

Exact adaptive confidence intervals for linear regression coefficients

Peter Hoff

Duke University,
e-mail: peter.hoff@duke.edu

Chaoyu Yu

University of Washington,
e-mail: chaoyuyu@uw.edu

Abstract: We propose an adaptive confidence interval procedure (CIP) for the coefficients in the normal linear regression model. This procedure has a frequentist coverage rate that is constant as a function of the model parameters, yet provides smaller intervals than the usual interval procedure, on average across regression coefficients. The proposed procedure is obtained by defining a class of CIPs that all have exact $1 - \alpha$ frequentist coverage, and then selecting from this class the procedure that minimizes a prior expected interval width. We describe an adaptive approach for estimating the prior distribution from the data, so that the potential risk of a poorly specified prior is reduced. The resulting adaptive confidence intervals maintain exact non-asymptotic $1 - \alpha$ coverage if two conditions are met - that the design matrix is full rank (which will be known) and that the errors are normally distributed (which can be checked empirically). No assumptions on the unknown parameters are necessary to maintain exact coverage. Additionally, in a “ p growing with n ” asymptotic scenario, this adaptive FAB procedure is asymptotically Bayes-optimal among $1 - \alpha$ frequentist CIPs.

MSC 2010 subject classifications: 62J05.

Keywords and phrases: Empirical Bayes, frequentist coverage, ridge regression, shrinkage, sparsity.

Received June 2018.

Contents

1	Introduction	95
2	Review of FAB intervals	97
3	FAB t -intervals for regression parameters	99
4	Approximate optimality	102
5	Alternative spending functions	104
5.1	Bounded-width adaptive FAB interval	104
5.2	Adaptive FAB regions for arbitrary prior distributions	105
6	Numerical examples	108
6.1	Motif regression	108
6.2	Motif regression simulation study	109
6.3	Diabetes progression	111

7 Discussion	112
Proofs	113
References	118

1. Introduction

Linear regression analyses routinely include point estimates and confidence intervals for the regression coefficients $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)$ of the linear model $\mathbf{y} \sim N_n(\mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{I})$. The most widely-used confidence interval procedure (CIP) for an element β_j of $\boldsymbol{\beta}$ is perhaps the usual t -interval centered around the ordinary least-squares (OLS) estimate $\hat{\beta}_j$. This interval is uniformly most accurate among CIPs that are derived from inversion of unbiased tests, and so it is called the uniformly most accurate unbiased (UMAU) CIP.

In this article we consider alternatives to the UMAU procedure that have *constant coverage*, that is, interval procedures $C_j(\mathbf{y})$ satisfying

$$\Pr(\beta_j \in C_j(\mathbf{y}) | \boldsymbol{\beta}, \sigma) = 1 - \alpha, \quad \forall (\boldsymbol{\beta}, \sigma) \in \mathbb{R}^p \times \mathbb{R}^+. \quad (1)$$

This property is what we normally think of as the usual frequentist definition of $1 - \alpha$ coverage - the random interval $C_j(\mathbf{y})$ covers the true value β_j with probability $1 - \alpha$, no matter what $\boldsymbol{\beta}$ and σ are. We introduce the term “constant coverage” to distinguish such intervals from other intervals whose coverage is bounded below by $1 - \alpha$ but varies with $(\boldsymbol{\beta}, \sigma^2)$, or so-called “frequentist intervals” whose coverage rate is only constant as a function of the parameters asymptotically. For example, the usual score interval for a coefficient in a logistic regression model has an actual finite-sample coverage rate that depends on the values of the parameters.

The UMAU interval procedure of course has constant coverage, and it also has an expected width that is constant for all values of $\boldsymbol{\beta}$. However, in many cases we have prior information that many of the elements of $\boldsymbol{\beta}$ may be close to a particular value, such as zero. In this case, we might prefer a CIP for β_j that has a smaller expected width for “likely” values of β_j in exchange for having wider intervals for values of β_j that are less likely. Specifically, if our prior information could be quantified in terms of a prior distribution with density $\pi(\boldsymbol{\beta})$, then arguably we would be interested in an interval procedure that minimizes the prior expected width

$$\mathbb{E}[|C_j(\mathbf{y})|] = \int \int |C_j(\mathbf{y})| p(\mathbf{y} | \boldsymbol{\beta}, \sigma) d\mathbf{y} \pi(\boldsymbol{\beta}) d\boldsymbol{\beta}$$

among all CIPs that satisfy the constant coverage property (1). Such a procedure would still be “frequentist” in that it would have $1 - \alpha$ constant coverage, but it would also be Bayes-optimal among frequentist procedures. We refer to such a statistical procedure as “frequentist, assisted by Bayes” or FAB.

The frequentist performance of such a FAB interval will be quite good if the true value of $\boldsymbol{\beta}$ has high prior probability. However, if an appropriate prior distribution is not known then the performance of a FAB interval using an

arbitrarily chosen prior could be quite poor. For this reason, we present a method for adaptively estimating a normal prior distribution for β from the data \mathbf{y} , and then using this estimated prior distribution to construct an approximately Bayes-optimal CIP for each regression coefficient $\beta_j, j = 1, \dots, p$. While such a procedure may be viewed as an empirical Bayes method, it differs from typical empirical Bayes procedures in that it satisfies the constant coverage condition (1) exactly and non-asymptotically under the normal regression model with $n > p$, while also being Bayes-optimal asymptotically as p and n increase to infinity. In contrast, standard Bayes and empirical Bayes procedures generally only maintain their coverage rates marginally with respect to the prior (Carlin and Gelfand, 1990; Yoshimori and Lahiri, 2014).

Our proposed adaptive CIP builds on the work of Pratt (1963), who obtained a Bayes-optimal frequentist confidence interval for the mean of a normal population with a known variance. Pratt’s interval has a shorter expected width than the UMAU interval for parameter values with high prior probability, but has an arbitrarily large width for values with low prior probability. To guard against a potential mismatch between the prior information and the true parameter value, Farchione and Kabaila (2008), Kabaila and Giri (2009) and Kabaila and Tissera (2014) developed confidence interval procedures that revert to the usual UMAU procedure when the point estimate is far from an *a priori* likely value. In this article, we instead guard against a potential mismatch between prior information and truth by adaptively estimating the prior distribution from the data itself.

Our confidence intervals are constructed by inverting classes of tests that vary as a function of the parameter in terms of how type I error is “spent.” Parameterization of a class of confidence intervals with such a “spending function” was first used by Stein (1962) and discussed further by Bartholomew (1971) to construct confidence intervals for parameters known to lie in particular subsets of the real line. This parameterization was also used by Puza and O’Neill (2006), who developed classes of asymmetric CIPs for one-sample problems.

In the next section we review Pratt’s FAB interval, obtain its representation in terms of a spending function, and discuss an extension developed in Yu and Hoff (2018) to accommodate an unknown variance. In Section 3 we extend these ideas to the case of interval estimation for a linear regression coefficient, and show how we may adaptively estimate a normal prior distribution for the elements of β . The resulting adaptive FAB confidence interval we propose maintains exact, non-asymptotic constant coverage. Additionally, since the accuracy of our adaptive estimate improves as n and p increase, in Section 4 we show that the adaptive FAB procedure is Bayes-optimal under this type of asymptotic regime. In Section 5 we develop new alternative model-based spending functions, including a modification of our adaptive FAB interval to guard against an outlying parameter value, as well as an adaptive FAB procedure derived from a “spike-and-slab” prior distribution. Section 6 includes several numerical examples illustrating the use of the adaptive FAB procedure, including analyses of two datasets and a small simulation study. A discussion follows in Section 7.

Several other authors have studied alternatives to UMAU intervals for regres-

sion parameters. O’Gorman (2001) developed a CIP based on a permutation test that adapts to non-normal error distributions, as opposed to adapting to small or sparse values of the regression coefficients. Kabaila and Tissera (2014) developed an improved CIP based on an adaptive estimate of the variance σ^2 . Their procedure depends on a user-specified spline function for which the constant coverage property must be checked numerically. In contrast, our proposed CIP is obtained by adaptively selecting from a class of constant-coverage intervals based on easy to obtain estimates of a few parameters. For our procedure, constant coverage follows by construction and does not need to be checked numerically. Lee et al. (2016) developed a procedure that has exact conditional coverage, given a model selection event and knowledge of σ^2 . However, for cases where σ^2 is unknown, their suggested modification uses a plug-in estimate of σ^2 and achieves exact coverage only asymptotically. Other authors (Bühlmann (2013), van de Geer et al. (2014), Zhang and Zhang (2014)) have considered confidence interval construction for sparse, high-dimensional regression, including the case that $p > n$. These approaches generally work by de-biasing sparse estimators of the regression coefficients. However, the coverage rates of these methods are asymptotic, and typically depend on the degree of sparsity of β . For example, in Section 4 we show that the finite-sample coverage of one such procedure can be very good in a sparse setting, but extremely poor if β is not sparse. In contrast, our proposed procedure maintains exact, non-asymptotic coverage if two conditions are met - that the design matrix is full rank (a condition that will be known), and that the errors are normally distributed (a condition that can be evaluated empirically). No assumptions on the unobserved parameters are necessary to maintain exact coverage.

2. Review of FAB intervals

Suppose $\hat{\theta}$ is normally distributed with unknown mean θ and known variance σ^2 . Then for any choice of $s \in [0, 1]$,

$$\begin{aligned} \Pr(z_{\alpha(1-s)} < (\theta - \hat{\theta})/\sigma < z_{1-\alpha s} | \theta) &= (1 - \alpha s) - \alpha(1 - s) \\ &= 1 - \alpha, \end{aligned}$$

where z_p denotes the p th quantile of the standard normal distribution. As noticed by Stein (1962), Bartholomew (1971) and Puza and O’Neill (2006), this implies that for any function $s : \mathbb{R} \rightarrow [0, 1]$ the set-valued function

$$C_s(\hat{\theta}) = \left\{ \theta : \hat{\theta} + \sigma z_{\alpha(1-s(\theta))} < \theta < \hat{\theta} + \sigma z_{1-\alpha s(\theta)} \right\} \quad (2)$$

is a $1 - \alpha$ confidence procedure, satisfying $\Pr(\theta \in C_s(\hat{\theta}) | \theta) = 1 - \alpha$. Puza and O’Neill (2006) referred to such a function s as a *tail function*, as it gets plugged into the standard normal quantile function. In this article, we refer to s as a *spending function*. This is because, as will be discussed, inversion of C_s yields a class of hypothesis tests where s determines how much type I error is “spent” in each part of the parameter space.

The usual procedure is $C_{1/2}(\hat{\theta})$, obtained from the constant spending function $s(\theta) = 1/2$. While $C_{1/2}$ is the uniformly most accurate unbiased (UMAU) confidence interval procedure (CIP), the lack of a uniformly most powerful test of $H_\theta : E[\hat{\theta}] = \theta$ versus $K_\theta : E[\hat{\theta}] \neq \theta$ means there are confidence procedures corresponding to collections of biased level- α tests that have smaller expected widths than the UMAU procedure for some regions of the parameter space. If prior information is available that θ is likely to be near some value μ , then we may be willing to incur wider intervals for θ -values far from μ in exchange for smaller intervals near μ . With this in mind, Pratt (1963) developed a Bayes-optimal $1 - \alpha$ CIP that minimizes the ‘‘Bayes width’’ or expected interval width averaged over values of both $\hat{\theta}$ and θ , where the latter averaging is done with respect to a $N(\mu, \tau^2)$ prior distribution for θ . The resulting CIP has $1 - \alpha$ frequentist coverage for each value of θ , but has lower expected width for values of θ near the prior mean (and wider expected widths elsewhere). We describe this interval as being ‘‘frequentist assisted by Bayes’’ or FAB. As shown in Yu and Hoff (2018), the spending function corresponding to Pratt’s FAB confidence interval is characterized as follows: If $\tau^2 > 0$, then

$$s(\theta) = g^{-1}(2\sigma(\theta - \mu)/\tau^2) \quad (3)$$

where $g(s) = \Phi^{-1}(\alpha s) - \Phi^{-1}(\alpha(1 - s))$ and Φ is the standard normal CDF. If $\tau^2 = 0$, then $s(\theta) = 1$ if $\theta > \mu$ and $s(\theta) < 0$ if $\theta < \mu$. The value of $s(\mu) \in [0, 1]$ does not affect the width of the confidence interval, but can affect whether or not μ is included in the interval or not (as an endpoint). We suggest taking $s(\mu)$ to be $1/2$ when $\tau^2 = 0$, as it is in the case that $\tau^2 > 0$.

Now consider confidence interval construction for θ in the more typical case that σ^2 is unknown. Suppose we will observe independent statistics $\hat{\theta}$ and $\hat{\sigma}^2$, where $\hat{\theta} \sim N(\theta, \sigma^2)$ and $q\hat{\sigma}^2/\sigma^2 \sim \chi_q^2$. Letting t_p be the p th quantile of the t -distribution with q degrees of freedom, any spending function $s : \mathbb{R} \rightarrow [0, 1]$ defines a class of acceptance regions

$$A_s(\theta) = \left\{ (\hat{\theta}, \hat{\sigma}) : \hat{\theta} + \hat{\sigma}t_{\alpha(1-s(\theta))} < \theta < \hat{\theta} + \hat{\sigma}t_{1-\alpha s(\theta)} \right\},$$

so that for each θ , $A_s(\theta)$ is the acceptance region of a level- α test. Inversion of this class of tests yields a CIP with exact $1 - \alpha$ constant coverage,

$$C_s(\hat{\theta}, \hat{\sigma}) = \left\{ \theta : \hat{\theta} + \hat{\sigma}t_{\alpha(1-s(\theta))} < \theta < \hat{\theta} + \hat{\sigma}t_{1-\alpha s(\theta)} \right\}. \quad (4)$$

Important properties of this CIP include the following:

Theorem 1. *If $\hat{\theta} \sim N(\theta, \sigma^2)$ and $q\hat{\sigma}^2/\sigma^2 \sim \chi_q^2$ are independent, then*

1. $\Pr(\theta \in C_s(\hat{\theta}, \hat{\sigma}) | \theta, \sigma) = 1 - \alpha \forall (\theta, \sigma) \in \mathbb{R} \times \mathbb{R}^+$, so the procedure has $1 - \alpha$ constant coverage;
2. If $s(\theta)$ is nondecreasing then $C_s(\hat{\theta}, \hat{\sigma})$ is an interval with probability 1;
3. If $s(\theta)$ is not nondecreasing then $C_s(\hat{\theta}, \hat{\sigma})$ is an interval with probability less than 1.

Items 1 and 2 were shown in Yu and Hoff (2018), and a proof of item 3 is in the appendix. This result says that every spending function corresponds to a $1 - \alpha$ frequentist confidence procedure, and every *nondecreasing* spending function corresponds to a $1 - \alpha$ confidence *interval* procedure. Yu and Hoff (2018) showed that the spending function (3) that corresponds to Pratt’s z -interval is strictly increasing. If such a nondecreasing spending function s is used, then the lower and upper endpoints of the interval, $\underline{\theta}$ and $\bar{\theta}$, are obtained by solving the equations

$$F\left(\frac{\underline{\theta} - \hat{\theta}}{\hat{\sigma}}\right) = \alpha(1 - s(\underline{\theta})), \quad F\left(\frac{\hat{\theta} - \bar{\theta}}{\hat{\sigma}}\right) = \alpha s(\bar{\theta}), \quad (5)$$

where F is the CDF of the t_q distribution. These equations can be solved using a zero-finding algorithm, and noting that $\underline{\theta} < \hat{\theta} + \hat{\sigma}t_\alpha$ and $\hat{\theta} + \hat{\sigma}t_{1-\alpha} < \bar{\theta}$. Furthermore, this implies that $\underline{\theta} < \hat{\theta} < \bar{\theta}$ as long as $\alpha < 1/2$.

The (frequentist) expected width of an interval based on Pratt’s spending function (3) is low when θ is near the prior mean μ but is increasing in $|\theta - \mu|$. To guard against a poor choice of μ , Farchione and Kabaila (2008) suggest an alternative confidence interval procedure that reduces to the UMAU interval when $\hat{\theta}$ is far from μ . Alternatively, in multiparameter settings, reasonable choices for μ and τ^2 may be obtained from the data itself. For example, estimates of some parameters might suggest plausible values for others. In this case, we may want to use a spending function $\tilde{s}(\theta)$ that minimizes a Bayes risk corresponding to a “prior” distribution that is adaptively estimated from the data. We refer to such a procedure as adaptive FAB. Fortunately, the results of Proposition 1 hold not just for fixed spending functions, but also those that are random but statistically independent of $\hat{\theta}$ and $\hat{\sigma}^2$:

Corollary 1. *If $\hat{\theta} \sim N(\theta, \sigma^2)$ and $q\hat{\sigma}^2/\sigma^2 \sim \chi_q^2$, and $\hat{\theta}$, $\hat{\sigma}^2$ and \tilde{s} are independent, then $\Pr(\theta \in C_{\tilde{s}}(\hat{\theta}, \hat{\sigma})|\theta, \sigma) = 1 - \alpha$.*

This result follows by conditioning on \tilde{s} : $\Pr(\theta \in C_{\tilde{s}}(\hat{\theta}, \hat{\sigma})|\theta, \sigma) = \mathbb{E}[\Pr(\theta \in C_{\tilde{s}}(\hat{\theta}, \hat{\sigma})|\tilde{s}, \theta, \sigma)|\theta, \sigma]$, but the inner conditional probability is $1 - \alpha$ since \tilde{s} and $(\hat{\theta}, \hat{\sigma})$ are independent and $C_{\tilde{s}}$ has $1 - \alpha$ coverage for each fixed \tilde{s} . Yu and Hoff (2018) made use of this fact to develop an adaptive FAB confidence interval procedure for the means of multiple normal populations. Their adaptive procedure for the mean θ of a given population is $C_{\tilde{s}}(\hat{\theta}, \hat{\sigma})$, where \tilde{s} is the spending function (3) with (μ, σ^2, τ^2) replaced by estimates using data from the other populations. This procedure provides exact $1 - \alpha$ confidence intervals for each population, and is asymptotically optimal in the case of the normal hierarchical model.

3. FAB t -intervals for regression parameters

We now show how the results discussed in the previous section may be used to construct adaptive frequentist confidence intervals for linear regression parameters. Under the normal linear regression model, the intervals we construct have

exact $1 - \alpha$ constant coverage and do not require asymptotic approximations or assumptions on the unknown parameter.

Consider the problem of constructing confidence intervals for the elements of an unknown vector $\boldsymbol{\beta} \in \mathbb{R}^p$ based on data $\mathbf{y} \in \mathbb{R}^n$ and $\mathbf{X} \in \mathbb{R}^{n \times p}$ from the normal linear regression model $\mathbf{y} \sim N_n(\mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{I})$. As is well known,

$$\begin{aligned}\hat{\boldsymbol{\beta}} &= (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} \sim N_p(\boldsymbol{\beta}, \sigma^2(\mathbf{X}^\top \mathbf{X})^{-1}) \\ \hat{\sigma}^2 &= \|\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}\|^2 / (n - p) \sim \frac{\sigma^2}{n-p} \chi_{n-p}^2,\end{aligned}$$

with $\hat{\boldsymbol{\beta}}$ and $\hat{\sigma}^2$ being independent. In particular, $(\hat{\beta}_j - \beta_j) / (w_j \hat{\sigma}) \sim t_{n-p}$, where w_j is the square-root of the j th diagonal entry of $(\mathbf{X}^\top \mathbf{X})^{-1}$. The UMAU confidence interval for β_j is

$$C(\hat{\beta}_j, \hat{\sigma}) = \left\{ \beta_j : \hat{\beta}_j + w_j \hat{\sigma} t_{\alpha/2} < \beta_j < \hat{\beta}_j + w_j \hat{\sigma} t_{1-\alpha/2} \right\}. \quad (6)$$

This interval has $1 - \alpha$ coverage probability and an expected width that is constant as a function of the true value of β_j .

Now suppose that prior information about $\boldsymbol{\beta}$ suggests that $\boldsymbol{\beta} \sim N_p(\mathbf{0}, \tau^2 \mathbf{I})$ for some value of τ^2 (other prior distributions will be discussed at the end of this section, and in Sections 5 and 6). If τ^2 and σ^2 were known, the Bayes-optimal CIP for β_j would be obtained simply by replacing $\hat{\theta}$ and σ in (2) and (3) with $\hat{\beta}_j$ and $w_j \sigma$, yielding

$$\begin{aligned}C_s(\hat{\beta}_j) &= \left\{ \beta_j : \hat{\beta}_j + w_j \sigma z_{\alpha(1-s(\beta_j))} < \beta_j < \hat{\beta}_j + w_j \sigma z_{1-\alpha s(\beta_j)} \right\} \\ s(\beta_j) &= g^{-1}(2w_j \sigma \beta_j / \tau^2).\end{aligned} \quad (7)$$

However, since τ^2 and σ^2 are unknown we alter this interval as follows:

- σ^2 is replaced by $\hat{\sigma}^2$, which is independent of $\hat{\beta}_j$ and satisfies $(n-p)\hat{\sigma}^2/\sigma^2 \sim \chi_{n-p}^2$;
- z -quantiles are replaced by the quantiles of the t_{n-p} distribution;
- $s(\beta_j)$ is replaced by $\tilde{s}(\beta_j) = g^{-1}(2w_j \tilde{\sigma} \beta_j / \tilde{\tau}^2)$, where $(\tilde{\tau}^2, \tilde{\sigma}^2)$ are independent of $(\hat{\beta}_j, \hat{\sigma}^2)$.

These modifications yield an adaptive FAB interval given by

$$C_{\tilde{s}}(\hat{\beta}_j, \hat{\sigma}) = \left\{ \beta_j : \hat{\beta}_j + w_j \hat{\sigma} t_{\alpha(1-\tilde{s}(\beta_j))} < \beta_j < \hat{\beta}_j + w_j \hat{\sigma} t_{1-\alpha \tilde{s}(\beta_j)} \right\}. \quad (8)$$

Such an interval satisfies the conditions of Corollary 1, thereby guaranteeing exact $1 - \alpha$ frequentist coverage, regardless of whether or not the values of $\boldsymbol{\beta}$ are approximately normally distributed, or if the estimates $\tilde{\tau}^2, \tilde{\sigma}^2$ are accurate. However, if the normal approximation and adaptive estimates are accurate, then we expect the resulting FAB interval (8) to be close to the ‘‘oracle’’ interval (7), which is Bayes-optimal and narrower on average than the UMAU procedure given by (6).

The approximate optimality of $C_{\tilde{s}}$ is considered more formally in the next section using an asymptotic argument. First, we discuss obtaining estimators $(\tilde{\tau}^2, \tilde{\sigma}^2)$ that are independent of $(\hat{\beta}_j, \hat{\sigma}^2)$ so that the conditions of Corollary 1 are met. Let $\mathbf{P}_X = \mathbf{X}(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}$ and $\mathbf{P}_0 = \mathbf{I} - \mathbf{P}_X$ be the projection matrices onto the space spanned by the columns of \mathbf{X} and the corresponding null space, respectively. Recall that the OLS estimate $\hat{\beta}_j$ is given by $\hat{\beta}_j = \mathbf{a}^\top \mathbf{y}$, where \mathbf{a} is the j th row of the matrix $(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top$. Let $\mathbf{P}_1 = \mathbf{a} \mathbf{a}^\top / \mathbf{a}^\top \mathbf{a}$ be the projection matrix associated with \mathbf{a} , and let $\mathbf{P}_2 = \mathbf{P}_X (\mathbf{I} - \mathbf{P}_1)$. We can decompose \mathbf{y} as

$$\begin{aligned} \mathbf{y} &= \mathbf{I} \mathbf{y} = (\mathbf{P}_0 + \mathbf{P}_X) \mathbf{y} \\ &= (\mathbf{P}_0 + \mathbf{P}_1 + \mathbf{P}_2) \mathbf{y} \\ &= \mathbf{P}_0 \mathbf{y} + \mathbf{P}_1 \mathbf{y} + \mathbf{P}_2 \mathbf{y} \equiv \mathbf{y}_0 + \mathbf{y}_1 + \mathbf{y}_2. \end{aligned}$$

Since $\mathbf{P}_k \mathbf{P}_l = \mathbf{0}$ for $k \neq l$, we have that \mathbf{y}_0 , \mathbf{y}_1 and \mathbf{y}_2 are statistically independent. Now the OLS estimate satisfies $\hat{\beta}_j = \mathbf{a}^\top \mathbf{P}_1 \mathbf{y} = \mathbf{a}^\top \mathbf{y}_1$, and $\hat{\sigma}^2 = \mathbf{y}_0^\top \mathbf{y}_0 / (n - p)$, and so both estimates are statistically independent of each other and the vector \mathbf{y}_2 . Therefore, any estimates $(\tilde{\tau}^2, \tilde{\sigma}^2)$ that are functions of \mathbf{y}_2 will be independent of $(\hat{\beta}_j, \hat{\sigma}^2)$ and so can be used to construct a spending function $\tilde{s}(\beta_j)$ that satisfies the conditions of Corollary 1.

To obtain such an estimate $(\tilde{\tau}^2, \tilde{\sigma}^2)$, let \mathbf{G}_2 be an orthonormal basis for the space spanned by \mathbf{P}_2 (for example, the matrix of eigenvectors of \mathbf{P}_2 that correspond to non-zero eigenvalues). Then $\mathbf{G}_2 \mathbf{G}_2^\top = \mathbf{P}_2$, $\mathbf{G}_2^\top \mathbf{G}_2 = \mathbf{I}_{p-1}$, and $\mathbf{z}_2 = \mathbf{G}_2^\top \mathbf{y}_2 = \mathbf{G}_2^\top \mathbf{y} \sim N_{p-1}(\mathbf{G}_2^\top \mathbf{X} \boldsymbol{\beta}, \sigma^2 \mathbf{I})$. Under the prior model $\boldsymbol{\beta} \sim N(\mathbf{0}, \tau^2 \mathbf{I})$, the marginal distribution for \mathbf{z}_2 is therefore

$$\mathbf{z}_2 \sim N_{p-1}(\mathbf{0}, \mathbf{X}_2 \mathbf{X}_2^\top \tau^2 + \sigma^2 \mathbf{I}), \quad (9)$$

where $\mathbf{X}_2 = \mathbf{G}_2^\top \mathbf{X}$. A variety of empirical Bayes estimates of (τ^2, σ^2) may be obtained from this marginal distribution. For example, noting that $E[\mathbf{z}_2^\top \mathbf{A}^\top \mathbf{A} \mathbf{z}_2] = \text{tr}(\mathbf{X}_2^\top \mathbf{A}^\top \mathbf{A} \mathbf{X}_2) \tau^2 + \text{tr}(\mathbf{A}^\top \mathbf{A}) \sigma^2$ for any matrix \mathbf{A} , unbiased moment estimates may be obtained by finding $(\tilde{\tau}^2, \tilde{\sigma}^2)$ that solve simultaneously two equations, given by $\mathbf{z}_2^\top \mathbf{A}^\top \mathbf{A} \mathbf{z}_2 = \text{tr}(\mathbf{X}_2^\top \mathbf{A}^\top \mathbf{A} \mathbf{X}_2) \tilde{\tau}^2 + \text{tr}(\mathbf{A}^\top \mathbf{A}) \tilde{\sigma}^2$, for two different values of \mathbf{A} . Alternatively, $(\tilde{\tau}^2, \tilde{\sigma}^2)$ may be taken to be the maximum likelihood estimate based on the marginal model (9). This estimate is discussed further in the next section.

To summarize, we have constructed statistics $\hat{\beta}_j, \hat{\sigma}^2, \tilde{\tau}^2, \tilde{\sigma}^2$ such that for each $(\boldsymbol{\beta}, \sigma^2) \in \mathbb{R}^p \times \mathbb{R}^+$,

- C1: $\hat{\beta}_j \sim N(\beta_j, w_j^2 \sigma^2)$, $(n - p) \hat{\sigma}^2 / \sigma^2 \sim \chi_{n-p}^2$, and $\hat{\beta}_j$ and $\hat{\sigma}^2$ are independent;
- C2: $(\tilde{\tau}^2, \tilde{\sigma}^2)$ are independent of $(\hat{\beta}_j, \hat{\sigma}^2)$.

Therefore the spending function $\tilde{s}(\beta_j) = g^{-1}(2w_j \tilde{\sigma} \beta_j / \tilde{\tau}^2)$ is independent of $(\hat{\beta}_j, \hat{\sigma}^2)$ and so the conditions of Corollary 1 are met. We summarize these results with the following theorem:

Theorem 2. *If $\mathbf{y} \sim N_n(\mathbf{X} \boldsymbol{\beta}, \sigma^2 \mathbf{I})$ and $(\hat{\beta}_j, \hat{\sigma}^2, \tilde{\tau}^2, \tilde{\sigma}^2)$ satisfy C1 and C2, then the adaptive FAB CIP given by (8) has $1 - \alpha$ coverage for every value of $\boldsymbol{\beta}$ and*

σ^2 , that is

$$\Pr(\beta_j \in C_{\tilde{s}}(\hat{\beta}_j, \hat{\sigma}) | \boldsymbol{\beta}, \sigma) = 1 - \alpha, \quad \forall (\boldsymbol{\beta}, \sigma) \in \mathbb{R}^p \times \mathbb{R}^+.$$

The $1 - \alpha$ frequentist coverage rate of this confidence interval does not rely on the prior model $\boldsymbol{\beta} \sim N(\mathbf{0}, \tau^2 \mathbf{I})$ being correct. Rather, the prior model just determines how \mathbf{y}_2 is used to construct the spending function $\tilde{s}(\beta_j)$. If it is plausible that the distribution of β_j -values is not centered around zero, then it makes sense to use a prior model with a non-zero mean, i.e. $\boldsymbol{\beta} \sim N_p(\boldsymbol{\mu} \mathbf{1}, \tau^2 \mathbf{I})$. If the regressor variables are related to each other in some known way, it may be preferable to consider more complex prior models. For example, if the regressor variables have spatial locations and their effects are likely to be smoothly varying in space, then one might consider a prior model of the form $\boldsymbol{\beta} \sim N_p(\mathbf{0}, \Sigma(\psi))$, where ψ is a low-dimensional parameter. Based on this model, the conditional mean and variance of β_j given \mathbf{y}_2 could be obtained and used to construct the spending function $\tilde{s}(\beta_j)$. Since this spending function only depends on \mathbf{y}_2 , it is statistically independent of $\hat{\beta}_j$ and $\hat{\sigma}^2$, and so exact $1 - \alpha$ frequentist coverage is maintained.

4. Approximate optimality

As discussed above, if $\beta_j \sim N(0, \tau^2)$ and σ^2 and τ^2 were known then the oracle FAB interval C_s given by (7) is Bayes-optimal in that it minimizes the prior expected interval width $\mathbb{E}[|C|]$ among procedures C that have $1 - \alpha$ frequentist coverage. This prior expected width is an expectation over both the estimate $\hat{\beta}_j$ and the value of β_j with respect to the $N(0, \tau^2)$ prior distribution.

The adaptive FAB interval $C_{\tilde{s}}$ given by (8) differs from the oracle FAB interval in three ways: the value of σ^2 has been replaced by $\hat{\sigma}^2$; the z -quantiles have been replaced by t -quantiles; and the spending function s that depends on (τ^2, σ^2) has been replaced by \tilde{s} that depends on $(\tilde{\tau}^2, \tilde{\sigma}^2)$. In this subsection we take $(\tilde{\tau}^2, \tilde{\sigma}^2)$ to be the maximizers of the likelihood given by the marginal model (9). The resulting interval still has $1 - \alpha$ frequentist coverage, but it is only an approximation to C_s , and so we must have $\mathbb{E}[|C_{\tilde{s}}|] > \mathbb{E}[|C_s|]$ since C_s is Bayes-optimal. However, if $n - p$ is large then the t -quantiles will be close to the corresponding z -quantiles, and we expect that $\hat{\sigma}^2 \approx \sigma^2$. If p is also large then under the prior $\boldsymbol{\beta} \sim N(0, \tau^2 \mathbf{I})$ we expect that $\tilde{\tau}^2 \approx \tau^2$ and $\tilde{\sigma}^2 \approx \sigma^2$. As a result, we should have $\tilde{s}(\beta_j) \approx s(\beta_j)$ and so we expect that $\mathbb{E}[|C_{\tilde{s}}|] \approx \mathbb{E}[|C_s|]$, that is, the FAB procedure will be approximately Bayes-optimal.

We investigate this more formally with an asymptotic comparison of the widths of the adaptive and oracle FAB procedures. We first obtain an asymptotic result for a single scalar parameter β_j , and then discuss the result in the context of the linear regression model. Consider a sequence of experiments indexed by n such that for each n we have statistics $(\hat{\beta}_j, \hat{\sigma}^2, \tilde{\tau}^2, \tilde{\sigma}^2)$ that satisfy coverage conditions C1 and C2 given above. Furthermore, suppose the following asymptotic conditions hold as $n \rightarrow \infty$:

- A1. $(n - p) \rightarrow \infty$ and $\sigma^2/n \rightarrow \sigma_\infty^2 > 0$;
 A2. $(\tilde{\tau}^2, \tilde{\sigma}^2/n) \rightarrow (\tau^2, \sigma_\infty^2)$ in probability.
 A3. $w_j^2 n \rightarrow w_0^2 > 0$;

We consider this case where σ^2 grows with n since otherwise, if σ^2 were fixed then the widths of the oracle FAB, adaptive FAB and UMAU intervals would all converge to zero at the same rate.

Lemma 1. *Under the conditions C1, C2, A1, A2 and A3, the width $|C_{\bar{s}}|$ of the FAB procedure (8) satisfies $E[|C_{\bar{s}}|] \rightarrow E[|C_s|]$ as $n \rightarrow \infty$, where C_s is the Bayes-optimal FAB procedure for the case that $\hat{\beta}_j \sim N(\beta_j, w_0^2 \sigma_\infty^2)$ and $\beta_j \sim N(0, \tau^2)$.*

A proof is in the appendix. The lemma says that under this asymptotic regime, the performance of the adaptive FAB interval is asymptotically equivalent to that of the oracle FAB interval: They both have $1 - \alpha$ frequentist coverage for each n , and the prior expected width of the FAB procedure approaches that of the oracle FAB interval as $n \rightarrow \infty$.

We now consider how this result applies to the linear regression model and the specific estimates $\hat{\beta}_j, \hat{\sigma}^2, \tilde{\tau}^2, \tilde{\sigma}^2$ described in the previous subsection. Consider a sequence of experiments indexed by n such that the following conditions hold:

- B1. For each n ,
- \mathbf{X} is full-rank;
 - $\mathbf{y} \sim N_n(\mathbf{X}\boldsymbol{\beta}, \sigma^2 \mathbf{I})$ with $\sigma^2 = n\sigma_\infty^2$;
 - $\boldsymbol{\beta} \sim N(\mathbf{0}, \tau^2 \mathbf{I})$.
- B2. $p/n \rightarrow c \in (0, 1)$ as $n \rightarrow \infty$.
 B3. The empirical distribution of the eigenvalues of $\mathbf{X}^\top \mathbf{X}/n$ is bounded uniformly in n , and converges in distribution to a non-degenerate limit as $n \rightarrow \infty$.

If conditions B1, B2 and B3 are met then the estimates $(\hat{\beta}_j, \hat{\sigma}^2, \tilde{\tau}^2, \tilde{\sigma}^2)$ defined in Section 3.1 satisfy the conditions C1, C2, A1 and A2 and so the FAB interval for β_j of any variable j satisfying condition A3 will satisfy the conditions of Theorem 1, and hence be asymptotically optimal. To see that this holds, first note that the definition of the model in condition B1 implies that $(\hat{\beta}_j, \hat{\sigma}^2, \tilde{\tau}^2, \tilde{\sigma}^2)$ satisfy the coverage conditions C1 and C2. Second, asymptotic condition A1 is met by the definition $\sigma^2 = n\sigma_\infty^2$ in B1 and that n is growing faster than p as assumed by B2. The remaining necessary result is the following:

Lemma 2. *Suppose B1, B2 and B3 hold, and that $(\tau^2, \sigma_\infty^2) \in \Theta$, a compact subset of $[0, \infty) \times (0, \infty)$. Let $(\tilde{\tau}^2, \tilde{\sigma}^2)$ be the maximizers over Θ of the likelihood given by the marginal model (9). Then $(\tilde{\tau}^2, \tilde{\sigma}^2/n) \rightarrow (\tau^2, \sigma_\infty^2)$ in probability as $n \rightarrow \infty$.*

This result is proven in the appendix. Putting Lemma 1 and Lemma 2 gives the following summary of the asymptotic behavior of $C_{\bar{s}}$:

Theorem 3. *Under the conditions of Lemma 2, for any variable j for which A3 holds, the FAB interval $C_{\bar{s}}(\hat{\beta}_j, \hat{\sigma}^2)$ has prior expected width $E[|C_{\bar{s}}|]$ that satisfies*

$E[|C_{\tilde{s}}|] \rightarrow E[|C_s|]$ as $n \rightarrow \infty$, where C_s is the Bayes-optimal FAB procedure for the case that $\hat{\beta}_j \sim N(\beta_j, w_0^2 \sigma_\infty^2)$ and $\beta_j \sim N(0, \tau^2)$.

This result makes precise the heuristic idea that if n and p are large, then the adaptive FAB interval should be nearly as good as the oracle FAB interval.

5. Alternative spending functions

5.1. Bounded-width adaptive FAB interval

As discussed above, the adaptive FAB interval based on a normal prior distribution is approximately optimal if the β_j 's are normally distributed with a common mean and variance. If the empirical distribution of the β_j 's is not approximately normal, then, while the adaptive FAB procedure will still have $1 - \alpha$ coverage for each β_j , some modifications to the procedure may be desired to achieve better expected interval width, particularly if there are a small number of outliers among β_1, \dots, β_p .

If it is known in advance that certain subvectors of β are likely to have very different magnitudes (for example, if β includes main effects as well as interactions), it makes sense to adaptively estimate a separate prior distribution for each subvector. This possibility is discussed further in Section 6.3.

If a single but unknown element of β , say β_1 , is an extreme outlier relative to the others, then the width improvement of the adaptive FAB intervals for β_2, \dots, β_p will be reduced, while the width for β_1 could be much larger than the UMAU interval. To see this, recall that the adaptive FAB interval for β_j is based on a $N(0, \tilde{\tau}^2)$ prior distribution, where $\tilde{\tau}^2$ is an estimate of the variability of $\beta_1, \dots, \beta_{j-1}, \beta_{j+1}, \dots, \beta_p$ around zero. If β_1 is an extreme outlier relative to the other β_j 's, then the estimate $\tilde{\tau}^2$ will be inflated relative to the average magnitude of the other β_j 's, making the adaptive FAB interval closer to the UMAU interval than if β_1 were smaller in magnitude. Of more concern may be the interval for the outlying value β_1 . Since the value of $\tilde{\tau}^2$ used in the prior for β_1 is estimating the magnitude of β_j 's that are much smaller than β_1 , the value of $\tilde{\tau}^2$ could be much smaller than β_1^2 , and so the adaptive FAB interval is expected to be larger than the UMAU interval, with width increasing as β_1 increases in magnitude.

Such an issue can be avoided by modifying the adaptive FAB interval so that its width relative to the UMAU interval is bounded. The modification is to construct an adaptive bounded-width FAB confidence interval $C_{\tilde{s}}$ via Equation 8 but with the estimated spending function $\tilde{s}(\beta_j)$ replaced by

$$\tilde{s}(\beta_j) = \max\{\gamma, \min\{\tilde{s}(\beta_j), 1 - \gamma\}\},$$

where γ is a user-specified value between 0 and 1/2. To see how this bounds the interval width, let β_l and β_u be the lower and upper endpoints of the interval, respectively, for a given value of $\hat{\beta}_j$. Recall that these endpoints are defined

as solutions to $\beta_l = \hat{\beta}_j + \hat{\sigma}t_{\alpha(1-\check{s}(\beta_l))}$ and $\beta_u = \hat{\beta}_j + \hat{\sigma}t_{1-\alpha\check{s}(\beta_u)}$. Therefore, the interval width as a function of $\hat{\beta}_j$ is given by $\hat{\sigma} \times (t_{1-\alpha\check{s}(\beta_u)} - t_{\alpha(1-\check{s}(\beta_l))})$. Suppose $\hat{\beta}_j$ is large and negative and so $\check{s}(\beta_l) = \gamma$, the minimum value of \check{s} . Then $\check{s}(\beta_u) \geq \gamma$ because \check{s} is nondecreasing, and so the interval width is less than or equal to $\hat{\sigma} \times (t_{1-\alpha\gamma} - t_{\alpha(1-\gamma)})$. Similarly, if $\hat{\beta}_j$ is very large and positive so that $\check{s}(\beta_u) = 1 - \gamma$, the maximum value of \check{s} , then the width is less than $\hat{\sigma} \times (t_{1-\alpha(1-\gamma)} - t_{\alpha\gamma})$, which is also equal to $\hat{\sigma} \times (t_{1-\alpha\gamma} - t_{\alpha(1-\gamma)})$. For in-between values of $\hat{\beta}_j$ for which neither $\check{s}(\beta_l) = \gamma$ nor $\check{s}(\beta_u) = 1 - \gamma$, this bounded-width FAB interval is the same as the standard FAB interval.

We illustrate this numerically in Figure 1, in the slightly simpler case of a 95% FAB z -interval based on $\hat{\beta} \sim N(\beta, 1)$ and the assumed prior distribution $\beta \sim N(0, 1)$. The plot on the left side of the figure shows the spending function corresponding to Pratt's z -interval given by Equation 3, plotted with a thick gray line. As $\hat{\beta} \rightarrow \infty$, the spending function at values of β near $\hat{\beta}$ converges to one, causing the quantile associated with the lower endpoint of Pratt's interval to decrease without bound. This property was mentioned in Pratt (1963), and also discussed for the special case of $\tau^2 = 0$ by Bartholomew (1971) and Farchione and Kabaila (2008). However, by using $\check{s}(\beta) = \max\{\gamma, \min\{\check{s}(\beta), 1 - \gamma\}\}$ to construct the interval, we are never using quantiles below $z_{\alpha\gamma}$ or above $z_{1-\alpha\gamma}$, and the width of the interval will be bounded. The second plot of Figure 1 shows the expected interval width of Pratt's FAB interval and its bounded-width modification, where the expectation is over $\hat{\beta} \sim N(\beta, 1)$. The original FAB interval achieves lower expected width than the UMAU and bounded-width FAB intervals over a set of β -values with high prior probability, but the width increases without bound as $|\beta| \rightarrow \infty$. In contrast, the bounded-width FAB interval still achieves improved expected width relative to the UMAU interval over a similar range of β -values, but has a bounded expected width as $|\beta| \rightarrow \infty$. The bound on the width can be made arbitrarily close to the width of the UMAU interval by letting $\gamma \rightarrow 1/2$, but doing so decreases the improvement relative to the UMAU interval.

While this bounded-width FAB interval can be used to bound the interval width in cases where there is concern that the prior variance τ^2 might be much smaller than the actual parameter value, in adaptive multiparameter settings this bounded interval procedure may have limited appeal over the non-bounded adaptive FAB version. First, if there are two or more outliers, then the estimate of $\tilde{\tau}^2$ used by each adaptive FAB interval will be based on at least one other outlier and so won't be exceedingly small. Second, in the case of a single large outlier, some practitioners may be willing to risk having a large interval for one parameter in exchange for narrow intervals for most other parameters.

5.2. Adaptive FAB regions for arbitrary prior distributions

When a normal model for the regression coefficients does not seem reasonable, an alternative to the adaptive FAB or bounded-width FAB procedures is to use a different family of prior distributions, perhaps one with heavy tails or one that




FIG 1. Comparison of spending functions and expected interval widths for the UMAU interval (solid black line), Pratt's FAB interval (thick gray line) and the bounded-width FAB interval (dashed black line) for the case that $\sigma^2 = \tau^2 = 1$, $\alpha = .05$ and $\gamma = .25$. The dashed gray line in the second plot is the bound on the width of the bounded-width FAB interval.

reflects potential sparsity among the β_j 's. Here we illustrate a general procedure for obtaining FAB confidence procedures from arbitrary prior distributions, and discuss "spike-and-slab" prior distributions as a special case.

First consider constructing a confidence procedure for a single parameter β , with estimator $\hat{\beta} \sim N(\beta, \sigma^2)$ where σ^2 is known. For every level- α confidence procedure there is a corresponding set-valued function $A(\beta_0)$ such that for each $\beta_0 \in \mathbb{R}$, $A(\beta_0)$ is the acceptance region of a level- α test of $H : \beta = \beta_0$ versus $K : \beta \neq \beta_0$. Given a prior distribution for β with density $\pi(\beta)$, the Bayes-optimal confidence region is obtained via $A_\pi(\beta_0)$, which for each value of β_0 is the acceptance region of the most powerful test of $H : \hat{\beta} \sim N(\beta_0, \sigma^2)$ versus $K_\pi : \beta \sim F_\pi$, where F_π has density

$$f_\pi(\hat{\beta}) = \int \sigma^{-1} \phi([\hat{\beta} - \beta]/\sigma) \pi(\beta) d\beta,$$

with ϕ being the standard normal density function. Here, F_π is the marginal distribution of $\hat{\beta}$, obtained by integrating the normal sampling distribution for $\hat{\beta}$ with respect to the prior distribution for β . By the Neyman-Pearson lemma, the most powerful level- α test of H versus K_π accepts when the likelihood ratio $L(\hat{\beta} : \beta_0) = f_\pi(\hat{\beta}) / (\sigma^{-1} \phi([\hat{\beta} - \beta_0]/\sigma))$ is less than some constant $c_\alpha(\beta_0)$. The acceptance regions of the tests of $H : \beta = \beta_0$ versus K_π are then, for each β_0 , given by $A_\pi(\beta_0) = \{\hat{\beta} : L(\hat{\beta} : \beta_0) < c_\alpha(\beta_0)\}$. If $A_\pi(\beta_0)$ is an interval for each β_0 , then we can write $A_\pi(\beta_0) = \{\hat{\beta} : l(\beta_0) < \hat{\beta} < u(\beta_0)\}$. We claim that such an acceptance region function corresponds to a spending function given by $s(\beta_0) = \Phi([l(\beta_0) - \beta_0]/\sigma) / \alpha$. To see this, recall that $A_\pi(\beta_0)$ is the acceptance region of a level- α test and $\hat{\beta} \sim N(\beta_0, \sigma^2)$ under the null hypothesis, so we must have $\alpha = 1 - \Phi([u(\beta_0) - \beta_0]/\sigma) + \Phi([l(\beta_0) - \beta_0]/\sigma)$, or equivalently, $1 - \Phi([u(\beta_0) - \beta_0]/\sigma) = \alpha(1 - s(\beta_0))$ and $\Phi([l(\beta_0) - \beta_0]/\sigma) = \alpha s(\beta_0)$ for some

$s(\beta_0) \in [0, 1]$. Solving for $u(\beta_0)$ and $l(\beta_0)$ gives $u(\beta_0) = \beta_0 + \sigma z_{1-\alpha(1-s(\beta_0))}$ and $l(\beta_0) = \beta_0 + \sigma z_{\alpha s(\beta_0)}$. Inverting the acceptance intervals and doing a bit of algebra gives the FAB confidence region as

$$C_\pi(\hat{\beta}) = \{\beta_0 : \hat{\beta} + \sigma z_{\alpha(1-s(\beta_0))} < \beta_0 < \hat{\beta} + \sigma z_{1-\alpha s(\beta_0)}\}.$$

A FAB t -region is obtained by replacing σ in the above formula with an estimate $\hat{\sigma}$, and replacing the z -quantiles with t -quantiles. An adaptive FAB interval is obtained by replacing the spending function $s(\beta_0)$ with an estimate \tilde{s} based on data that are statistically independent of $\hat{\beta}$ and $\hat{\sigma}$.

In theory, this procedure could be applied to any class of prior distributions to obtain an adaptive FAB confidence procedure. In practice, finding and inverting the acceptance region can be numerically challenging. The relatively simple form of Pratt's FAB interval given by (3) results from f_π having a closed-form expression (it is a normal density). In general, if there is not a closed-form expression for f_π , then the interval will require numerical integration to construct.

One class of prior distributions for which f_π is available are finite mixtures of normal distributions. A special case of such distributions are the so-called "spike and slab" prior distributions (Mitchell and Beauchamp, 1988), under which β is equal to a $N(0, \tau^2)$ random variable with probability λ , and is otherwise equal to zero. Under this prior distribution, the marginal distribution F_π of $\hat{\beta}$ is a mixture of a $N(0, \sigma^2)$ distribution and a $N(0, \tau^2 + \sigma^2)$ distribution, with probabilities $1 - \lambda$ and λ respectively. The log likelihood ratio for testing $H : \hat{\beta} \sim N(\beta_0, \sigma^2)$ versus $K_\pi : \hat{\beta} \sim F_\pi$ can be shown to be equal to

$$l(\hat{\beta} : \beta_0) = -\hat{\beta}\beta_0 + \log(ae^{b\hat{\beta}^2/2} + 1) + k,$$

where $a = \frac{\lambda}{1-\lambda}(\frac{\sigma^2}{\sigma^2+\tau^2})^{1/2}$, $b = \sigma^{-2}\frac{\tau^2}{\sigma^2+\tau^2}$ and k is constant in $\hat{\beta}$. This is convex in $\hat{\beta}$ and so the acceptance region for each β_0 is an interval. Based on the discussion above, the spending function corresponding to this prior distribution can be obtained by numerically finding $l(\beta_0)$ for each $\beta_0 \in \mathbb{R}$ and setting $s(\beta_0) = \Phi([l(\beta_0) - \beta_0]/\sigma)/\alpha$.

The spending function corresponding to one such spike and slab prior distribution (with $\tau^2 = 25$, $\lambda = .10$, $\alpha = .05$ and $\sigma^2 = 1$) is displayed in the left panel of Figure 2. For these values of the parameters the spending function is not non-decreasing, and so by item 3 of Theorem 1, the FAB t -region may not be an interval. To illustrate this, the middle panel of Figure 2 plots the acceptance region and confidence interval of the FAB t -test with 5 degrees of freedom and conditional on the standard deviation estimate $\hat{\sigma}$ being equal to 1. The acceptance region, obtained by looking at the range of $\hat{\beta}$ values between the two solid lines for each β , is an interval. However, the corresponding confidence procedure does not yield an interval for each value of $\hat{\beta}$. For values of $\hat{\beta}$ between about 8.18 and 9.52 (and between about -9.52 and -8.18), the confidence region consists of two disconnected intervals, one that contains $\hat{\beta}$ and another that is close to the prior mean (zero).




FIG 2. The spending function (left panel) and confidence region (right panel) corresponding to a spike and slab prior distribution. Dashed lines in the right plot indicate the bounds of the UMAU t -interval.

Overall, the FAB region corresponding to this spike and slab prior is narrower than the UMAU interval in a range near zero, and similar to the UMAU interval for moderately large values of $\hat{\beta}$ (e.g. on the scale of τ). However, the procedure has poor performance for values in-between. This mirrors to some extent some known results about shrinkage estimators obtained from sparsity-inducing prior distributions or penalties. For example, Leeb and Pötscher (2008) generalize some results on Hodge's estimator (LeCam, 1953) to show how certain sparse estimators of linear regression coefficients have poor performance not at zero, but nearby.

6. Numerical examples

In this section we illustrate the adaptive FAB procedure numerically, and show how it can be modified to accommodate different adaptation strategies. For example, in the next subsection we use an empirically estimated prior distribution that is not centered around zero, thereby providing improved performance if most of the effects are of a common sign. In the following subsection, we show how adaptation may be done separately for different groups of parameters, such as main effects and interactions. We also provide a simulation study that illustrates how a CIP that adapts to sparsity may have very poor coverage if the regression parameter is not actually sparse, whereas the adaptive FAB procedure maintains constant coverage for all parameter values.

6.1. Motif regression

Conlon et al. (2003) measured the binding intensity of a protein to each of $n = 287$ DNA segments, and related each intensity to scores measuring abundance

of the DNA segment in $p = 195$ genetic motifs. These data were also used as an example by Meinshausen et al. (2009), among others.

Assuming a normal linear regression model for the centered and scaled data, the usual unbiased estimate of σ^2 is 0.77, and the usual standard errors for the OLS regression coefficients range from 0.12 to 0.85 with a mean of 0.30. On the other hand, empirical Bayes estimates of μ and τ^2 under the prior $\beta \sim N_p(\mu\mathbf{1}, \tau^2\mathbf{I})$ are around 0.004 and 0.001 respectively, ($\hat{\tau} \approx 0.036$) suggesting that the true values of the elements of β are highly concentrated around zero.

We constructed 95% FAB confidence intervals for the effects of the $p = 195$ genetic motifs, using the adaptive FAB procedure described in Section 3.1 except under a $N_p(\mu\mathbf{1}, \tau^2\mathbf{I})$ distribution for β . This is to allow for the possibility that the distribution of true effects is not centered around zero, which seems reasonable for this particular dataset where it is expected that abundance has either a positive or negligible effect on binding intensity. In the analysis that follows, for each coefficient j , values of $(\tilde{\mu}, \tilde{\tau}^2, \tilde{\sigma}^2)$ are estimated from the j -specific vector \mathbf{y}_2 defined in Section 3.1, thereby ensuring that $\hat{\beta}_j$ is independent of $(\tilde{\mu}, \tilde{\tau}^2, \tilde{\sigma}^2)$ and constant coverage of the FAB confidence interval for each β_j is maintained.

The intervals are shown graphically in Figure 3, along with the UMAU intervals for comparison. The FAB intervals are shorter than the UMAU intervals for 189 of the 195 effects (97%), with relative widths ranging from 0.83 to 1.11, and being 0.85 on average across effects. The number of “significant” effects identified by the two procedures is similar: twelve of the FAB CIs and eleven of the UMAU CIs do not contain zero. However, the two procedures identify somewhat different significant motifs: seven motifs are identified by both procedures, all with positive OLS effect estimates. The FAB procedure identifies an additional five motifs all with positive effect estimates, whereas the four additional motifs identified by the UMAU procedure all have negative effect estimates.

We also computed the bounded-width FAB intervals described in Section 5.1, where the γ parameter was set to 0.25, which ensures that the t -quantiles used to construct the intervals are bounded by 0.0125 for the lower endpoint and 0.9875 for the upper. This results in shorter intervals than the standard adaptive FAB intervals for four of the parameters, with the maximum improvement being about 8%. This comes at a cost of being wider than the standard adaptive FAB interval for 191 of the parameters, and being larger by about 8% on average across all parameters.

6.2. Motif regression simulation study

Zhang and Zhang (2014) developed a confidence interval procedure for sparse parameters in high-dimensional normal linear regression models. When applied to the motif dataset, this low dimensional projection (LDP) procedure produces intervals that are narrower than the FAB intervals for all regression coefficients, with relative widths ranging from 0.27 to 0.71, and being about half as wide on average across coefficients. However, unlike the FAB and UMAU procedures,




FIG 3. 95% confidence intervals for motif regression effects. UMAU intervals are thick gray lines, FAB intervals are thin black lines.

the actual coverage rates of LDP intervals are guaranteed to achieve their nominal rates only asymptotically, and only if certain sparsity conditions on β are met.

To compare the performance of the UMAU, FAB and LDP procedures we constructed two related simulation studies based on the motif binding dataset described in the previous subsection. In each study, we obtained estimates (β_0, σ_0^2) from the real data \mathbf{y} and \mathbf{X} , and used these estimates to simulate new response vectors $\mathbf{y}^{(k)} \sim N(\mathbf{X}\beta_0, \sigma_0^2\mathbf{I})$ independently for $k = 1, \dots, 5000$. For each response vector $\mathbf{y}^{(k)}$ we construct UMAU, FAB and LDP confidence intervals for each of the $p = 195$ regression coefficients. These intervals are used to obtain Monte Carlo approximations to the finite-sample coverage rates of the LDP procedure, as well as approximations to the expected interval widths of the UMAU, FAB and LDP procedures.

In the first of these two simulation studies we simulated 5000 datasets from the model $\mathbf{y}^{(k)} \sim N_n(\mathbf{X}\beta_0, \sigma_0^2\mathbf{I})$, where \mathbf{X} is the original design matrix and β_0 is the lasso estimate from the original data, using an empirical Bayes estimate of the L_1 -penalty parameter. This resulted in a sparse β_0 -vector with 176 of the 195 coefficients being identically zero, so in the context of this simulation study, the “truth” is highly sparse. The value of σ_0^2 used to simulate the data was the usual unbiased estimate from the original data. We computed the UMAU, FAB and LDP confidence intervals for each of the 5000 simulated datasets. The widths of the FAB and LDP intervals were 85% and 43% of the UMAU interval widths respectively, on average across datasets and parameters. The empirical coverage rates of the nominal 95% LDP intervals ranged between 93.8 and 96.1 percent. There was some evidence that the coverage rates were not exactly 95%: Exact level-.05 binomial tests rejected the hypothesis that the coverage rates were 95% for 71 of the 195 regression parameters (36%). All of these 71 parameters had true values of 0, and the empirical coverage rates of 64 of these 71 parameters were larger than 95%, suggesting that LDP intervals




FIG 4. Relative interval widths (left panel) and coverage rates (right panel) for the motif simulation study with non-sparse β . Error bars for the coverage rates are Clopper-Pearson 95% intervals.

slightly overcover β_j when it is zero. However, in general the coverage rates of the LDP procedure were very close to the nominal rates, in this case where the truth is sparse.

The second simulation study was the same as the first except the value β_0 used to generate the simulated data was the OLS estimate from the original data, and so in this case the “true” regression model is not sparse. On average across the 5000 simulated datasets and 195 parameters, the widths of the FAB and LDP intervals were 88% and 54% of the UMAU interval widths respectively, similar to the results from the first study. These relative widths are shown in the left panel of Figure 4. However, the coverage rates for the LDP intervals were generally far from their nominal levels: Based on exact binomial tests, coverage rates for 183 of the 195 parameters were significantly different from 95% (at level 0.05). As shown in the right panel of Figure 4, the LDP intervals generally overcover parameter values near zero, and greatly undercover parameters larger in magnitude. For comparison, the empirical coverage rates of the FAB intervals are also shown. These rates show no evidence of deviation from the nominal rates, as should be the case - the FAB intervals have exact 95% coverage for each component of β by construction.

6.3. Diabetes progression

Efron et al. (2004) considered parameter estimation for a model of diabetes progression from data on ten explanatory variables from each of $n = 442$ subjects. The expected progression of a subject was assumed to be a linear function of the linear, quadratic and two-way interaction effects of the ten variables, resulting in a linear model with $p = 64$ regressors total (the binary sex variable does not

have a separate quadratic effect).

We expect that main effects will be larger than quadratic effects and two-way interactions. For this reason, we obtain adaptive intervals separately for these three types of parameters, so that the spending function \tilde{s} used to obtain the confidence interval for the effect of a given regressor is obtained adaptively from the estimated effects of regressors in the same category. This can easily be done as follows: Write the design matrix \mathbf{X} as $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3]$, where \mathbf{X}_1 , \mathbf{X}_2 , \mathbf{X}_3 are the design matrices corresponding to the main effects, quadratic effects and two-way interactions, respectively, and let $\boldsymbol{\beta}^\top = [\boldsymbol{\beta}_1^\top, \boldsymbol{\beta}_2^\top, \boldsymbol{\beta}_3^\top]$ be the corresponding partition of $\boldsymbol{\beta}$. To obtain the FAB CIs for the main effects, we let \mathbf{G} be an orthonormal basis for the null space of $[\mathbf{X}_2, \mathbf{X}_3]$. Letting $\tilde{\mathbf{y}} = \mathbf{G}^\top \mathbf{y}$ and $\tilde{\mathbf{X}} = \mathbf{G}^\top \mathbf{X}$, we have $\tilde{\mathbf{y}} \sim N_{n-p_2-p_3}(\tilde{\mathbf{X}}\boldsymbol{\beta}_1, \sigma^2\mathbf{I})$. We can then apply the FAB CI procedure to $(\tilde{\mathbf{y}}, \tilde{\mathbf{X}})$ to obtain intervals that adapt to the magnitude of $\boldsymbol{\beta}_1$ (and not to the magnitudes of $\boldsymbol{\beta}_2$ and $\boldsymbol{\beta}_3$). Adaptive confidence intervals for $\boldsymbol{\beta}_2$ and $\boldsymbol{\beta}_3$ can be obtained analogously.

In the analysis that follows we use an adaptively estimated $N(0, \tau^2)$ prior distribution for each coefficient. Recall that our FAB procedure generates an empirical Bayes estimate $\tilde{\tau}^2$ of $\tau^2 = \text{Var}[\beta_j]$ for each coefficient j that is statistically independent of the OLS estimate $\hat{\beta}_j$. For the main effects the values of $\tilde{\tau}$ ranged between 0.19 and 0.21, with a mean of 0.20, and were larger than the standard errors of the OLS coefficients except for those of four somewhat co-linear predictors. In contrast, values of $\tilde{\tau}$ for the quadratic and interaction terms were all less than 0.03, and were all less than the corresponding standard errors.

We computed the adaptive FAB interval for each regression coefficient using these coefficient-specific estimates of τ^2 . The FAB intervals are as narrow or narrower than all but three of the corresponding UMAU intervals, with the relative interval widths ranging from 0.84 to 1.0003, and being 0.86 on average. The FAB CIs for the main effects are essentially the same as the UMAU CIs, whereas the FAB CIs for the quadratic and interaction terms are all narrower than the corresponding UMAU intervals, by about 16% on average. This example illustrates some flexibility of the FAB procedure, in that the adaptation for a particular parameter may be based on a subset of the data information that is deemed most relevant for that parameter.

7. Discussion

We have constructed a class of $1 - \alpha$ confidence interval procedures (CIPs) for individual regression coefficients of the normal regression model $\mathbf{y} \sim N_n(\mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{I})$. Each member of this class corresponds to a spending function $s : \mathbb{R} \rightarrow [0, 1]$. Under the regression model, every member of the class has constant $1 - \alpha$ coverage for all possible values of $\boldsymbol{\beta}$, σ^2 and full-rank design matrices \mathbf{X} . We have described a method of adaptively selecting the spending function so that the across-parameter average interval width is reduced, and the $1 - \alpha$ coverage rate is maintained for each regression coefficient. The coverage guarantee is non-

asymptotic, does not rely on β being sparse and does not rely on conditions on the design matrix (other than it being full rank). However, under some assumptions on the distribution of the elements of β and the design matrix, the adaptive technique we propose is asymptotically optimal as both n and p increase.

The spending function $s(\beta)$ that we adaptively estimate from the data is based on a normal prior distribution for the elements of β . As such, we expect our procedure to provide the most improvement when the empirical distribution of β_1, \dots, β_p is approximately normal. If instead we suspect that β is sparse, it may seem preferable to base the adaptation on other families of prior distributions, such as Laplace or “spike and slab” distributions. Some numerical work not presented here suggests that FAB intervals obtained using the Laplace family of priors are in practice similar to those obtained with normal priors. However, FAB procedures based on spike and slab priors do seem more efficient but also present a problem: The spending function for a spike and slab prior is not generally nondecreasing, and so by Theorem 1 the corresponding confidence region may not be an interval. We suspect that non-interval confidence regions have limited appeal in practice, but even if they were of interest they present the numerical challenge of identifying multiple disconnected sets of parameter values to include in the region.

Adaptive FAB intervals for linear regression coefficients may be computed using the R-package `fabCI`. Complete replication code for the numerical examples in this article is available at the first author’s website. This research was partially supported by NSF grant DMS-1505136.

Proofs

Proof of Theorem 1

Items 1 and 2 of the theorem were shown in Yu and Hoff (2018). To prove Item 3, suppose $s(\theta)$ is not nondecreasing so that there exists $\theta_1 < \theta_2$ with $s(\theta_1) > s(\theta_2)$. We will show that there are a range of values of $(\hat{\theta}, \hat{\sigma})$ with $\hat{\theta} < \theta_1$ for which θ_2 (and $\hat{\theta}$) are in $C_s(\hat{\theta}, \hat{\sigma})$ but θ_1 is not. Let $\underline{t}_j = t_{\alpha(1-s(\theta_j))}$ and $\bar{t}_j = t_{1-\alpha s(\theta_j)}$ for $j = 1, 2$. Both $t_{\alpha(1-s)}$ and $t_{1-\alpha s}$ are decreasing in s so $\underline{t}_1 < \underline{t}_2 < \bar{t}_1 < \bar{t}_2$. For θ_2 to be in the confidence region and θ_1 not to be, we need $\hat{\theta} + \hat{\sigma}\underline{t}_2 < \theta_2 < \hat{\theta} + \hat{\sigma}\bar{t}_2$ and $\theta_1 > \hat{\theta} + \hat{\sigma}\bar{t}_1$, or equivalently

$$\bar{t}_1 < (\theta_1 - \hat{\theta})/\hat{\sigma} \tag{10}$$

$$\underline{t}_2 < (\theta_2 - \hat{\theta})/\hat{\sigma} < \bar{t}_2. \tag{11}$$

The set of values $(\hat{\theta}, \hat{\sigma})$ for which this holds has positive Lebesgue measure on $(-\infty, \theta_1) \times (0, \infty)$. For example, both $(\theta_1 - \hat{\theta})/\hat{\sigma}$ and $(\theta_2 - \hat{\theta})/\hat{\sigma}$ can be made simultaneously arbitrarily close to any number between $\bar{t}_1 \wedge \underline{t}_2$ and \bar{t}_2 by taking $\hat{\sigma}$ and $-\hat{\theta}$ to be sufficiently large. The probability of observing values of $(\hat{\theta}, \hat{\sigma})$

satisfying (10) and (11) that yield a non-interval confidence region is therefore greater than zero.

Proof of Lemma 1

For notational convenience, in this proof we write β_j and w_j as β and w , and write the α -quantile of the t_q distribution as $t(\alpha)$, suppressing the index that denotes the degrees of freedom. We begin the proof of Lemma 2 with another lemma:

Lemma 3. *The width $|C_{\bar{s}}|$ of $C_{\bar{s}}$ satisfies $|C_{\bar{s}}| < |\hat{\beta}| + w\hat{\sigma}(|t(\alpha/2)| + |t(1-\alpha/2)|)$.*

Proof. Recall that the endpoints $\underline{\beta}$ and $\bar{\beta}$ of $C_{\bar{s}}$ are solutions to

$$\begin{aligned}\hat{\beta} &= \bar{\beta} - w\hat{\sigma}t(1 - \alpha\tilde{s}(\bar{\beta})) \\ \hat{\beta} &= \underline{\beta} - w\hat{\sigma}t(\alpha(1 - \tilde{s}(\underline{\beta})))\end{aligned}$$

Here $\tilde{s}(\beta)$ is defined as $\tilde{s}(\beta) = g^{-1}(2w\tilde{\sigma}\beta/\tilde{\tau}^2)$, where $g(s) = \Phi^{-1}(\alpha s) - \Phi^{-1}(\alpha(1-s))$. At the upper endpoint, we have $\tilde{s}(\bar{\beta}) = F((\hat{\beta} - \bar{\beta})/(w\hat{\sigma}))/\alpha$, where F is the CDF of the t_q -distribution. When $\bar{\beta} > 0$, we have $\tilde{s}(\bar{\beta}) > g^{-1}(0) = 1/2$. Thus $\bar{\beta} < \hat{\beta} - w\hat{\sigma}t(\alpha/2)$. Also, $g^{-1}(2w\tilde{\sigma}\bar{\beta}/\tilde{\tau}^2) < 1$, so $\bar{\beta} > \hat{\beta} - w\hat{\sigma}t(\alpha)$. When $\bar{\beta} < 0$, $\hat{\beta} - w\hat{\sigma}t(\alpha/2) < \bar{\beta}$. This implies that

$$\begin{aligned}\hat{\beta} - w\hat{\sigma}t(\alpha) &< \bar{\beta} < \hat{\beta} - w\hat{\sigma}t(\alpha/2) && \text{if } \bar{\beta} > 0 \\ \hat{\beta} - w\hat{\sigma}t(\alpha/2) &< \bar{\beta} < 0 && \text{if } \bar{\beta} < 0.\end{aligned}$$

Similarly we have

$$\begin{aligned}0 < \underline{\beta} < \hat{\beta} - w\hat{\sigma}t(1 - \alpha/2) && \text{if } \underline{\beta} > 0 \\ \hat{\beta} - w\hat{\sigma}t(1 - \alpha/2) < \underline{\beta} < \hat{\beta} - w\hat{\sigma}t(1 - \alpha) && \text{if } \underline{\beta} < 0.\end{aligned}$$

Therefore

$$|C_{\bar{s}}| = \bar{\beta} - \underline{\beta} < |\hat{\beta}| + w\hat{\sigma}(|t(\alpha/2)| + |t(1 - \alpha/2)|).$$

□

Now we prove Lemma 1. We denote the endpoints of the oracle CIP C_s as $\bar{\beta}$ and $\underline{\beta}$, which are the solutions to

$$\begin{aligned}\bar{\beta} - w_0\sigma_\infty\Phi^{-1}(1 - \alpha s(\bar{\beta})) &= \hat{\beta} \\ \underline{\beta} - w_0\sigma_\infty\Phi^{-1}(\alpha(1 - s(\underline{\beta}))) &= \hat{\beta}.\end{aligned}$$

We denote the endpoints of $C_{\bar{s}}$ as $\bar{\beta}^n$ and $\underline{\beta}^n$, which are the solutions to

$$\begin{aligned}\bar{\beta}^n - w\hat{\sigma}t(1 - \alpha\tilde{s}(\bar{\beta}^n)) &= \hat{\beta}^n \\ \underline{\beta}^n - w\hat{\sigma}t(\alpha(1 - \tilde{s}(\underline{\beta}^n))) &= \hat{\beta}^n.\end{aligned}$$

We first prove that $|C_{\hat{s}}| - |C_s| = (\bar{\beta}^n - \bar{\beta}) + (\underline{\beta} - \underline{\beta}^n) \xrightarrow{P} 0$ as $n \rightarrow \infty$ for each fixed $\hat{\beta}$. We can write the upper endpoints as $\bar{\beta}^n = G_n(w\hat{\sigma}, w\tilde{\sigma}, \tilde{\tau}^2, \hat{\beta}^n)$, and $\bar{\beta} = G(w_0\sigma_\infty, w_0\sigma_\infty, \tau^2, \hat{\beta})$, where G and G_n are continuous functions of their parameters. The functions G and G_n are different in that the former is based on z -quantiles, while the latter uses t -quantiles. We have

$$\begin{aligned} |\bar{\beta}^n - \bar{\beta}| &= |G_n(w\hat{\sigma}, w\tilde{\sigma}, \tilde{\tau}^2, \hat{\beta}^n) - G(w_0\sigma_\infty, w_0\sigma_\infty, \tau^2, \hat{\beta})| \\ &\leq |G_n(w\hat{\sigma}, w\tilde{\sigma}, \tilde{\tau}^2, \hat{\beta}^n) - G(w\hat{\sigma}, w\tilde{\sigma}, \tilde{\tau}^2, \hat{\beta}^n)| \\ &\quad + |G(w\hat{\sigma}, w\tilde{\sigma}, \tilde{\tau}^2, \hat{\beta}^n) - G(w_0\sigma_\infty, w_0\sigma_\infty, \tau^2, \hat{\beta})|. \end{aligned} \quad (12)$$

The second term converges to zero in probability because $(w\hat{\sigma}, w\tilde{\sigma}, \tilde{\tau}^2, \hat{\beta}^n) \xrightarrow{P} (w_0\sigma_\infty, w_0\sigma_\infty, \tau^2, \hat{\beta})$. Elaborating on the convergence of the first term, note that G_n is a monotone sequence of continuous functions: Given G_{n_1} and G_{n_2} where $n_2 > n_1$, and suppose the corresponding degrees-of-freedom of the t -quantiles are q_1 and q_2 . We have $q_2 \geq q_1$, thus $t_{q_2}(1 - \alpha\tilde{s}) \leq t_{q_1}(1 - \alpha\tilde{s})$. Hence $\hat{\beta}^n - w\hat{\sigma}t_{q_2}(1 - \alpha\tilde{s}(\hat{\beta}^n)) \geq \hat{\beta}^n - w\hat{\sigma}t_{q_1}(1 - \alpha\tilde{s}(\hat{\beta}^n))$. Therefore $G_{q_2}(w\hat{\sigma}^2, w\tilde{\sigma}^2, \tilde{\tau}^2, \hat{\beta}^n) \leq G_{q_1}(w\hat{\sigma}^2, w\tilde{\sigma}^2, \tilde{\tau}^2, \hat{\beta}^n)$, and so by Dini's theorem, $G_n \rightarrow G$ uniformly on a compact set of $(s^2\hat{\sigma}^2, s^2\tilde{\sigma}^2, \tilde{\tau}^2, \hat{\beta}^n)$ values. Since $(s^2\hat{\sigma}^2, s^2\tilde{\sigma}^2, \tilde{\tau}^2, \hat{\beta}^n)$ converges in probability to $(v^2\sigma_\infty^2, v^2\sigma_\infty^2, \tau^2, \hat{\beta})$, for arbitrary $\epsilon > 0$ and $\delta > 0$, there exists a number $N(\epsilon, \delta)$ such that when $n > N(\epsilon, \delta)$, $|w\hat{\sigma} - w_0\sigma_\infty| \leq \delta$, $|w\tilde{\sigma} - w_0\sigma_\infty| \leq \delta$, $|\tilde{\tau}^2 - \tau^2| \leq \delta$ and $|\hat{\beta}^n - \hat{\beta}| \leq \delta$ with probability at least $1 - \epsilon$. Therefore, for arbitrary $\eta > 0$,

$$\lim_{n \rightarrow \infty} P(|G_n(w\hat{\sigma}, w\tilde{\sigma}, \tilde{\tau}^2, \hat{\beta}^n) - G(w\hat{\sigma}, w\tilde{\sigma}, \tilde{\tau}^2, \hat{\beta})| < \eta) > 1 - \epsilon.$$

Since ϵ is arbitrary, we conclude that the first term in (12) converges to zero in probability.

Now we show the expected width converges to the oracle width by integrating over $\hat{\beta}$. This is done by first showing $|C_{\hat{s}}|$ is uniformly integrable and then applying Vitali's theorem. By the previous lemma we know that

$$|C_{\hat{s}}| < |\hat{\beta}^n| + w\hat{\sigma}(|t(\alpha/2)| + |t(1 - \alpha/2)|).$$

Note that $|t(\alpha/2)| + |t(1 - \alpha/2)| < |t_1(\alpha/2)| + |t_1(1 - \alpha/2)| = c_1 < \infty$, where t_1 is the t -quantile with one degree of freedom. We now show $|C_{\hat{s}}|$ is L_2 -bounded. We have

$$|C_{\hat{s}}|^2 < |\hat{\beta}^n|^2 + c_1^2 w^2 \hat{\sigma}^2 + 2|\hat{\beta}^n|c_1 w\hat{\sigma}.$$

Here $E[|\hat{\beta}^n|^2] = \beta^2 + w^2\sigma^2$ and $E[w^2\hat{\sigma}^2] = w^2\sigma^2$. Since $w^2\sigma^2 \rightarrow w_0^2\sigma_\infty^2 < \infty$, thus $E[|\hat{\beta}^n|^2]$ and $E[w^2\hat{\sigma}^2]$ are both bounded for all n . Similarly, $E[w\hat{\sigma}]$ and $E[|\hat{\beta}^n|]$ are also bounded. Therefore, it is easy to see that $|C_{\hat{s}}|$ is L_2 -bounded, which implies that $|C_{\hat{s}}|$ is uniformly integrable. By Vitali's Theorem, $\lim_{n \rightarrow \infty} E[|C_{\hat{s}}|] = E[|C_s|] = 0$.

Proof of Lemma 2

We first prove a consistency result for a notationally simpler model, and then discuss how the result applies to the marginal model (9). The simpler model is $\mathbf{y} \sim N_n(\mathbf{0}, \Lambda\tau^2 + \mathbf{I}\sigma^2)$ where Λ is a known diagonal matrix with positive entries. Let $\theta_0 = (\tau_0^2, \sigma_0^2)$ be the true value of $\theta = (\tau^2, \sigma^2)$, and let $q_i(\theta) = y_i^2/(\lambda_i\tau^2 + \sigma^2) - \log \frac{\lambda_i\tau_0^2 + \sigma_0^2}{\lambda_i\tau^2 + \sigma^2}$, which is -2 times the contribution of y_i to the log likelihood plus a constant. The MLE is therefore given by $\hat{\theta} = \arg \min Q_n(\theta)$, where $Q_n(\theta) = \sum_{i=1}^n q_i(\theta)/n$.

Lemma 4. *Assume $\theta_0 \in \Theta$, a compact subset of $[0, \infty) \times (0, \infty)$, and let $\hat{\theta}_n = \arg \min_{\Theta} Q_n(\theta)$. For each n assume that F_n , the empirical distribution of the diagonal entries of Λ , has support on $[0, \bar{\lambda}]$, where $\bar{\lambda}$ is fixed for all n . If F_n converges weakly to a nondegenerate distribution F_0 as $n \rightarrow \infty$ then $\hat{\theta} \xrightarrow{P} \theta_0$ as $n \rightarrow \infty$.*

We prove consistency of $\hat{\theta}$ in three steps: First, we show that $Q_n(\theta) - E[Q_n(\theta)]$ converges uniformly to zero as $n \rightarrow \infty$. Second, we show that this implies that as a function of $\theta \in \Theta$, Q_n converges uniformly to a function Q_0 . Third, we show that Q_0 is uniquely minimized at θ_0 . Consistency of $\hat{\theta}$ follows from these latter two results (see, for example, Theorem 2.1 of Newey and McFadden (1994)).

For each n the expectation of $Q_n(\theta) - E[Q_n(\theta)]$ is zero and the variance is given by

$$\text{Var}[Q_n(\theta)] = \frac{2}{n} \left(\frac{1}{n} \sum_{i=1}^n \left(\frac{\lambda_i\tau_0^2 + \sigma_0^2}{\lambda_i\tau^2 + \sigma^2} \right)^2 \right) = \frac{2}{n} \left(\int \left(\frac{\lambda\tau_0^2 + \sigma_0^2}{\lambda\tau^2 + \sigma^2} \right)^2 dF_n(\lambda) \right).$$

The integrand is a bounded function of λ on any bounded set $[0, \bar{\lambda}]$ as long as $\sigma^2 > 0$. Therefore, if F_n converges weakly to a distribution F_0 with such bounded support, then for all $\theta = (\tau^2, \sigma^2) \in [0, \infty) \times (0, \infty)$ the integral in the parentheses converges to a finite limit and the variance of $Q_n(\theta) - E[Q_n(\theta)]$ converges to zero. Thus $Q_n(\theta) - E[Q_n(\theta)] \xrightarrow{P} 0$ as $n \rightarrow \infty$. To show that this convergence is uniform we use Theorem 3 of Andrews (1992). Using basic calculus, it can be shown that the functions $q_i(\theta)$ satisfy the Lipschitz condition $|q_i(\theta) - q_i(\theta')| < g(y_i, \lambda_i) \times \|\theta - \theta'\|$ where $g(y_i, \lambda_i) = (1 + y_i^2/\epsilon)(\lambda_i + 1)/\epsilon$ and $\epsilon = \min_{\Theta} \sigma^2$. Andrews' Theorem 5 says that if (i) Θ is compact, (ii) $Q_n \xrightarrow{P} 0$ for each $\theta \in \Theta$, and (iii) $\sup_{n \geq 1} \sum E[g(y_i, \lambda_i)]/n < \infty$, then $\sup_{\Theta} |Q_n(\theta)| \xrightarrow{P} 0$. This last condition holds in this case because $E[g(y_i, \lambda_i)]$ is quadratic in λ_i , the values of which are all bounded by assumption. Therefore, $Q_n(\theta)$ converges uniformly in probability to zero as $n \rightarrow \infty$.

Now consider the limiting value of $E[Q_n(\theta)]$. We have

$$\begin{aligned} \lim_{n \rightarrow \infty} E[Q_n(\theta)] &= \lim_{n \rightarrow \infty} \int \left(\frac{\lambda\tau_0^2 + \sigma_0^2}{\lambda\tau^2 + \sigma^2} - \log \frac{\lambda\tau_0^2 + \sigma_0^2}{\lambda\tau^2 + \sigma^2} \right) dF_n(\lambda) \\ &= \int \left(\frac{\lambda\tau_0^2 + \sigma_0^2}{\lambda\tau^2 + \sigma^2} - \log \frac{\lambda\tau_0^2 + \sigma_0^2}{\lambda\tau^2 + \sigma^2} \right) dF_0(\lambda) \equiv Q_0(\theta), \end{aligned}$$

for all values of θ for which the integrand is a bounded function of λ . On $\lambda \in [0, \bar{\lambda}]$, the ratio $(\lambda\tau_0^2 + \sigma_0^2)/(\lambda\tau^2 + \sigma^2)$ is bounded between $\min\{\frac{\bar{\lambda}\tau_0^2 + \sigma_0^2}{\lambda\tau^2 + \sigma^2}, \frac{\sigma_0^2}{\sigma^2}\}$ and $\max\{\frac{\bar{\lambda}\tau_0^2 + \sigma_0^2}{\lambda\tau^2 + \sigma^2}, \frac{\sigma_0^2}{\sigma^2}\}$, and so the integrand is bounded in $\lambda \in [0, \bar{\lambda}]$ for all $(\tau^2, \sigma^2) \in [0, \infty) \times (0, \infty)$. Furthermore, the convergence of $E[Q_n(\theta)]$ to $Q_0(\theta)$ is uniform, since the integrand is bounded as a function of θ on the compact set Θ (Ranga Rao (1962)). Together with the uniform convergence in probability of $Q_n(\theta) - E[Q_n(\theta)]$ to zero, this implies uniform convergence in probability of Q_n to Q_0 .

Finally we show that $Q_0(\theta)$ has a unique minimizing value at θ_0 . After computing the gradient of $Q_0(\theta)$, it is easily shown that a critical point (τ^2, σ^2) must satisfy

$$\int \lambda^k \frac{\lambda\tau^2 + \sigma^2}{(\lambda\tau^2 + \sigma^2)^2} dF_0(\lambda) = \int \lambda^k \frac{\lambda\tau_0^2 + \sigma_0^2}{(\lambda\tau^2 + \sigma^2)^2} dF_0(\lambda)$$

for $k \in \{0, 1\}$. Rearranging, it can be shown that a critical point satisfies

$$\begin{pmatrix} m_1 & 1 \\ m_2 & m_1 \end{pmatrix} \begin{pmatrix} \tau^2 \\ \sigma^2 \end{pmatrix} = \begin{pmatrix} m_1 & 1 \\ m_2 & m_1 \end{pmatrix} \begin{pmatrix} \tau_0^2 \\ \sigma_0^2 \end{pmatrix} \quad (13)$$

where $m_k(\tau^2, \sigma^2)$, $k \in \{1, 2\}$ is given by

$$m_k(\tau^2, \sigma^2) = \frac{\int \lambda^k (\lambda\tau^2 + \sigma^2)^{-2} dF_0(\lambda)}{\int (\lambda\tau^2 + \sigma^2)^{-2} dF_0(\lambda)}.$$

For each $(\tau^2, \sigma^2) \in [0, 1) \times (0, \infty)$, m_1 and m_2 are the first and second moments of λ under a probability measure having density with respect to F_0 proportional to $(\lambda\tau^2 + \sigma^2)^{-2}$. If F_0 is not degenerate, then $m_2 > m_1^2$, and so the determinant of the matrix in (13) is non-zero. Therefore, the matrix is invertible and (τ_0^2, σ_0^2) is the only solution to (13). The matrix of second derivatives of Q_0 is given by

$$\frac{\partial^2 Q_0}{\partial\theta\partial\theta^\top} = \int \begin{pmatrix} \lambda^2 & \lambda \\ \lambda & 1 \end{pmatrix} \frac{2(\lambda\tau_0^2 + \sigma_0^2) - (\lambda\tau^2 + \sigma^2)}{\lambda\tau^2 + \sigma^2} \times (\lambda\tau^2 + \sigma^2)^{-2} dF_0(\lambda).$$

At the critical point (τ_0^2, σ_0^2) this simplifies to the expectation of the matrix in the integrand with respect to the probability measure with density proportional to $(\lambda\tau^2 + \sigma^2)^{-2}$ with respect to F_0 . Again, if F_0 is not degenerate then the expectation of this matrix, and hence the Hessian of Q_0 , is strictly positive definite. The critical point is a local minimum, and since it is the only critical point of the continuous function Q_0 , it is the unique minimizer. This completes the proof of Lemma 4.

To see how this applies to the properties of the empirical Bayes estimates $(\tilde{\tau}^2, \tilde{\sigma}^2)$ of (τ^2, σ^2) based on the marginal model (9), let \mathbf{U} be the $(p-1) \times (p-1)$ matrix of left singular vectors of \mathbf{X}_2 , and let $n\mathbf{\Lambda}_2$ be the diagonal matrix of the squared singular values. Then $\mathbf{U}_2^\top \mathbf{z}_2 / \sqrt{n} \sim N_{p-1}(\mathbf{0}, \mathbf{\Lambda}_2\tau^2 + \mathbf{I}\sigma_\infty^2)$, and so the properties of the MLE of $(\tau^2, \sigma_\infty^2)$ based on \mathbf{z}_2 will be the same as those of (τ^2, σ^2) in Lemma 4 if $\mathbf{\Lambda}_2$ satisfies the assumption of the Lemma. To see that it

does, recall that the assumption of Lemma 2 was that the empirical distribution of the eigenvalues of $\mathbf{X}^\top \mathbf{X}/n$ is uniformly bounded and converges weakly to a non-degenerate distribution with finite support. For a given n , let $\gamma_1, \dots, \gamma_p$ be the eigenvalues of $\mathbf{X}^\top \mathbf{X}/n$. Since $\mathbf{X}_2^\top \mathbf{X}_2/n$ is a compression of $\mathbf{X}^\top \mathbf{X}/n$, by the Cauchy interlacing theorem we have $\gamma_1 \leq \lambda_1 \leq \gamma_2 \leq \dots \leq \gamma_{p-1} \leq \lambda_{p-1} \leq \gamma_p$. Therefore, if the values of $\{\gamma_1, \dots, \gamma_p\}$ are bounded uniformly in p and have an empirical distribution that converges to a nondegenerate limit, then the same properties hold for the values of $\{\lambda_1, \dots, \lambda_{p-1}\}$.

References

- Andrews, D. W. K. (1992). Generic uniform convergence. *Econometric Theory* 8(2), 241–257. [MR1179512](#)
- Bartholomew, D. J. (1971). A comparison of frequentist and Bayesian approaches to inference with prior knowledge. pp. 417–434. With comments by G. A. Barnard, R. J. Buehler, D. R. Cox, V. P. Godambe, I. J. Good, W. J. Hall, J. A. Hartigan, O. Kempthorne and J. W. Pratt and a reply by the author. [MR0378156](#)
- Bühlmann, P. (2013). Statistical significance in high-dimensional linear models. *Bernoulli* 19(4), 1212–1242. [MR3102549](#)
- Carlin, B. P. and A. E. Gelfand (1990). Approaches for empirical Bayes confidence intervals. *J. Amer. Statist. Assoc.* 85(409), 105–114. [MR1137356](#)
- Conlon, E. M., X. S. Liu, J. D. Lieb, and J. S. Liu (2003). Integrating regulatory motif discovery and genome-wide expression analysis. *Proceedings of the National Academy of Sciences* 100(6), 3339–3344.
- Efron, B., T. Hastie, I. Johnstone, and R. Tibshirani (2004). Least angle regression. *Ann. Statist.* 32(2), 407–499. With discussion, and a rejoinder by the authors. [MR2060166](#)
- Farchione, D. and P. Kabaila (2008). Confidence intervals for the normal mean utilizing prior information. *Statist. Probab. Lett.* 78(9), 1094–1100. [MR2422965](#)
- Kabaila, P. and K. Giri (2009). Confidence intervals in regression utilizing prior information. *J. Statist. Plann. Inference* 139(10), 3419–3429. [MR2549091](#)
- Kabaila, P. and D. Tissera (2014). Confidence intervals in regression that utilize uncertain prior information about a vector parameter. *Australian & New Zealand Journal of Statistics* 56(4), 371–383. [MR3300167](#)
- LeCam, L. (1953). On some asymptotic properties of maximum likelihood estimates and related Bayes' estimates. *Univ. California Publ. Statist.* 1, 277–329. [MR0054913](#)
- Lee, J. D., D. L. Sun, Y. Sun, and J. E. Taylor (2016). Exact post-selection inference, with application to the lasso. *Ann. Statist.* 44(3), 907–927. [MR3485948](#)
- Leeb, H. and B. M. Pötscher (2008). Sparse estimators and the oracle property, or the return of Hodges' estimator. *J. Econometrics* 142(1), 201–211. [MR2394290](#)
- Meinshausen, N., L. Meier, and P. Bühlmann (2009). p -values for high-

- dimensional regression. *J. Amer. Statist. Assoc.* 104(488), 1671–1681. [MR2750584](#)
- Mitchell, T. J. and J. J. Beauchamp (1988). Bayesian variable selection in linear regression. *J. Amer. Statist. Assoc.* 83(404), 1023–1036. With comments by James Berger and C. L. Mallows and with a reply by the authors. [MR0997578](#)
- Newey, W. K. and D. McFadden (1994). Large sample estimation and hypothesis testing. In *Handbook of econometrics, Vol. IV*, Volume 2 of *Handbooks in Econom.*, pp. 2111–2245. North-Holland, Amsterdam. [MR1315971](#)
- O’Gorman, T. W. (2001). Using adaptive weighted least squares to reduce the lengths of confidence intervals. *Canad. J. Statist.* 29(3), 459–471. [MR1872647](#)
- Pratt, J. W. (1963). Shorter confidence intervals for the mean of a normal distribution with known variance. *The Annals of Mathematical Statistics* 34(2), 574–586. [MR0148150](#)
- Puza, B. and T. O’Neill (2006). Interval estimation via tail functions. *Canadian Journal of Statistics* 34(2), 299–310. [MR2324011](#)
- Ranga Rao, R. (1962). Relations between weak and uniform convergence of measures with applications. *Ann. Math. Statist.* 33, 659–680. [MR0137809](#)
- Stein, C. M. (1962). Confidence sets for the mean of a multivariate normal distribution. *J. Roy. Statist. Soc. Ser. B* 24, 265–296. [MR0148184](#)
- van de Geer, S., P. Bühlmann, Y. Ritov, and R. Dezeure (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *Ann. Statist.* 42(3), 1166–1202. [MR3224285](#)
- Yoshimori, M. and P. Lahiri (2014). A second-order efficient empirical Bayes confidence interval. *Ann. Statist.* 42(4), 1–29. [MR3226156](#)
- Yu, C. and P. D. Hoff (2018). Adaptive multigroup confidence intervals with constant coverage. *Biometrika* 105(2), 319–335. [MR3804405](#)
- Zhang, C.-H. and S. S. Zhang (2014). Confidence intervals for low dimensional parameters in high dimensional linear models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* 76(1), 217–242. [MR3153940](#)