Expected length of post-model-selection confidence intervals conditional on polyhedral constraints

Danijel Kivaranovic

Department of Statistics and Operations Research, University of Vienna,

and

Hannes Leeb

Department of Statistics and Operations Research, University of Vienna

January 25, 2018

Abstract

Valid inference after model selection is currently a very active area of research. The polyhedral method, pioneered by Lee et al. (2016), allows for valid inference after model selection if the model selection event can be described by polyhedral constraints. In that reference, the method is exemplified by constructing two valid confidence intervals when the Lasso estimator is used to select a model. We here study the expected length of these intervals. For one of these confidence intervals, that is easier to compute, we find that its expected length is always infinite. For the other of these confidence intervals, whose computation is more demanding, we give a necessary and sufficient condition for its expected length to be infinite. In simulations, we find that this condition is typically satisfied.

Keywords: Lasso, inference after model selection, hypothesis test.

1 Introduction

Lee et al. (2016) recently introduced a new technique for valid inference after model selection, the so-called polyhedral method. Using this method, and using the Lasso for model selection in linear regression, Lee et al. (2016) derived two new confidence sets that are valid conditional on the outcome of the model selection step. More precisely, let \hat{m} denote the model containing those regressors that correspond to non-zero coefficients of the Lasso estimator, and let \hat{s} denote the sign-vector of those non-zero Lasso coefficients. Then Lee et al. (2016) constructed confidence intervals $[L_{\hat{m},\hat{s}}, U_{\hat{m},\hat{s}}]$ and $[L_{\hat{m}}, U_{\hat{m}}]$ whose coverage probability is $1 - \alpha$, conditional on the events $\{\hat{m} = m, \hat{s} = s\}$ and $\{\hat{m} = m\}$, respectively (provided that the probability of the conditioning event is positive). The computational effort in constructing these intervals is considerably lighter for $[L_{\hat{m},\hat{s}}, U_{\hat{m},\hat{s}}]$. In simulations, Lee et al. (2016) noted that this latter interval can be quite long in some cases; cf. Figure 10 in that reference. We here analyze the (conditional) expected length of these intervals.

1.1 Overview of findings

Throughout, we use the same setting and assumptions as Lee et al. (2016). In particular, we assume that the response vector is distributed as $N(\mu, \sigma^2 I_n)$ with unknown mean $\mu \in \mathbb{R}^n$ and known variance $\sigma^2 > 0$ (our results carry over to the unknown-variance case; see the end of Section 3), and that the non-stochastic regressor matrix has columns in general position. Write $\mathbb{P}_{\mu,\sigma^2}$ and $\mathbb{E}_{\mu,\sigma^2}$ for the probability measure and the expectation operator, respectively, corresponding to $N(\mu, \sigma^2 I_n)$.

For the interval $[L_{\hat{m},\hat{s}}, U_{\hat{m},\hat{s}}]$, we find the following: Fix a non-empty model m, a sign-vector s, as well as $\mu \in \mathbb{R}^n$ and $\sigma^2 > 0$. If $\mathbb{P}_{\mu,\sigma^2}(\hat{m} = m, \hat{s} = s) > 0$, then

$$\mathbb{E}_{\mu,\sigma^2} \left[U_{\hat{m},\hat{s}} - L_{\hat{m},\hat{s}} \right| \hat{m} = m, \hat{s} = s \right] = \infty.$$
 (1)

Obviously, this statement continues to hold if the event $\hat{m} = m, \hat{s} = s$ is replaced by the larger event $\hat{m} = m$ throughout (because there are only finitely many possible values for the sign-vector s). And this statement continues to hold if the condition $\mathbb{P}_{\mu,\sigma^2}(\hat{m} = m, \hat{s} = s) > 0$ is dropped and the conditional expectation in (1) is replaced by the unconditional one.

For the interval $[L_{\hat{m}}, U_{\hat{m}}]$, we derive a necessary and sufficient condition for its expected length to be infinite, conditional on the event $\hat{m} = m$. That condition depends on the regressor matrix, on the model m and also on a linear contrast that defines the quantity of interest, and is daunting to verify in all but the most basic examples. We also provide a sufficient condition for infinite expected length that is easy to verify. In simulations, we find that this sufficient condition for infinite expected length is typically satisfied when the model m excludes a significant portion of all the available regressors (e.g., if the selected model is 'sparse'). And even if the model m is not sparse, we find that this condition is still satisfied for a sizable fraction of the linear contrasts that define the quantity of interest. See Table 1 and the attending discussion for more detail.

The methods developed in this paper can also be used if the Lasso, as the model selector, is replaced by any other procedure that allows for application of the polyhedral method. In particular, we see that confidence intervals based on the polyhedral method in Gaussian regression can have infinite expected length. Our findings suggest that the expected length of confidence intervals based on the polyhedral method should be closely scrutinized, in Gaussian regression but also in non-Gaussian settings and in other variations of the polyhedral method.

The rest of the paper is organized as follows: We conclude this section by discussing a number of related results that put our findings in context. Section 2 describes the confidence intervals of Lee et al. (2016) in detail and introduces some notation. Section 3 contains

two core results, Propositions 1 and 2, which entail our main findings, the simulation study mentioned earlier, as well as a discussion of the unknown variance case. The appendix contains the proofs of our core results and some auxiliary lemmas.

1.2 Context and related results

There are currently several exciting ongoing developments based on the polyhedral method, not least because it proved to be applicable to more complicated settings, and there are several generalization of this framework. See, among others, Tibshirani et al. (2016), Taylor & Tibshirani (2017), Tian & Taylor (2015). Certain optimality results of the method of Lee et al. (2016) are given in Fithian et al. (2017). Using a different approach, Berk et al. (2013) proposed the so-called PoSI-intervals which are unconditionally valid. A benefit of the PoSI-intervals is that they are valid after selection with any possible model selector, instead of a particular one like the Lasso; however, as a consequence, the PoSI-intervals are typically very conservative (that is, the actual coverage probability is above the nominal level). Nonetheless, Bachoc et al. (2016) showed in a Monte Carlo simulation that, in certain scenarios, the PoSI-intervals can be shorter than the intervals of Lee et al. (2016). The results of the present paper are based on the first author's master's thesis.

It is important to note that all confidence sets discussed so far are non-standard, in the sense that the parameter to be covered is not the true parameter in an underlying correct model (or components thereof), but instead is a model-dependent quantity of interest. (See Section 2 for details and the references in the preceding paragraph for more extensive discussions.) An advantage of this non-standard approach is that it does not rely on the assumption that any of the candidate models is correct. Valid inference for an underlying true parameter is a more challenging task, as demonstrated by the impossibility results in Leeb & Pötscher (2006 a, b, 2008). There are several proposals of valid confidence intervals after model selection (in the sense that the actual coverage probability of the true parameter is at or above the nominal level) but these are rather large compared to the standard confidence intervals from the full model (supposing that one can fit the full model); see Pötscher (2009), Pötscher & Schneider (2010), Schneider (2016). In fact, Leeb & Kabaila (2017) showed that the usual confidence interval obtained by fitting the full model is admissible also in the unknown variance case; therefore, one cannot obtain uniformly smaller valid confidence sets for a component of the true parameter by any other method.

2 Assumptions and confidence intervals

Let Y denote the $N(\mu, \sigma^2 I_n)$ -distributed response vector, $n \ge 1$, where $\mu \in \mathbb{R}^n$ is unknown and $\sigma^2 > 0$ is known. Let $X = (x_1, \ldots, x_p)$, $p \ge 1$, with $x_i \in \mathbb{R}^n$ for each $i = 1, \ldots, p$, be the non-stochastic $n \times p$ regressor matrix. We assume that the columns of X are in general position (this mild assumption is further discussed in the following paragraph). The full model $\{1, \ldots, p\}$ is denoted by m_F . All subsets of the full model are collected in \mathcal{M} , that is, $\mathcal{M} = \{m : m \subseteq m_F\}$. The cardinality of a model m is denoted by |m|. For any $m = \{i_1, \ldots, i_k\} \in \mathcal{M} \setminus \emptyset$ with $i_1 < \cdots < i_k$, we set $X_m = (x_{i_1}, \ldots, x_{i_k})$. Analogously, for any vector $v \in \mathbb{R}^p$, we set $v_m = (v_{i_1}, \ldots, v_{i_k})'$. If m is the empty model, then X_m is to be interpreted as the zero vector in \mathbb{R}^n and v_m as 0.

The Lasso estimator, denoted by $\hat{\beta}(y)$, is a minimizer of the least squares problem with an additional penalty on the absolute size of the regression coefficients (Frank & Friedman 1993, Tibshirani 1996):

$$\min_{\boldsymbol{\beta}\in\mathbb{R}^p} \frac{1}{2} \|\boldsymbol{y} - \boldsymbol{X}\boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1, \quad \boldsymbol{y}\in\mathbb{R}^n, \ \lambda > 0.$$

The Lasso has the property that some coefficients of $\hat{\beta}(y)$ are zero with positive probability. A minimizer of the Lasso objective function always exists, but it is not necessarily unique.

Uniqueness of $\hat{\beta}(y)$ is guaranteed here by our assumption that the columns of X are in general position (Tibshirani 2013). This assumption is relatively mild; e.g., if the entries of the matrix X are drawn from a (joint) distribution that has a Lebesgue density, then the columns of X are in general position with probability 1 (Tibshirani 2013). The model $\hat{m}(y)$ selected by the Lasso and the sign-vector $\hat{s}(y)$ of non-zero Lasso coefficients can now formally be defined through

$$\hat{m}(y) = \left\{ j : \hat{\beta}_j(y) \neq 0 \right\}$$
 and $\hat{s}(y) = \operatorname{sign}\left(\hat{\beta}_{\hat{m}(y)}(y)\right).$

Recall that \mathcal{M} denotes the set of all possible submodels and set $\mathcal{S}_m = \{-1, 1\}^{|m|}$ for each $m \in \mathcal{M}$. For later use we also denote by \mathcal{M}^+ and \mathcal{S}_m^+ the collection of non-empty models and the collection of corresponding sign-vectors, that occur with positive probability, i.e.,

$$\mathcal{M}^{+} = \{ m \in \mathcal{M} \setminus \emptyset : \mathbb{P}_{\mu,\sigma^{2}}(\hat{m}(Y) = m) > 0 \},$$
$$\mathcal{S}_{m}^{+} = \{ s \in \mathcal{S}_{m} : \mathbb{P}_{\mu,\sigma^{2}}(\hat{m}(Y) = m, \hat{s}(Y) = s) > 0 \} \qquad (m \in \mathcal{M} \setminus \emptyset).$$

These sets do not depend on μ and σ^2 as the measure $\mathbb{P}_{\mu,\sigma^2}$ is equivalent to Lebesgue measure with respect to null sets. Also, our assumption that the columns of X are in general position guarantees that \mathcal{M}^+ only contains models m for which X_m has columnrank m (Tibshirani 2013).

Inference is focused on a non-standard, model dependent, quantity of interest. Throughout the following, fix $m \in \mathcal{M}^+$ and let

$$\beta^m = \mathbb{E}_{\mu,\sigma^2} \left[(X'_m X_m)^{-1} X'_m Y \right] = (X'_m X_m)^{-1} X'_m \mu$$

For $\gamma^m \in \mathbb{R}^{|m|} \setminus \{0\}$, the goal is to construct a confidence interval for $\gamma^{m'}\beta^m$ with conditional coverage probability $1 - \alpha$ on the event $\{\hat{m} = m\}$. Clearly, the quantity of interest can also be written as $\gamma^{m'}\beta^m = \eta^{m'}\mu$ for $\eta^m = X_m(X'_mX_m)^{-1}\gamma^m$. For later use, write P_{η^m} for the orthogonal projection on the space spanned by η^m .

At the core of the polyhedral method lies the observation that the event where $\hat{m} = m$ and where $\hat{s} = s$ describes a convex polytope in sample space \mathbb{R}^n (up to a Lebesgue null set): For each $m \in \mathcal{M}^+$ and each $s \in \mathcal{S}_m$, we have

$$\{y: \ \hat{m}(y) = m, \hat{s}(y) = s\} \stackrel{\text{a.s.}}{=} \{y: \ A_{m,s}y < b_{m,s}\},$$
(2)

cf. Theorem 3.3 in Lee et al. (2016) (explicit formulas for the matrix $A_{m,s}$ and the vector $b_{m,s}$ are also repeated in Appendix C in our notation). Fix $z \in \mathbb{R}^n$ orthogonal to η^m . Then the set of y satisfying $(I_n - P_{\eta^m})y = z$ and $A_{m,s}y < b$ is either empty or a line segment. In either case, that set can be written as $\{z + \eta^m w : \mathcal{V}_{m,s}^-(z) < w < \mathcal{V}_{m,s}^+(z)\}$. The endpoints satisfy $-\infty \leq \mathcal{V}_{m,s}^-(z) \leq \mathcal{V}_{m,s}^+(z) \leq \infty$ (see Lemma 4.1 of Lee et al. 2016; formulas for these quantities are also given in Appendix C in our notation). Now decompose Y into the sum of two independent Gaussians $P_{\eta^m}Y$ and $(I_n - P_{\eta^m})Y$, where the first one is a linear function of $\eta^{m'}Y \sim N(\eta^{m'}\mu, \sigma^2\eta^{m'}\eta^m)$. With this, the conditional distribution of $\eta^{m'}Y$, conditional on the event that $\hat{m}(Y) = m$, $\hat{s}(Y) = s$ and $(I_n - P_{\eta^m})(Y) = z$, is the conditional $N(\eta^{m'}\mu, \sigma^2\eta^{m'}\eta^m)$ -distribution, conditional on the set $(\mathcal{V}_{m,s}^-(z), \mathcal{V}_{m,s}^+(z))$ (in the sense that the latter conditional distribution is a regular conditional distribution if one starts with the conditional distribution of $\eta^{m'}Y$ given $\hat{m} = m$ and $\hat{s} = s$ – which is always well-defined – and if one then conditions on the random variable $(I_n - P_{\eta^m})Y)$.

To use these observations for the construction of confidence sets, consider first the conditional distribution of a random variable $W \sim N(\theta, \varsigma^2)$ conditional on the event $W \in T$, where $\theta \in \mathbb{R}$, where $\varsigma^2 > 0$, and where $T \subseteq \mathbb{R}$ is the union of finitely many non-empty open intervals. Write $F_{\theta,\varsigma^2}^T(\cdot)$ for the cumulative distribution function (c.d.f) of W given $W \in T$. The corresponding law can be viewed as a 'truncated normal' distribution and will be denoted by $TN(\theta, \varsigma^2, T)$ in the following. A confidence interval for θ with coverage probability $1 - \alpha$ conditional on the event $W \in T$ is obtained by the usual method of

collecting all values θ_0 for which a hypothesis test of $H_0: \theta = \theta_0$ against $H_1: \theta \neq \theta_0$ does not reject. In particular, for $w \in \mathbb{R}$, define L(w) and U(w) through

$$F_{L(w),\varsigma^2}^T(w) = 1 - \frac{\alpha}{2}$$
 and $F_{U(w),\varsigma^2}^T(w) = \frac{\alpha}{2}$

(these quantities are well-defined in view of Lemma A.2). With this, we have $P(\theta \in [L(W), U(W)]|W \in T) = 1 - \alpha$, irrespective of $\theta \in \mathbb{R}$.

Fix $m \in \mathcal{M}^+$ and $s \in \mathcal{S}_m^+$, and let $\sigma_m^2 = \sigma^2 \eta^{m\prime} \eta^m$ and $T_{m,s}(z) = (\mathcal{V}_{m,s}^-(z), \mathcal{V}_{m,s}^+(z))$ for z orthogonal to η^m . With this, we have

$$\eta^{m'}Y \Big| \{ \hat{m} = m, \hat{s} = s, (I_n - P_{\eta^m})Y = z \} \sim TN(\eta^{m'}\mu, \sigma_m^2, T_{m,s}(z)).$$

for each $z \in \{(I_n - P_{\eta^m})y : A_{m,s}y < b_{m,s}\}$. Now define $L_{m,s}(y)$ and $U_{m,s}(y)$ through

$$F_{L_{m,s}(y),\sigma_m^2}^{T_{m,s}((I_n - P_{\eta^m})y)}(\eta^{m'}y) = 1 - \frac{\alpha}{2} \quad \text{and} \quad F_{U_{m,s}(y),\sigma_m^2}^{T_{m,s}((I_n - P_{\eta^m})y)}(\eta^{m'}y) = \frac{\alpha}{2}$$

for each y so that $A_{m,s}y < b_{m,s}$. Then, by construction, the random interval $[L_{m,s}(Y), U_{m,s}(Y)]$ covers $\eta^{m'}\mu$ with probability $1 - \alpha$ conditional on the event that $\hat{m} = m$ and $\hat{s} = s$.

In a similar fashion, fix $m \in \mathcal{M}^+$, set $T_m(z) = \bigcup_{s \in \mathcal{S}_m^+} T_{m,s}(z)$ for z orthogonal to η^m , and define $L_m(y)$ and $U_m(y)$ through

$$F_{L_m(y),\sigma_m^2}^{T_m((I_n - P_\eta m)y)}(\eta^{m'}y) = 1 - \frac{\alpha}{2} \quad \text{and} \quad F_{U_m(y),\sigma_m^2}^{T_m((I_n - P_\eta m)y)}(\eta^{m'}y) = \frac{\alpha}{2}.$$

Again by construction, the random interval $[L_m(Y), U_m(Y)]$ covers $\eta^{m'}\mu$ with probability $1 - \alpha$ conditional on the event that $\hat{m} = m$.

Remark. (i) If $\tilde{m} = \tilde{m}(y)$ is any other model selection procedure, so that the event $\{y : \tilde{m} = m\}$ is the union of a finite number of polyhedra (up to null sets), then the polyhedral method can be applied to obtain a confidence set for $\eta^{m'}\mu$ with conditional coverage probability $1 - \alpha$, conditional on the event $\{\tilde{m} = m\}$, if that event has positive

probability. Indeed, for such a model selection procedure, the arguments following (1) also apply, mutatis mutandis.

(ii) So far, we have defined confidence intervals only on the events $W \in T$, $m \in \mathcal{M}^+$ and $s_m \in \mathcal{S}_m^+$, and $m \in \mathcal{M}^+$, respectively. In the remaining cases, the interval endpoints (and the corresponding quantity of interest) can be chosen arbitrarily (measurable) without affecting our results. It is easy to choose constructions so that one obtains meaningful confidence intervals that are defined everywhere in sample space.

(iii) In Theorem 3.3 of Lee et al. (2016), relation (2) is stated as an equality, not as an equality up to null sets, and with the right-hand side replaced by $\{y : A_{m,s}y \leq b_{m,s}\}$ (in our notation). Because (2) differs from this only on a Lebesgue null set, the difference is inconsequential for the purpose of the present paper. The statement in Lee et al. (2016) is based on the fact that \hat{m} was defined as the equicorrelation set (Tibshirani 2013) in that paper. But if \hat{m} is the equicorrelation set, then there can exist vectors $y \in \{\hat{m} = m\}$ such that some coefficients of $\hat{\beta}(y)$ are zero, which clashes with the idea that \hat{m} contains those variables whose Lasso coefficients are non-zero. However, for any $m \in \mathcal{M}^+$, the set of such ys is a Lebesgue null set.

3 Core results

We first analyze the simple confidence set [L(W), U(W)] that was introduced in the preceding section, which covers θ with probability $1 - \alpha$, conditional on $W \in T$, where $W \sim N(\theta, \varsigma^2)$. By assumption, T is of the form $T = \bigcup_{i=1}^{K} (a_i, b_i)$ where $K < \infty$ and $-\infty \leq a_1 < b_1 < \cdots < a_K < b_K \leq \infty$. Figure 1 exemplifies the length of [L(w), U(w)]when T is bounded (left panel) and when T is unbounded (right panel). The dashed line is the length of the standard (unconditional) confidence interval for θ . In the left panel, we see that the length of [L(w), U(w)] diverges as w approaches the far left or the far right boundary point of the truncation set (i.e., -3 and 3). On the other hand, in the right panel we see that the length of [L(w), U(w)] is bounded and converges to the length of the standard interval as $|w| \to \infty$.



Figure 1: Length of the interval [L(w), U(w)] for the case where $T = (-3, -2) \cup (-1, 1) \cup (2, 3)$ (left panel) and the case where $T = (-\infty, -2) \cup (-1, 1) \cup (2, \infty)$ (right panel). In both cases, we took $\varsigma^2 = 1$ and $\alpha = 0.05$.

Write $\Phi(w)$ and $\phi(w)$ for the c.d.f. and p.d.f. of the standard normal distribution, respectively, where we adopt the usual convention that $\Phi(-\infty) = 0$ and $\Phi(\infty) = 1$.

Proposition 1. If T is bounded either from above or from below, then

$$E[U(W) - L(W)|W \in T] = \infty.$$

If T is unbounded from above and from below, then

$$\frac{U(W) - L(W)}{\varsigma} \stackrel{a.s.}{\leq} 2\Phi^{-1}(1 - p_*\alpha/2) \\ \leq 2\Phi^{-1}(1 - \alpha/2) + \frac{a_K - b_1}{\varsigma},$$

where $p_* = \inf_{\vartheta \in \mathbb{R}} P(N(\vartheta, \varsigma^2) \in T)$ and where $a_K - b_1$ is to be interpreted as 0 in case K = 1. [The first inequality trivially continues to hold if T is bounded, as then $p_* = 0$.]

Intuitively, one expects confidence intervals to be wide if one conditions on a bounded set because extreme values cannot be observed on a bounded set and the confidence intervals have to take this into account. We find that the conditional expected length is infinite in this case. If, for example, T is bounded from below, i.e., if $-\infty < a_1$, then first statement in the proposition follows from two facts: First, the length of U(w) - L(w) behaves like $1/(w - a_1)$ as w approaches a_1 from above; and, second, the p.d.f. of the truncated normal distribution at w is bounded away from 0 zero as w approaches a_1 from above. See the proof in Section B for a more detailed version of this argument. On the other hand, if the truncation set is unbounded, extreme values are observable and confidence intervals, therefore, do not have to be extremely wide. The second upper bound provided by the proposition for that case will be useful later.

We see that the boundedness of the truncation set T is critical for the interval length. When the Lasso is used as a model selector, this prompts the question whether the truncation sets $T_{m,s}(z)$ and $T_m(z)$ are bounded or not, because the intervals $[L_{m,s}(y), U_{m,s}(y)]$ and $[L_m(y), U_m(y)]$ are obtained from conditional normal distributions with truncation sets $T_{m,s}((I_n - P_{\eta^m})y)$ and $T_m((I_n - P_{\eta^m})y)$, respectively. For $m \in \mathcal{M}^+$, $s \in \mathcal{S}_m^+$, and z orthogonal to η^m , recall that $T_{m,s}(z) = (\mathcal{V}_{m,s}^-(z), \mathcal{V}_{m,s}^+(z))$, and that $T_m(z)$ is the union of these intervals over $s \in \mathcal{S}_m^+$. Write $[\eta^m]^{\perp}$ for the orthogonal complement of the span of η^m .

Proposition 2. For each $m \in \mathcal{M}^+$ and each $s \in \mathcal{S}_m$, we have

$$\forall z \in [\eta^m]^{\perp}: \quad -\infty < \mathcal{V}^-_{m,s}(z) \qquad or \quad \forall z \in [\eta^m]^{\perp}: \quad \mathcal{V}^+_{m,s}(z) < \infty$$

or both.

For the confidence interval $[L_{\hat{m},\hat{s}}(Y), U_{\hat{m},\hat{s}}(Y)]$, the statement in (1) now follows immediately: If m is a non-empty model and s is a sign-vector so that the event $\{\hat{m} = m, \hat{s} = s\}$ has positive probability, then $m \in \mathcal{M}^+$ and $s \in \mathcal{S}_m^+$. Now Proposition 2 entails that $T_{m,s}((I_n - P_{\eta^m})Y)$ is almost surely bounded on the event $\{\hat{m} = m, \hat{s} = s\}$, and Proposition 1 entails that (1) holds.

For the confidence interval $[L_{\hat{m}}(Y), U_{\hat{m}}(Y)]$, we obtain that its conditional expected length is finite, conditional on $\hat{m} = m$ with $m \in \mathcal{M}^+$, if and only if its corresponding truncation set $T_m(Y)$ is almost surely unbounded from above and from below on that event. More precisely, for $m \in \mathcal{M}^+$, we have

$$\mathbb{E}_{\mu,\sigma^2}[U_{\hat{m}}(Y) - L_{\hat{m}}(Y)|\hat{m} = m] = \infty$$
(3)

if and only if there exists a $s \in \mathcal{S}_m^+$ and a vector y satisfying $A_{m,s}y < b_{m,s}$, so that

$$T_m((I_n - P_{\eta^m})y)$$
 is bounded from above or from below. (4)

Before proving this equivalence, recall that $T_m((I_n - P_{\eta^m})y)$ is the union of the intervals $(\mathcal{V}_{m,s}^-((I_n - P_{\eta^m})y), \mathcal{V}_{m,s}^+((I_n - P_{\eta^m})y))$ with $s \in \mathcal{S}_m^+$. Inspection of the explicit formulas for the interval endpoints given in Appendix C now immediately reveals the following: The lower endpoint $\mathcal{V}_{m,s}^-((I_n - P_{\eta^m})y)$ is either constant equal to $-\infty$ on the set $\{y : A_{m,s}y < b_{m,s}\}$, or it is the minimum of a finite number of linear functions of y (and hence finite and continuous) on that set. Similarly the upper endpoint $\mathcal{V}_{m,s}^+((I_n - P_{\eta^m})y)$ is either constant

equal to ∞ on that set, or it is the maximum of a finite number of linear functions of y (and hence finite and continuous) on that set.

To prove the equivalence, we first assume, for some s and y with $s \in S_m^+$ and $A_{m,s}y < b_{m,s}$, that the set in (4) is bounded from above (the case of boundedness from below is similar). Then there is an open neighborhood O of y, so that each point $w \in O$ satisfies $A_{m,s}w < b_{m,s}$ and also so that $T_m((I_n - P_{\eta^m})w)$ is bounded from above. Because O has positive Lebesgue measure, (3) now follows from Proposition 1. To prove the converse, assume for each $s \in S_m^+$ and each y satisfying $A_{m,s}y < b_{m,s}$ that $T_m((I_n - P_{\eta^m})y)$ is unbounded from above and from below. Because the sets $\{y : A_{m,s}y < b_{m,s}\}$ for $s \in S_m^+$ are disjoint by construction, the same is true for the sets $T_{m,s}((I_n - P_{\eta^m})y)$ for $s \in S_m^+$. Using Proposition 1, we then obtain that $U_{\hat{m}}(Y) - L_{\hat{m}}(Y)$ is bounded by a linear function of

$$\max\{\mathcal{V}_{m,s}^{-}((I_n - P_{\eta^m})Y) : s \in \mathcal{S}_m^+\} - \min\{\mathcal{V}_{m,s}^{+}((I_n - P_{\eta^m})Y) : s \in \mathcal{S}_m^+\}$$

Lebesgue almost everywhere on the event $\{\hat{m} = m\}$. (The maximum and the minimum in the preceding display correspond to a_K and b_1 , respectively, in Proposition 1.) It remains to show that the expression in the preceding display has finite conditional expectation on the event $\{\hat{m} = m\}$. But this expression is the maximum of a finite number of Gaussians minus the minimum of a finite number of Gaussians. Its unconditional expectation, and hence also its conditional expectation on the event $\{\hat{m} = m\}$, is finite.

In order to infer (3) from (4), that latter condition needs to be checked for every point yin a union of polyhedra. While this is easy in some simple examples like, say, the situation depicted in Figure 1 of Lee et al. (2016), searching over polyhedra in \mathbb{R}^n is hard in general. In our simulations, we therefore use a simpler sufficient condition that implies (3): After observing the data, i.e., after observing a particular value y^* of Y, and hence also observing $\hat{m}(y^*) = m$ and $\hat{s}(y^*) = s$, we check whether $T_m((I_n - P_{\eta^m})y^*)$ is bounded from above or

from below (and also whether $A_{m,s}y^* < b_{m,s}$, which, if satisfied, entails that $m \in \mathcal{M}^+$ and that $s \in \mathcal{S}_m^+$). If this is the case, then it follows, ex post, that (3) holds. Note that these computations occur naturally during the computation of $[L_m(y^*), U_m(y^*)]$ and can hence be performed as a safety precaution with little extra effort.

Remark. If \tilde{m} is any other model selection procedure, so that the event $\{y : \tilde{m} = m\}$ is the union of a finite number of polyhedra (up to null sets), then the polyhedral method can be applied to obtain a confidence set for $\eta^{m'}\mu$ with conditional coverage probability $1 - \alpha$, conditional on the event $\{\tilde{m} = m\}$ if that event has positive probability. Clearly, for such a model selection procedure, an equivalence similar to (3)–(4) holds. Indeed, the derivation of this equivalence relies on Proposition 1 but not on the Lasso-specific Proposition 2.

3.1 Simulation results

To investigate whether condition (4) is restrictive or not, we perform an exploratory simulation exercise consisting of repeated samples of size n = 100 in the nine scenarios corresponding to the rows in Table 1, which cover all combinations of the cases p = 20, p = 50and p = 200 (i.e., p small, moderate, large) with the cases $\lambda = 0.1$, $\lambda = 1$ and $\lambda = 10$ (i.e., λ small, moderate, large). For each of the nine scenarios, we generate an $n \times p$ regressor matrix X whose rows follow a p-variate Gaussian distribution with mean zero, so that the diagonal elements of the covariance matrix all equal 1 and the off-diagonal elements all equal 0.2. We then generate an n-vector y^* whose entries are i.i.d. standard Gaussians, compute the Lasso estimator $\hat{\beta}(y^*)$ and the resulting selected model $m = \hat{m}(y^*)$ (if $m = \emptyset$, this process is repeated with a newly generated vector y^*). Finally, we generate 2000 directions γ_j^m that are i.i.d. uniform on the unit-sphere in \mathbb{R}^m and set $\eta_j^m = X_m(X'_m X_m)^{-1} \gamma_j^m$, $1 \le j \le 2000$. For each η_j^m we now check if the sufficient condition outlined in the preceding

paragraph is satisfied with η_j^m replacing η^m . If it is satisfied, the corresponding confidence set $[L_{\hat{m}}(Y), U_{\hat{m}}(Y)]$ (with η_j^m replacing η^m) is guaranteed to have infinite expected length conditional on the event that $\hat{m} = m$. The fraction of indices $j, 1 \leq j \leq 2000$, for which this is the case, together with the number of parameters in the selected model are displayed in the cells of Table 1, for 50 independent replications of y^* in each scenario (row). In each of the nine scenarios, and for each of the 50 replications in each, this fraction estimates a lower bound for the probability that the confidence interval $[L_{\hat{m}}(Y), U_{\hat{m}}(Y)]$ has infinite expected length conditional on $\hat{m} = m$ if the direction of interest, i.e., γ^m or, equivalently, η^m , is chosen at random.

As soon as the selected model excludes more than a few variables, we see that the conditional expected length of $[L_{\hat{m}}(Y), U_{\hat{m}}(Y)]$ is guaranteed to be infinite in a substantial number of cases. In particular, this always occurs if p > n. (Also keep in mind that we only check a sufficient condition for infinite expected length, not a necessary one.) Also, within each row in the table, the number of cases with infinite conditional expected length is roughly increasing as the number of parameters in the selected model decreases. Beyond these observations, there appears to be no simple relation between the number of selected variables in the model and the percentage of cases where the interval has infinite conditional expected length. We also stress here that a simulation study can not be exhaustive and that other simulation scenarios will give different results.

3.2 The unknown variance case

Suppose here that $\sigma^2 > 0$ is unknown and that $\hat{\sigma}^2$ is an estimator for σ^2 . Fix $m \in \mathcal{M}^+$ and $s \in \mathcal{S}_m^+$. Note that the set $A_{m,s}y < b_{m,s}$ does not depend on σ^2 and hence also $\mathcal{V}_{m,s}^-((I_n - P_{\eta^m})y)$ and $V_{m,s}^-((I_n - P_{\eta^m})y)$ do not depend on σ^2 . For each $\varsigma^2 > 0$ and for each y so that $A_{m,s}y < b_{m,s}$ define $L_{m,s}(y,\varsigma^2)$, $U_{m,s}(y,\varsigma^2)$, $L_m(y,\varsigma^2)$, and $U_m(y,\varsigma^2)$ like

<i>p</i>	λ	y_1^*	y_2^*	y_3^*		y_{48}^{*}	y_{49}^{*}	y_{50}^{*}
20	0.1	0%	0%	0%	•••	80%	100%	100%
		(20)	(20)	(20)	•••	(19)	(19)	(19)
	1	0%	0%	0%	•••	100%	100%	100%
		(20)	(20)	(20)	•••	(16)	(16)	(14)
	10	49%	60%	72%	•••	100%	100%	100%
		(4)	(7)	(8)	•••	(3)	(3)	(3)
50	0.1	0%	0%	0%	•••	100%	100%	100%
		(50)	(50)	(50)	•••	(48)	(47)	(47)
	1	96%	99%	100%	•••	100%	100%	100%
		(47)	(44)	(47)		(39)	(38)	(37)
	10	100%	100%	100%	•••	100%	100%	100%
		(18)	(18)	(17)	•••	(7)	(4)	(3)
200	0.1	100%	100%	100%	•••	100%	100%	100%
		(100)	(100)	(100)	•••	(97)	(97)	(95)
	1	100%	100%	100%	•••	100%	100%	100%
		(96)	(94)	(94)		(84)	(83)	(82)
	10	100%	100%	100%	•••	100%	100%	100%
		(41)	(37)	(36)	•••	(21)	(20)	(19)

Table 1: Percentage of cases where η^m is such that the confidence interval $[L_{\hat{m}}(Y), U_{\hat{m}}(Y)]$ for $\eta^{m'}\mu$ is guaranteed to have infinite expected length conditional on $\hat{m} = m$, with $m = \hat{m}(y_i^*)$ and, in parentheses, the number of parameters in the model, i.e., |m|. The entries in each row are ordered to improve readability, first by percentage (increasing) and second by number of parameters (decreasing).

 $L_{m,s}(y)$, $U_{m,s}(y)$, $L_m(y)$, and $U_m(y)$ in Section 2 with ς^2 replacing σ^2 in the formulas. (Note that, say, $L_{m,s}(y)$ depends on σ^2 through $\sigma_m^2 = \sigma^2 \eta^{m'} \eta^m$.) The asymptotic coverage probability of the intervals $[L_{m,s}(Y, \hat{\sigma}^2), U_{m,s}(Y, \hat{\sigma}^2)]$ and $[L_m(Y, \hat{\sigma}^2), U_m(Y, \hat{\sigma}^2)]$, conditional on the events $\{\hat{m} = m, \hat{s} = s\}$ and $\{\hat{m} = m\}$, respectively, is discussed in Lee et al. (2016).

If $\hat{\sigma}^2$ is independent of $\eta^{m'}Y$ and positive with positive probability, then it is easy to see that (1) continues to hold with $[L_{m,s}(Y,\hat{\sigma}^2), U_{m,s}(Y,\hat{\sigma}^2)]$ replacing $[L_{m,s}(Y), U_{m,s}(Y)]$ for each $m \in \mathcal{M}^+$ and each $s \in \mathcal{S}_m^+$. And if, in addition, $\hat{\sigma}^2$ has finite mean conditional on the event $\{\hat{m} = m\}$ for $m \in \mathcal{M}^+$, then it is elementary to verify that the equivalence (3)–(4) continues to hold with $[L_m(Y,\hat{\sigma}^2), U_m(Y,\hat{\sigma}^2)]$ replacing $[L_m(Y), U_m(Y)]$ (upon repeating the arguments following (3)–(4) and upon using the finite conditional mean of $\hat{\sigma}^2$ in the last step).

In the case where p < n, the usual variance estimator $||Y - X(X'X)^{-1}X'Y||^2/(n-p)$ is independent of $\eta^{m'}Y$, is positive with probability 1, and has finite unconditional (and hence also conditional) mean. For variance estimators in the case where $p \ge n$, we refer to Lee et al. (2016) and the references therein.

Appendix A Auxiliary results

In this section, we collect some properties of functions like $F_{\theta,\varsigma^2}^T(w)$ that will be needed in the proofs of Proposition 1 and Proposition 2. The following result will be used repeatedly in the following and is easily verified using L'Hospital's method.

Lemma A.1. For all a, b with $-\infty \le a < b \le \infty$, the following holds:

$$\lim_{\theta \to \infty} \frac{\Phi \left(a - \theta \right)}{\Phi \left(b - \theta \right)} = 0$$

Write $F_{\theta,\varsigma^2}^T(w)$ and $f_{\theta,\varsigma^2}^T(w)$ for the c.d.f. and p.d.f. of the $TN(\theta,\varsigma^2,T)$ -distribution, where $T = \bigcup_{i=1}^K (a_i, b_i)$ with $-\infty \leq a_1 < b_1 < a_2 < \cdots < a_K < b_K \leq \infty$. For $w \in T$ and for k so that $a_k < w < b_k$, we have

$$F_{\theta,\varsigma^2}^T(w) = \frac{\Phi\left(\frac{w-\theta}{\varsigma}\right) - \Phi\left(\frac{a_k-\theta}{\varsigma}\right) + \sum_{i=1}^{k-1} \Phi\left(\frac{b_i-\theta}{\varsigma}\right) - \Phi\left(\frac{a_i-\theta}{\varsigma}\right)}{\sum_{i=1}^{K} \Phi\left(\frac{b_i-\theta}{\varsigma}\right) - \Phi\left(\frac{a_i-\theta}{\varsigma}\right)};$$

if k = 1, the sum in the numerator is to be interpreted as 0. And for w as above, the density $f_{\theta,\varsigma^2}^T(w)$ is equal to $\phi((w - \theta)/\varsigma)/\varsigma$ divided by the denominator in the preceding display.

Lemma A.2. For each fixed $w \in T$, $F_{\theta,\varsigma^2}^T(w)$ is continuous and strictly decreasing in θ , and

$$\lim_{\theta \to -\infty} F_{\theta,\varsigma^2}^T(w) = 1 \quad and \quad \lim_{\theta \to \infty} F_{\theta,\varsigma^2}^T(w) = 0.$$

Proof. Continuity is obvious and monotonicity has been shown in Lee et al. (2016) for the case where T is a single interval, i.e., K = 1; it is easy to adapt that argument to also cover the case K > 1. Next consider the formula for $F_{\theta,\varsigma^2}^T(w)$. As $\theta \to \infty$, Lemma A.1 implies that the leading term in the numerator is $\Phi((w-\theta)/\varsigma)$ while the leading term in the denominator is $\Phi((b_K - \theta)/\varsigma)$. Using Lemma A.1 again gives $\lim_{\theta\to\infty} F_{\theta,\varsigma^2}^T(w) = 0$. Finally, it is easy to see that $F_{\theta,\varsigma^2}^T(w) = 1 - F_{-\theta,\varsigma^2}^{-T}(-w)$ (upon using the relation $\Phi(t) = 1 - \Phi(-t)$ and a little algebra). With this, we also obtain that $\lim_{\theta\to-\infty} F_{\theta,\varsigma^2}^T(w) = 1$.

For $\gamma \in (0, 1)$ and $w \in T$, define $Q_{\gamma}(w)$ through

$$F^T_{Q_{\gamma}(w),\varsigma^2}(w) \quad = \quad \gamma$$

Lemma A.2 ensures that $Q_{\gamma}(w)$ is well-defined. Note that $L(w) = Q_{1-\alpha/2}(w)$ and $U(w) = Q_{\alpha/2}(w)$.

Lemma A.3. For fixed $w \in T$, $Q_{\gamma}(w)$ is strictly decreasing in γ on (0, 1). And for fixed $\gamma \in (0, 1)$, $Q_{\gamma}(w)$ is continuous and strictly increasing in $w \in T$ so that $\lim_{w \searrow a_1} Q_{\gamma}(w) = -\infty$ and $\lim_{w \nearrow b_K} Q_{\gamma}(w) = \infty$.

Proof. Fix $w \in T$. Strict monotonicity of $Q_{\gamma}(w)$ in γ follows from strict monotonicity of $F_{\theta,\varsigma^2}^T(w)$ in θ ; cf. Lemma A.2.

Fix $\gamma \in (0, 1)$ throughout the following. To show that $Q_{\gamma}(\cdot)$ is strictly increasing on T, fix $w, w' \in T$ with w < w'. We get

$$\gamma \quad = \quad F^T_{Q_{\gamma}(w),\varsigma^2}(w) \quad < \quad F^T_{Q_{\gamma}(w),\varsigma^2}(w'),$$

where the inequality holds because the density of $F_{Q_{\gamma}(w),\varsigma^2}^T(\cdot)$ is positive on T. The definition of $Q_{\gamma}(w')$ and Lemma A.2 entail that $Q_{\gamma}(w) < Q_{\gamma}(w')$.

To show that $Q_{\gamma}(\cdot)$ is continuous on T, we first note that $F_{\theta,\varsigma^2}^T(w)$ is continuous in $(\theta, w) \in \mathbb{R} \times T$ (which is easy to see from the formula for $F_{\theta,\varsigma^2}^T(w)$ given after Lemma A.1). Now fix $w \in T$. Because $Q_{\gamma}(\cdot)$ is monotone, it suffices to show that $Q_{\gamma}(w_n) \to Q_{\gamma}(w)$ for any increasing sequence w_n in T converging to w from below, and for any decreasing sequence w_n converging to w from above. If the w_n increase towards w from below, the sequence $Q_{\gamma}(w_n)$ is increasing and bounded by $Q_{\gamma}(w)$ from above, so that $Q_{\gamma}(w_n)$ converges to a finite limit \overline{Q} . With this, and because $F_{\theta,\varsigma^2}^T(w)$ is continuous in (θ, w) , it follows that

$$\lim_{n} F_{Q_{\gamma}(w_{n}),\varsigma^{2}}^{T}(w_{n}) = F_{\overline{Q},\varsigma^{2}}^{T}(w).$$

In the preceding display, the sequence on the left-hand side is constant equal to γ by definition of $Q_{\gamma}(w_n)$, so that $F_{\overline{Q},\varsigma^2}^T(w) = \gamma$. It follows that $\overline{Q} = Q_{\gamma}(w)$. If the w_n decrease towards w from above, a similar argument applies.

To show that $\lim_{w \nearrow b_K} Q_{\gamma}(w) = \infty$, let $w_n, n \ge 1$, be an increasing sequence in T that converges to b_K . It follows that $Q_{\gamma}(w_n)$ converges to a (not necessarily finite) limit \overline{Q} as

 $n \to \infty$. If $\overline{Q} < \infty$, we get for each $b < b_K$ that

$$\liminf_{n} F_{Q_{\gamma}(w_n),\varsigma^2}^T(w_n) \ge \liminf_{n} F_{Q_{\gamma}(w_n),\varsigma^2}^T(b) = F_{\overline{Q},\varsigma^2}^T(b).$$

In this display, the inequality holds because $F_{Q_{\gamma}(w_n),\varsigma^2}^T(\cdot)$ is a c.d.f., and the equality holds because $F_{\theta,\varsigma^2}^T(b)$ is continuous in θ . As this holds for each $b < b_K$, we obtain that $\liminf_n F_{Q_{\gamma}(w_n),\varsigma^2}^T(w_n) = 1$. But in this equality, the left-hand side equals γ – a contradiction. By similar arguments, it also follows that $\lim_{w \to a_1} Q_{\gamma}(w) = -\infty$.

Lemma A.4. The function $Q_{\gamma}(\cdot)$ satisfies

$$\lim_{w \nearrow b_K} (b_K - w) Q_{\gamma}(w) = -\varsigma^2 \log(\gamma) \qquad \text{if } b_K < \infty \text{ and}$$
$$\lim_{w \searrow a_1} (a_1 - w) Q_{\gamma}(w) = -\varsigma^2 \log(1 - \gamma) \qquad \text{if } a_1 > -\infty.$$

Proof. As both statements follow from similar arguments, we only give the details for the first one. As w approaches b_k from below, $Q_{\gamma}(w)$ converges to ∞ by Lemma A.3. This observation, the fact that $F_{Q_{\gamma}(w),\varsigma^2}^T(w) = \gamma$ holds for each w, and Lemma A.1 together imply that

$$\lim_{w \nearrow b_k} \frac{\Phi\left(\frac{w - Q_{\gamma}(w)}{\varsigma}\right)}{\Phi\left(\frac{b_k - Q_{\gamma}(w)}{\varsigma}\right)} = \gamma$$

Because $\Phi(-x)/(\phi(x)/x) \to 1$ as $x \to \infty$ (cf. Feller 1957, Lemma VII.1.2.), we get that

$$\lim_{w \nearrow b_k} \frac{\phi\left(\frac{w - Q_{\gamma}(w)}{\varsigma}\right)}{\phi\left(\frac{b_k - Q_{\gamma}(w)}{\varsigma}\right)} = \gamma.$$

The claim now follows by plugging-in the formula for $\phi(\cdot)$ on the left-hand side, simplifying, and then taking the logarithm of both sides.

Appendix B Proof of Proposition 1

Proof of the first statement in Proposition 2. Assume that $b_K < \infty$ (the case where $a_1 > -\infty$ is treated similarly). Lemma A.4 entails that $\lim_{w \nearrow b_K} (b_K - w)(U(w) - L(w)) = \varsigma^2 C$, where $C = \log((1 - \alpha/2)/(\alpha/2))$ is positive. Hence, there exists a constant $\epsilon > 0$ so that

$$U(w) - L(w) > \frac{1}{2} \frac{\varsigma^2 C}{b_K - w}$$

whenever $w \in (b_K - \epsilon, b_K) \cap T$. Set $B = \inf\{f_{\theta,\varsigma^2}^T(w) : w \in (b_K - \epsilon, b_K) \cap T\}$. For $w \in T$, $f_{\theta,\varsigma^2}^T(w)$ is a Gaussian density divided by a constant scaling factor, so that B > 0. Because $U(w) - L(w) \ge 0$ in view of Lemma A.3, we obtain that

$$\mathbb{E}_{\theta,\varsigma^2}[U(W) - L(W)|W \in T] \geq \frac{\varsigma^2 BC}{2} \int_{(b_K - \epsilon) \cap T} \frac{1}{b_K - w} dw = \infty.$$

Proof of the first inequality in Proposition 2. Define $R_{\gamma}(w)$ through $\Phi((w-R_{\gamma}(w))/\varsigma) = \gamma$, i.e, $R_{\gamma}(w) = w - \varsigma \Phi^{-1}(\gamma)$ Then, on the one hand, we have

$$F_{R_{\gamma}(w),\varsigma^{2}}^{T}(w) = \frac{P(N(R_{\gamma}(w),\varsigma^{2}) \leq w, N(R_{\gamma}(w),\varsigma^{2}) \in T)}{P(N(R_{\gamma}(w),\varsigma^{2}) \in T)}$$
$$\leq \frac{P(N(R_{\gamma}(w),\varsigma^{2}) \leq w)}{\inf_{\vartheta} P(N(\vartheta,\varsigma^{2}) \in T)} = \frac{\gamma}{p_{*}},$$

while, on the other,

$$F_{R_{\gamma}(w),\varsigma^{2}}^{T}(w) \geq \frac{P(N(R_{\gamma}(w),\varsigma^{2}) \leq w) - P(N(R_{\gamma}(w),\varsigma^{2}) \notin T)}{P(N(R_{\gamma}(w),\varsigma^{2}) \in T)}$$

$$\geq \inf_{\vartheta} \frac{P(N(R_{\gamma}(w),\varsigma^{2}) \leq w) - 1 + P(N(\vartheta,\varsigma^{2}) \in T)}{P(N(\vartheta,\varsigma^{2}) \in T)}$$

$$= \frac{\gamma - 1 + p_{*}}{p_{*}}.$$

The inequalities in the preceding two displays, together with the fact that $F_{\theta,\varsigma^2}^T(w)$ is decreasing in θ , imply that

$$R_{1-p_*(1-\gamma)}(w) \leq Q_{\gamma}(w) \leq R_{p_*\gamma}(w).$$

(Note that the inequality in the third-to-last display continues to hold with $p_*\gamma$ replacing γ ; in that case, the upper bound reduces to γ . And, similarly, the inequality in the second-to-last display continues to hold with $1 - p_*(1 - \gamma)$ replacing γ , in which case the lower bound reduces to γ). In particular, we get that $U(w) = Q_{\alpha/2}(w) \leq R_{p_*\alpha/2}(w) = w - \varsigma \Phi^{-1}(p_*\alpha/2)$ and that $L(w) = Q_{1-\alpha/2}(w) \geq R_{1-p_*\alpha/2}(w) = w - \varsigma \Phi^{-1}(1 - p_*\alpha/2)$. The last two inequalities, and the symmetry of $\Phi(\cdot)$ around zero, imply the second inequality in the proposition.

Proof of the second inequality in Proposition 2. Note that $p_* \geq p_\circ = \inf_{\vartheta} P(N(\vartheta, \varsigma^2) < b_1 \text{ or } N(\vartheta, \varsigma^2) > a_K)$, because T is unbounded above and below. Setting $\delta = (a_K - b_1)/(2\varsigma)$, we note that $\delta \geq 0$ and that it is elementary to verify that $p_\circ = 2\Phi(-\delta)$. Because $\Phi^{-1}(1 - p_*\alpha/2) \leq \Phi^{-1}(1 - p_\circ\alpha/2)$, the inequality will follow if we can show that $\Phi^{-1}(1 - p_\circ\alpha/2) \leq \Phi^{-1}(1 - \alpha/2) + \delta$ or, equivalently, that $\Phi^{-1}(p_\circ\alpha/2) \geq \Phi^{-1}(\alpha/2) - \delta$. Because $\Phi(\cdot)$ is strictly increasing, this is equivalent to

$$p_{\circ}\alpha/2 = \Phi(-\delta)\alpha \ge \Phi(\Phi^{-1}(\alpha/2) - \delta)$$

To this end, we set $f(\delta) = \alpha \Phi(-\delta)/\Phi(\Phi^{-1}(\alpha/2) - \delta)$ and show that $f(\delta) \ge 1$ for $\delta \ge 0$. Because f(0) = 1, it suffices to show that $f'(\delta)$ is non-negative for $\delta > 0$. The derivative can be written as a fraction with positive denominator and with numerator equal to

$$-\alpha\phi(-\delta)\Phi(\Phi^{-1}(\alpha/2)-\delta) + \alpha\Phi(-\delta)\phi(\Phi^{-1}(\alpha/2)-\delta).$$

The expression in the preceding display is non-negative if and only if

$$\frac{\Phi(-\delta)}{\phi(-\delta)} \geq \frac{\Phi(\Phi^{-1}(\alpha/2) - \delta)}{\phi(\Phi^{-1}(\alpha/2) - \delta)}.$$

This will follow if the function $g(x) = \Phi(-x)/\phi(x)$ is decreasing in $x \ge 0$. The derivative g'(x) can be written as a fraction with positive denominator and with numerator equal to

$$-\phi(x)^{2} + x\Phi(-x)\phi(x) = x\phi(x)\left(\Phi(-x) - \frac{\phi(x)}{x}\right)$$

Using the well-known inequality $\Phi(-x) \leq \phi(x)/x$ for x > 0 (Feller 1957, Lemma VII.1.2.), we see that the expression in the preceding display is non-positive for x > 0.

Appendix C Proof of Proposition 2

From Lee et al. (2016), we recall the formulas for the expressions on the right-hand side of (2), namely $A_{m,s} = (A_{m,s}^0, A_{m,s}^1)'$ and $b_{m,s} = (b_{m,s}^0, b_{m,s}^1)'$ with $A_{m,s}^0$ and $b_{m,s}^0$ given by

$$\frac{1}{\lambda} \begin{pmatrix} X'_{m^c}(I_n - P_{X_m}) \\ -X'_{m^c}(I_n - P_{X_m}) \end{pmatrix} \text{ and } \begin{pmatrix} \iota - X'_{m^c} X_m (X'_m X_m)^{-1} s \\ \iota + X'_{m^c} X_m (X'_m X_m)^{-1} s \end{pmatrix},$$

respectively, and with $A_{m,s}^1 = -\operatorname{diag}(s)(X'_m X_m)^{-1}X'_m$ and $b_{m,s}^1 = -\lambda \operatorname{diag}(s)(X'_m X_m)^{-1}s$ (in the preceding display, P_{X_m} denotes the orthogonal projection matrix onto the column space spanned by X_m and ι denotes an appropriate vector of ones). Moreover, it is easy to see that the set $\{y : A_{m,s}y < b_{m,s}\}$ can be written as $\{y : \text{ for } z = (I_p - P_{X_m})y, \text{ we have } \mathcal{V}_{m,s}^-(z) < \eta^{m'}y < \mathcal{V}_{m,s}^+(z), \mathcal{V}_{m,s}^0(z) > 0\}$, where

$$\begin{aligned}
\mathcal{V}_{m,s}^{-}(z) &= \max\left(\{(b_{m,s} - A_{m,s}z)_{i}/(A_{m,s}c^{m})_{i} : (A_{m,s}c^{m})_{i} < 0\} \cup \{-\infty\}\right), \\
\mathcal{V}_{m,s}^{+}(z) &= \min\left(\{(b_{m,s} - A_{m,s}z)_{i}/(A_{m,s}c^{m})_{i} : (A_{m,s}c^{m})_{i} > 0\} \cup \{\infty\}\right), \\
\mathcal{V}_{m,s}^{0}(z) &= \min\left(\{(b_{m,s} - A_{m,s}z)_{i} : (A_{m,s}c^{m})_{i} = 0\} \cup \{\infty\}\right)
\end{aligned}$$

with $c^m = \eta^m / \|\eta^m\|^2$; cf. also Lee et al. (2016).

Proof of Proposition 2. Set $I_{-} = \{i : (A_{m,s}c^{m})_{i} < 0\}$ and $I_{+} = \{i : (A_{m,s}c^{m})_{i} > 0\}$. In view of the formulas of $\mathcal{V}_{m,s}^{-}(z)$ and $\mathcal{V}_{m,s}^{+}(z)$ given earlier, it suffices to show that either I_{-} or I_{+} is non-empty. Conversely, assume that $I_{-} = I_{+} = \emptyset$. Then $A_{m,s}c^{m} = 0$ and hence also $A_{m,s}^{1}c^{m} = 0$. Using the explicit formula for $A_{m,s}^{1}$ and the definition of η^{m} , i.e., $\eta^{m} = X_{m}(X'_{m}X_{m})^{-1}\gamma^{m}$, it follows that $\gamma^{m} = 0$, which contradicts our assumption that $\gamma^{m} \in \mathbb{R}^{|m|} \setminus \{0\}$.

References

- Bachoc, F., Leeb, H. & Pötscher, B. M. (2016), 'Valid confidence intervals for post-modelselection predictors', arXiv preprint arXiv:1412.4605.
- Berk, R., Brown, L., Buja, A., Zhang, K. & Zhao, L. (2013), 'Valid post-selection inference', Ann. Statist. 41, 802–837.
- Feller, W. (1957), An Introduction to Probability Theory and its Applications, Vol. 1, 2nd edn, Wiley, New York, NY.
- Fithian, W., Sun, D. & Taylor, J. (2017), 'Optimal inference after model selection', arXiv preprint arXiv:1410.2597.
- Frank, I. E. & Friedman, J. H. (1993), 'A statistical view of some chemometrics regression tools', *Technometrics* 35, 109–135.
- Lee, J. D., Sun, D. L., Sun, Y. & Taylor, J. E. (2016), 'Exact post-selection inference, with application to the Lasso', Ann. Statist. 44, 907–927.
- Leeb, H. & Kabaila, P. (2017), 'Admissibility of the usual confidence set for the mean of a

univariate or bivariate normal population: The unknown variance case', J. R. Stat. Soc. Ser. B Stat. Methodol. **79**, 801–813.

- Leeb, H. & Pötscher, B. M. (2006a), 'Can one estimate the conditional distribution of post-model-selection estimators?', Ann. Statist. 34, 2554–2591.
- Leeb, H. & Pötscher, B. M. (2006b), 'Performance limits for estimators of the risk or distribution of shrinkage-type estimators, and some general lower risk-bound results', *Econometric Theory* 22, 69–97.
- Leeb, H. & Pötscher, B. M. (2008), 'Can one estimate the unconditional distribution of post-model-selection estimators?', *Econometric Theory* 24, 338–376.
- Pötscher, B. M. (2009), 'Confidence sets based on sparse estimators are necessarily large', Sankhyā Ser. A 71, 1–18.
- Pötscher, B. M. & Schneider, U. (2010), 'Confidence sets based on penalized maximum likelihood estimators in Gaussian regression', *Electron. J. Statist.* 4, 334–360.
- Schneider, U. (2016), 'Confidence sets based on thresholding estimators in high-dimensional Gaussian regression models', *Econometric Rev.* 35, 1412–1455.
- Taylor, J. & Tibshirani, R. (2017), 'Post-selection inference for l_1 -penalized likelihood models', *Canad. J. Statist.* forthcoming.
- Tian, X. & Taylor, J. E. (2015), 'Selective inference with a randomized response', arXiv preprint arXiv:1507.06739.
- Tibshirani, R. (1996), 'Regression shrinkage and selection via the Lasso', J. R. Stat. Soc. Ser. B Stat. Methodol. 58, 267–288.

- Tibshirani, R. J. (2013), 'The Lasso problem and uniqueness', *Electron. J. Statist.* 7, 1456–1490.
- Tibshirani, R. J., Taylor, J., Lockhart, R. & Tibshirani, R. (2016), 'Exact post-selection inference for sequential regression procedures', J. Amer. Statist. Assoc. 111, 600–620.