

On inference in high-dimensional regression

H. S. Battey and N. Reid *

November 11, 2022

Abstract

This paper develops an approach to inference in a linear regression model when the number of potential explanatory variables is larger than the sample size. The approach treats each regression coefficient in turn as the interest parameter, the remaining coefficients being nuisance parameters, and seeks an optimal interest-respecting transformation, inducing sparsity on the relevant blocks of the notional Fisher information matrix. The induced sparsity is exploited through a marginal least squares analysis for each variable, as in a factorial experiment, thereby avoiding penalization. One parameterization of the problem is found to be particularly convenient, both computationally and mathematically. In particular, it permits an analytic solution to the optimal transformation problem, facilitating theoretical analysis and comparison to other work. In contrast to regularized regression such as the lasso and its extensions, neither adjustment for selection nor rescaling of the explanatory variables is needed, ensuring the physical interpretation of regression coefficients is retained. Recommended usage is within a broader set of inferential statements, so as to reflect uncertainty over the model as well as over the parameters. The considerations involved in extending the work to other regression models are briefly discussed.

Some key words: conditionality; confidence sets of models; factorial contrasts; fixed design; inducement of sparsity; nuisance parameters; parameter orthogonalization.

1 Introduction

We consider inference for the coefficient parameters of a high-dimensional linear regression model in which the number of potential explanatory variables p is larger than the sample size n . In this context, the debiased lasso (Zhang and Zhang, 2014; van de Geer et al., 2014) and decorrelated score (Ning and Liu, 2017) require two sparsity assumptions: a natural one on the parameter vector, and a less natural one on the inverse Fisher information matrix. The key contribution of the present paper is an approach that induces sparsity on relevant

* Heather Battey (h.battey@imperial.ac.uk). Department of Mathematics, Imperial College London, 180 Queen's Gate, London SW7 2AZ, UK.

* Nancy Reid (nancym.reid@utoronto.ca). Department of Statistical Sciences, University of Toronto, 700 University Ave, 9th Floor, Toronto, Ontario M5G 1Z5, Canada.

blocks of the notional Fisher information matrix, removing the second assumption without affecting interpretation of the parameter of interest. The induced sparsity is exploited through a marginal least squares analysis for each variable, as in a factorial experiment. Since regularized regression is not used, there is no need to standardize the columns of the covariate matrix to have unit length. Thus, the physical interpretation of coefficients is retained.

Primary inspiration for the proposal came from the analysis of matched comparison problems, in which it is sometimes feasible to eliminate a large number of nuisance parameters by simple operations on the sample space. There are also connections to Cox and Reid (1987) as discussed in §2. Section 4 provides a more detailed comparison to the debiased lasso and related proposals.

The remainder of the present section focusses on the broader inferential strategy. Usual usage of the debiased lasso and its extensions, as described for example in Bühlmann et al. (2014), involves hypothesis tests for all the parameters followed by a correction for multiple testing. The conclusion is typically that the null hypothesis of no effect is rejected for a very small number of variables; the joint explanatory power of sets of variables tends not to be assessed. The present paper instead proposes confidence intervals for the regression coefficient parameters as an adjunct to confidence sets of models (Cox and Battey, 2017; Battey and Cox, 2018), as detailed in §5. The latter papers emphasized the construction of statistically indistinguishable sets of variables based on their fit to the data, but did not consider inferential statements about the individual parameters or their associated compatibility with the models in the confidence set. This strategy is illustrated on a set of gene-expression data from Bühlmann et al. (2014).

Several quite different approaches to inference in high-dimensional regression have been proposed. In Lockhart et al. (2014), Lee et al. (2014), and Tian and Taylor (2018) inferential statements are corrected for variable selection by conditioning on the region of the sample space that led to the selection being made. Sample-splitting is a simple alternative that avoids detailed characterization of the selection event. The efficiency loss of sample-splitting was calculated in a simple example by Cox (1975) and several more elaborate splitting strategies have been proposed in Rasines and Young (2021) and Leiner et al. (2021). For a different perspective see Zhao et al. (2021). The approach advocated by Berk et al. (2013), and further discussed in Leeb et al. (2015), aims instead for inferential statements that are universally valid under all possible model selection procedures; this leads to highly conservative confidence regions.

In contrast with much of the earlier literature on this topic, the analysis of §3 is conditional on X , as would be dictated by ancillarity when $p < n$. In addition to ancillarity, there are substantive reasons for preferring a conditional analysis. In observational studies X , although random, is frequently realized in advance of the response variable Y , in which case it may be possible to interpret the effect of X on Y as causal. For example X may be information on the genome or epi-genome, and Y a measure of some condition expected to be caused by changes in X . However, even in cases where the values of $(Y_i, X_i)_{i=1}^n$ are realized contemporaneously,

it is appropriate to condition on X for studying its association with the outcome variable Y . Otherwise it would be possible that an inference justified on the basis of its unconditional properties performed poorly for any given sample of X .

2 Approximate orthogonalization

Let x_i be a vector of measurements on p covariates for individual i . Outcomes are assumed to be generated from a linear regression model

$$Y_i = x_i^T \beta + \varepsilon_i, \quad i = 1, \dots, n,$$

where $(\varepsilon_i)_{i=1}^n$ are independent with mean zero and variance $\tau > 0$, and $\beta = (\beta_1, \dots, \beta_p)^T$ is of dimension $p \gg n$, but satisfying the sparsity condition $\|\beta\|_0 = \sum_{u=1}^p \mathbb{1}\{\beta_u \neq 0\} = s \ll p$. The set of signal indices is $\mathcal{S} = \{v : \beta_v \neq 0\}$. The arbitrary designated interest parameter is β_v , so that the vector β_{-v} of nuisance parameters has elements $\{\beta_1, \dots, \beta_p\} \setminus \beta_v$. Let $Y = (Y_1, \dots, Y_n)^T$ and let X be the matrix whose i th row is x_i^T . The column of X corresponding to the interest parameter β_v is written $x_v = (x_{1,v}, \dots, x_{n,v})^T$, where indices u, v, w, \dots from the end of the alphabet distinguish arbitrary columns of X from arbitrary transposed rows. The remaining columns are those of the $n \times (p-1)$ dimensional matrix, X_{-v} , whose i th row is $x_{i,-v}^T$. On writing $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)^T$, the linear regression model is

$$Y = X\beta + \varepsilon = x_v\beta_v + X_{-v}\beta_{-v} + \varepsilon. \quad (1)$$

If the columns in X_{-v} are orthogonal to x_v then a simple regression of Y on x_v estimates β_v without bias, this being the motivation for factorial experiments and more elaborate experimental designs. In fact only the columns corresponding to signal variables need be orthogonal to x_v , but it is not known which columns these may be.

Since the parameter β in (1) is unchanged by premultiplication of both sides by an $n \times n$ matrix, for each interest parameter β_v a matrix A^v is sought such that the columns \tilde{x}_v^v of the transformed covariate matrix $\tilde{X}^v = A^v X$ are as orthogonal as possible to \tilde{x}_v^v . The parameter of interest β_v is then estimated by simple linear regression of $\tilde{Y}^v = A^v Y$ on \tilde{x}_v^v . The associated least squares estimator is

$$\tilde{\beta}_v = (\tilde{x}_v^{vT} \tilde{x}_v^v)^{-1} \tilde{x}_v^{vT} \tilde{Y}^v = (x_v^T A^{vT} A^v x_v)^{-1} x_v^T A^{vT} A^v Y. \quad (2)$$

Under model (1)

$$\begin{aligned} \mathbb{E}(\tilde{\beta}_v) &= \beta_v + \sum_{w \in \mathcal{S}} \vartheta_w \beta_w =: \beta_v + b_v, \\ \text{var}(\tilde{\beta}_v) &= (\tilde{x}_v^{vT} \tilde{x}_v^v)^{-1} \tilde{x}_v^{vT} A^v \mathbb{E}(\varepsilon \varepsilon^T) A^{vT} \tilde{x}_v^v (\tilde{x}_v^{vT} \tilde{x}_v^v)^{-1} =: \tau \sigma_{vv}, \end{aligned} \quad (3)$$

where

$$\begin{aligned}\vartheta_w &= \vartheta_w(A^v) &= (\tilde{x}_v^T \tilde{x}_v)^{-1} \tilde{x}_v^T \tilde{x}_w, \\ \sigma_{vv} &= \sigma_{vv}(A^v) &= (\tilde{x}_v^T \tilde{x}_v)^{-2} \tilde{x}_v^T A^v A^{vT} \tilde{x}_v,\end{aligned}\tag{4}$$

and by the Cauchy-Schwarz inequality, $b_v^2 \leq \|\beta_{-v}\|_2^2 \sum_{w \in \mathcal{S}} \vartheta_w^2$. Since $\|\beta_{-v}\|_2$ and \mathcal{S} are unknown, this suggests choosing A^v to minimize $\sigma_{vv}(A^v) + \sum_{w \neq v} \vartheta_w^2(A^v)$, an upper bound on the observable components of the mean squared error. The relevant objects in terms of $q_v = A^{vT} A^v x_v$ are

$$\begin{aligned}\tilde{\beta}_v &= (q_v^T x_v)^{-1} q_v^T Y, \\ \sigma_{vv} &= (q_v^T x_v)^{-2} (q_v^T q_v), \\ \vartheta_w^2 &= (q_v^T x_v)^{-2} (q_v^T x_w)^2.\end{aligned}\tag{5}$$

On noting that

$$(q_v^T x_w)^2 = q_v^T x_w x_w^T q_v,$$

the sum of the variance $\tau \sigma_{vv}(q_v)$ and squared potential biasing terms $\sum_{w \neq v} \vartheta_w^2(q_v)$ is minimized when q_v solves the simple unconstrained optimization problem

$$\operatorname{argmin}_{q \in \mathbb{R}^n} (q^T x_v)^{-2} q^T (I_n + \sum_{w \neq v} x_w x_w^T) q.$$

This estimator, which weights squared bias and variance equally, is a natural choice. However, a generalization is helpful for gaining geometric insight into the problem and for comparison with other work. Thus define q_v by

$$q_v \in \mathcal{Q}_v^{(\delta)} = \operatorname{argmin}_{q \in \mathbb{R}^n} (q^T x_v)^{-2} q^T (\delta I_n + X_{-v} X_{-v}^T) q,\tag{6}$$

for some fixed $\delta > 0$ to be chosen. In §3, such solutions will be shown to exist in closed form under a condition on X defined below at Proposition 1. If $q_v \in \mathcal{Q}_v^{(\delta)}$, then any non-zero scalar multiple of q_v is also in $\mathcal{Q}_v^{(\delta)}$, as can be seen from the form of the objective function. Hereafter, q_v refers to an arbitrary solution to the minimization problem of equation (6) unless otherwise specified.

The connection to parameter orthogonalization (Cox and Reid, 1987) is that, for Gaussian linear regression, both approaches seek to induce sparsity on the relevant row and column of the Fisher information matrix. With β_v identified as the interest parameter, a so-called interest-respecting orthogonal parameterization of model (1) is (β_v, η) , where

$$\eta := X_{-v}^T x_v \beta_v + X_{-v}^T X_{-v} \beta_{-v},\tag{7}$$

or alternatively, if $(X_{-v}^T X_{-v})^{-1}$ exists, (β_v, ϕ) , where

$$\phi := (X_{-v}^T X_{-v})^{-1} \eta = \beta_{-v} + (X_{-v}^T X_{-v})^{-1} X_{-v}^T x_v \beta_v, \quad (8)$$

as orthogonal parameters are only defined up to linear transformation. Expressing (1) in either of parameterizations (7) or (8) requires $p < n$. This suggests an alternative to (8) in which $(X_{-v}^T X_{-v})^{-1}$ is replaced by a regularized inverse $(\delta I_p + X_{-v}^T X_{-v})^{-1}$, an approach discussed briefly in §7.

Section 3 derives some theoretical aspects of the above approximate orthogonalization procedure, particularly the interpretation of q_v , which clarifies the relationship between the present proposal and that of Zhang and Zhang (2014). Proofs of the main results are given in Appendix A.

Remark 1. For any two parameters β_u and β_v , the covariance between $\tilde{\beta}_u$ and $\tilde{\beta}_v$ constructed according to equations (2) and (6), is

$$\text{cov}(\tilde{\beta}_u, \tilde{\beta}_v) = \tau (q_u^T x_u)^{-1} q_u^T q_v (q_v^T x_v)^{-1} =: \tau \sigma_{uv}, \quad (9)$$

which is known up to τ .

Remark 2. There is a simple transformation of X , \hat{X}^v , say, whose columns are exactly orthogonal to a given one, \hat{x}_v^v say. This is $\hat{X}^v = \hat{P}^v X + C^v$ where $\hat{P}^v = I_n - (x_v^T x_v)^{-1} x_v x_v^T$ projects the columns of X into the subspace of \mathbb{R}^n orthogonal to x_v , and C^v is a matrix of zeros apart from the column $c_v^v = \hat{x}_v^v = x_v$, which replaces the single column of zeros produced by $\hat{P}^v X$. Since this transformation to \hat{X}^v is not of the form $A^v X$, it is not interest-respecting. An interest-respecting version could be obtained by projection of \hat{X}^v , in a suitable matrix norm, into the space of matrices of the form $A X$, where $A \in \mathbb{R}^{n \times n}$. This is however much less convenient than (6), both computationally and theoretically.

A nominal α -level confidence interval for β_v is

$$\tilde{\mathcal{C}}_v(\alpha) = [\tilde{\beta}_v - z_{1-\alpha/2}(\tau \sigma_{vv})^{1/2}, \tilde{\beta}_v + z_{1-\alpha/2}(\tau \sigma_{vv})^{1/2}], \quad (10)$$

where $\Phi(z_{1-\alpha}) = \alpha$. As shown in §3, the confidence interval (10) is valid as $n \rightarrow \infty$, to the extent that q_v is successful in orthogonalizing \tilde{x}_v^v to \tilde{x}_w^v for all $w \in \mathcal{S}$. The qualification here is because the objective function in (6) does not ensure \tilde{x}_v^v is exactly orthogonal to \tilde{x}_w^v for all $w \in \mathcal{S}$, only that the cumulative non-orthogonality to \tilde{x}_w^v over $w \neq v$ is small. For some v this non-orthogonality will concentrate relatively more on the unknown signal indices, resulting in larger bias for those parameters. However, as shown in Proposition 3, this bias decays rapidly with the sample size.

One could in principle attempt to identify positions where the signal variables are most likely to be and adjust the objective function (6) accordingly, a suggestion to which we return in §7. A variant of this idea has been proposed recently by Li et al. (2021). Since their

procedure is based on the debiased lasso (Zhang and Zhang, 2014; van de Geer et al., 2014) the resulting inference is not scale-invariant, and the theoretical guarantees entail a similar condition of sparsity of the inverse Fisher information matrix, which in their random design setting corresponds to sparsity of the inverse covariance matrix of the covariates.

3 Inferential properties

In setting up the objective function

$$m(q) = (q^T x_v)^{-2} q^T (\delta I_n + X_{-v} X_{-v}^T) q \quad (11)$$

in equation (6), $\delta > 0$ was introduced to enable comparison to other work. However, the procedure is remarkably robust to this choice, a phenomenon discussed in further detail following Proposition 3.

Let $(\delta I_n + X_{-v} X_{-v}^T) = M_\delta$, and define

$$L_\delta = M_\delta - \{x_v^T M_\delta^{-1} x_v\}^{-1} x_v x_v^T. \quad (12)$$

Proposition 1 specifies the set of solutions $\mathcal{Q}_v^{(\delta)}$ from (6) in closed form.

Proposition 1. *Suppose $\delta > 0$ is such that the eigenvalues of L_δ are non-negative. Then any q_v of the form*

$$q_v = a(\delta I_n + X_{-v} X_{-v}^T)^{-1} x_v \quad (13)$$

is a minimizer of (11), where a is any non-zero real number. If the eigenvalue condition is violated, then such a q_v is a saddlepoint of (11).

Modulo multiplication by a positive scalar, L_δ is the matrix field $\nabla \nabla^T m$ evaluated at any q_v of the form (13). Let $z = \{x_v^T M_\delta^{-1} x_v\}^{-1/2} x_v$ so that $L_\delta = M_\delta - z z^T$. Then

$$\det(L_\delta) = (1 - z^T M_\delta^{-1} z) \det(M_\delta) = 0,$$

where the first equality is the so-called matrix determinant lemma for determinants of rank one perturbations and the last equality follows as $z^T M_\delta^{-1} z = 1$. Thus, the matrix L_δ has at least one zero eigenvalue, showing that, if the condition of Proposition 1 is satisfied, m is locally weakly convex at any q_v of the form (13) but not strictly convex (e.g. Boyd and Vandenberghe, 2004). Let

$$P^v(a, \delta) = a(\delta I_n + X_{-v} X_{-v}^T)^{-1}. \quad (14)$$

so that $q_v = P^v(a, \delta) x_v$. Although any choice of a yields a minimizer of (11), the choice $a = \delta$ improves interpretability of (14) in view of equation (15) of Proposition 2. It is also of some theoretical value to consider $a = \delta \rightarrow 0^+$.

Proposition 2. Set $a = \delta$ and define X_{-v}^+ to be the Moore-Penrose pseudo-inverse of X_{-v} given in (36) of Appendix A.2. Then

$$P^v(\delta, \delta) = I_n - X_{-v}(\delta I_{p-1} + X_{-v}^T X_{-v})^{-1} X_{-v}^T \quad (15)$$

and

$$P^v := \lim_{\delta \rightarrow 0^+} P^v(\delta, \delta) = I_n - X_{-v} X_{-v}^+ \quad (16)$$

is an orthogonal projection operator that maps any vector in \mathbb{R}^n into the left null space of X_{-v} or, equivalently, the kernel of X_{-v}^T , $\ker(X_{-v}^T) = \{u \in \mathbb{R}^n : X_{-v}^T u = 0\} \subset \mathbb{R}^n$.

This limiting case only serves to supply insight and is of low practical relevance when $p > n$. When $p < n$, $X_{-v}^+ = (X_{-v}^T X_{-v})^{-1} X_{-v}^T$ so that $P^v = I_n - X_{-v} (X_{-v}^T X_{-v})^{-1} X_{-v}^T$ and (15) can be viewed as a ridge-regularized version of this projection (Hoerl and Kennard, 1970). See §4 for further discussion.

From (5), (6), and Proposition 4 below, $\tau^{-1/2} S_n(\beta_v)$ is an asymptotically pivotal quantity, where

$$S_n(\beta_v) := \frac{(q_v^T x_v)(\tilde{\beta}_v - \beta_v - b_v)}{(q_v^T q_v)^{1/2}}. \quad (17)$$

However the confidence intervals in (10) ignore the bias term $(q_v^T x_v)b_v/(\tau q_v^T q_v)^{1/2}$. While, for sufficiently small δ , q_v effectively minimizes the observable part of an upper bound on b_v^2 , it is inevitable that $|b_v|$ will be larger for some components β_v than for others, due to lack of knowledge of \mathcal{S} . Proposition 3 establishes the behaviour of the bias b_v in (3) and the quantity $(q_v^T x_v)b_v/(\tau q_v^T q_v)^{1/2}$ as a function of the sample size. The simulations in §6.2 are consistent with these conclusions.

By Proposition 2, $P^v(\delta, \delta)z$ is a regularized projection of $z \in \mathbb{R}^n$ into the kernel of X_{-v}^T , so by construction the elements of $X_{-v}^T P^v(\delta, \delta)z$ are small and bounded away from zero provided that δ is. To avoid explicitly quantifying the dependence on n , p and δ , write $X_{-v}^T P^v(\delta, \delta)x_v \asymp h(n, p, \delta)$.

Proposition 3. Let $\delta > 0$ be bounded away from zero. Provided that the non-zero elements of β_{-v} are bounded, the bias terms in (17) satisfy

$$b_v = O(s/n), \quad \text{and} \quad \frac{(q_v^T x_v)b_v}{(\tau q_v^T q_v)^{1/2}} = O\left(\frac{s}{n^{1/2}} \cdot \frac{h(n, p, \delta)^{1/2}}{\delta^{1/2}}\right).$$

Proposition 3 suggests that one sensible choice of δ would solve the fixed point equation $\delta \asymp h(n, p, \delta)$, so that

$$\frac{(q_v^T x_v)b_v}{(\tau q_v^T q_v)^{1/2}} = O\left(\frac{s}{n^{1/2}}\right).$$

We have found empirically that the stability of the conclusions to the choice of δ is remarkably high. Such stability can always be checked using the data at hand.

In applying Proposition 3, confidence intervals are based on the asymptotically pivotal

quantity $V_n^{-1/2}U_n(\beta_v)$, where V_n is an estimator of τ discussed below and

$$U_n(\beta_v) := \frac{(q_v^T x_v)(\tilde{\beta}_v - \beta_v)}{(q_v^T q_v)^{1/2}}. \quad (18)$$

The validity of this is established via Proposition 3, combined with the limiting distribution of $S_n(\beta_v)$ derived below. We also see that, for fixed sample size, the bias plays a relatively larger role for components whose associated variances are smaller. This is visually apparent from Figure 3 of §6.2.

Proposition 4. *Suppose that $(\varepsilon_i)_{i=1}^n$ are independent mean zero random variables of variance τ and bounded moments of order $2 + \gamma$ for some $\gamma > 0$. Then*

$$\sup_{z \in \mathbb{R}} |\text{pr}\{\tau^{-1/2}S_n(\beta_v) \leq z\} - \Phi(z)| \rightarrow 0, \text{ as } n \rightarrow \infty,$$

where Φ is the standard normal cumulative distribution function. If the above condition holds with $\gamma = 1$,

$$\sup_{z \in \mathbb{R}} |\text{pr}\{\tau^{-1/2}S_n(\beta_v) \leq z\} - \Phi(z)| \leq Ce_n, \quad (19)$$

where $C \in (0.4097, 0.5606]$, $(q_{j,v})_{j=1}^n$ are the elements of q_v , and

$$e_n = \frac{\sum_{j=1}^n |q_{j,v}|^3 \mathbb{E}|\varepsilon_j|^3}{(n\tau)^{3/2} (q_v^T q_v)^{3/4}}.$$

With any consistent estimator V_n of the error variance, the limiting normal approximation to the Studentized statistic $V_n^{-1/2}S_n(\beta_v)$ remains valid by Slutsky's theorem. Proposition 5 provides more precise quantification.

Proposition 5. *Suppose that the conditions of Proposition 4 hold with $\gamma = 1$ and let V_n be a consistent estimator of τ . On writing*

$$\text{pr}(|V_n^{-1/2} - \tau^{-1/2}| > t_n) \asymp r_n(t_n), \quad (20)$$

where t_n and $r_n(t_n)$ are both decreasing in n , the normal approximation to the distribution of $V_n^{-1/2}S_n(\beta_v)$ has uniform error of order

$$\sup_{z \in \mathbb{R}} |\text{pr}\{V_n^{-1/2}S_n(\beta_v) \leq z\} - \Phi(z)| \leq c \max\{e_n, 1 - \Phi(g_n), t_n g_n, r_n(t_n)\}, \quad (21)$$

where c is a constant that we do not quantify, g_n is any nondecreasing sequence such that $t_n g_n = o(1)$, and e_n is defined in Proposition 4.

While it is unclear which choice of g_n delivers the sharpest bound in (21), a reasonably tight bound is obtained by choosing g_n such that $1 - \Phi(g_n)$ is roughly the same order as e_n .

By Mill's ratio

$$\frac{1 - \Phi(x)}{\phi(x)} \in \left(\frac{1}{x + 1/x}, \frac{1}{x} \right), \quad x > 0,$$

where ϕ is the standard normal density function. Thus, using the approximation $e_n = g_n^{-1}\phi(g_n)$ and solving approximately for g_n using Lagrange's inversion formula (e.g. de Bruijn, 1981, p.25–29), we obtain

$$g_n = \left([\log(e_n^{-2}) - \log\{\log(e_n^{-2})\}] + O[\log\{\log(e_n^{-2})\}/\log(e_n^{-2})] \right)^{1/2},$$

which is a very slowly growing function of n . Thus, asymptotically, the confidence interval based on the infeasible statistic $V_n^{-1/2}S_n(\beta_v)$ has nominal coverage probability $1 - \alpha$, and for any fixed n the error in this approximation is roughly of order $\max\{e_n, t_n g_n, r_n(t_n)\}$.

4 The debiased lasso

Our proposal is related to inferential methods based on inverting the Karush-Kuhn-Tucker conditions associated with a lasso solution. In one respect, the most general procedure is that of van de Geer et al. (2014), which covers penalized likelihood inference. Their debiased lasso estimator, $\widehat{\beta}^d$ say, reduces essentially to the proposal of Zhang and Zhang (2014) in the case of the linear model. Zhang and Zhang (2014) modified the score equation for the least squares estimator to $z_v^T(Y - x_v\beta_v) = 0$, with z_v to be chosen. By comparison the estimating equation that defines $\widetilde{\beta}_v$ is

$$\widetilde{x}_v^{vT}(\widetilde{Y}^v - \widetilde{x}_v^v\beta_v) = q_v^T(Y - x_v\beta_v) = 0, \quad \text{with } q_v = A^{vT}A^v x_v. \quad (22)$$

Zhang and Zhang (2014) suggested a bias correction based on a pilot estimator $\widehat{\beta}_w^{(\text{init})}$ of each nuisance parameter β_w :

$$\frac{z_v^T Y}{z_v^T x_v} - \sum_{w \neq v} \frac{\widehat{\beta}_w^{(\text{init})} z_v^T x_w}{z_v^T x_v} = \beta_v + \frac{z_v^T \varepsilon}{z_v^T x_v}.$$

Thus the bias is controlled by

$$\left(\frac{\max_{w \neq v} |z_v^T x_w|}{|z_v^T x_v|} \right) \|\widehat{\beta}_{-v}^{(\text{init})} - \beta_{-v}\|_1 \quad (23)$$

and Zhang and Zhang (2014) recommended that $\widehat{\beta}_{-v}^{(\text{init})}$ be taken as the lasso estimator, so that $\|\widehat{\beta}_{-v}^{(\text{init})} - \beta_{-v}\|_1$ is small with high probability under regularity conditions. They also recommended that z_v be chosen using the lasso, motivated by ordinary least squares when $p < n$, where the score equation $z_v^T(Y - x_v\beta_v) = 0$ holds for the least squares estimator with

$$z_v = x_v^\perp := (I_n - X_{-v}(X_{-v}^T X_{-v})^{-1} X_{-v})x_v,$$

i.e., the projection of x_v on the orthogonal complement of the space spanned by the columns of X_{-v} . Equation (15) in Proposition 2 shows the parallels between the two approaches, although they were derived from different perspectives. In particular, when $a = \delta$ our operations on the sample space, designed to minimize a feasible upper bound on the mean squared error, recover the ridge regression residual vector for z_v , rather than the nodewise lasso residuals from an ℓ_1 -constrained linear regression of x_v on X_{-v} recommended by Zhang and Zhang (2014) and van de Geer et al. (2014).

While there is no obvious direct analogue of our quantities b_v or $(q_v^T x_v) b_v / (\tau q_v^T q_v)^{1/2}$, van de Geer et al. (2014) show that their Studentized test statistic is the sum of a standard normally distributed random variable and the quantity $\tilde{\Delta}_v$, satisfying

$$\text{pr}\left(|\tilde{\Delta}_v| \geq 8n \left(\frac{s\lambda_v\lambda}{\tau\phi_0^2 \|x_v - X_{-v}\hat{\gamma}_v\|_2} \right)\right) \leq 2\exp(-t^2).$$

Here, ϕ_0^2 is the so-called compatibility constant of the lasso, $\hat{\gamma}_v$ is the nodewise lasso estimator of the coefficient vector from a regression of x_v on X_{-v} with regularization parameter λ_v , and λ (chosen as a function of $t > 0$) is the lasso tuning parameter associated with the regression of Y on X . They also show that their debiased lasso estimator satisfies $\sqrt{n}(\hat{\beta}^d - \beta) = W + \Delta$ where W is a normal random variable of zero mean and Δ satisfies $\|\Delta\|_\infty = o_{\text{pr}}(1)$. This theoretical guarantee is established under an assumption of row-sparsity of the inverse Fisher information matrix. Our result of Proposition 3 shows that $b_v = O(s/n)$ without the extra sparsity condition (beyond sparsity of β) and in a Gaussian setting induces sparsity on the relevant row and column of the Fisher information matrix, as in Cox and Reid (1987). Thus our approach seeks to achieve by construction what is implicitly assumed in van de Geer et al. (2014) and Zhang and Zhang (2014). Inducement of sparsity cannot be accomplished simultaneously for all parameters, but is achievable by considering each coefficient in turn as the interest parameter.

Another important advantage of our approach is that it does not require rescaling the columns of X to unit length, so the physical interpretation of the estimate and confidence intervals is maintained. Neither the lasso estimator nor its debiased counterpart is invariant to componentwise rescaling, and simulations in §6.1 indicate that the approach is rather conservative after transformation to the original scale.

Javanmard and Montanari (2014) also consider a version of the debiased lasso estimator for the linear model. Their estimator is of a similar form to that of van de Geer et al. (2014) and Zhang and Zhang (2014), namely

$$\hat{\beta}^{\text{JM}} = \hat{\beta}_{\text{lasso}} + n^{-1} M X^T (Y - X \hat{\beta}_{\text{lasso}}),$$

where, in a similar spirit to our proposal, M is chosen to minimize the variance of the resulting estimator subject to a constraint on the bias. Specifically, $M = (m_1, \dots, m_p)^T$, where m_v

solves

$$m_v = \underset{m}{\operatorname{argmin}} m^T (n^{-1} \sum_i X_i X_i^T) m \quad \text{s.t.} \quad \|(n^{-1} \sum_i X_i X_i^T) m - e_v\|_\infty \leq \mu,$$

with e_v the zero vector with a 1 in the v th position. Thus M can be viewed as a regularized inverse of $n^{-1} \sum_i X_i X_i^T$. A key difficulty with this proposal is that the tuning parameter μ needs to decay sufficiently quickly in order that the bias of the estimator (the analogue of our b_v) may be controlled. Javanmard and Montanari (2014) state in their Algorithm 1 that if any of the constraints is violated, then M should be set to the identity matrix. This seems problematic since the constraint is used to derive the theoretical guarantees. We found that it was often violated in examples we tried.

5 Confidence sets of models

5.1 Background

For sparse high-dimensional regression problems, Cox and Battey (2017) emphasised that several or even many low-dimensional models are likely to fit the data indistinguishably well. They argued that while an arbitrary choice between them would be reasonable for prediction, for subject-matter understanding, it is more appropriate to report a confidence set of models. Our suggested use for the estimates and confidence intervals obtained as described in §2 is as an adjunct to such confidence sets, the two procedures being performed in isolation, thereby avoiding issues of post-selection inference.

The first phase in the construction of confidence sets of models is to identify a set of retained explanatory variables, $\widehat{\mathcal{S}}$, where $|\widehat{\mathcal{S}}|$ is reasonably large, but considerably less than n . The model that includes all variables in $\widehat{\mathcal{S}}$ is called the encompassing model. A confidence set of models, \mathcal{M} , is obtained by comparing the fit of the encompassing model to the fit of any model using $s^\#$ or fewer variables from $\widehat{\mathcal{S}}$, and retaining any such models that are statistically indistinguishable from the encompassing one. The choice of $s^\#$ is arbitrary and based on an assumption of sparsity.

The confidence intervals in §2 can be used to refine this large confidence set, \mathcal{M} , to a smaller set $\mathcal{R} \subset \mathcal{M}$, by removing from \mathcal{M} any models which have coefficient estimates that are not included in the relevant confidence interval (10) (with τ replaced by V_n). Specifically, if any of the estimated coefficients from fitting a model in \mathcal{M} exceed such confidence limits, this casts doubt on the plausibility of that model. If none of the models in \mathcal{M} is consistent with the collection of confidence intervals, then this suggests that $|\widehat{\mathcal{S}}|$ is too small.

We assume that \mathcal{M} is exactly calibrated, i.e. that $\operatorname{pr}(\mathcal{S} \in \mathcal{M}) = 1 - \vartheta$. Asymptotic calibration of \mathcal{M} is established under certain conditions in Lewis and Battey (2022). As discussed in §3, while the individual confidence intervals are all calibrated at their nominal levels for sufficiently large n , small-sample bias affects inference on some parameters more than others. The following calculations indicate the implications of this bias on the small-sample coverage properties of \mathcal{R} .

Let $q = q(u)$ be the proportion of confidence intervals whose coverage probability is below some acceptable threshold, $1 - u$ say. Since poor calibration is unrelated to whether a variable is in \mathcal{S} , the indices of signal variables can be treated as a simple random sample of size s from $\{1, \dots, p\}$. Thus, the number of poorly calibrated intervals among these s is binomial of index s and parameter q , and an approximate lower bound on the probability that no signal variable is excluded from \mathcal{M} is $((1 - q)(1 - u))^s$.

The analysis of artificially generated responses and real covariate data in §6.1 have $q = 23/4088 = 0.0056$ intervals with coverage smaller than $1 - u = 0.95$ when the nominal level of our intervals is 0.01, $q = 118/4088 = 0.0289$ with coverage smaller than $1 - u = 0.98$ at nominal level 0.005 and $q = 20/4088 = 0.0049$ with coverage smaller than $1 - u = 0.99$ at nominal level 0.001 (see Table 3). Thus, for $s = 5$, and $\text{pr}(\mathcal{S} \in \mathcal{M}) = 0.99$, the coverage probability of \mathcal{R} using the 0.001-level intervals is approximately $((1 - q)(1 - u))^5(1 - \vartheta) = 0.9186$.

The above arguments would not work for the debiased lasso, as simulations in §6.2 show that, while the coverage probabilities are high on average, the debiased lasso intervals are more systematically miscalibrated for signal variables.

Whether the refined set \mathcal{R} is more useful than a version of \mathcal{M} whose coverage matches \mathcal{R} depends on q and its associated value of u . This can be assessed through an analysis similar to that in §6.1.

5.2 Illustration

Bühlmann et al. (2014) describe an investigation to study the association of the response variable, vitamin B_2 production, with the logarithmic expression levels of 4088 genes. In the sample of size $n = 71$ the correlation among columns of X ranges from -0.926 to 0.991 , with a mean correlation 0.029 . See Bühlmann et al. (2014) for original references.

The first analysis is the estimation of each of the regression coefficients and their confidence limits at nominal level $\alpha = 0.001$ using the approximate orthogonalization approach of §2. A separate analysis was carried out to choose the encompassing model $\hat{\mathcal{S}}$ and a confidence set of models \mathcal{M} . The variables in $\hat{\mathcal{S}}$ were identified through a series of model fits based on partially balanced incomplete block arrangements (Yates, 1936; Cox and Battey, 2017). Then each possible low-dimensional sub-model of $\hat{\mathcal{S}}$ was tested against the encompassing model using an F -test at level $\vartheta = 0.01$. Models not rejected by the F -test comprise the confidence set \mathcal{M} .

The confidence intervals constructed as described in §2 for the 22 variables in $\hat{\mathcal{S}}$ are reported in Table 1 for $\alpha = 0.05$ and $\alpha = 0.001$. These have been ordered according to the proportion of models from \mathcal{M} to which they belong. For comparison, the lasso and the elastic net, fitted to all 71 observations and tuned to retain 22 variables, find nine and fourteen of the same variables. These are indicated in the first column of Table 1. Further explanation of $\hat{\mathcal{S}}$ and \mathcal{M} is provided in Appendix B, including details of error variance estimation and how asymptotic calibration of \mathcal{M} is ensured.

There are 34555 models in \mathcal{M} , but only 773 of these give estimates of the regression

variable index v	proportion of models in \mathcal{M}	$\tilde{\beta}_v$	lower limit (0.05)	upper limit (0.05)	lower limit (0.001)	upper limit (0.001)
2138	0.218	-0.062	-0.382	0.259	-0.599	0.476
2564 ^{L,E}	0.218	-1.481	-1.801	-1.160	-2.018	-0.943
1516 ^{L,E}	0.218	0.343	0.022	0.663	-0.195	0.880
1503 ^{L,E}	0.217	-0.325	-0.646	-0.0050	-0.863	0.212
1639 ^{L,E}	0.217	-0.406	-0.726	-0.086	-0.944	0.132
1603	0.217	-1.048	-1.368	-0.728	-1.586	-0.510
4008 ^E	0.217	-0.366	-0.686	-0.046	-0.903	0.172
4002 ^{L,E}	0.216	-0.505	-0.825	-0.185	-1.043	0.033
1069 ^E	0.216	-0.398	-0.718	-0.078	-0.936	0.140
1436 ^E	0.215	-0.463	-0.783	-0.143	-1.001	0.075
3291	0.215	-0.640	-0.960	-0.320	-1.178	-0.102
978	0.214	-0.259	-0.580	0.061	-0.797	0.278
3514 ^{L,E}	0.214	1.373	1.053	1.694	0.836	1.911
1297 ^{L,E}	0.213	0.219	-0.102	0.539	-0.319	0.756
1285	0.213	0.172	-0.148	0.493	-0.366	0.710
3808 ^E	0.213	0.677	0.356	0.997	0.139	1.214
1423	0.212	0.043	-0.277	0.363	-0.495	0.581
1278 ^{L,E}	0.212	0.147	-0.173	0.467	-0.391	0.685
403	0.211	0.902	0.582	1.44	0.365	1.323
1290	0.211	0.189	-0.131	0.510	-0.348	0.727
1303 ^E	0.211	0.187	-0.133	0.507	-0.351	0.725
1312 ^{L,E}	0.209	0.490	0.169	0.810	0.048	1.027

Table 1: Indices for variables in $\hat{\mathcal{S}}$, ordered by prevalence in \mathcal{M} , plus 95% and 99.9% confidence limits for the associated regression coefficients. Superscripts L and E indicate that the variable was also found by the lasso and the elastic net fitted to the full sample.

coefficients that are compatible with our nominal 99.9% confidence intervals for the individual parameters in Table 1. Let \mathcal{R} denote this set of 773 models. Since the decision to report a model in \mathcal{R} depends on the simultaneous correctness of the associated confidence intervals, a Bonferroni-type correction should be applied as in §5.1. In fact, to account for a degree of finite-sample miscalibration of our intervals due to the term b_v , we use the argument of §5.1 with q and the corresponding value of u estimated from Table 3. Thus the effective coverage probability of \mathcal{R} assuming, optimistically, that the coverage probability for \mathcal{M} is exactly calibrated at its nominal level, is approximately 0.9186. For comparison, the size of the model confidence set based only on a likelihood ratio test at nominal level $1 - 0.9186 = 0.0814$ is 15084, illustrating the advantage of our intervals as an adjunct to confidence sets of models. Any choice between the 773 models in \mathcal{R} would require either additional data or subject matter expertise. Table 2 reports central regions of the confidence “distribution” of models (in the sense of Fisher, 1930; Cox, 1958) obtained as a nested sequence of confidence sets at different levels. The full confidence set of models is reported in the supplementary material.

	2138	2564	1516	1503	1639	1603	4008	4002	1069	1436	3291	978	3514	1297	1285	3808	1423	1278	403	1290	1303	1312	
1	-	-	-	-	-	-	-	-	-	-	-	*	-	*	-	-	-	-	-	-	*	-	-
2	-	-	-	-	-	-	-	-	*	-	-	-	-	-	*	-	-	-	-	-	-	-	-
3	-	-	-	-	-	-	-	-	*	-	-	-	-	*	-	-	-	-	-	-	-	-	-
4	-	-	-	-	-	-	-	-	*	-	-	-	-	-	-	-	*	-	*	-	-	-	-
5	-	-	-	-	-	-	-	-	*	-	-	-	-	-	*	-	-	*	-	-	-	-	-
6	-	-	-	-	-	-	-	-	*	-	-	*	-	-	*	-	-	-	*	-	-	-	-
7	-	-	-	-	-	-	-	-	-	-	-	*	-	*	*	-	-	-	-	-	-	-	-
8	-	-	-	-	-	-	-	-	-	-	-	*	-	*	*	-	*	-	-	-	-	-	-
9	-	-	-	*	*	-	-	-	-	-	-	*	-	-	-	*	-	-	-	-	-	-	-
10	-	-	-	-	-	-	-	-	*	-	-	-	-	*	-	-	-	-	-	-	-	*	-
11	-	-	-	*	-	-	-	-	*	-	-	-	*	-	-	-	-	-	-	-	-	-	-
12	-	-	-	*	-	-	-	-	-	*	-	-	*	-	-	-	-	-	-	-	-	-	-
13	-	-	-	*	-	*	*	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
14	-	-	-	-	-	-	-	-	*	-	-	-	-	*	-	-	-	-	-	-	*	-	-
15	-	-	-	*	*	-	-	-	-	-	-	-	-	-	-	*	-	-	-	-	-	-	-
16	-	-	-	-	-	-	-	-	*	-	-	-	-	-	-	-	*	-	-	-	-	-	*
17	-	-	-	-	-	-	-	-	*	-	-	-	-	*	*	-	-	*	-	-	-	-	-
18	-	-	*	*	-	-	-	-	-	-	-	-	-	-	-	*	-	-	-	-	-	-	*
19	-	-	-	-	-	-	-	-	*	-	-	*	-	*	*	*	-	-	-	-	-	-	-
20	-	-	*	-	-	*	-	*	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
21	-	-	-	-	-	-	-	-	*	-	-	-	-	*	*	*	-	-	-	-	-	-	-
22	-	-	-	-	-	-	-	-	*	-	-	-	-	*	-	-	-	-	-	-	*	*	-
23	-	-	-	-	-	-	-	-	*	-	-	-	-	*	-	-	*	-	-	-	-	-	*
24	-	-	-	*	*	-	-	-	*	*	-	-	-	-	-	*	-	-	-	-	-	-	-
25	-	-	*	*	-	-	-	-	-	*	-	-	-	-	-	*	-	-	-	-	-	-	*
26	-	-	-	-	-	-	-	-	*	-	-	-	-	*	-	-	*	*	-	-	-	-	-
27	-	-	-	-	*	*	-	*	-	*	-	-	-	-	-	-	-	-	-	-	-	-	-
28	-	-	-	-	-	-	-	-	*	-	-	*	-	*	*	-	-	-	*	-	-	-	-
29	-	-	-	*	-	-	-	-	*	-	-	*	-	-	-	-	-	-	*	-	-	-	*
30	-	-	-	*	-	-	-	-	*	-	-	*	-	-	*	-	-	-	-	-	-	-	*
31	-	-	-	*	-	-	-	-	*	-	-	*	-	-	-	-	-	-	-	-	*	-	*
32	-	-	-	-	-	-	-	-	*	-	-	-	-	*	*	-	*	-	-	-	-	-	*
33	-	-	-	*	*	-	-	-	*	-	*	-	-	-	-	*	-	-	-	-	-	-	-
34	-	-	*	*	-	-	-	-	-	*	-	-	-	*	-	*	-	-	-	-	-	-	-

Table 2: Central region of a confidence distribution of models, obtained from successively smaller values of α . Shading corresponds to (from dark to light): $\alpha \in \{0.5, 0.4, 0.3, 0.2\}$. The full confidence set of models \mathcal{R} with coverage probability 0.9186 is reported as supplementary material.

6 Simulations

6.1 Using covariate data from Bühlmann et al. (2014)

In this section we use the covariate data from Section 5.2 to study the coverage and bias in a moderately realistic example, and compare the results to the debiased lasso. Simulations of this nature are a helpful adjunct to the analysis above, to assess the security of the results.

We drew indices for five signal variables at random from the index set $\{1, \dots, p\}$, giving $\mathcal{S} = \{313, 689, 724, 1747, 2470\}$, fixed across 1000 ensuing Monte Carlo replications. In each, an artificial outcome was generated as $Y = X\beta + \varepsilon$, where ε is a vector of independent standard normally distributed random variables. The signal strength was 1 for all signal variables. The assumption of known error variance used here was relaxed in §5.2. Confidence intervals at the $\alpha \in \{0.05, 0.01, 0.005, 0.001\}$ nominal levels for each of the 4088 coefficients were constructed as described in §2, and the coverage probabilities and mean lengths for each coefficient were estimated as simulation averages. Modal and median simulated coverage probabilities over the 4088 variables and the median of the simulated mean lengths are reported in Table 3. Detailed summary information is depicted in Figure 1, where the coverage probabilities are plotted against $|b_v|$ and mean length. Yellow stars indicate signal variables. As expected, intervals with higher coverage (1, right) are on average longer, and the actual coverage decreases as the bias increases, although relatively slowly.

nominal level	modal coverage	median coverage	median length	$1 - u$			
				0.9	0.95	0.98	0.99
0.05	0.951	0.942	3.32	0.939	0.228	0	0
0.01	0.989	0.987	4.37	1	0.994	0.839	0.288
0.005	0.989	0.987	4.77	1	0.999	0.971	0.798
0.001	0.995	0.993	5.59	1	1	1	0.995

Table 3: Summary statistics over the 4088 coefficients of the simulated coverage probabilities and simulated mean lengths for each confidence interval. The last four columns are the values of $1 - q$ used in the calculation of §5.1.

For comparison, we also recorded the coverage properties and lengths for the debiased lasso estimator $\hat{\beta}^d$ (van de Geer et al., 2014) which, for the linear model, is essentially the same estimator as that of Zhang and Zhang (2014). Since lasso-based procedures require the columns of X to be on the same scale, we standardized them to unit length before constructing the debiased lasso in the usual way. In order to compare them to our approach, we then converted the conclusions back to the scale of interest by dividing the value of each entry of $\hat{\beta}^d$ by the length of the corresponding column of the original X matrix. Estimates of standard errors were adjusted by dividing them by the square root of this length. While more appropriate than a debiased lasso analysis on the original scale, the rescaling operation should not be expected to be adequate, as neither Wald-based inference nor the debiased lasso is invariant to rescaling.




Figure 1: Real covariate data with artificially generated outcomes: simulated coverage probabilities plotted against the absolute bias $|b_v|$ of the estimator with the the mean length of the confidence interval over Monte Carlo replications indicated by the shading. Nominal levels are 0.05 (left) and 0.01 (right)

Summary statistics are reported in Table 4 and Figure 2. This demonstrates a degree of over-coverage, with the lengths of each interval tending to be larger than is necessary to achieve the nominal level (by comparison with Table 3). Section 6.2 gives a further comparison to the debiased lasso when the explanatory variables are all of similar magnitudes, so that no rescaling is required. The empirical coverage probabilities in that idealized setting are close to nominal, although the procedure appears to be systematically miscalibrated for signal variables, which is problematic for the type of analysis discussed in §5.1.

nominal level	modal coverage	median coverage	median length	$1 - u$			
				0.9	0.95	0.98	0.99
0.05	1	0.998	4.29	1	0.996	0.964	0.889
0.01	1	1	5.64	1	1	0.998	0.992

Table 4: Analogue of first two rows of Table 3 for debiased lasso-based confidence intervals.

We close this comparison with a comment on computation. Construction of our confidence intervals for a single component β_v only entails construction of q_v , which is essentially instantaneous because of its closed-form solution. The computational burden associated with construction of a single interval is therefore approximately independent of p . By contrast, specification of the debiased lasso estimator and its variance estimate for β_v entails construction of the whole matrix $\hat{\Theta}_{\text{Lasso}}$ from equation (8) of van de Geer et al. (2014). This involves solving p nodewise lasso regressions, whose tuning parameters are ideally chosen by cross validation as in the experiments above. Specifically, there are $p + 1000$ tuning parameters chosen by cross-validation involved in the generation of Figure 2, compared to one tuning parameter fixed at




Figure 2: Analogue of Figure 1 for debiased lasso-based confidence intervals.

$\delta = 1$ across Monte Carlo replicates for Figure 1, since the performance of our approach was found to be stable with respect to δ , provided that a value extremely close to zero was not used. The computation of the confidence set \mathcal{M} of §5 is computationally more intensive as this entails checking all sub-models of dimension $\leq s^\#$ of the encompassing model for their compatibility with the data. With $s^\# = 5$, computation can be performed within a minute on a standard computer for $|\hat{\mathcal{S}}| = 22$ but computation time increases rapidly with $|\hat{\mathcal{S}}|$, and to a lesser extent with $s^\#$.

6.2 A factorial experiment

In this section we assess how coverage properties and lengths of the proposed intervals are affected by p , n and the correlation between columns of X . In particular, we seek to understand the small-sample behaviour of the procedure in various circumstances.

Our Monte Carlo experiment was designed as follows. An X matrix was obtained by generating each row x_i^T for $i = 1, \dots, n$ from a normal distribution of mean zero and covariance matrix $\rho 1_p 1_p^T + (1 - \rho)I$, and the columns were each centered to have sample mean zero. The matrix X was fixed across subsequent Monte Carlo replications, as was a sparse vector β , consisting of ones in the last five entries and zeros elsewhere. The values of n , p and ρ were taken at all 2^3 combinations of high and low levels as in Table 5, the value of n/p being the same in two of the four configurations at high and low values of ρ . In each replication, a vector ε of independent standard normally distributed errors was generated anew and a vector of outcomes constructed as $Y = X\beta + \varepsilon$. Each of the first 1000 entries of β was treated in turn as the interest parameter and the coverage properties and lengths, obtained as in §6.1, are reported in Table 5. These set $a = \delta = 1$ in formula (13). Simulations (not reported here) for $a = 1$ and $\delta = n^{-1}$, produced identical numbers to those in Table 5 up to three significant figures.

Table 6 reports point estimates of the main effects of ρ , n and p , from logistic regression of

ρ	n	p	modal coverage	median coverage	proportion with coverage > 0.9	median length	median ϑ_w^2	95th p.c. ϑ_w^2
0.9	70	2450	0.987	0.980	0.922	1.975	0.0065	0.056
0.9	70	1225	0.990	0.980	0.924	2.000	0.0063	0.055
0.9	35	2450	0.986	0.979	0.898	2.786	0.0133	0.118
0.9	35	1225	0.988	0.981	0.937	2.801	0.0133	0.118
0.1	70	2450	0.989	0.870	0.454	0.663	0.0065	0.056
0.1	70	1225	0.987	0.906	0.508	0.671	0.0063	0.054
0.1	35	2450	0.991	0.836	0.415	0.936	0.0135	0.120
0.1	35	1225	0.988	0.917	0.538	0.947	0.0133	0.120

Table 5: Summary statistics for simulated coverage probabilities and mean lengths of $\alpha = 0.01$ confidence intervals for the first 1000 coefficients (null variables only). The last two columns are the median and 95th percentile of the $1000 \times (p - 1)$ values of ϑ_w^2 for the first 1000 coefficients.

the coverage probabilities, and linear regression of the length, using the values in Table 5. For the coverage probabilities, the multiplicative effects are reported, with estimates close to one indicating a negligible effect on the coverage. The modal coverage probabilities are essentially unaffected by ρ , n and p , as would be expected, as for most of the coefficients the approximate orthogonalization is effective and the statistic used is close to the pivotal one. At these small values of n , an increase from $n = 35$ to $n = 70$ reduces the length of the intervals but appears to have a negligible effect on the median coverage probability, while increasing p has no effect on the length but multiplies the odds associated with the median coverage probabilities by 0.624. This is to be expected: in optimizing the composite non-orthogonality over more variables, the signal variables play a relatively smaller role in the optimization problem. By far the largest effect on the median coverage is due to ρ . The median length of the confidence intervals also increases with the correlation, reflecting expression (5) for the variance of $\tilde{\beta}_v$. See Table 7 for further analysis on the role of n .

Detailed summary information is in Figure 3, where we plot coverage probabilities in each of the eight settings as a function of the absolute bias $|b_v|$. The z axis indicated by the colour is the mean length of the intervals over simulations. Coverage probabilities for signal variables and their associated values of $|b_v|$ are represented as yellow stars. As expected in view of §3, the coverage probabilities associated with the $\rho = 0.1$ settings decay more sharply with absolute bias than with $\rho = 0.9$. The decrease in $|b_v|$ for increasing n , as quantified in Proposition 3, is also reflected in these plots.

Proposition 3 indicates that the procedure is well calibrated asymptotically for all β_v . We confirm this empirically by taking the worst case from Table 5 ($p = 2450$ and $\rho = 0.1$) and considering the performance over a sequence of n values.

The experiments suggest that while the procedure is valuable in settings with high correlations between columns of X , it is relatively less secure in the almost uncorrelated case when

response variable	estimated main effect		
	ρ	n	p
modal coverage	0.917	1.000	1.000
median coverage	6.540	1.094	0.624
median length	1.586	-0.790	-0.015

Table 6: Estimated main effects of ρ , n and p on the properties of our $\alpha = 0.01$ nominal confidence intervals. Estimated effects of coverage are multiplicative on the odds scale.

	$n = 6^2$	$n = 12^2$	$n = 18^2$	$n = 24^2$	$n = 30^2$
median coverage	0.801	0.886	0.892	0.922	0.951
proportion > 0.9	0.377	0.478	0.492	0.561	0.655
median length	0.925	0.466	0.324	0.257	0.227

Table 7: Properties of confidence intervals for the $\rho = 0.1$, $p = 2450$ setting with increasing n . Nominal level $\alpha = 0.01$.

the sample size is small. For any given data, we suggest performing preliminary checks as in §6.1 as an indication of the reliability of the method on the data at hand.

Performance of the debiased lasso on the same data is summarized in Table 8 and Figure 4. Since the equicorrelation matrix used to generate X does not have associated with it a sparse inverse, a key assumption for the theoretical guarantees of van de Geer et al. (2014) is violated. Their approach nevertheless performs favourably for variables with no real effect (see Table 8 which concerns null coefficients only), although it appears to be more systematically miscalibrated for signal variables, as illustrated in Figure 4.




Figure 3: Coverage probabilities (vertical axis) for each β_v plotted against the associated absolute bias $|b_v|$ (horizontal axis) and the mean length of the confidence interval (colour axis). High and low values of p are 2450 and 1225 and those of n are 70 and 35. Yellow stars correspond to signal variables.

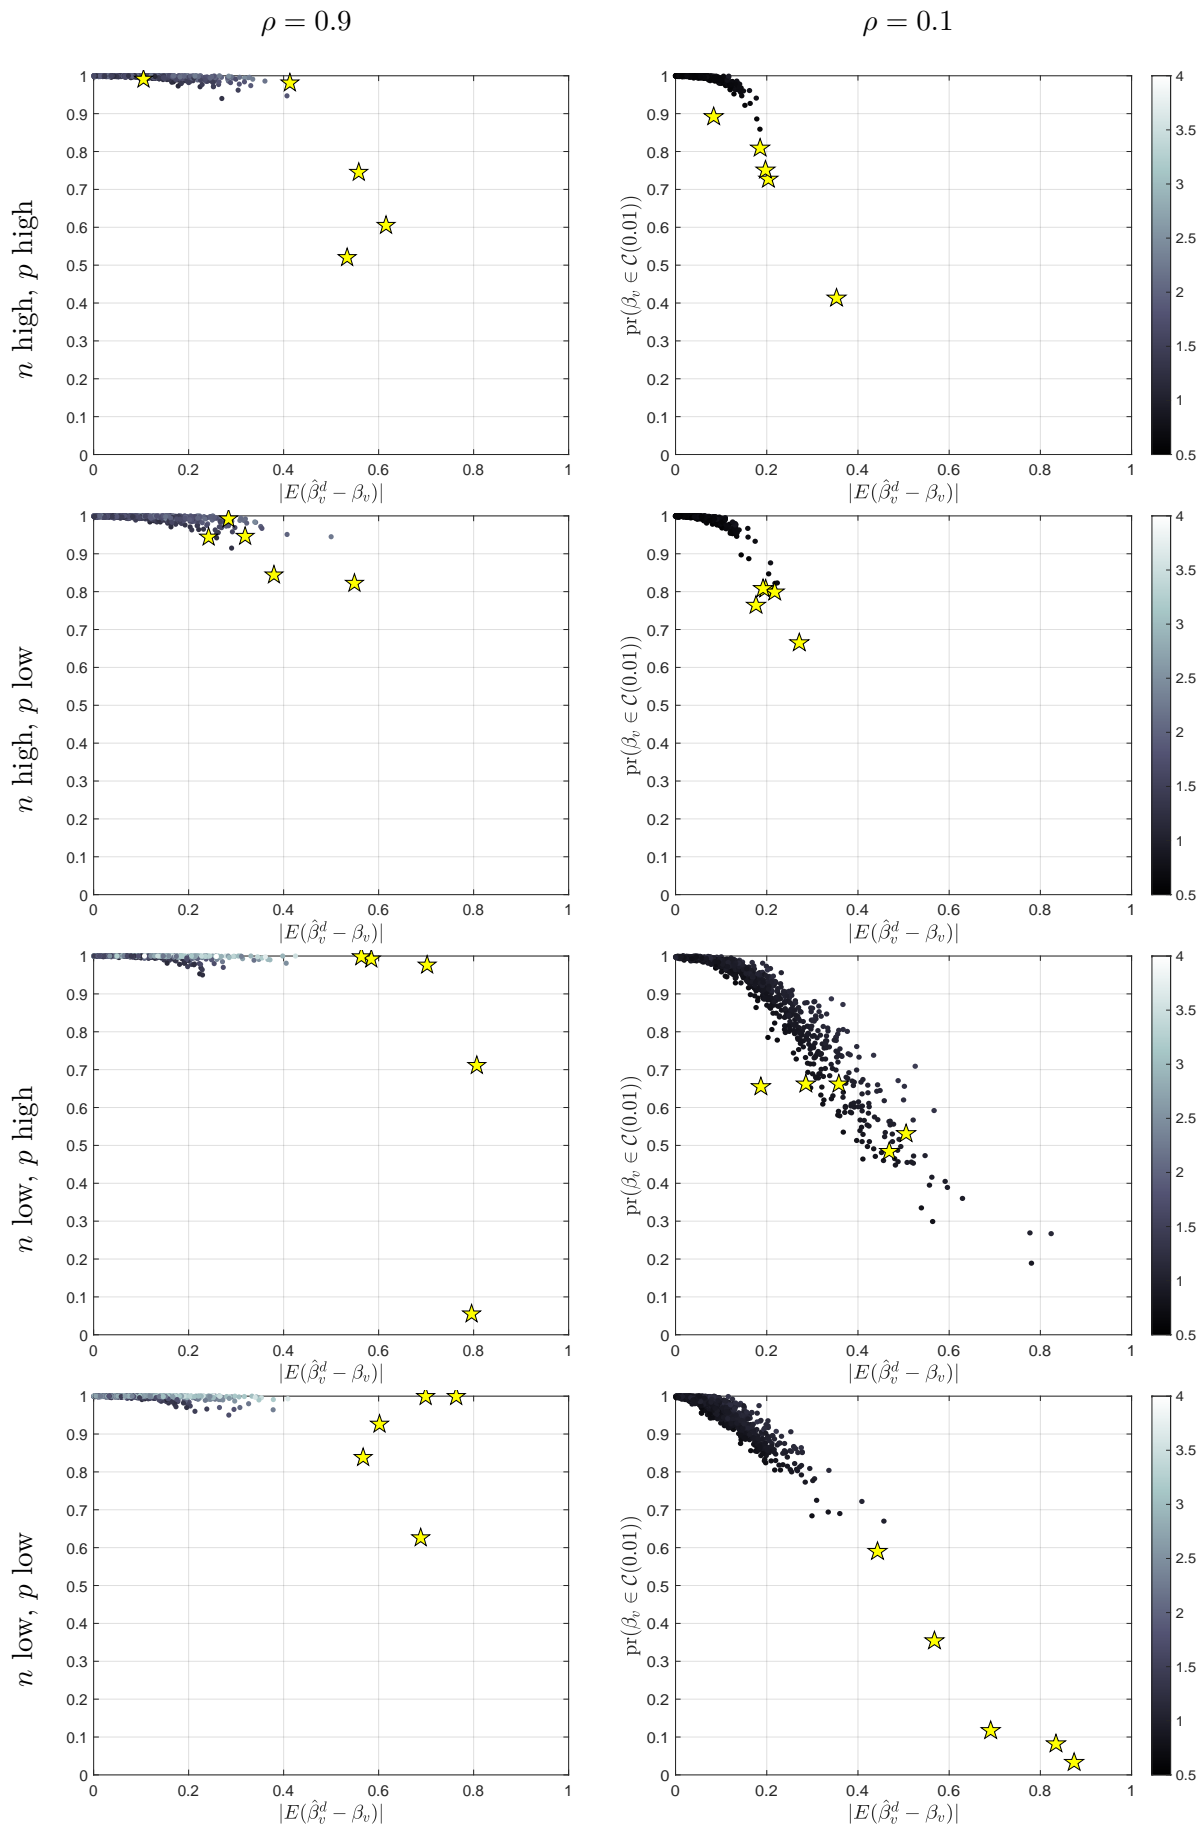


Figure 4: Coverage probabilities of debiased lasso confidence intervals (vertical axis) for each β_v plotted against the absolute value of the Monte Carlo average of $\hat{\beta}_v^d - \beta_v$ (horizontal axis) and the mean length of the confidence interval (colour axis). Yellow stars correspond to signal variables.

ρ	n	p	modal coverage	median coverage	proportion with coverage > 0.9	median length
0.9	70	2450	1	0.999	1	1.174
0.9	70	1225	1	0.999	1	1.754
0.9	35	2450	1	1	1	2.176
0.9	35	1225	1	1	1	2.398
0.1	70	2450	1	0.999	0.998	0.647
0.1	70	1225	1	0.999	0.993	0.641
0.1	35	2450	0.998	0.966	0.682	0.924
0.1	35	1225	0.999	0.981	0.868	0.930

Table 8: Modal and median Monte Carlo coverage probabilities and the median of the Monte Carlo mean length of confidence intervals across the first 1000 coefficients (corresponding to null variables only) for the debiased lasso with tuning parameter chosen by cross-validation. The nominal coverage probability is 0.99.

7 Concluding remarks and some open problems

We have proposed a new asymptotically unbiased estimator of each component of β in a linear regression model. Our estimator has a simple closed form formula and effectively no tuning parameters, since the stability with respect to δ was found to be high. Importantly, the estimator does not require all columns of X to be on the same scale. Theoretical conclusions associated with the procedure are confirmed by simulation in §6.2.

We used these estimates and associated confidence intervals to propose an inferential framework based on refined confidence sets of models. The logical argument used in §5 seems new and is of the following form. Suppose for a possible contradiction that model m is true. This leads to standard least squares estimates $\hat{\beta}(m)$ of the low-dimensional vector $\beta(m)$ associated with model m . Under model m , $\hat{\beta}(m)$ is unbiased and consistent for $\beta(m)$. Let β_v be an arbitrary entry of $\beta(m)$. To the extent that the orthogonalization makes $b_v = 0$ our confidence interval for β_v is unbiased and consistent for any model. Thus extremity of any entry of $\hat{\beta}(m)$, as calibrated by the associated limits of the corresponding confidence interval, contradicts validity of m . Our intervals are asymptotically calibrated by Proposition 3 but allowance for finite-sample miscalibration is discussed in §5.1 and §5.2. In contrast, the usual proposed use of the debiased lasso requires adjustments to the confidence limits on account of the large number of comparisons.

An alternative to the orthogonalization described in §2 is to first estimate the set of signal variables, and modify (6) to

$$q_v(\hat{\mathcal{S}}) \in \left\{ \underset{q \in \mathbb{R}^n}{\operatorname{argmin}} (q^T x_v)^{-2} q^T (\delta I_n + \sum_{w \in \hat{\mathcal{S}} \setminus \{v\}} x_w x_w^T) q \right\}, \quad (24)$$

where $q_v = q_v(\{1, \dots, p\})$. This proposal was assessed in simulations (not reported) when $\hat{\mathcal{S}}$

was constructed using the lasso. The coverage properties depended strongly on the lasso tuning parameter. A tuning parameter could always be found to yield a median coverage probability close to nominal. However, the coverage probabilities corresponding to the signal variables was appreciably reduced.

In §5.2 we used a simple estimator of the error variance, obtained as part of the construction of calibrated confidence sets of models; see Appendix B. A similar approach, following Fan et al. (2012), could be based on cross-validation. In factorial experiments estimates of the variance are obtained from the estimates of the null effects: if $(\varepsilon_i)_{i=1}^n$ are normally distributed and the approximate orthogonalization via equation 6 was completely effective, then

$$(q_v^T q_v)^{-1} (q_v^T x_v)^2 \tilde{\beta}_v^2 \quad (25)$$

provide estimates of τ . The composite of p such statistics is, however, contaminated by those corresponding to signal variables and by those for which the optimization problem (6) fails to orthogonalize \tilde{x}_v^v to \tilde{x}_w^v for one or more $w \in \mathcal{S}$. One might however construct a robust estimator of variance from something like a trimmed mean of these p estimates.

The extent to which the ideas of §2 apply beyond the linear regression model remains unclear. In a generalized linear model with canonical link we can write the log-likelihood function in terms of an interest parameter β_v as

$$\ell(\psi, \lambda) = \phi^{-1} \sum_{i=1}^n \left\{ Y_i(x_{i,v}\beta_v + x_{i,-v}^T\beta_{-v}) - K(x_{i,v}\beta_v + x_{i,-v}^T\beta_{-v}) \right\}, \quad (26)$$

and the function that is linear in the true parameter β is

$$\eta_i := g(\mu_i) = \{(\partial/\partial\eta)K\}^{-1}\{\mathbb{E}(Y_i)\} = x_i^T\beta, \quad (27)$$

with $\mu_i = \mathbb{E}(Y_i)$. An estimate $\hat{\eta} = (\hat{\eta}_1, \dots, \hat{\eta}_n)^T$ can be expressed as

$$\hat{\eta} = X\beta + \varepsilon = x_v\beta_v + X_{-v}\beta_{-v} + \varepsilon, \quad (28)$$

where ε captures the error in estimating η . In the linear model $\hat{\eta} = Y$ is an estimator of $X\beta$ with error satisfying $\mathbb{E}(\varepsilon) = 0$, so that (28) is a generalization. Following §2 we might seek a matrix A^v that approximately orthogonalizes the v th column, \tilde{x}_v^v , of $\tilde{X}^v = A^v X$ to its remaining columns, thereby justifying a marginal regression of $A^v \hat{\eta}$ on \tilde{x}_v^v to provide inference on β_v . The same reparameterization arguments of §2 apply. The challenge is to specify an estimator $\hat{\eta}$, perhaps highly inefficient, but whose properties are sufficiently well understood.

An alternative generalization might be based on the ideas of Cox and Reid (1987). The possibility for this in the linear model is apparent from equation (8), where the factor multiplying β_v is the least squares estimate of the coefficient vector in a regression of x_v on X_{-v} . The introduction of δI_n would accommodate $p > n$ by replacing this least squares estimate by a ridge regression estimate (Hoerl and Kennard, 1970) and would yield a parameter approxi-

mately orthogonal to β_v when $p > n$. Further analysis is needed to understand how far this idea generalizes, as our preliminary calculations for a logistic regression setting (not reported) suggested some difficulties. To be useful, any subsequent operation would have to exploit the induced sparsity in the relevant column of the Fisher information matrix similarly to Ning and Liu (2017), who assumed such sparsity in the inverse of this matrix without preliminary parameter-orthogonalizing steps.

Acknowledgements. We are grateful to Yanbo Tang and to three anonymous referees for important comments on an earlier draft. Implementation of the debiased lasso in §6.1 and §6.2 was based on highly efficient and publicly available R code written by Dr. Rajen Shah. The work was partially supported by the UK Engineering and Physical Sciences Research Council (EP/T01864X/1) and the Natural Sciences and Engineering Council of Canada.

A Proofs for Section 3

A.1 Proof of Proposition 1

Proof. Let $M_\delta = (\delta I_n + X_{-v} X_{-v}^T)$ and write the objective function in equation (6) as $m = gh$, where $g(q) = (q^T x_v)^{-2}$ and $h(q) = q^T M_\delta q$. The associated gradient field evaluated at q is $\nabla m = \nabla(gh) = g\nabla h + h\nabla g$, or more explicitly

$$\nabla m = 2(q^T x_v)^{-2} M_\delta q - 2(q^T x_v)^{-3} (q^T M_\delta q) x_v. \quad (29)$$

Any stationary point, q_v^* say, solves

$$(q_v^{*T} x_v) M_\delta q_v^* = (q_v^{*T} M_\delta q_v^*) x_v. \quad (30)$$

Since the right hand side belongs to the one-dimensional subspace of \mathbb{R}^n spanned by x_v , we infer by inspection of the left hand side that $(\delta I_n + X_{-v} X_{-v}^T) q_v^*$ spans the same subspace. Equivalently, there exists an $a \in \mathbb{R} \setminus \{0\}$ such that $M_\delta q_v^* = a x_v$. The matrix M_δ is invertible for any values of n and p and any $\delta > 0$ bounded away from zero, since the minimum eigenvalue of the Gramian matrix $X_{-v} X_{-v}^T$ is zero and for any eigenvalue-eigenvector pair for $X_{-v} X_{-v}^T$, (γ_k, ξ_k) say,

$$(\delta I_n + X_{-v} X_{-v}^T) \xi_k = (\delta + \gamma_k) \xi_k.$$

Thus write

$$q_v^* = a(\delta I_n + X_{-v} X_{-v}^T)^{-1} x_v. \quad (31)$$

Upon substitution into equation (30) we verify that any such choice of a supplies a valid solution to (30).

This set of solutions consists of local minimizers of m if and only if $\nabla \nabla^T m(q_v^*)$ is non-negative definite for any q_v^* of the form given in equation (31). Equivalently $z^T \{\nabla \nabla^T m(q_v^*)\} z \geq$

0 for any $z \in \mathbb{R}^n$. Otherwise the solutions q_v^* represent other critical points. Let $z \in \mathbb{R}^n$. The corresponding directional derivative, $\nabla_z m(q) := z^T (\nabla m)$, of m at q is

$$\nabla_z m(q) = 2(q^T x_v)^{-2} z^T M_\delta q - 2(q^T x_v)^{-3} (q^T M_\delta q) z^T x_v. \quad (32)$$

For $u, z \in \mathbb{R}^n$,

$$\begin{aligned} \nabla_u \nabla_z m(q) &= u^T \{ \nabla (z^T \nabla m) \} = -4(q^T x_v)^{-3} (z^T M_\delta q) (u^T x_v) + 2(q^T x_v)^{-2} u^T M_\delta z \\ &\quad + 6(q^T x_v)^{-4} (q^T M_\delta q) (z^T x_v) (u^T x_v) - 4(q^T x_v)^{-3} (u^T M_\delta q) (z^T x_v), \end{aligned}$$

so that,

$$\nabla_z \nabla_z m(q_v^*) = 2 \{ x_v^T M_\delta^{-1} x_v \}^{-2} z^T [M_\delta - \{ x_v^T M_\delta^{-1} x_v \}^{-1} x_v x_v^T] z. \quad (33)$$

The conclusion follows because, for any function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, $z^T \{ \nabla \nabla^T f(q) \} z = \nabla_z \nabla_z f(q)$. \square

A.2 Proof of Proposition 2

Proof. In view of the well known matrix identity

$$(A + BCD)^{-1} = A^{-1} - A^{-1}B(C^{-1} + DA^{-1}B)^{-1}DA^{-1},$$

we see that

$$P^v(a, \delta) = a(\delta I_n + X_{-v} X_{-v}^T)^{-1} = a\delta^{-1} (I_n - X_{-v} (\delta I_{p-1} + X_{-v}^T X_{-v})^{-1} X_{-v}^T). \quad (34)$$

On setting $a = \delta$ and taking the limit as $\delta \rightarrow 0^+$,

$$P^v = I_n - X_{-v} \lim_{\delta \rightarrow 0^+} (\delta I_{p-1} + X_{-v}^T X_{-v})^{-1} X_{-v}^T = I_n - X_{-v} X_{-v}^+, \quad (35)$$

where X_{-v}^+ is the Moore-Penrose pseudo-inverse of X_{-v} , the unique matrix satisfying the equations (Penrose, 1955):

$$\begin{aligned} X_{-v} X_{-v}^+ X_{-v} &= X_{-v} & (X_{-v} X_{-v}^+)^T &= X_{-v} X_{-v}^+ \\ X_{-v}^+ X_{-v} X_{-v}^+ &= X_{-v}^+ & (X_{-v}^+ X_{-v})^T &= X_{-v}^+ X_{-v}. \end{aligned} \quad (36)$$

It follows from these equations that $I_n - X_{-v} X_{-v}^+$ is symmetric and equal to its square, from which we conclude that it is an orthogonal projection operator into the kernel of X_{-v}^T . \square

A.3 Proof of Proposition 3

Proof. Recall that $b_v = \sum_{w \in \mathcal{S}} \vartheta_w \beta_w$, where $\vartheta_w = (q_v^T x_v)^{-1} q_v^T x_w$. When $(p-1) > n$, there exists a $(p-1)$ -dimensional vector b , say, $(p-1) - n$ entries of which are zero, such that

$x_v = X_{-v}b$. Thus, on substituting the expression (13) for q_v ,

$$\begin{aligned} b_v &= \{x_v^T(\delta I_n + X_{-v}X_{-v}^T)^{-1}x_v\}^{-1}\beta_{-v}^T X_{-v}^T(\delta I_n + X_{-v}X_{-v}^T)^{-1}x_v \\ &= \{b^T X_{-v}^T(\delta I_n + X_{-v}X_{-v}^T)^{-1}x_v\}^{-1}\beta_{-v}^T X_{-v}^T(\delta I_n + X_{-v}X_{-v}^T)^{-1}x_v. \end{aligned} \quad (37)$$

For any $(p-1)$ -dimensional vector z whose entries are bounded away from zero and of order $h(n, p, \delta)$, to be discussed below, $b^T z \asymp n \cdot h(n, p, \delta)$. Similarly, by sparsity of β , $\beta^T z \asymp s \cdot h(n, p, \delta)$. Since δ is bounded away from zero, so are the entries of $X_{-v}^T \delta (\delta I_n + X_{-v}X_{-v}^T)^{-1}x_v = X_{-v}^T P^v(\delta, \delta)x_v$ by equation (15). It follows on multiplying and dividing (37) by δ that $b_v = O(s/n)$.

It remains to consider the behaviour of

$$\begin{aligned} \frac{q_v^T q_v}{(q_v^T x_v)^2} &= \frac{x_v^T(\delta I_n + X_{-v}X_{-v}^T)^{-1}(\delta I_n + X_{-v}X_{-v}^T)^{-1}x_v}{(x_v^T(\delta I_n + X_{-v}X_{-v}^T)^{-1}x_v)^2} \\ &= \frac{\delta x_v^T \delta (\delta I_n + X_{-v}X_{-v}^T)^{-1}(\delta I_n + X_{-v}X_{-v}^T)^{-1}x_v}{(x_v^T \delta (\delta I_n + X_{-v}X_{-v}^T)^{-1}x_v)^2}. \end{aligned}$$

With b the same vector as defined above, the last line of the previous equation can be written as

$$\frac{\delta b^T X_{-v}^T P^v(\delta, \delta)z}{(b^T X_{-v}^T P^v(\delta, \delta)x_v)^2} \asymp \frac{\delta n h(n, p, \delta)}{(n h(n, p, \delta))^2},$$

where $z = (\delta I_n + X_{-v}X_{-v}^T)^{-1}x_v$ and the rate of convergence arises because both $X_{-v}^T P^v(\delta, \delta)z$ and $X_{-v}^T P^v(\delta, \delta)x_v$ have entries of order $h(n, p, \delta)$. It follows that the term ignored by using $U_n(\beta_v)$ in place of $S_n(\beta_v)$ is

$$\frac{(q_v^T x_v)b_v}{(\tau q_v^T q_v)^{1/2}} = O\left(\frac{s}{n^{1/2}} \cdot \frac{h(n, p, \delta)^{1/2}}{\delta^{1/2}}\right).$$

□

A.4 Proof of Proposition 4

Proof. Let $q_{j,v}$ denote the j th entry of q_v and write $S_n(\beta_v)$ as $\sum_{j=1}^n \xi_j$, where

$$\xi_j := (q_v^T q_v)^{-1/2} q_{j,v} \varepsilon_j.$$

We have $\mathbb{E}(\xi_j) = 0$ and $\text{var}(\sum_j \xi_j) = \tau$. To establish a Lindeberg-Feller central limit theorem (e.g. Kallenberg, 1997, Theorem 4.12) for $\sum_j \xi_j$, we need only verify the Lindeberg condition, i.e., for all $\eta > 0$,

$$\lim_{n \rightarrow \infty} \tau^{-1} \sum_{j=1}^n \mathbb{E}(\xi_j^2 \mathbb{I}\{|\xi_j| > \eta \tau^{1/2}\}) = 0.$$

Let c_n denote a lower bound on $q_v^T q_v$. Note that c_n is bounded away from zero because otherwise the evaluation of the objective function in (6) at q_v is (positive and) unbounded,

contradicting the definition of q_v as a minimizer. Let $\zeta_n = \|q_v\|_{\max}$. Then $|\xi_j| \leq c_n^{-1}\zeta_n|\varepsilon_j|$ and the event $\{|\xi_j| > \eta\tau^{1/2}\}$ is contained in the event $\{|\varepsilon_j| > \eta\tau^{1/2}c_n\zeta_n^{-1}\}$. It follows that

$$\begin{aligned} \tau^{-1} \sum_{j=1}^n \mathbb{E}(\xi_j^2 \mathbb{I}\{|\xi_j| > \eta\tau^{1/2}\}) &\leq \tau^{-1} \sum_{j=1}^n \mathbb{E}(\xi_j^2 \mathbb{I}\{|\varepsilon_j| > \eta\tau^{1/2}c_n\zeta_n^{-1}\}) \\ &= \tau^{-1} (q_v^T q_v)^{-1} \sum_{j=1}^n q_{j,v}^2 \mathbb{E}(\varepsilon_j^2 \mathbb{I}\{|\varepsilon_j| > \eta\tau^{1/2}c_n\zeta_n^{-1}\}) \\ &= \tau^{-1} \mathbb{E}(\varepsilon_j^2 \mathbb{I}\{|\varepsilon_j| > \eta\tau^{1/2}c_n\zeta_n^{-1}\}). \end{aligned}$$

Let $\delta = \eta\tau^{1/2}c_n\zeta_n^{-1}$. For any $\gamma > 0$,

$$\mathbb{E}(\varepsilon_j^2 \mathbb{I}\{|\varepsilon_j| > \delta\}) \leq \mathbb{E}(\varepsilon_j^2 \delta^{-\gamma} |\varepsilon_j|^\gamma \mathbb{I}\{|\varepsilon_j| > \delta\}) \leq \delta^{-\gamma} \mathbb{E}(|\varepsilon_j|^{2+\gamma}).$$

Thus, since $\mathbb{E}(|\varepsilon_j|^{2+\gamma}) < \infty$ and $\zeta_n^\gamma = \|q_v\|_{\max}^\gamma = o(\|q_v\|_2^\gamma) \asymp c_n^\gamma$,

$$\lim_{n \rightarrow \infty} \tau^{-1} \sum_{j=1}^n \mathbb{E}(\xi_j^2 \mathbb{I}\{|\xi_j| > \eta\tau^{1/2}\}) = \lim_{n \rightarrow \infty} (\eta\tau^{1/2}c_n\zeta_n^{-1})^{-\gamma} \mathbb{E}(|\varepsilon_j|^{2+\gamma}) = 0$$

and the Lindeberg condition is verified. That $\|q_v\|_{\max}^\gamma = o(\|q_v\|_2^\gamma)$ for any $\gamma > 0$ follows from (13), which ensures that the number of non-zero elements of q_v increases with n .

Equation (19) is an application of the Berry-Esseen Theorem for non-identically distributed random variables (e.g. Petrov, 1995, Chapter 5) with the constant C improved by Tyurin (2010). \square

A.5 Proof of Proposition 5

Proof. The second term in the decomposition

$$\begin{aligned} \sup_{z \in \mathbb{R}} |\text{pr}\{V_n^{-1/2}S_n(\beta_v) \leq z\} - \Phi(z)| &\leq \sup_{z \in \mathbb{R}} |\text{pr}\{V_n^{-1/2}S_n(\beta_v) \leq z\} - \text{pr}\{\tau^{-1/2}S_n(\beta_v) \leq z\}| \\ &+ \sup_{z \in \mathbb{R}} |\text{pr}\{\tau^{-1/2}S_n(\beta_v) \leq z\} - \Phi(z)| =: \text{I} + \text{II}, \end{aligned} \quad (38)$$

is bounded by Ce_n using Proposition 4.

For any $t > 0$, define the events

$$\begin{aligned} \mathcal{E}_1(t) &:= \{(\tau^{-1/2} - V_n^{-1/2})S_n(\beta_v) \leq t\}, \\ \mathcal{E}_2(t) &:= \{|(V_n^{-1/2} - \tau^{-1/2})S_n(\beta_v)| \leq t\}. \end{aligned}$$

On the event $\{V_n^{-1/2}S_n(\beta_v) \leq z\} \cap \mathcal{E}_1(t)$,

$$\tau^{-1/2}S_n(\beta_v) - t \leq V_n^{-1/2}S_n(\beta_v) \leq z.$$

Thus $\{V_n^{-1/2}S_n(\beta_v) \leq z\} \cap \mathcal{E}_1(t) \subseteq \{\tau^{-1/2}S_n(\beta_v) \leq z + t\}$, and it follows on writing

$$\{V_n^{-1/2}S_n(\beta_v)\} = \left\{ \{V_n^{-1/2}S_n(\beta_v)\} \cap \mathcal{E}_1(t) \right\} \cup \left\{ \{V_n^{-1/2}S_n(\beta_v)\} \cap \mathcal{E}_1^c(t) \right\}$$

that

$$\text{pr}\{V_n^{-1/2}S_n(\beta_v) \leq z\} \leq \text{pr}\{\tau^{-1/2}S_n(\beta_v) \leq z + t\} + \text{pr}\{\mathcal{E}_1^c(t)\}, \quad (39)$$

where $\mathcal{E}_1^c(t)$ denotes the complement of $\mathcal{E}_1(t)$. Later t will be replaced by a decreasing function of n , thus control over I in (38) requires an asymptotically matching lower bound on $\text{pr}\{V_n^{-1/2}S_n(\beta_v) \leq z\}$ as $t \rightarrow 0$. To this end, consider

$$V_n^{-1/2}S_n(\beta_v) \leq \tau^{-1/2}S_n(\beta_v) + |(V_n^{-1/2} - \tau^{-1/2})S_n(\beta_v)|,$$

so

$$\{V_n^{-1/2}S_n(\beta_v) \leq z\} \supseteq \{\tau^{-1/2}S_n(\beta_v) + |(V_n^{-1/2} - \tau^{-1/2})S_n(\beta_v)| \leq z\} =: \mathcal{A}.$$

Furthermore,

$$\mathcal{A} \supseteq \mathcal{A} \cap \mathcal{E}_2(t) \supseteq \{\tau^{-1/2}S_n(\beta_v) \leq z - t\} \cap \mathcal{E}_2(t),$$

so that

$$\text{pr}\{V_n^{-1/2}S_n(\beta_v)\} \geq \text{pr}\{\tau^{-1/2}S_n(\beta_v) \leq z - t\} \text{pr}\{\mathcal{E}_2(t)\} \quad (40)$$

and equations (39) and (40) together give

$$\begin{aligned} & \sup_{z \in \mathbb{R}} |\text{pr}\{V_n^{-1/2}S_n(\beta_v) \leq z\} - \text{pr}\{\tau^{-1/2}S_n(\beta_v) \leq z\}| \quad (41) \\ & \leq \sup_{z': |z-z'| \leq t} \max \left\{ |\text{pr}\{\tau^{-1/2}S_n(\beta_v) \leq z'\} \text{pr}\{\mathcal{E}_2(t)\} - \text{pr}\{\tau^{-1/2}S_n(\beta_v) \leq z\}|, \right. \\ & \quad \left. |\text{pr}\{\tau^{-1/2}S_n(\beta_v) \leq z'\} - \text{pr}\{\tau^{-1/2}S_n(\beta_v) \leq z\}| + \text{pr}\{\mathcal{E}_1^c(t)\} \right\} \end{aligned}$$

for any $t > 0$. Additionally, by Proposition 4 and Taylor series expansion of the standard normal distribution function,

$$\begin{aligned} & \sup_{z, z': |z-z'| \leq t} |\text{pr}\{\tau^{-1/2}S_n(\beta_v) \leq z'\} - \text{pr}\{\tau^{-1/2}S_n(\beta_v) \leq z\}| \\ & \asymp \sup_{z, z': |z-z'| \leq t} |\Phi(z') - \Phi(z)| + e_n \asymp t + e_n. \end{aligned}$$

As a function of $t \rightarrow 0$, we have

$$\text{pr}\{\mathcal{E}_1^c(t)\} \asymp \text{pr}\{|(V_n^{-1/2} - \tau^{-1/2})S_n(\beta_v)| \leq t\} = 1 - \text{pr}\{\mathcal{E}_2(t)\}$$

so that $\text{pr}\{\mathcal{E}_2(t)\} \rightarrow 1$ at the same rate in t as $\text{pr}\{\mathcal{E}_1^c(t)\} \rightarrow 0$. Let d_n be a decreasing function and g_n a non-decreasing function, both to be specified. For two arbitrary constants c_1 and c_2 ,

consider the event

$$\mathcal{D} := \{|(V_n^{-1/2} - \tau^{-1/2})| |S_n(\beta_v)| > c_1 c_2 d_n g_n\} = \{\mathcal{D} \cap \mathcal{B}\} \cup \{\mathcal{D} \cap \mathcal{B}^c\},$$

where

$$\mathcal{B} = \{|S_n(\beta_v)| \leq c_2 g_n\} = \{|\tau^{-1/2} S_n(\beta_v)| \leq \tau^{-1/2} c_2 g_n\}.$$

On \mathcal{B} , $|S_n(\beta_v)/c_2 g_n| \leq 1$, yielding $\mathcal{D} \cap \mathcal{B} \subseteq \{|V_n^{-1/2} - \tau^{-1/2}| > c_1 d_n\}$. Thus

$$\text{pr}\{|(V_n^{-1/2} - \tau^{-1/2}) S_n(\beta_v)| > c_1 c_2 d_n g_n\} \leq \text{pr}(|V_n^{-1/2} - \tau^{-1/2}| > c_1 d_n) + \text{pr}(\mathcal{B}^c), \quad (42)$$

where $\text{pr}(\mathcal{B}^c) \asymp \{1 - \Phi(g_n)\} + e_n$ by Proposition 4. The conclusion follows by combining (38), (41) and (42). \square

B Details of the analysis of §5.2

The cross correlations of all potential explanatory variables were examined and any adjacent variables whose correlations exceeded 0.95 were treated as a single variable for the purpose of constructing $\widehat{\mathcal{S}}$, with the first of each such pair being retained in the analysis. Among the 4088 variables initially present, 61 were discarded on these grounds, leaving 4027 variables.

We implemented a refinement of Cox and Battey's method, as Battey and Cox (2018) demonstrated that using all the available data to identify first the encompassing model $\widehat{\mathcal{S}}$ and then the confidence set of models \mathcal{M} leads to lower than nominal coverage for the confidence set of models. Thus we used a sample splitting procedure, which serves to give a better calibrated set \mathcal{M} but also enables unbiased estimation of error variance for construction of the confidence intervals for components of β . The sample of size $n = 71$ observations was randomly partitioned into subsets, \mathcal{I}_1 and \mathcal{I}_2 , say, of sizes $n_1 = 35$ and $n_2 = 36$. In the terminology of Cox and Battey (2017), the encompassing model was identified using \mathcal{I}_1 for the cube reduction and \mathcal{I}_2 for the square reduction. The confidence set of models and estimate of error variance was constructed using \mathcal{I}_1 .

To improve the stability of this sample-splitting procedure, we repeated it 1000 times, re-randomizing the positions of the variable indices in the cube each time. The subset $\widehat{\mathcal{S}}$ was taken to be all those variables that survived the two-stage procedure in at least half of these randomizations. The confidence intervals constructed as described in §2 for the 22 variables in $\widehat{\mathcal{S}}$ are reported in Table 1. These have been ordered according to the proportion of models from \mathcal{M} to which they belong. For comparison, the lasso and the elastic net, fitted to all 71 observations and tuned to retain 22 variables, find nine and fourteen of the same variables. These are indicated in the first column of Table 1.

From the variables constituting the models in Table 2, three were the first in a string of two or more variables whose empirical correlations exceeded 0.95 (see the first paragraph of this subsection). As far as \mathcal{M} is concerned, such variables are essentially interchangeable. The

affected variables are 1278 (paired with 1279), 1285 (paired with 1286, 1287, 1288 and 1289) and 4002 (paired with 4003, 4004, 4005 and 4006).

References

- BATTEY, H. S. and COX, D. R. (2018). Large numbers of explanatory variables: a probabilistic assessment. *Proc. R. Soc. Lond. A.*, 474, 20170631
- BERK, R., BROWN, L., BUJA, A., ZHANG, K. and ZHAO, L. (2013). Valid post-selection inference. *Ann. Statist.*, 41, 802–837.
- BOYD, S. and VANDENBERGHE, L. (2004). *Convex Optimization*. Cambridge University Press, Cambridge, UK.
- DE BRUIJN, N. G. (1981). *Asymptotic Methods in Analysis*. Corrected reprint of the third edition. Dover Publications, New York.
- BÜHLMANN, P., KALISCH, M. and MEIR, L. (2014). High-dimensional statistics with a view toward applications in biology. *Annu. Rev. Stat. Appl.*, 1, 255–278.
- COX, D. R. (1958). Some problems connected with statistical inference. *Ann. Math. Statist.*, 29, 357–372.
- COX, D. R. (1975). A note on data-splitting for the evaluation of significance levels. *Biometrika*, 62, 441–444.
- COX, D. R. and BATTEY, H. S. (2017). Large numbers of explanatory variables, a semi-descriptive analysis. *Proc. Nat. Acad. Sci.*, 114, 8592–8595.
- COX, D. R. and REID, N. (1987). Parameter orthogonality and approximate conditional inference (with discussion). *J. R. Statist. Soc. B*, 49, 1–39.
- FAN, J., GUO, S. and HAO, N. (2012). Variance estimation using refitted cross-validation in ultrahigh dimensional regression. *J. R. Statist. Soc. B*, 74, 37–65.
- FISHER, R. A. (1930). Inverse probability. *Proc. Cam. Phil. Soc*, 26, 528–535.
- HOERL, A. E. and KENNARD, R. W. (1970). Ridge regression: biased estimation for nonorthogonal problems. *Technometrics*, 12, 55–67.
- JAVANMARD, A. and MONTANARI, A. (2014). Confidence intervals and hypothesis testing for high-dimensional regression. *J. Mach. Learning. Res.*, 15, 2869–2909.
- KALLENBERG, O. (1997). *Foundations of Modern Probability*. Springer, Berlin.
- LI, Y., ZHENG, Z., ZHOU, J. and ZHU, Z. (2021). High-dimensional inference via hybrid orthogonalization. *arXiv:2111.1339v1*.

- LEE, J. D., SUN, D. L., SUN, Y. and TAYLOR, J. E. (2016). Exact post-selection inference, with application to the lasso. *Ann. Statist.*, 44, 907–927.
- LEEB, H. and PÖTSCHER, B. M. and EWALD, K. (2015). On various confidence intervals post-model-selection. *Statist. Sci.*, 30, 216–227.
- LEINER, J., DUAN, B., WASSERMAN, L. and RAMDAS, A. (2021). Data blurring: sample splitting a single sample. *arXiv:2112.11079*.
- LEWIS, R. M. and BATTEY, H. S. (2022). Some theoretical elucidation of Cox reduction and confidence sets of models. *Preprint*.
- LOCKHART, R., TAYLOR, J. E., TIBSHIRANI, R. J. and TIBSHIRANI, R. (2014). A significance test for the lasso. *Ann. Statist.*, 42, 413–468.
- NING, Y. and LIU, H. (2017). A general theory of hypothesis tests and confidence regions for sparse high dimensional models. *Ann. Statist.*, 45, 158–195.
- PENROSE, R. (1955). A generalized inverse for matrices. *Proc. Cambridge Philos. Soc.*, 51, 406–413.
- PETROV, V. V. (1995). *Limit Theorems of Probability Theory*. Oxford University Press, Oxford.
- RASINES, D. G. and YOUNG, G. A. (2021). Splitting strategies for post-selection inference. *arXiv:2102.02159*.
- TIAN, X. and TAYLOR, J. E. (2018). Selective inference with a randomized response. *Ann. Statist.*, 46, 679–710.
- TYURIN, I. S. (2010). An improvement of upper estimates of the constants in the Lyapunov theorem. *Russian Mathematical Surveys*, 65, 201–202.
- VAN DE GEER, S., BÜHLMANN, P., RITOV, Y. and DEZEURE, R. (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *Ann. Statist.*, 42, 1166–1202.
- YATES, F. (1936). A new method of arranging variety trials involving a large number of varieties. *J. Agric. Sci.*, 26, 424–455.
- ZHANG, C-H. and ZHANG, S. S. (2014). Confidence intervals for low dimensional parameters in high dimensional linear models. *J. R. Statist. Soc. B*, 76, 217–242.
- ZHAO, S., WITTEN, D. and SHAJAIE, A. (2021). In defence of the indefensible: a very naïve approach to high-dimensional inference. *Statist. Sci.*, 36, 562–577.