eta & 0 & \frac{\partial^2}{\partial \alpha^2} Pm_\eta \end{pmatrix}.$$

Consequently, the determinant of this matrix is given by

$$\det V_\eta = \frac{\partial^2}{\partial \vartheta^2} Pm_\eta \times \det \begin{pmatrix} \frac{\partial^2}{\partial \theta_\gamma^2} Pm_\eta & \frac{\partial^2}{\partial \vartheta \partial \alpha} Pm_\eta \\ \frac{\partial^2}{\partial \vartheta \partial \alpha} Pm_\eta & \frac{\partial^2}{\partial \alpha^2} Pm_\eta \end{pmatrix}.$$

The determinant on the right-hand side of this expression, evaluated at η_γ^* , is non-zero since the Pm_η is concave with respect to (θ_γ, α) , as shown in Proposition 3. Moreover, the fact that the first derivative $\frac{\partial}{\partial \vartheta} Pm_\eta = 0$ at $\eta = \eta_\gamma^*$ together with the fact that η_γ^* is the unique maximizer implies that $\frac{\partial^2}{\partial \vartheta^2} Pm_\eta \neq 0$. Consequently, the matrix of second derivatives of Pm_η is nonsingular at η_γ^* . The asymptotic normality result follows by Theorem 5.23 from van der Vaart (1998).

Two-piece Laplace errors ($k = 2$). First, we note that under assumptions A1–A4, where $j = 2$ in A4 throughout, Proposition 3 implies the existence and uniqueness of η_γ^* . The gradient of $m_\eta(y_1, x_1)$, which is given by (i) in Proposition 2 (with $n = 1$), is bounded for almost all $\eta \in \Gamma$ and for each (y_1, x_1) , due to the compactness of Γ . Now, a direct application of the Minkowski inequality implies that $\|\dot{m}_\eta(y_1, x_1)\|$ is upper bounded almost surely by the sum of the absolute values of the entries of $\dot{m}_\eta(y_1, x_1)$. Let us now define $K(y_1, x_1) = \sup_{\eta \in \mathcal{B}_{\eta_\gamma^*}} \|\dot{m}_\eta(y_1, x_1)\|$, where $\mathcal{B}_{\eta_\gamma^*} \subset \Gamma$ is any neighborhood of η_γ^* , whose projection over θ_γ coincides with $\mathcal{B}_{\theta_\gamma^*}$. Thus, from the expression of $\dot{m}_\eta(y_1, x_1)$ together with assumption A4, it follows that

$$\int K(y_1, x_1)^2 dS_0(y_1|x_1) d\Psi(x_1) < \infty,$$

Then, by using the mean value theorem and the Cauchy-Schwartz inequality, it follows that for $\eta_1, \eta_2 \in \mathcal{B}_{\eta_\star}$, with probability 1,

$$\begin{aligned} |m_{\eta_1}(y_1, x_1) - m_{\eta_2}(y_1, x_1)| &= |\dot{m}_{\eta_\star}(y_1, x_1)^T(\eta_1 - \eta_2)| \\ &\leq \|\dot{m}_{\eta_\star}(y_1, x_1)\| \cdot \|\eta_1 - \eta_2\| \\ &\leq K(y_1, x_1) \cdot \|\eta_1 - \eta_2\|, \end{aligned}$$

where $\eta_\star = (1 - c)\eta_1 + c\eta_2$, for some $c \in (0, 1)$.

Now, for each x_1 :

$$\begin{aligned} Pm_{\eta|x_1} = \mathbb{E}[m_\eta|x_1] &= -\log(2) - \frac{1}{2}\log(\vartheta) - \frac{1}{\sqrt{\vartheta}(1+\alpha)} \int_{-\infty}^{x_1^T\theta_\gamma} S_0(y_1|x_1)dy_1 \\ &\quad - \frac{1}{\sqrt{\vartheta}(1-\alpha)} \int_{x_1^T\theta_\gamma}^{\infty} 1 - S_0(y_1|x_1)dy_1. \end{aligned}$$

Then, the gradient of $Pm_{\eta|x_1}$ is given by

$$\begin{aligned} \frac{\partial}{\partial\theta_\gamma} Pm_{\eta|x_1} &= -\frac{x_1 S_0(x_1^T\theta_\gamma|x_1)}{\sqrt{\vartheta}(1+\alpha)} + \frac{x_1[1 - S_0(x_1^T\theta_\gamma|x_1)]}{\sqrt{\vartheta}(1-\alpha)}, \\ \frac{\partial}{\partial\vartheta} Pm_{\eta|x_1} &= -\frac{1}{2\vartheta} + \frac{I_1}{2\vartheta^{3/2}(1+\alpha)} + \frac{I_2}{2\vartheta^{3/2}(1-\alpha)}, \\ \frac{\partial}{\partial\alpha} Pm_{\eta|x_1} &= \frac{I_1}{\sqrt{\vartheta}(1+\alpha)^2} - \frac{I_2}{\sqrt{\vartheta}(1-\alpha)^2}, \end{aligned}$$

where $I_1 = \int_{-\infty}^{x_1^T\theta_\gamma} S_0(y_1|x_1)dy$, and $I_2 = \int_{x_1^T\theta_\gamma}^{\infty} 1 - S_0(y_1|x_1)dy$, which are finite by assumption A4. Then, the second derivative matrix is given by

$$\begin{aligned} \frac{\partial^2}{\partial\theta_\gamma^2} Pm_{\eta|x_1} &= -\frac{2x_1x_1^T s_0(x_1^T\theta_\gamma|x_1)}{\sqrt{\vartheta}(1-\alpha^2)}, \\ \frac{\partial^2}{\partial\vartheta^2} Pm_{\eta|x_1} &= \frac{1}{2\vartheta^2} - \frac{3I_1}{4\vartheta^{5/2}(1+\alpha)} - \frac{3I_2}{4\vartheta^{5/2}(1-\alpha)}, \\ \frac{\partial^2}{\partial\alpha^2} Pm_{\eta|x_1} &= -\frac{2I_1}{\sqrt{\vartheta}(1+\alpha)^3} - \frac{2I_2}{\sqrt{\vartheta}(1-\alpha)^3}, \\ \frac{\partial^2}{\partial\vartheta\partial\theta_\gamma} Pm_{\eta|x_1} &= \frac{x_1 S_0(x_1^T\theta_\gamma|x_1)}{2\vartheta^{3/2}(1+\alpha)} - \frac{x_1[1 - S_0(x_1^T\theta_\gamma|x_1)]}{2\vartheta^{3/2}(1-\alpha)}, \\ \frac{\partial^2}{\partial\alpha\partial\theta_\gamma} Pm_{\eta|x_1} &= \frac{x_1 S_0(x_1^T\theta_\gamma|x_1)}{\sqrt{\vartheta}(1+\alpha)^2} + \frac{x_1[1 - S_0(x_1^T\theta_\gamma|x_1)]}{\sqrt{\vartheta}(1-\alpha)^2}, \\ \frac{\partial^2}{\partial\vartheta\partial\alpha} Pm_{\eta|x_1} &= -\frac{I_1}{2\vartheta^{3/2}(1+\alpha)^2} + \frac{I_2}{2\vartheta^{3/2}(1-\alpha)^2}. \end{aligned}$$

These entries are finite for all $\eta \in \Gamma$ by assumption A4. Note that $Pm_\eta = \mathbb{E}[Pm_{\eta|x_1}]$, where the expectation is taken over x_1 . Assumptions A1–A4 together with Proposition 3,

imply that Pm_η is finite and that this expectation is concave and has a unique maximum at η_γ^* . From assumption A5,

$$\begin{aligned} \left. \frac{\partial}{\partial \theta_\gamma} Pm_\eta \right|_{\eta=\eta_\gamma^*} &= \mathbb{E} \left[\left. \frac{\partial}{\partial \theta_\gamma} Pm_{\eta|x_1} \right] \right|_{\eta=\eta_\gamma^*} = 0, \\ \left. \frac{\partial}{\partial \alpha} Pm_\eta \right|_{\eta=\eta_\gamma^*} &= \mathbb{E} \left[\left. \frac{\partial}{\partial \alpha} Pm_{\eta|x_1} \right] \right|_{\eta=\eta_\gamma^*} = 0, \end{aligned}$$

which in turn imply that $\frac{\partial^2}{\partial \vartheta \partial \theta_\gamma} Pm_\eta = 0$ and $\frac{\partial^2}{\partial \vartheta \partial \alpha} Pm_\eta = 0$ at $\eta = \eta_\gamma^*$. Thus, it follows that the matrix of second derivatives evaluated at η_γ^* has the structure:

$$V_\eta = \begin{pmatrix} \frac{\partial^2}{\partial \theta_\gamma^2} Pm_\eta & 0 & \frac{\partial^2}{\partial \vartheta \partial \alpha} Pm_\eta \\ 0 & \frac{\partial^2}{\partial \vartheta^2} Pm_\eta & 0 \\ \frac{\partial^2}{\partial \vartheta \partial \alpha} Pm_\eta & 0 & \frac{\partial^2}{\partial \alpha^2} Pm_\eta \end{pmatrix}.$$

Consequently, the determinant of this matrix is given by

$$\det V_\eta = \frac{\partial^2}{\partial \vartheta^2} Pm_\eta \times \det \begin{pmatrix} \frac{\partial^2}{\partial \theta_\gamma^2} Pm_\eta & \frac{\partial^2}{\partial \vartheta \partial \alpha} Pm_\eta \\ \frac{\partial^2}{\partial \vartheta \partial \alpha} Pm_\eta & \frac{\partial^2}{\partial \alpha^2} Pm_\eta \end{pmatrix}.$$

The determinant on the right-hand side of this expression, evaluated at η_γ^* , is non-zero since the Pm_η is concave with respect to (θ_γ, α) , as shown in Proposition 3. Moreover, the fact that the first derivative $\frac{\partial}{\partial \vartheta} Pm_\eta = 0$ at $\eta = \eta_\gamma^*$ together with the fact that η_γ^* is the unique maximizer implies that $\frac{\partial^2}{\partial \vartheta^2} Pm_\eta \neq 0$. Consequently, the matrix of second derivatives of Pm_η is nonsingular at η_γ^* . The asymptotic normality result follows by Theorem 5.23 from van der Vaart (1998).

9.5. Proof of Corollary 5. The result when $\epsilon_i \sim L(0, \vartheta)$ follows directly from Pollard (1991) Theorem 1, hence it suffices to find the expression for f_0 under each assumed residual distribution. The median for a general two-piece distribution with a mode at 0 is given by $\sqrt{\vartheta}(1 + \alpha)F^{-1}\left(\frac{1}{2(1+\alpha)}\right)$ if $\alpha > 0$ and $\sqrt{\vartheta}(1 - \alpha)F^{-1}\left(\frac{(1-2\alpha)}{2(1-\alpha)}\right)$ if $\alpha \leq 0$, where $F(\cdot)$ is the cdf of the standard underlying distribution with mode 0, $\vartheta = 1$ (Arellano-Valle et al. (2005), Expression (9)).

When $\epsilon_i \sim N(0, \vartheta)$ we have $m = 0$ and hence $f_0 = N(0; 0, \vartheta) = 1/(\sqrt{2\pi\vartheta})$. When $\epsilon_i \sim \text{AN}(0, \vartheta, \alpha)$ we have $m = \sqrt{\vartheta}(1 + \alpha)\Phi^{-1}(0.5/(1 + \alpha))$ if $\alpha > 0$ and $m = \sqrt{\vartheta}(1 - \alpha)\Phi^{-1}(0.5(1 - 2\alpha)/(1 - \alpha))$ if $\alpha < 0$, where $\Phi^{-1}(\cdot)$ is the inverse standard cdf, and

hence $f_0 = \exp \left\{ -\frac{1}{2} \left(\Phi^{-1} \left(\frac{0.5}{1+|\alpha|} \right) \right)^2 \right\} \frac{1}{\sqrt{2\pi\vartheta}}$. For the Laplace and Asymmetric Laplace, we note that the inverse cdf of the standard Laplace distribution evaluated at a quantile $q \in [0, 1]$ is $F^{-1}(q) = \log(2q)$ if $q < 0.5$ and $F^{-1}(q) = -\log(2(1-q))$ if $q \geq 0.5$. When $\epsilon_i \sim L(0, \vartheta)$ we have $m = 0$ and $f_0 = 1/(2\sqrt{\vartheta})$. Finally, when $\epsilon_i \sim \text{AL}(0, \vartheta, \alpha)$ we have $m = -\sqrt{\vartheta}(1+\alpha)\log(1+\alpha)$ if $\alpha > 0$ and $m = \sqrt{\vartheta}(1-\alpha)\log(1-\alpha)$ if $\alpha < 0$, from which it follows that $f_0 = \frac{1}{2\sqrt{\vartheta}} \exp \{-\log(1+|\alpha|)\} = \frac{1}{2\sqrt{\vartheta}(1+|\alpha|)}$.

The results for the true Normal model follows by using classic asymptotic results on least square estimators (see *e.g.* Newey and Powell (1987))

9.6. Proof of Proposition 6. We provide the proof for the asymmetric Normal and asymmetric Laplace ($\alpha \neq 0$), their symmetric counterparts follow as particular cases. $\eta_\gamma = (\theta_\gamma, \vartheta_\gamma, \alpha_\gamma)$ denotes the parameter vector under model γ , $\hat{\eta}_\gamma$ the MLE and $\tilde{\eta}_\gamma$ the posterior mode for a given observed (y, X) . Further, $M_k(\eta_\gamma) = E(\log L_k(\eta_\gamma))$ where the expectation is with respect to the data-generating truth and $\eta_\gamma^* = \arg \max_{\eta \in \Gamma_\gamma} M_k(\eta_\gamma)$ is the optimal parameter value under γ . We wish to characterize the asymptotic behaviour of the Bayes factors given by a Laplace approximation to the integrated likelihood

$$(25) \quad \frac{\hat{p}(y | \gamma)}{\hat{p}(y | \gamma^*)} = e^{\log L_k(\tilde{\eta}_\gamma) - \log L_k(\tilde{\eta}_{\gamma^*})} \times \frac{p(\tilde{\eta}_\gamma | \gamma)}{p(\tilde{\eta}_{\gamma^*} | \gamma^*)} \times (2\pi)^{\frac{p_\gamma - p_{\gamma^*}}{2}} \times \frac{|H_k(\tilde{\eta}_{\gamma^*})|^{\frac{1}{2}}}{|H_k(\tilde{\eta}_\gamma)|^{\frac{1}{2}}}$$

when (y, X) arise from the data-generating model in Condition A1, which may differ from the assumed model. We note that the third term on the right hand side is $O_p(1)$ given that p_γ is fixed and the fourth term is $O_p(n^{(p_{\gamma^*} - p_\gamma)/2})$, since by Proposition 3 we have that $\tilde{\eta}_\gamma = \hat{\eta}_\gamma + o_p(1) \xrightarrow{P} \eta_\gamma^*$ and thus by the continuous mapping theorem $n^{-1}H_k(\tilde{\eta}_\gamma) \xrightarrow{P} H_k(\tilde{\eta}_\gamma)$.

The strategy is to first show that when $M_k(\eta_\gamma^*) - M_k(\eta_{\gamma^*}^*) < 0$ (*i.e.* $\gamma^* \not\subset \gamma$) the first term of the right hand decreases exponentially with n and the second term is $O_p(1)$, giving that $\log(\hat{p}(y | \gamma)/\hat{p}(y | \gamma^*)) = O_p(n)$ as desired. Subsequently we shall show that when $M_k(\eta_\gamma^*) - M_k(\eta_{\gamma^*}^*) = 0$ (the case $\gamma^* \subset \gamma$) the first terms is essentially the likelihood ratio test statistic and is $O_p(1)$, whereas analogously to the results in Johnson and Rossell (2010); Rossell and Telesca (2015) the second term is $O_p(1)$ under local priors, but it is $O_p(n^{p_{\gamma^*} - p_\gamma})$ under the pMOM prior and $O_p(e^{-\sqrt{n}})$ under the eMOM prior, which again gives the desired result.

Consider first the case when $\gamma^* \not\subset \gamma$, which implies $M_k(\eta_\gamma^*) - M_k(\eta_{\gamma^*}^*) < 0$. Then by continuity of $p(\eta_\gamma | \gamma)$ we have that $p(\tilde{\eta}_\gamma | \gamma) \xrightarrow{P} p(\eta_\gamma^* | \gamma) > 0$, thus $p(\tilde{\eta}_\gamma | \gamma) = O_p(1)$ and analogously $p(\tilde{\eta}_{\gamma^*} | \gamma^*) = O_p(1)$. The law of large numbers and uniform convergence of $\log L_k$ to its expected value shown in Proposition 3 give that

$$(26) \quad \frac{1}{n}(\log L_k(\tilde{\eta}_\gamma) - \log L_k(\tilde{\eta}_{\gamma^*})) \xrightarrow{P} (M_k(\eta_\gamma^*) - M_k(\eta_{\gamma^*}^*)),$$

and thus $(\log L_k(\tilde{\eta}_\gamma) - \log L_k(\tilde{\eta}_{\gamma^*})) = nO_p(1)$ as we wanted to prove.

Next consider the case when $\gamma^* \subset \gamma$, which implies $M_k(\eta_\gamma^*) - M_k(\eta_{\gamma^*}^*) = 0$. Since $\tilde{\eta}_\gamma \xrightarrow{P} \eta_\gamma^*$ by Proposition 3 we have that under a local prior

$$(27) \quad \frac{p(\tilde{\eta}_\gamma | \gamma)}{p(\tilde{\eta}_{\gamma^*} | \gamma^*)} \xrightarrow{P} \frac{p(\eta_\gamma | \gamma)}{p(\eta_{\gamma^*} | \gamma^*)} > 0,$$

thus the left hand side is $O_p(1)$. Under a non-local prior we still have $p(\eta_{\gamma^*} | \gamma^*) > 0$ but in contrast $p(\eta_\gamma | \gamma) = 0$, thus it is necessary to characterize the rate at which the latter term vanishes. Briefly, following the proof of Theorem 1 in Koenker and Bassett (1982) the fact that $\log L_k$ converges uniformly to its expectation (see the proof of our Proposition 3) and consistency of $\tilde{\eta}_\gamma \xrightarrow{P} \eta_\gamma^*$ give that $\log L_k$ can be approximated by a quadratic function plus a term that is $o_p(1)$. Then the argument leading to Rossell and Telesca (2015) Proposition 2(i) gives that $\tilde{\theta}_{\gamma j} - \hat{\theta}_{\gamma j} = O_p(1/n)$ and thus $\tilde{\theta}_{\gamma j} = O_p(n^{-1/2})$ under the pMOM prior p_M , whereas $\tilde{\theta}_{\gamma j} = O_p(n^{-1/4})$ under the peMOM prior p_E . It follows that $\pi_M(\tilde{\theta}_\gamma) = O_p(1) \prod_{\theta_{\gamma j}^* \neq 0} \tilde{\theta}_{\gamma j}^2 = O_p(n^{-(p_\gamma - p_{\gamma^*})})$, and $\pi_E(\tilde{\eta}) = O_p(1) \prod_{\theta_{\gamma j}^* \neq 0} e^{O_p(1)/\tilde{\theta}_{\gamma j}^2} = O_p(e^{-\sqrt{n}})$, as desired.

To conclude the proof it only remains to show that $\log L_k(\tilde{\eta}_\gamma) - \log L_k(\tilde{\eta}_{\gamma^*}) = \lambda(y) + o_p(1) = O_p(1)$, where $\lambda(y) = \log L_k(\hat{\eta}_\gamma) - \log L_k(\hat{\eta}_{\gamma^*})$ is the likelihood ratio statistic. To see this, the strategy is to show that $\lambda(y) = \lambda(y; \vartheta_\gamma^*) + o_p(1)$ where $\lambda(y; \vartheta_\gamma^*) = \log L_k(\hat{\theta}_\gamma, \vartheta_\gamma^*, \hat{\alpha}_\gamma) - \log L_k(\hat{\theta}_{\gamma^*}, \vartheta_\gamma^*, \hat{\alpha}_{\gamma^*})$ is the likelihood ratio obtained by plugging in the oracle $\vartheta_\gamma^* = \vartheta_{\gamma^*}^*$, for which classical results establishing its asymptotic distribution exist. Straightforward algebra shows that for $k = 1$ the MLE must satisfy

$$\hat{\vartheta}_\gamma = \frac{1}{n} \left(\sum_{i \in A(\theta)} \frac{(y_i - x_i^T \hat{\theta}_\gamma)^2}{(1 + \alpha)} \sum_{i \notin A(\theta)} \frac{(y_i - x_i^T \hat{\theta}_\gamma)^2}{(1 - \alpha)} \right),$$

whereas for $k = 2$ it satisfies

$$\hat{\vartheta}_\gamma^{\frac{1}{2}} = \frac{1}{n} \left(\sum_{i \in A(\theta)} \frac{|y_i - x_i^T \hat{\theta}_\gamma|}{(1 + \alpha)} \sum_{i \notin A(\theta)} \frac{|y_i - x_i^T \hat{\theta}_\gamma|}{(1 - \alpha)} \right),$$

thus by plugging $\hat{\vartheta}_\gamma$ into the likelihoods and following Koenker and Machado (1999) Section 2.1 one obtains

$$(28) \quad \begin{aligned} \lambda(y) &= -\frac{n}{2} \log \left(\frac{\hat{\vartheta}_\gamma}{\hat{\vartheta}_{\gamma^*}} \right) = -\frac{n}{2} \log \left(1 + \frac{\hat{\vartheta}_\gamma - \hat{\vartheta}_{\gamma^*}}{\hat{\vartheta}_{\gamma^*}} \right) = -\frac{n}{2} \frac{\hat{\vartheta}_\gamma - \hat{\vartheta}_{\gamma^*}}{\hat{\vartheta}_{\gamma^*}} + o_p(1) \\ &= -\frac{1}{2} \frac{\hat{\vartheta}_\gamma - \hat{\vartheta}_{\gamma^*}}{\vartheta_{\gamma^*}^*} + o_p(1) = \lambda(y; \vartheta_\gamma^*) + o_p(1), \end{aligned}$$

since $n^{-1} \hat{\vartheta}_{\gamma^*} \xrightarrow{P} \vartheta_{\gamma^*}^*$ by Proposition 3.

Now following the proof of Theorem 1 in Koenker and Bassett (1982), given that $\log L_k$ converges uniformly to its expected value and $\hat{\alpha}_\gamma \xrightarrow{P} \alpha_\gamma^*$, $\hat{\alpha}_{\gamma^*} \xrightarrow{P} \alpha_{\gamma^*}^*$ where $\alpha_\gamma^* = \alpha_{\gamma^*}^*$, we obtain $\lambda(y; \vartheta_\gamma^*) = \lambda(y; \vartheta_\gamma^*, \alpha_\gamma^*) + o_p(1)$ where $\lambda(y; \vartheta_\gamma^*, \alpha_\gamma^*) = \log L_k(\hat{\theta}_\gamma, \vartheta_\gamma^*, \alpha_\gamma^*) -$

$\log L_k(\hat{\theta}_{\gamma^*}, \vartheta_{\gamma^*}^*, \alpha_{\gamma^*}^*)$ is the likelihood ratio test for fixed $(\vartheta_{\gamma^*}^*, \alpha_{\gamma^*}^*)$. Note that the symmetric Normal and Laplace correspond to $\hat{\alpha}_{\gamma} = 0$ and thus in that case $\lambda(y; \vartheta_{\gamma}^*) = \lambda(y; \vartheta_{\gamma^*}^*, \alpha_{\gamma^*}^*)$ exactly. To conclude the proof it suffices to show that $\lambda(y; \vartheta_{\gamma^*}^*, \alpha_{\gamma^*}^*) = O_p(1)$. When $k = 1$ the log-likelihood $\log L_1$ satisfies the regularity conditions in Schrader and Hettmansperger (1980), which are borrowed from Huber (1973) and essentially require the gradient of $\log L_k$ to be bounded and continuous, which it is by Proposition 1 and the assumption that θ_{γ} takes values in a compact space. Hence from Schrader and Hettmansperger (1980) we have that after appropriate rescaling by a constant $\lambda(y; \vartheta_{\gamma^*}^*; \alpha_{\gamma^*}^*)$ is asymptotically chi-square distributed with degrees of freedom $p_{\gamma} - p_{\gamma^*}$, and thus is $O_p(1)$ as we required. When $k = 2$ the conditions in Schrader and Hettmansperger (1980) are not satisfied due to gradient discontinuity, however then $\lambda(y; \vartheta_{\gamma^*}^*; \alpha_{\gamma^*}^*)$ corresponds to the likelihood ratio test for quantile regression that Koenker and Bassett (1982) showed to be asymptotically $\chi_{p_{\gamma} - p_{\gamma^*}}^2$ (after rescaling by a constant) precisely under our Conditions A2-A3. This concludes the proof.

10. APPROXIMATIONS TO THE INTEGRATED LIKELIHOOD

For ease of notation we drop the subindex k denoting the set of active variables and let $\theta = (\theta_1, \dots, \theta_{|k|})$ be their coefficients. Both the Laplace and Importance Sampling approximations require maximizing and evaluating the hessian of $h_l(\theta, \vartheta, \tilde{\alpha}) = \log L(\theta, \vartheta, \tilde{\alpha}) + \log p(\theta, \vartheta, \tilde{\alpha})$, where $L(\cdot)$ and $p(\cdot)$ are the appropriate likelihood and prior density. Denote by $g_l(\theta, \vartheta, \tilde{\alpha})$ the gradient of $h_l(\cdot)$ and by $H_l(\theta, \vartheta, \tilde{\alpha})$ its hessian, Algorithm 10 finds the posterior mode.

Posterior mode via Newton-Raphson

- (1) Initialize $(\theta^{(0)}, \tilde{\vartheta}^{(0)}, \tilde{\alpha}^{(0)}) = (\hat{\theta}, \log(\hat{\vartheta}), \text{atanh}(\hat{\alpha}))$ where $(\hat{\theta}, \hat{\vartheta}, \hat{\alpha})$ is the MLE given by Algorithm 4.2. Set $t = 1$ and repeat Steps 2-3 until e is below some small tolerance (default 10^{-5}).
 - (2) Update $(\theta^{(t)}, \tilde{\vartheta}^{(t)}, \tilde{\alpha}^{(t)}) =$
 $(\theta^{(t-1)}, \tilde{\vartheta}^{(t-1)}, \tilde{\alpha}^{(t-1)}) - H_l^{-1}(\theta^{(t-1)}, \tilde{\vartheta}^{(t-1)}, \tilde{\alpha}^{(t-1)}) g_l(\theta^{(t-1)}, \tilde{\vartheta}^{(t-1)}, \tilde{\alpha}^{(t-1)})$.
 - (3) Compute $e = \|(\theta^{(t)}, \tilde{\vartheta}^{(t)}, \tilde{\alpha}^{(t)}) - (\theta^{(t-1)}, \tilde{\vartheta}^{(t-1)}, \tilde{\alpha}^{(t-1)})\|^\infty$ where $\|\mathbf{z}\|^\infty$ is the largest element of \mathbf{z} in absolute value. Set $t = t + 1$.
-

As usual in the event that $(\theta^{(t)}, \tilde{\vartheta}^{(t)}, \tilde{\alpha}^{(t)})$ does not increase $h_l(\cdot)$ Step 2 can be adjusted by adding a constant λ to the diagonal of $H_l(\cdot)$, which for large λ gives the direction of the gradient and is guaranteed to decrease $h_l(\cdot)$. However we observed that this is extremely rare in practice. Usually the simple Newton step increases $h_l(\cdot)$ at each iteration and converges to the maximum in a few iterations.

Both $g_l(\cdot)$ and $H_l(\cdot)$ are the sum of a term coming from the log-likelihood plus a term coming from the log-prior density. The exact expressions are given below separately.

As an alternative to Algorithm 10 we also provide Algorithm 10 based on Coordinate Descent (*i.e.* successive univariate optimization). Note that the Newton steps to update

θ_j and α are in the direction of the gradient and are hence guaranteed to increase the objective function for small enough λ . Step 2 takes advantage of the fact that the maximizer with respect to $\tilde{\vartheta}$ for fixed (θ, α) is available in closed form.

Posterior mode via CDA

- (1) Initialize $\theta^{(0)}$ to the least squares estimate, $\alpha^{(0)} = 0$, $t = 0$.
- (2) For the MOM prior set $\tilde{\vartheta}^{(t)} = \log(s/(n + p + 3a_{\vartheta}))$, where

$$s = \left(b_{\vartheta} + \theta^{(t)T} \theta^{(t)} + \sum_{i \in A(\theta)} \frac{(y_i - x_i^T \theta^{(t)})^2}{(1 + \alpha^{(t)})^2} + \sum_{i \notin A(\theta)} \frac{(y_i - x_i^T \theta^{(t)})^2}{(1 - \alpha^{(t)})^2} \right).$$

For eMOM and iMOM use a Newton-Raphson step.

- (3) For $j = 1, \dots, p$
 - (a) Set $\lambda = 1$ and $\theta^* = \theta_j^{(t-1)} - \lambda g^*/h^*$, where g^* and h^* are the first and second derivatives of $f(\theta_j) = \log L_1(\theta_1^{(t-1)}, \dots, \theta_{j-1}^{(t-1)}, \theta_j, \theta_{j+1}^{(t-1)}, \dots, \theta_p^{(t-1)}, \vartheta^{(t)}, \alpha) + \log p(\theta_j | \vartheta)$ evaluated at $\theta_j = \theta_j^{(t-1)}$.
 - (b) If $f(\theta^*) > f(\theta_j^{(t-1)})$ set $\theta_j^{(t)} = \theta^*$, else set $\lambda = 0.5\lambda$ and repeat Step 3-(1).
 - (4) Let $\tilde{\alpha}^* = \tilde{\alpha}^{(t-1)} - \lambda g^*/h^*$, where g^* and h^* are the first and second derivatives of $f(\tilde{\alpha}) = \log L_1(\theta^{(t)}, \vartheta^{(t)}, \tilde{\alpha}) + \log p(\tilde{\alpha})$ at $\tilde{\alpha} = \tilde{\alpha}^{(t-1)}$. If $f(\tilde{\alpha}^*) > f(\tilde{\alpha}^{(t-1)})$ set $\alpha^{(t)} = \tilde{\alpha}^*$, else set $\lambda = 0.5\lambda$ and repeat Step 4.
 - (5) Compute $e = \max |(\theta^{(t)}, \tilde{\vartheta}^{(t)}, \tilde{\alpha}^{(t)}) - (\theta^{(t-1)}, \tilde{\vartheta}^{(t-1)}, \tilde{\alpha}^{(t-1)})|$. If $e < 10^{-5}$ stop, else set $t = t + 1$ and go back to Step 1.
-

10.1. Derivatives of the log-likelihood.

10.1.1. *Two-piece Normal.* Under the re-parameterization $\tilde{\vartheta} = \log(\vartheta)$, $\tilde{\alpha} = \operatorname{atanh}(\alpha)$ the two-piece Normal log-likelihood (3) has gradient

$$\begin{pmatrix} \frac{1}{\exp(\tilde{\vartheta})} X^T W (y - X\theta) \\ -\frac{n}{2} + \frac{1}{2\exp(\tilde{\vartheta})} (y - X\theta)^T W (y - X\theta) \\ \frac{1}{2\exp(\tilde{\vartheta})} (y - X\theta)^T W^* (y - X\theta) \end{pmatrix},$$

where as usual $W = \operatorname{diag}(w)$, $w_i = [1 + \tanh(\tilde{\alpha})]^{-2}$ if $i \in A(\theta)$ and $w_i = [1 - \tanh(\tilde{\alpha})]^{-2}$ if $i \notin A(\theta)$, and $W^* = \operatorname{diag}(w^*)$ with $w_i^* = -\frac{2\operatorname{sech}^2(\tilde{\alpha})}{(\tanh(\tilde{\alpha})+1)^3}$ if $i \in A(\theta)$ and $w_i^* = \frac{2\operatorname{sech}^2(\tilde{\alpha})}{(1-\tanh(\tilde{\alpha}))^3}$ if $i \notin A(\theta)$. Its Hessian is given by

$$(29) \quad -e^{-\tilde{\vartheta}} \begin{pmatrix} X^T W X & X^T W (y - X\theta) & X^T W^* (y - X\theta) \\ \frac{1}{2} (y - X\theta)^T W (y - X\theta) & -\frac{1}{2} (y - X\theta)^T W^* (y - X\theta) & \frac{1}{2} (y - X\theta)^T W^{**} (y - X\theta) \end{pmatrix},$$

where $W^{**} = \operatorname{diag}(w^{**})$ with $w_i^{**} = 2e^{-4\tilde{\alpha}} (e^{2\tilde{\alpha}} + 2)$ if $i \in A(\theta)$ and $w_i^{**} = 2e^{2\tilde{\alpha}} + 4e^{4\tilde{\alpha}}$ if $i \notin A(\theta)$.

10.1.2. *Two-piece Laplace*. The asymmetric Laplace $\log L_2(\theta, \tilde{\vartheta}, \tilde{\alpha})$ where $\tilde{\vartheta} = \log(\vartheta)$, $\tilde{\alpha} = \operatorname{atanh}(\alpha)$ has gradient

$$\begin{pmatrix} -e^{-\tilde{\vartheta}/2} X^T \bar{w} \\ -\frac{n}{2} + \frac{1}{2} e^{-\tilde{\vartheta}/2} w^T |y - X\theta| \\ e^{-\tilde{\vartheta}/2} |y - X\theta|^T \bar{w}^* \end{pmatrix},$$

and hessian

$$(30) \quad e^{-\tilde{\vartheta}/2} \times \begin{pmatrix} 0 & \frac{1}{2} X^T \bar{w} & X^T w^* \\ \frac{1}{2} \bar{w}^T X & -\frac{1}{4} w^T |y - X\theta| & -\frac{1}{2} |y - X\theta|^T \bar{w}^* \\ (X^T w^*)^T & -\frac{1}{2} |y - X\theta|^T \bar{w}^* & -2 |y - X\theta|^T w^* \end{pmatrix},$$

where $w_i = \bar{w}_i = (1 + \alpha)^{-1}$, $w_i^* = \bar{w}_i^* = e^{-2\alpha}$ if $i \in A(\theta)$, and $w_i = (1 - \alpha)^{-1}$, $\bar{w}_i = -w_i$, $w_i^* = e^{2\alpha}$, $\bar{w}_i^* = -w_i^*$ if $i \notin A(\theta)$. Naturally, symmetric Laplace errors are the particular case $\alpha = 0$ and give $w_i = w_i^* = 1$.

10.1.3. *Expected two-piece Laplace log-likelihood*. We derive $\bar{L}_2 = E(\log L_2(\eta))$ where $\eta = (\theta, \vartheta, \alpha)$ and its derivatives under the data-generating model $y_i = x_i^T \theta_0 + \epsilon_i$ for some $\theta_0 \in \mathbb{R}^p$ where ϵ_i are independent across $i = 1, \dots, n$ and arise from an arbitrary probability density function $s_0(y_i|x_i)$. After some algebra and noting that $\epsilon_i = y_i - x_i^T \theta_0$ gives

$$\begin{aligned} \bar{L}_2 = \int \log L_2(\eta) s_0(\epsilon|x) d\epsilon &= -n \log(2) - \frac{n}{2} \log(\vartheta) - \sum_{i=1}^n \frac{1}{\sqrt{\vartheta}(1+\alpha)} \int_{-\infty}^{x_i^T(\theta-\theta_0)} S_0(\epsilon_i) d\epsilon_i \\ &\quad - \sum_{i=1}^n \frac{1}{\sqrt{\vartheta}(1-\alpha)} \int_{x_i^T(\theta-\theta_0)}^{\infty} (1 - S_0(\epsilon_i)) d\epsilon_i, \end{aligned}$$

where $S_0(\epsilon_i) = S_0(\epsilon_i|0)$ is the cumulative probability function associated to $s_0(\epsilon_i) = s_0(\epsilon_i|0)$, where 0 indicates a zero covariate vector. Then taking derivatives we obtain

$$\begin{aligned} \frac{\partial}{\partial \theta} \bar{L}_2 &= \sum_{i=1}^n -\frac{x_i S_0(x_i^T(\theta - \theta_0))}{\sqrt{\vartheta}(1+\alpha)} + \frac{x_i [1 - S_0(x_i^T(\theta - \theta_0))]}{\sqrt{\vartheta}(1-\alpha)}, \\ \frac{\partial}{\partial \vartheta} \bar{L}_2 &= \sum_{i=1}^n -\frac{1}{2\vartheta} + \frac{I_{i1}}{2\vartheta^{3/2}(1+\alpha)} + \frac{I_{i2}}{2\vartheta^{3/2}(1-\alpha)}, \\ \frac{\partial}{\partial \alpha} \bar{L}_2 &= \sum_{i=1}^n \frac{I_{i1}}{\sqrt{\vartheta}(1+\alpha)^2} - \frac{I_{i2}}{\sqrt{\vartheta}(1-\alpha)^2}, \end{aligned}$$

where $I_{i1} = \int_{-\infty}^{x_i^T(\theta-\theta_0)} S_0(\epsilon_i) d\epsilon_i$, $I_{i2} = \int_{x_i^T(\theta-\theta_0)}^{\infty} (1 - S_0(\epsilon_i)) d\epsilon_i$. The second derivatives are

$$\begin{aligned}
\frac{\partial^2}{\partial\theta^2}\bar{L}_2 &= -\sum_{i=1}^n \frac{2x_i x_i^T s_0(x_i^T(\theta - \theta_0))}{\sqrt{\vartheta}(1 - \alpha^2)}, \\
\frac{\partial^2}{\partial\vartheta^2}\bar{L}_2 &= \sum_{i=1}^n \frac{1}{2\vartheta^2} - \frac{3I_{i1}}{4\vartheta^{5/2}(1 + \alpha)} - \frac{3I_{i2}}{4\vartheta^{5/2}(1 - \alpha)}, \\
\frac{\partial^2}{\partial\alpha^2}\bar{L}_2 &= -\sum_{i=1}^n \frac{2I_{i1}}{\sqrt{\vartheta}(1 + \alpha)^3} - \frac{2I_{i2}}{\sqrt{\vartheta}(1 - \alpha)^3}, \\
\frac{\partial^2}{\partial\vartheta\partial\theta}\bar{L}_2 &= \sum_{i=1}^n \frac{x_i S_0(x_i^T(\theta - \theta_0))}{2\vartheta^{3/2}(1 + \alpha)} - \frac{x_i[1 - S_0(x_i^T(\theta - \theta_0))]}{2\vartheta^{3/2}(1 - \alpha)}, \\
\frac{\partial^2}{\partial\alpha\partial\theta}\bar{L}_2 &= \sum_{i=1}^n \frac{x_i S_0(x_i^T(\theta - \theta_0))}{\sqrt{\vartheta}(1 + \alpha)^2} + \frac{x_i[1 - S_0(x_i^T(\theta - \theta_0))]}{\sqrt{\vartheta}(1 - \alpha)^2}, \\
\frac{\partial^2}{\partial\vartheta\partial\alpha}\bar{L}_2 &= -\sum_{i=1}^n \frac{I_{i1}}{2\vartheta^{3/2}(1 + \alpha)^2} + \frac{I_{i2}}{2\vartheta^{3/2}(1 - \alpha)^2}.
\end{aligned}$$

Simple inspection reveals that $(\partial/\partial\theta)\bar{L}_2 = 0$ implies $(\partial^2/\partial\theta\partial\vartheta)\bar{L}_2 = 0$, and likewise $(\partial/\partial\alpha)\bar{L}_2 = 0$ implies $(\partial^2/\partial\theta\partial\alpha)\bar{L}_2 = 0$. Since the maximum likelihood estimator $(\hat{\theta}, \hat{\vartheta}, \hat{\alpha})$ converges in probability to the maximizer of \bar{L}_2 , these second derivatives evaluated at $(\hat{\theta}, \hat{\vartheta}, \hat{\alpha})$ also converge in probability to 0.

We wish to find an asymptotic expression for the remaining second derivatives evaluated at $(\hat{\theta}, \hat{\vartheta}, \hat{\alpha})$ when the data-generating truth is $\epsilon_i \sim \text{AL}(x_i^T\theta_0, \vartheta_0, \alpha_0)$ for some $(\theta_0, \vartheta_0, \alpha_0)$. Given that $(\hat{\theta}, \hat{\vartheta}, \hat{\alpha}) \xrightarrow{P} (\theta_0, \vartheta_0, \alpha_0)$, the expressions above require evaluating the density of an asymmetric Laplace $s_0(0) = 1/(2\sqrt{\vartheta_0})$ and its cumulative probability function $S_0(0) = (1 + \alpha_0)/2$. Similarly, direct integration gives $I_{i1} = \sqrt{\vartheta_0}(1 + \alpha_0)^2/2$ and $I_{i2} = \sqrt{\vartheta_0}(1 - \alpha_0)^2/2$.

$$\begin{aligned}
\frac{\partial^2}{\partial\theta^2}\bar{L}_2 &\xrightarrow{P} -X^T X \frac{1}{\vartheta_0(1 - \alpha_0^2)}, \\
\frac{\partial^2}{\partial\vartheta^2}\bar{L}_2 &\xrightarrow{P} \frac{n}{2\vartheta_0^2} - \frac{3n(1 + \alpha_0)}{8\vartheta_0^2} - \frac{3(1 - \alpha_0)}{8\vartheta_0^2} = -\frac{n}{4\vartheta_0^2}, \\
\frac{\partial^2}{\partial\alpha^2}\bar{L}_2 &\xrightarrow{P} -\frac{n}{1 + \alpha_0} - \frac{n}{1 - \alpha_0} = -\frac{2n}{1 - \alpha_0^2}, \\
\frac{\partial^2}{\partial\alpha\partial\theta}\bar{L}_2 &\xrightarrow{P} \frac{n\bar{x}}{\sqrt{\vartheta_0}} \left(\frac{1}{2(1 + \alpha_0)} + \frac{1}{2(1 - \alpha_0)} \right) = \frac{n\bar{x}}{\sqrt{\vartheta_0}(1 - \alpha_0^2)}.
\end{aligned}$$

10.2. Derivatives of the log-prior density. The log-prior density is $\log p(\theta, \tilde{\vartheta}) = \log p(\theta \mid \tilde{\vartheta}) + \log p(\tilde{\vartheta})$ when $\tilde{\alpha} = 0$ under the assumed model and $\log p(\theta, \tilde{\vartheta}, \tilde{\alpha}) = \log p(\theta, \tilde{\vartheta}) + \log p(\tilde{\alpha})$ when $\tilde{\alpha} \neq 0$, where $p(\theta \mid \tilde{\vartheta})$ and $p(\tilde{\alpha})$ are the pMOM, piMOM

or peMOM priors and $p(\tilde{\vartheta}) = \text{IG}(e^{\tilde{\vartheta}}; a_{\vartheta}/2, b_{\vartheta}/2)e^{\tilde{\vartheta}}$. For ease of notation let θ^{-a} be the vector with elements θ_i^{-a} for $i = 1, \dots, |k|$.

10.2.1. pmOM prior. Straightforward algebra gives

$$\begin{aligned} \nabla \log p_M(\theta, \tilde{\vartheta}, \tilde{\alpha}) &= \begin{pmatrix} 2\theta^{-1} - \theta e^{-\tilde{\vartheta}}/g_{\theta} \\ -\frac{3|k|+a_{\vartheta}}{2} + (\theta^T \theta/g_{\theta} + b_{\vartheta})e^{-\tilde{\vartheta}}/2 \\ 2\tilde{\alpha}^{-1} - \tilde{\alpha}g_{\alpha}^{-1} \end{pmatrix}, \\ \nabla^2 \log p_M(\theta, \tilde{\vartheta}, \tilde{\alpha}) &= \begin{pmatrix} \text{diag}(-2\theta^{-2} - e^{-\tilde{\vartheta}}/g_{\theta}) & \theta e^{-\tilde{\vartheta}}/g_{\theta} & 0 \\ \theta^T e^{-\tilde{\vartheta}}/g_{\theta} & -e^{-\tilde{\vartheta}}(\theta^T \theta/g_{\theta} + b_{\vartheta})/2 & 0 \\ 0 & 0 & -2\tilde{\alpha}^{-2} - g_{\alpha}^{-1} \end{pmatrix}, \end{aligned}$$

10.2.2. piMOM prior. We obtain

$$\begin{aligned} \nabla \log p_I(\theta, \tilde{\vartheta}, \tilde{\alpha}) &= \begin{pmatrix} -2\theta^{-1} + 2g_{\theta}e^{\tilde{\vartheta}}\theta^{-3} \\ (|k| - a_{\vartheta})/2 + b_{\vartheta}e^{-\tilde{\vartheta}}/2 - g_{\theta}e^{\tilde{\vartheta}}\sum_i \theta_i^{-2} \\ -2\tilde{\alpha}^{-1} - 2g_{\alpha}\tilde{\alpha}^{-3} \end{pmatrix}, \\ \nabla^2 \log p_I(\theta, \tilde{\vartheta}, \tilde{\alpha}) &= \begin{pmatrix} \text{diag}(2\theta^{-2} - 6g_{\theta}e^{\tilde{\vartheta}}\theta^{-4}) & 2g_{\theta}e^{\tilde{\vartheta}}\theta^{-3} & 0 \\ (-2g_{\theta}e^{\tilde{\vartheta}}\theta^{-3})^T & -b_{\vartheta}e^{-\tilde{\vartheta}}/2 - e^{\tilde{\vartheta}}g_{\theta}\sum_i \theta_i^{-2} & 0 \\ 0 & 0 & 2\tilde{\alpha}^{-2} + 6g_{\alpha}\tilde{\alpha}^{-4} \end{pmatrix}. \end{aligned}$$

10.2.3. peMOM prior. We obtain

$$\nabla \log p_E(\theta, \tilde{\vartheta}, \tilde{\alpha}) = \begin{pmatrix} 2g_{\theta}e^{\tilde{\vartheta}}\theta^{-3} - \theta e^{-\tilde{\vartheta}}g_{\theta}^{-1} \\ -(|k| + a_{\vartheta})/2 + (b_{\vartheta} + \theta^T \theta/g_{\theta})e^{-\tilde{\vartheta}}/2 - g_{\theta}e^{\tilde{\vartheta}}\sum_i \theta_i^{-2} \\ 2g_{\alpha}\tilde{\alpha}^{-3} - \tilde{\alpha}g_{\alpha}^{-1} \end{pmatrix},$$

and $\nabla^2 \log p_E(\theta, \tilde{\vartheta}, \tilde{\alpha}) =$

$$\begin{pmatrix} \text{diag}(-6g_{\theta}e^{\tilde{\vartheta}}\theta^{-4} - e^{-\tilde{\vartheta}}g_{\theta}^{-1}) & 2g_{\theta}e^{\tilde{\vartheta}}\theta^{-3} + \theta e^{-\tilde{\vartheta}}g_{\theta}^{-1} & 0 \\ (2g_{\theta}e^{\tilde{\vartheta}}\theta^{-3} + \theta e^{-\tilde{\vartheta}}g_{\theta}^{-1})^T & -(b_{\vartheta} + \theta^T \theta/g_{\theta})e^{-\tilde{\vartheta}}/2 - e^{\tilde{\vartheta}}g_{\theta}\sum_i \theta_i^{-2} & 0 \\ 0 & 0 & -6g_{\alpha}\tilde{\alpha}^{-4} - g_{\alpha}^{-1} \end{pmatrix}.$$

10.3. Quadratic approximation to asymmetric Laplace log-likelihood. The goal is to approximate the curvature of the one-dimensional function $f(\lambda) = \log L_2(\theta_{\lambda}, \tilde{\vartheta}, \tilde{\alpha})$ around $\lambda = 0$, where $\theta_{\lambda} = (\hat{\theta}_1, \dots, \hat{\theta}_{j-1}, \hat{\theta}_j + \lambda, \hat{\theta}_{j+1}, \dots, \hat{\theta}_p)$ is fixed to the maximum likelihood estimator except for the j^{th} regression parameter, which is a function of $\lambda \in \mathbb{R}$. Given that $f(0)$ is known and that its derivative at $\lambda = 0$ is 0 ($\hat{\theta}$ is a maximum) we seek $h_j^* < 0$ such that $f(\lambda) - f(0) \approx 0.5h_j^*\lambda^2$. Our strategy is to evaluate $f(\lambda_k)$ on a grid λ_k for $k = 1, \dots, K$ and use the least-squares estimate $h_j^* = 2 \sum_{k=1}^K \lambda_k^2 (f(\lambda_k) - f(0)) / \sum_{k=1}^K \lambda_k^4$,

where the form of $\log L_2$ gives the simple expression

$$f(\lambda_k) - f(0) = -\frac{1}{\sqrt{\hat{\vartheta}}} \sum_{i=1}^n |r_i - \lambda_k x_{ij}| \left(\frac{\mathbb{I}(r_i \leq \lambda_k x_{ij})}{1 + \hat{\alpha}} + \frac{\mathbb{I}(r_i > \lambda_k x_{ij})}{1 - \hat{\alpha}} \right),$$

and $r_i = y_i - x_i^T \hat{\theta}$. Once h_1^*, \dots, h_p^* have been obtained we let $D = \text{diag}(h_1^*/\bar{h}_{11}, \dots, h_p^*/\bar{h}_{pp})$ where $\bar{H} = (X^T X)/(\hat{\vartheta}(1 - \hat{\alpha}^2))$ is the asymptotic hessian under asymmetric Laplace errors, and we approximate the hessian of $\log L_2(\theta, \hat{\vartheta}, \hat{\alpha})$ around $\theta = \hat{\theta}$ with $H^* = D^{\frac{1}{2}} \bar{H} D^{\frac{1}{2}}$. The construction ensures that the diagonal elements in H^* are h_1^*, \dots, h_p^* , i.e. the quadratic approximation matches the actual curvature of $\log L_2$ along each canonical axis. From Section 4 the correlation structure borrowed from \bar{H} remains asymptotically valid as long as the residuals are independent and identically distributed, however in our experience the approximation usually suffices for practical purposes even when these assumptions is violated.

The problem has been thus reduced to choosing the grid $\lambda_1, \dots, \lambda_K$. One naive option is to take the n points of non-differentiability $\lambda = r_i/x_{ij}$, however by the nature of least squares this strategy tends to approximate better $f(\lambda)$ for large λ^2 and we are interested in local approximations around $\lambda = 0$, further evaluating $f(\lambda)$ at n points requires $O(n^2)$ operations for each $j = 1, \dots, p$ and is thus computationally costly. Instead we evaluate $f(\lambda)$ only at the $K = 2$ points given by the endpoints of the asymptotic 95% confidence interval $\lambda = \{-1.96\bar{v}_j, 1.96\bar{v}_j\}$ where \bar{v}_j is the j^{th} diagonal element in \bar{H}^{-1} . This simple strategy ensures that the approximation holds locally around $\lambda = 0$ in the sense of having non-negligible likelihood, requires only $O(n)$ operations and we have observed to deliver reasonably accurate approximations in practice. Our approximation is similar in spirit to the rank-based score test inversion used to obtain confidence intervals in quantile regression, which has been amply described to deliver fairly precise intervals, with the important difference that rank inversion requires an ordering of observations that scales poorly with p and n .

Supplementary Figure 5 shows an example with the likelihood L_2 (scaled to $(0, 1)$) and the two quadratic approximations based on the asymptotic covariance and its least-squares adjustment for an intercept-only model ($p = 1$) and $n = 200$. When residuals were truly generated from an asymmetric Laplace (left panel) the two quadratic approximations were essentially identical, however under truly normally distributed residuals the asymptotic covariance over-estimated the curvature.

11. SUPPLEMENTARY RESULTS

11.1. Simulation study. We assessed the sensitivity of the results of the $p = 6$ simulation study in Section 6.1 of the main paper to the prior on the asymmetry coefficient by setting g_α such that $P(|\alpha| > 0.1) = 0.99$. Supplementary Table 3 summarizes the inference on the error distribution and Supplementary Figure 6 the marginal variable inclusion probabilities. The latter were virtually identical to those in Figure 2 obtained under g_α such that $P(|\alpha| > 0.2) = 0.99$, showing that variable inclusion is robust to moderate changes in g_α .




FIGURE 5. Quadratic approximation to L_2 (solid grey) with $p = 1, n = 200$ from asymptotic covariance (dotted black) and least-squares adjustment (solid black). Left: $\epsilon_i \sim \text{AL}(0, 2, -0.5)$; Right: $\epsilon_i \sim N(0, 2)$.

Truth	Average $p(\gamma_{p+1}, \gamma_{p+2} y)$			
	$\gamma_{p+1} = \gamma_{p+2} = 0$	$\gamma_{p+1} = 1, \gamma_{p+2} = 0$	$\gamma_{p+1} = 0, \gamma_{p+2} = 1$	$\gamma_{p+1} = \gamma_{p+2} = 0$
$p = 6, g_\alpha = 0.357, \text{Laplace } p(\gamma y)$				
$N(0, 2)$	0.91	0.02	0.06	0.00
$\text{AN}(0, 2, -0.5)$	0.11	0.81	0.01	0.06
$L(0, 2)$	0.14	0.00	0.84	0.02
$\text{AL}(0, 2, -0.5)$	0.02	0.12	0.01	0.85
$p = 6, g_\alpha = 0.357, \text{Monte Carlo } p(\gamma y)$				
$N(0, 2)$	0.91	0.02	0.06	0.00
$\text{AN}(0, 2, -0.5)$	0.11	0.81	0.01	0.07
$L(0, 2)$	0.12	0.01	0.85	0.02
$\text{AL}(0, 2, -0.5)$	0.02	0.12	0.01	0.85
$p = 6, g_\alpha = 0.087, \text{Laplace } p(\gamma y)$				
$N(0, 2)$	0.87	0.07	0.06	0.01
$\text{AN}(0, 2, -0.5)$	0.07	0.86	0.01	0.07
$L(0, 2)$	0.13	0.01	0.79	0.07
$\text{AL}(0, 2, -0.5)$	0.01	0.13	0.01	0.85

TABLE 3. Simulation study for $p = 6$. Posterior probability of the 4 error distributions under $\vartheta = 2, \theta = (0, 0.5, 1, 1.5, \dots, 0), n = 100, \rho_{ij} = 0.5$.

Truth	Average $p(\gamma_{p+1}, \gamma_{p+2} y)$			
	$\gamma_{p+1} = \gamma_{p+2} = 0$	$\gamma_{p+1} = 1, \gamma_{p+2} = 0$	$\gamma_{p+1} = 0, \gamma_{p+2} = 1$	$\gamma_{p+1} = \gamma_{p+2} = 0$
$p = 101, \vartheta = 1$				
$N(0, 2)$	0.91	0.01	0.08	0.00
$AN(0, 2, -0.5)$	0.03	0.86	0.00	0.11
$L(0, 2)$	0.15	0.01	0.83	0.02
$AL(0, 2, -0.5)$	0.00	0.13	0.01	0.86
$p = 101, \vartheta = 2$				
$N(0, 2)$	0.89	0.01	0.10	0.00
$AN(0, 2, -0.5)$	0.02	0.89	0.00	0.09
$L(0, 2)$	0.15	0.01	0.82	0.02
$AL(0, 2, -0.5)$	0.00	0.16	0.01	0.83
$p = 501, \vartheta = 1$				
$N(0, 2)$	0.85	0.00	0.14	0.00
$AN(0, 2, -0.5)$	0.01	0.85	0.01	0.14
$L(0, 2)$	0.18	0.00	0.80	0.02
$AL(0, 2, -0.5)$	0.00	0.15	0.00	0.84
$p = 501, \vartheta = 2$				
$N(0, 2)$	0.83	0.00	0.16	0.00
$AN(0, 2, -0.5)$	0.00	0.87	0.00	0.12
$L(0, 2)$	0.19	0.00	0.79	0.01
$AL(0, 2, -0.5)$	0.00	0.22	0.00	0.77

TABLE 4. Simulation study for $p = 101, 501$. Posterior probability of the 4 error distributions under $g_\alpha = 0.357$, $\theta = (0, 0.5, 1, 1.5, \dots, 0)$, $n = 100$, $\rho_{ij} = 0.5$. Laplace approximation to $p(y | \gamma)$ was used.

We also assessed the accuracy of the Laplace approximations to the integrated likelihood $p(y | \gamma)$ by comparing the results with those obtained with the importance sampling estimates with $B = 10,000$ draws described in Section 5 of the main paper. Supplementary Figure 7 displays the results for $g_\alpha = 0.357$. These are extremely similar to those based on Laplace approximation in Figure 2.

Supplementary Figure 10 shows analogous results for $p = 100$, with $g_\alpha = 0.357$ and $p(y | \gamma)$ estimated via Laplace approximations.

11.2. DLD data. Supplementary Table 7 shows the six genes with largest marginal inclusion probabilities $p(\gamma_j = 1 | y)$ when conditioning on Normal errors and when inferring the error distribution. The figures were similar for the four top genes, but the Normal model assigned somewhat higher probability to FBXL19 substantially lower probability to MTMR1.

	$p = 100$				$p = 500$			
	$p(\gamma_0 y)$	$p(\hat{\gamma} = \gamma_0)$	FP	TP	$p(\gamma_0 y)$	$p(\hat{\gamma} = \gamma_0)$	FP	TP
Truly $\epsilon \sim N(0, 1)$								
Normal	0.46	0.63	0.1	2.7	0.26	0.37	0.2	2.4
Two-piece Normal	0.43	0.63	0.2	2.7	0.24	0.38	0.3	2.4
Laplace	0.26	0.42	0.5	2.6	0.12	0.19	0.8	2.3
Two-piece Laplace	0.23	0.39	0.7	2.6	0.12	0.21	0.9	2.3
Inferred	0.45	0.62	0.2	2.7	0.25	0.37	0.2	2.4
LASSO-LS		0.00	12.4	3.0		0.00	20.4	2.9
LASSO-LAD		0.00	14.7	2.9		0.00	27.8	2.7
SCAD		0.07	4.2	2.9		0.01	7.3	2.8
Truly $\epsilon \sim AN(0, 1, -0.5)$								
Normal	0.38	0.55	0.2	2.6	0.21	0.34	0.5	2.4
Two-piece Normal	0.59	0.73	0.1	2.8	0.40	0.55	0.4	2.6
Laplace	0.20	0.35	0.7	2.5	0.07	0.14	1.2	2.4
Two-piece Laplace	0.33	0.48	0.5	2.7	0.18	0.32	1.1	2.5
Inferred	0.57	0.72	0.1	2.8	0.38	0.52	0.4	2.6
LASSO-LS		0.00	12.4	3.0		0.00	21.9	2.9
LASSO-LAD		0.00	13.8	2.8		0.00	27.9	2.7
SCAD		0.07	4.0	2.9		0.03	7.3	2.8
Truly $\epsilon \sim L(0, 1)$								
Normal	0.11	0.14	0.3	2.0	0.03	0.02	0.6	1.6
Two-piece Normal	0.11	0.15	0.3	2.1	0.04	0.04	1.1	1.7
Laplace	0.29	0.38	0.2	2.4	0.13	0.19	0.4	2.0
Two-piece Laplace	0.28	0.35	0.3	2.4	0.12	0.18	0.5	2.0
Inferred	0.28	0.38	0.2	2.4	0.12	0.18	0.4	2.0
LASSO-LS		0.00	11.3	2.8		0.00	21.4	2.5
LASSO-LAD		0.01	13.1	2.9		0.00	23.9	2.6
SCAD		0.02	5.0	2.7		0.00	9.0	2.4
Truly $\epsilon \sim AL(0, -0.5)$								
Normal	0.07	0.10	0.4	1.9	0.02	0.02	1.1	1.5
Two-piece Normal	0.21	0.27	0.2	2.2	0.11	0.15	0.3	2.0
Laplace	0.16	0.19	0.4	2.1	0.05	0.07	0.7	1.8
Two-piece Laplace	0.43	0.51	0.2	2.5	0.27	0.34	0.4	2.3
Inferred	0.41	0.48	0.2	2.5	0.25	0.33	0.4	2.2
LASSO-LS		0.00	11.6	2.8		0.00	20.1	2.5
LASSO-LAD		0.00	12.5	2.6		0.00	24.7	2.3
SCAD		0.01	5.2	2.6		0.01	9.4	2.3

TABLE 5. Simulation results under $\vartheta = 1$. γ_0 : true predictors. $\hat{\gamma}$: selected variables. CC: number of correctly classified variables ($\sum_{j=1}^p \mathbb{I}(\hat{\gamma}_j = \gamma_{0j})$). FP: number of false positives; TP: number of true positives.

	$p = 100$				$p = 500$			
	$p(\gamma_0 y)$	$p(\hat{\gamma} = \gamma_0)$	FP	TP	$p(\gamma_0 y)$	$p(\hat{\gamma} = \gamma_0)$	FP	TP
Truly $\epsilon \sim N(0, 1)$								
Normal	0.01	0.01	0.4	1.2	0.00	0.00	0.8	0.9
Two-piece Normal	0.01	0.01	0.5	1.2	0.00	0.00	0.9	0.8
Laplace	0.00	0.00	0.7	1.1	0.00	0.00	1.0	0.8
Two-piece Laplace	0.00	0.01	0.8	1.1	0.00	0.00	1.1	0.8
Inferred	0.01	0.01	0.5	1.2	0.00	0.00	0.7	0.9
LASSO-LS		0.00	11.9	2.5		0.00	18.0	2.0
LASSO-LAD		0.00	12.6	2.0		0.00	23.1	1.6
SCAD		0.00	6.3	2.3		0.01	10.4	1.8
Truly $\epsilon \sim AN(0, 1, -0.5)$								
Normal	0.00	0.00	0.5	1.2	0.00	0.00	0.7	0.9
Two-piece Normal	0.01	0.01	0.4	1.4	0.00	0.01	0.7	1.1
Laplace	0.00	0.00	0.9	1.0	0.00	0.00	1.4	0.7
Two-piece Laplace	0.01	0.01	0.7	1.2	0.00	0.00	1.5	1.0
Inferred	0.01	0.01	0.4	1.4	0.00	0.01	0.9	1.0
LASSO-LS		0.00	11.0	2.4		0.00	19.4	1.9
LASSO-LAD		0.00	12.7	1.8		0.00	22.6	1.4
SCAD		0.00	6.1	2.1		0.00	10.1	1.8
Truly $\epsilon \sim L(0, 1)$								
Normal	0.01	0.01	0.4	1.3	0.00	0.00	0.8	0.9
Two-piece Normal	0.01	0.01	0.5	1.3	0.00	0.00	0.9	1.0
Laplace	0.05	0.06	0.4	1.7	0.01	0.01	0.7	1.2
Two-piece Laplace	0.05	0.07	0.4	1.7	0.01	0.01	0.8	1.2
Inferred	0.04	0.04	0.3	1.7	0.01	0.01	0.7	1.2
LASSO-LS		0.00	10.8	2.5		0.00	20.4	2.0
LASSO-LAD		0.01	12.7	2.5		0.00	21.7	2.0
SCAD		0.00	5.9	2.2		0.00	10.3	1.8
Truly $\epsilon \sim AL(0, -0.5)$								
Normal	0.00	0.00	0.5	1.1	0.00	0.00	0.9	0.8
Two-piece Normal	0.02	0.01	0.4	1.5	0.01	0.01	0.6	1.2
Laplace	0.02	0.01	0.6	1.3	0.00	0.01	0.8	1.0
Two-piece Laplace	0.09	0.12	0.3	1.9	0.04	0.05	0.7	1.5
Inferred	0.09	0.10	0.3	1.8	0.04	0.05	0.6	1.4
LASSO-LS		0.00	10.9	2.3		0.00	18.0	1.8
LASSO-LAD		0.00	11.8	2.1		0.00	21.0	1.6
SCAD		0.01	5.7	2.1		0.00	10.2	1.6

TABLE 6. Simulation results under $\vartheta = 2$. γ_0 : true predictors. $\hat{\gamma}$: selected variables. CC: number of correctly classified variables ($\sum_{j=1}^p \mathbb{I}(\hat{\gamma}_j = \gamma_{0j})$). FP: number of false positives; TP: number of true positives.




FIGURE 6. Sensitivity analysis with $g_\alpha = 0.087$. $P(\theta_i \neq 0 \mid y)$ for $p = 5$, $\vartheta = 2$, $\theta = (0.5, 1, 1.5, 0, 0)$, $n = 100$, $\rho_{ij} = 0.5$. Black circles show the mean.

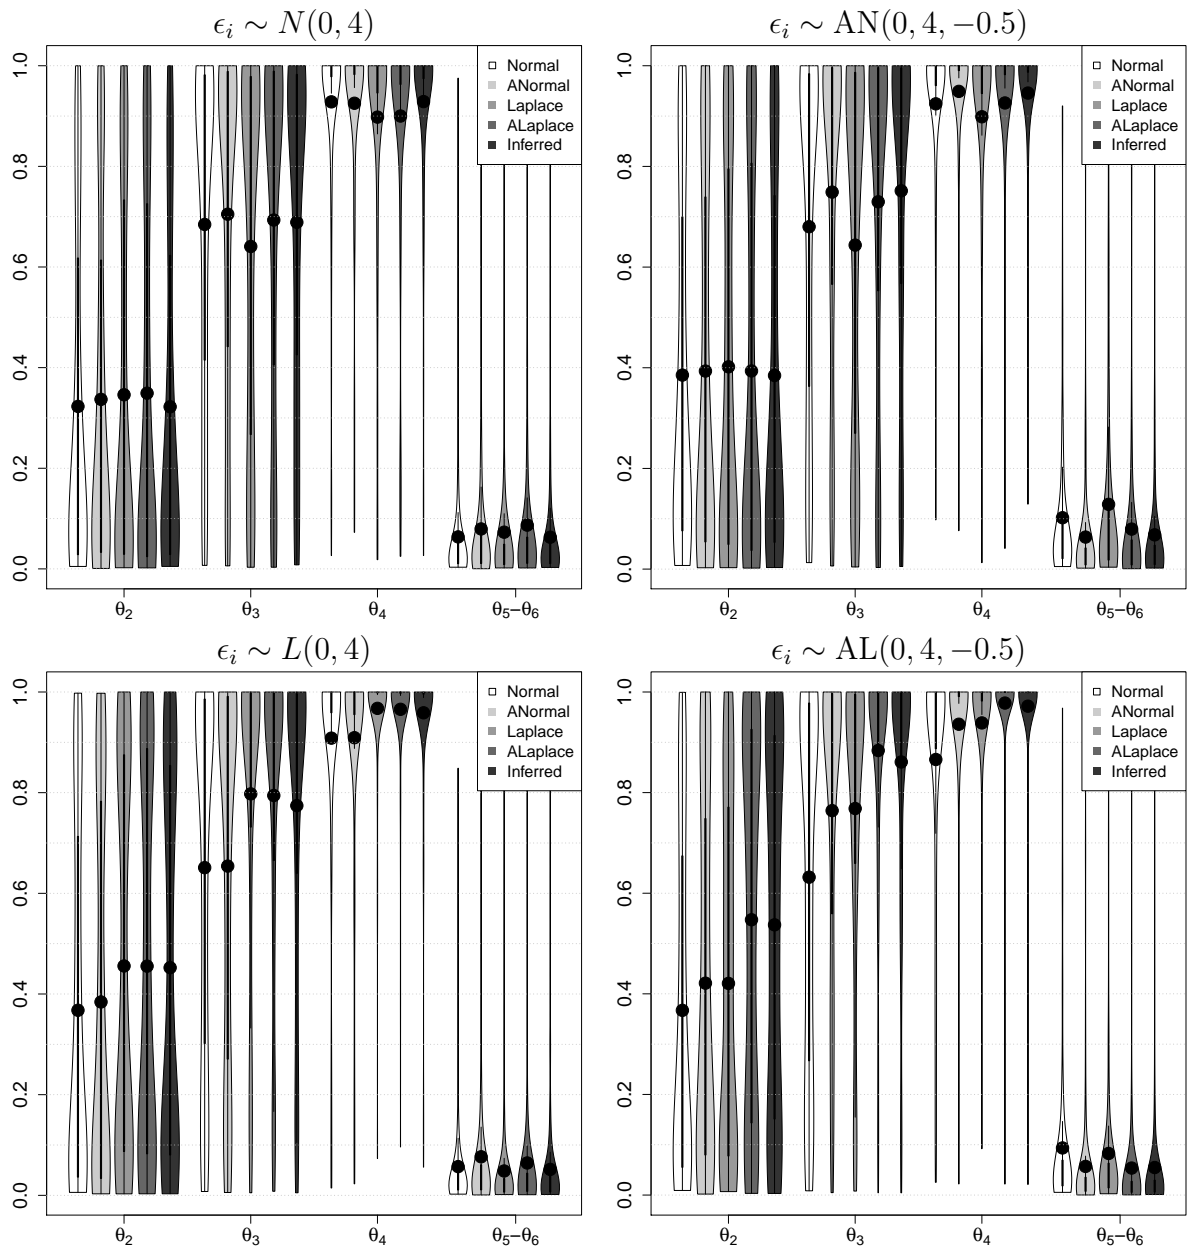


FIGURE 7. Monte Carlo estimates ($B = 10,000$) under $g_\alpha = 0.357$. $P(\theta_i \neq 0 \mid y)$ for $p = 5$, $\vartheta = 2$, $\theta = (0.5, 1, 1.5, 0, 0)$, $n = 100$, $\rho_{ij} = 0.5$. Black circles show the mean.

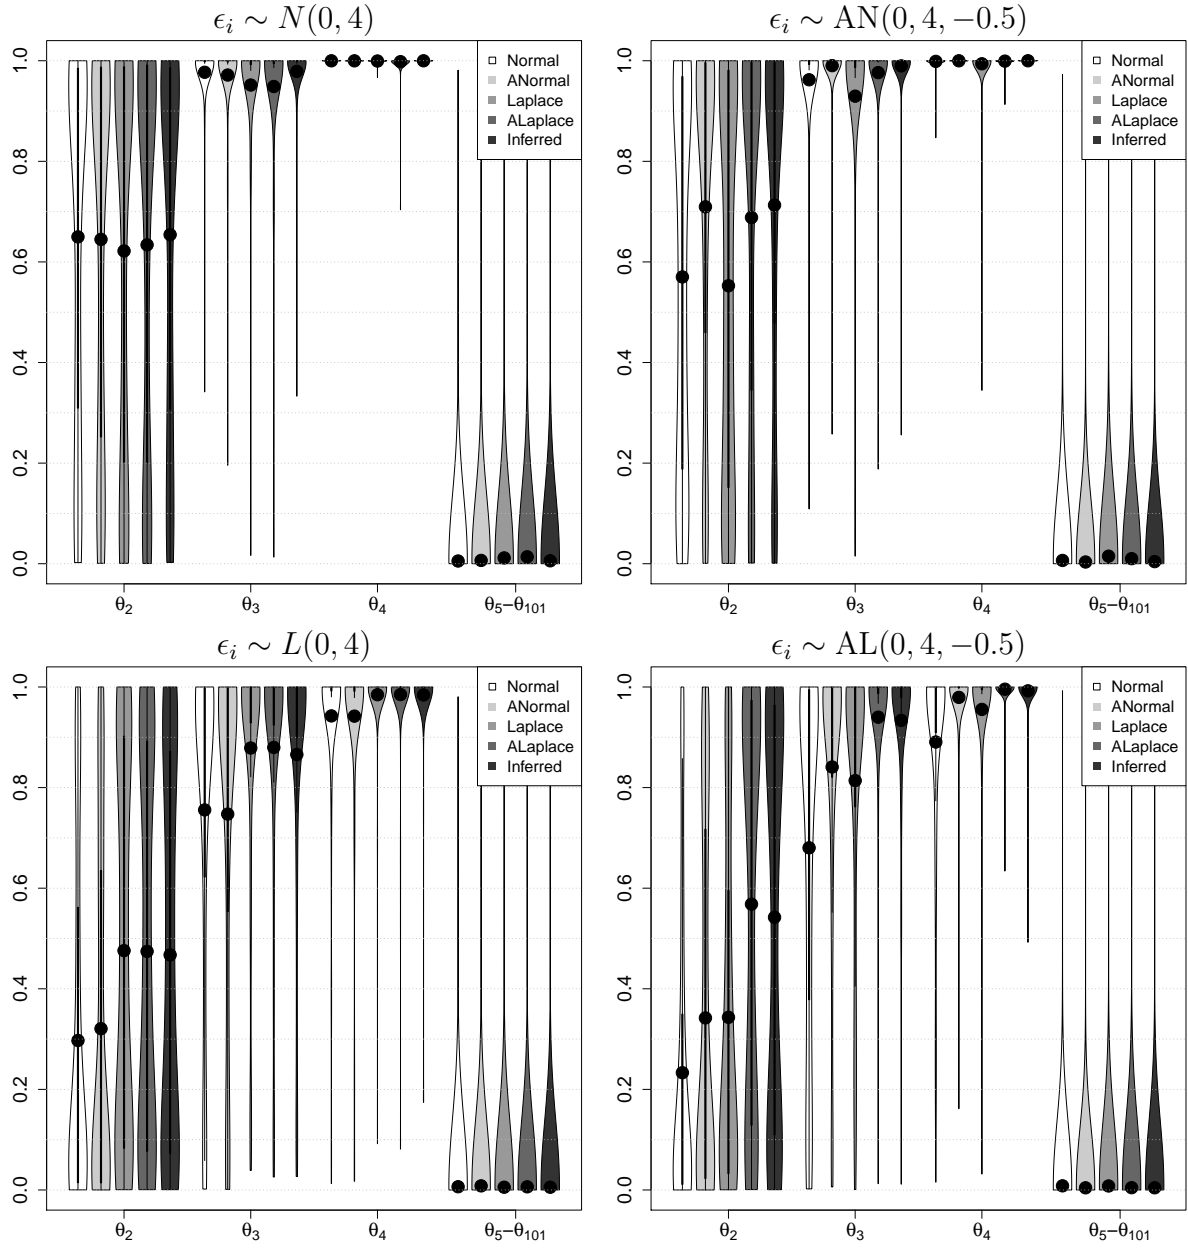


FIGURE 8. $P(\theta_i \neq 0 \mid y)$ for $p = 100$, $\vartheta = 1$, $\theta = (0, 0.5, 1, 1.5, 0, \dots, 0)$, $n = 100$, $\rho_{ij} = 0.5$. Black circles show the mean.

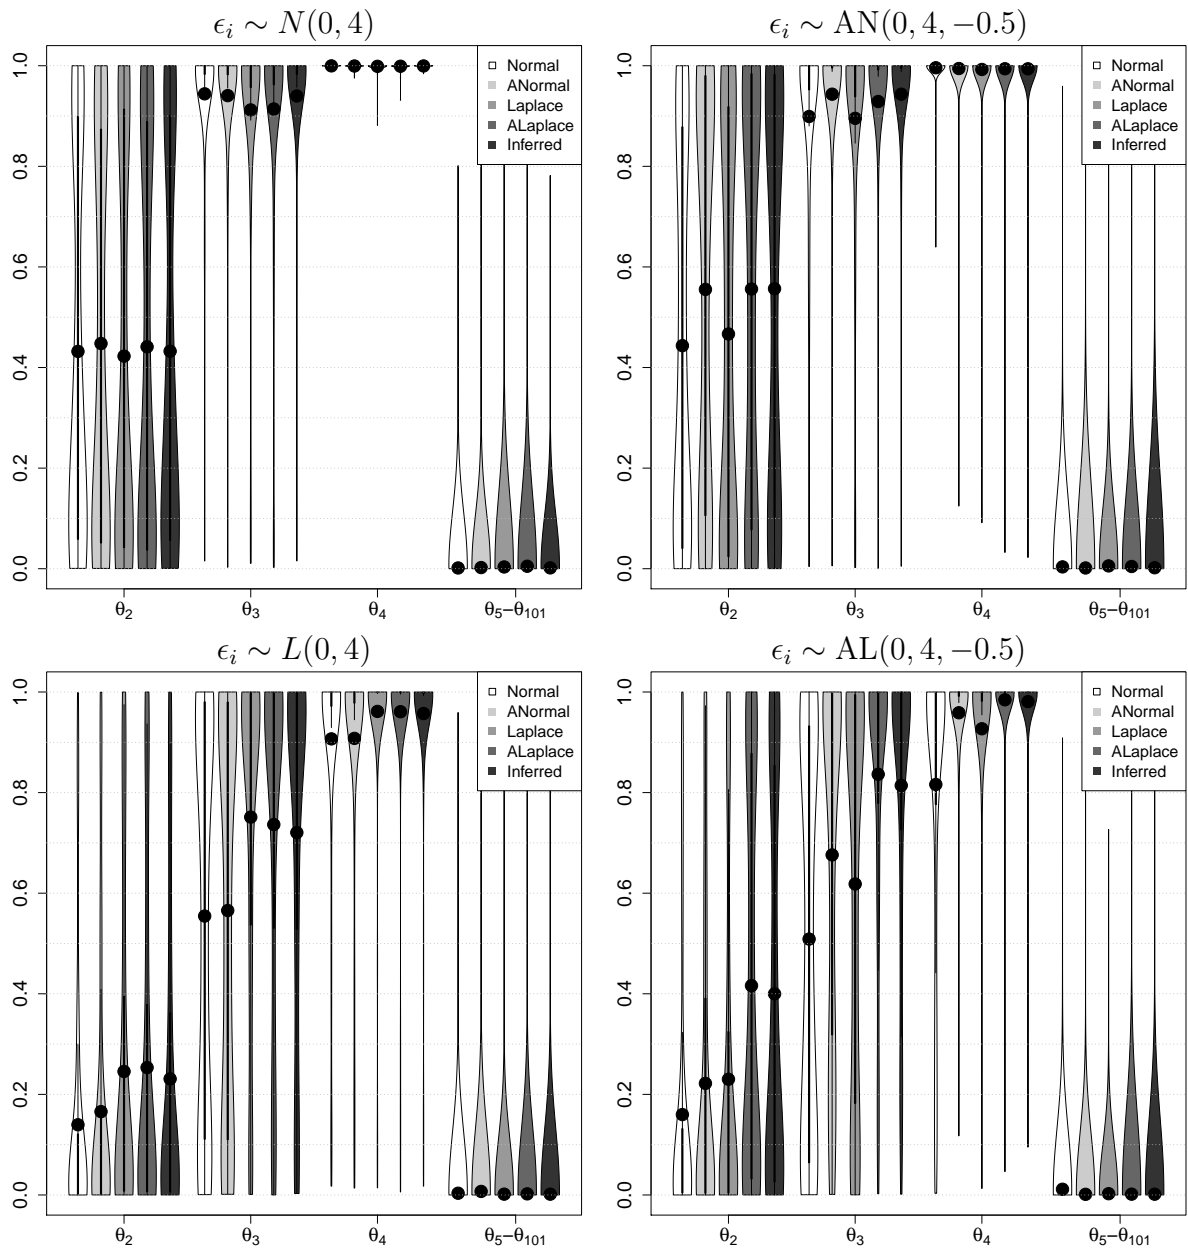


FIGURE 9. $P(\theta_i \neq 0 \mid y)$ for $p = 500$, $\vartheta = 1$, $\theta = (0, 0.5, 1, 1.5, 0, \dots, 0)$, $n = 100$, $\rho_{ij} = 0.5$. Black circles show the mean.

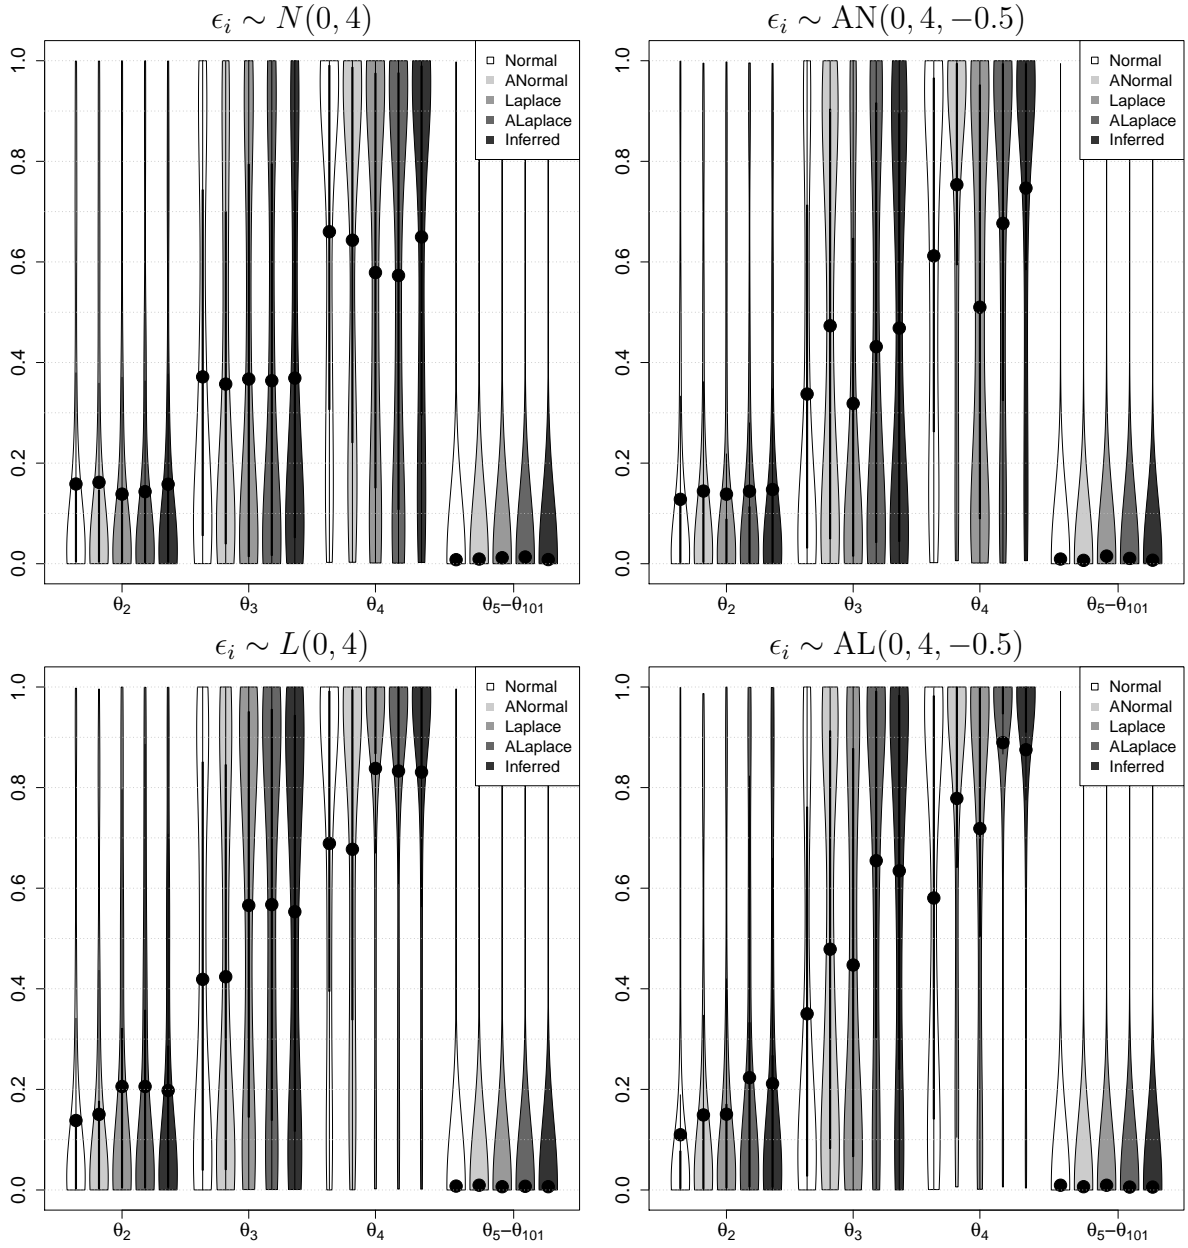


FIGURE 10. $P(\theta_i \neq 0 | y)$ for $p = 100$, $\vartheta = 2$, $\theta = (0, 0.5, 1, 1.5, 0, \dots, 0)$, $n = 100$, $\rho_{ij} = 0.5$. Black circles show the mean.

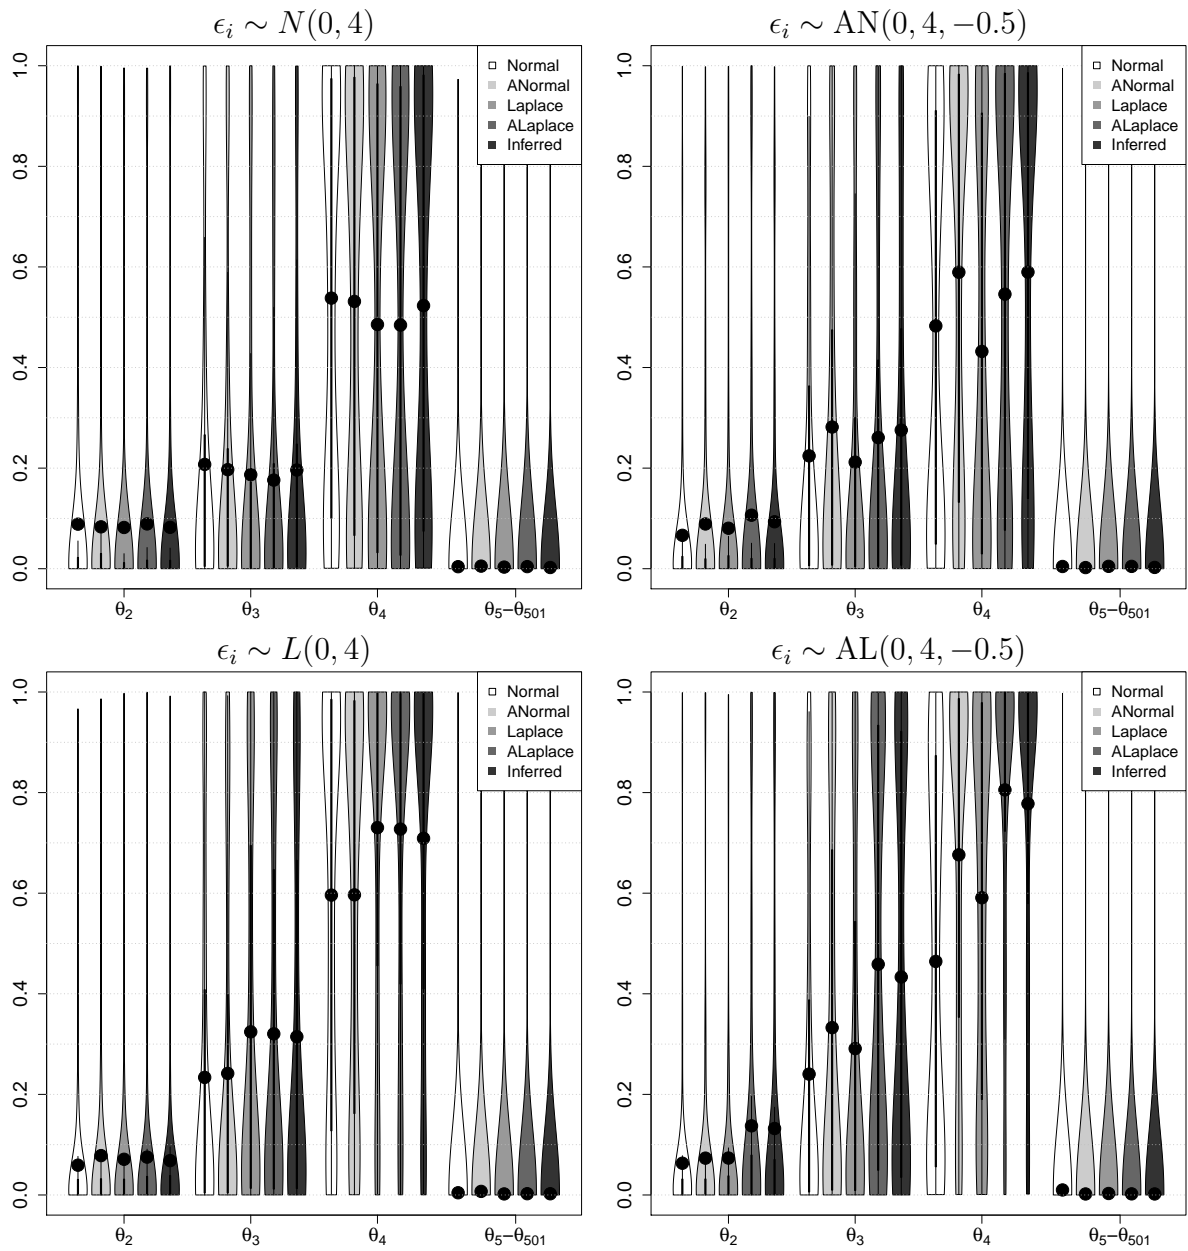


FIGURE 11. $P(\theta_i \neq 0 | y)$ for $p = 500$, $\vartheta = 2$, $\theta = (0, 0.5, 1, 1.5, 0, \dots, 0)$, $n = 100$, $\rho_{ij} = 0.5$. Black circles show the mean.

Gene symbol	Normal	Inferred
C6orf226	1.000	1.000
ECH1	1.000	1.000
CSF2RA	1.000	1.000
RRP1B	0.972	0.999
FBXL19	0.990	0.609
MTMR1	0.144	0.443
SLC35B4	0.324	0.395
RAB3GAP2	0.051	0.035

TABLE 7. Six genes with largest $p(\gamma_j = 1 \mid y)$ in the DLD dataset under assumed normality and inferred error distribution.

REFERENCES

- A. Alfons, C. Croux, and S. Gelper. Sparse least trimmed squares regression for analyzing high-dimensional large data sets. *The Annals of Applied Statistics*, 7(1):226–248, 2013.
- R.B. Arellano-Valle, H.W. Gómez, and F.A. Quintana. Statistical inference for a general class of asymmetric distributions. *Journal of statistical planning and inference*, 128(2):427–443, 2005.
- O. Arslan. Weighted LAD-LASSO method for robust parameter estimation and variable selection in regression. *Computational Statistics and Data Analysis*, 56(6):1952–1965, 2012.
- M.M. Barbieri and J.O. Berger. Optimal predictive model selection. *The Annals of Statistics*, 32(3):870–897, 2004.
- M.J. Bayarri, J.O. Berger, A. Forte, and G. Garcia-Donato. Criteria for Bayesian model choice with application to variable selection. *The Annals of statistics*, 40(3):1550–1577, 2012.
- P. Breheny and J. Huang. Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Annals of Applied Statistics*, 5(1):232–253, 2011.
- A. Calon, E. Espinet, S. Palomo-Ponce, D.V.F. Tauriello, M. Iglesias, M.V. Céspedes, M. Sevillano, C. Nadal, P. Jung, X.H.-F. Zhang, D. Byrom, A. Riera, D. Rossell, R. Mangués, J. Massagué, E. Sancho, and E. Batlle. Dependency of colorectal cancer on a tgf-beta-driven programme in stromal cells for metastasis initiation. *Cancer Cell*, 22(5):571–584, 2012.
- M. Chae, L. Lin, and D.B. Dunson. Bayesian sparse linear regression with unknown symmetric error. *arXiv*, 1608.02143:1–34, 2016.
- Y. Chung and D.B. Dunson. Nonparametric bayes conditional distribution modeling with variable selection. *Journal of the American Statistical Association*, 104(488):1646–1660, 2009.
- H. Dette, C. Ley, and F.J. Rubio. Natural (non-) informative priors for skew-symmetric distributions. *arXiv preprint arXiv:1605.02880*, 2016.
- F. Eicker. Asymptotic normality and consistency of the least squares estimator for families of linear regressions. *Annals of Mathematical Statistics*, 34(2):447–456, 1964.
- J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360, 2001.
- J. Fan, Y. Fan, and E. Barut. Adaptive robust variable selection. *The Annals of Applied Statistics*, 42(1):324–351, 2014.
- C. Fernández and M.F.J. Steel. On Bayesian modeling of fat tails and skewness. *Journal of the American Statistical Association*, 93(441):359–371, 1998.
- I. Gijbels and I. Vrinssen. Robust nonnegative garrote variable selection in linear regression. *Computational Statistics and Data Analysis*, 85:1–22, 2015.
- R. Gottardo and A.E. Raftery. Bayesian robust transformation and variable selection: a unified approach. *The Canadian Journal of Statistics*, 37(3):361–380, 2007.
- P. Grünwald and T. van Ommen. Inconsistency of Bayesian inference for misspecified linear models, and a proposal for repairing it. *arXiv*, 1412.3730:1–70, 2014.

- P.J. Huber. Robust regression: asymptotics, conjectures and Monte Carlo. *The Annals of Statistics*, 1(5):799–821, 1973.
- V.E. Johnson. On numerical aspects of Bayesian model selection in high and ultrahigh-dimensional settings. *Bayesian Analysis*, 8(4):741, 2013.
- V.E. Johnson and D. Rossell. Prior densities for default Bayesian hypothesis tests. *Journal of the Royal Statistical Society B*, 72(2):143–170, 2010.
- V.E. Johnson and D. Rossell. Bayesian model selection in high-dimensional settings. *Journal of the American Statistical Association*, 24(498):649–660, 2012.
- A. C. Kimber. Methods for the two-piece normal distribution. *Communications in Statistics - Theory and Methods*, 14(1):235–245, 1985.
- K. Knight. Asymptotics for l1-estimators of regression parameters under heteroscedasticity. *The Canadian Journal of Statistics*, 27(3):497–507, 1999.
- M. Kocherginsky, X. He, and Y. Mu. Practical confidence intervals for regression quantiles. *Journal of Computational and Graphical Statistics*, 14(1):41–55, 2005.
- G. Koenker and J.A.F. Machado. Goodness of fit and related inference processes for quantile regression. *Journal of the American Statistical Association*, 94(448):1296–1310, 1999.
- R. Koenker. Confidence intervals for regression quantiles. In *Proceedings of the 5th Prague Symposium on Asymptotic Statistics*, pages 349–359. Springer-Verlag, 1994.
- R. Koenker. *Quantile regression*. Cambridge University Press, Cambridge, 2005.
- R. Koenker and G. Bassett. Tests of linear hypotheses and l1 estimation. *Econometrica*, 50(6):1577–1584, 1982.
- S. Kundu and D.B. Dunson. Bayes variable selection in semiparametric linear models. *Journal of the American Statistical Association*, 109(505):437–447, 2014.
- S. Lambert-Lacroix. Robust regression through the huber’s criterion and adaptive lasso penalty. *Electronic Journal of Statistics*, 5:1015–1053, 2011.
- K. Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, 2(2):164–168, 1944.
- P.-L. Loh. Statistical consistency and asymptotic normality for high-dimensional robust M-estimators. *arXiv*, 1501.00312v1:1–56, 2015.
- H. Mallick and Y. Nengjun. Bayesian methods for high dimensional linear models. *Journal of biometrics & biostatistics*, 1:005, 2013.
- D. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal on Applied Mathematics*, 11(2):431–441, 1963.
- S. Mendelson. Learning without concentration for general loss functions. *ArXiv*, 1410.3192, 2014.
- G.S. Mudholkar and A.D. Hutson. The epsilon-skew-normal distribution for analyzing near-normal data. *Journal of statistical planning and inference*, 83(2):291–309, 2000.
- W. K. Newey and J. L. Powell. Asymmetric least squares estimation and testing. *Econometrica*, 55(4):819–847, 1987.
- D. Pollard. Asymptotics for least absolute deviation regression estimators. *Econometric Theory*, 7(2):186–199, 1991.
- R.T. Rockafellar. *Convex analysis*. Princeton university press, Princeton, 2015.

- D. Rossell and D. Telesca. Non-local priors for high-dimensional estimation. *Journal of the American Statistical Association*, (in press):1–33, 2015.
- D. Rossell, D. Telesca, and V.E. Johnson. High-dimensional Bayesian classifiers using non-local priors. In *Statistical Models for Data Analysis XV*, pages 305–314. Springer, 2013.
- D. Rossell, J.D. Cook, D. Telesca, and P. Roebuck. *mombf: Moment and Inverse Moment Bayes Factors*, 2016. URL <https://CRAN.R-project.org/package=mombf>. R package version 1.8.1.
- F.J. Rubio and M.G. Genton. Bayesian linear regression with skew-symmetric error distributions with applications to survival analysis. *Statistics in Medicine*, 35(4):2441–2454, 2016.
- F.J. Rubio and M.F.J. Steel. Inference in two-piece location-scale models with jeffreys priors (with discussion). *Bayesian Analysis*, 9(1):1–22, 2014.
- F.J. Rubio and K. Yu. Flexible objective Bayesian linear regression with applications in survival analysis. *Journal of Applied Statistics*, 2016; in press.
- R.M. Schrader and T.P. Hettmansperger. Robust analysis of variance upon a likelihood ratio criterion. *Biometrika*, 67(1):93–101, 1980.
- J.G. Scott and J.O. Berger. Bayes and empirical Bayes multiplicity adjustment in the variable selection problem. *The Annals of Statistics*, 38(5):2587–2619, 2010.
- M. Shin, A. Bhattacharya, and V.E. Johnson. Scalable Bayesian variable selection using nonlocal prior densities in ultrahigh-dimensional settings. *Texas A&M University (technical report)*, pages 1–33, 2015.
- D.C. Sorensen. Newton’s method with a model trust region modification. *SIAM Journal of Numerical Analysis*, 19(2):409–426, 1982.
- M.S. Srivastava. On fixed width confidence bounds for regression parameters. *The Annals of Mathematical Statistics*, 42(4):1403–1411, 1971.
- R. Tibshirani. Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society, B*, 58(1):267–288, 1996.
- A.W. van der Vaart. *Asymptotic statistics*. Cambridge University Press, New York, 1998.
- K.F. Wallis. The two-piece normal, binormal, or double Gaussian distribution: its origin and rediscoveries. *Statistical Science*, 29(1):106–112, 2014.
- H. Wang, G. Li, and G. Jiang. Robust regression shrinkage and consistent variable selection through the LAD-LASSO. *Journal of Business & Economic Statistics*, 25(3):347–355, 2007.
- L. Wang and R. Li. Weighted Wilcoxon-type smoothly clipped absolute deviation method. *Biometrics*, 65(2):564–571, 2009.
- L. Wang, Y. Tang, D. Sinha, D. Pati, and S. Lipsitz. Bayesian variable selection for skewed heteroscedastic response. *arXiv preprint arXiv:1602.09100*, 2016.
- Y. Yan and A. Kottas. A new family of error distributions for Bayesian quantile regression. Technical report, University of California Santa Cruz, 2015.
- K. Yu, C.W.S. Chen, C. Reed, and D.B. Dunson. Bayesian variable selection in quantile regression. *Statistics and its interface*, 6(2):261–274, 2013.

T. Yuan, X. Huang, M. Woodcock, M. Du, R. Dittmar, Y. Wang, S. Tsai, M. Kohli, L. Boardman, T. Patel, and L. Wang. Plasma extracellular rna profiles in healthy and cancer patients. *Scientific reports*, 6:1–11, 2016.

UNIVERSITAT POMPEU FABRA, DEPARTMENT OF BUSINESS AND ECONOMICS, BARCELONA (SPAIN)

LONDON SCHOOL OF HYGIENE & TROPICAL MEDICINE, LONDON (UK)