

DISCUSSION: “A SIGNIFICANCE TEST FOR THE LASSO”

BY JIANQING FAN¹ AND ZHENG TRACY KE²

Princeton University

We wholeheartedly congratulate Lockhart, Taylor, Tibshirani and Tibshirani on the stimulating paper, which provides insights into statistical inference based on the lasso solution path. The authors proposed novel covariance statistics for testing the significance of predictor variables as they enter the active set, which formalizes the data-adaptive test based on the lasso path. The observation that “shrinkage” balances “adaptivity” to yield to an asymptotic $\text{Exp}(1)$ null distribution is inspiring, and the mathematical analysis is delicate and intriguing.

Adopting the notation from the paper under discussion, the main results are that the covariance statistics (Theorem 1)

$$(1) \quad (T_{k_0+1}, T_{k_0+2}, \dots, T_{k_0+d}) \xrightarrow{d} (\text{Exp}(1), \text{Exp}(1/2), \dots, \text{Exp}(1/d))$$

for orthogonal designs, and under the global null model (Theorem 2), $T_1 \xrightarrow{d} \text{Exp}(1)$, and under the general model (Theorem 3), $P(T_{k_0+1} \geq t) \leq \exp(-t) + o(1)$. These remarkable results are derived under a number of critical assumptions such as the normality, the sure screening [borrowing the terminology of Fan and Lv (2008)] or model selection consistency of the lasso path. As pointed out in Fan and Li (2001), lasso introduces biases that are hard to account for. This together with the popularity of lasso give rise to the importance of this work, which results in informal statistical inference for the lasso. We welcome the opportunity to make a few comments.

1. Asymptotic null distributions. A natural question is how accurate the approximation (1) is and whether it holds for more general design matrices. We illustrate this using a small-scale numerical study. We take the same settings as in Section 5.2 (Table 2) by considering the global null true model with four types of design matrices: orthogonal, equal correlation, $AR(1)$ and block diagonal, where the parameter $\rho = 0.8$. We fix $n = 100$ and $p = 10$ and 50 . When $p = 50$, the marginal distributions of $\{T_1, T_2, T_3\}$ are very close to the theoretical ones given by (1). However, when $p = 10$, the approximation is not accurate for the “equal correlation” and “ $AR(1)$ ” designs. Figure 1 depicts the results for $p = 10$. The accuracies for the “orthogonal” and “block diagonal” designs are reasonable (omitted) and the accuracy for T_3 is in general worse than those for T_1 and T_2 .

Received December 2013.

¹Supported by NIH Grants R01-GM072611 and R01-GM100474 and NSF Grant DMS-12-06464.

²Supported by NIH Grant R01-GM072611 and NSF Grant DMS-12-06464.

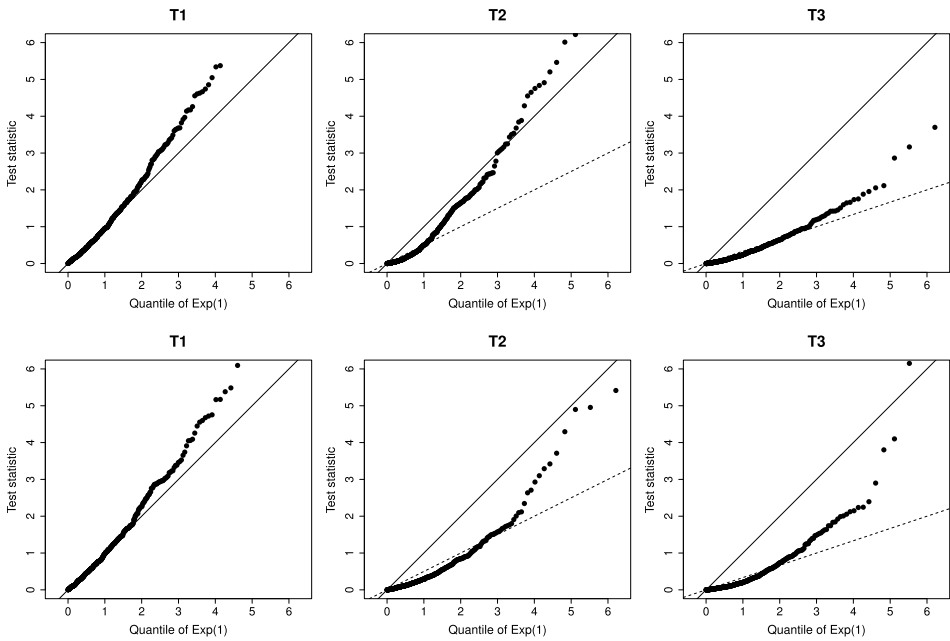


FIG. 1. *Quantile–quantile plots of the covariance test statistics versus their theoretical distributions under the global true null model with “equal correlation” design (top panel) and “AR(1)” design (bottom panel) for $n = 100$ and $p = 10$, based on 500 simulations.*

To check the bivariate behavior of the covariance statistics T_1 , T_2 and T_3 , we transform the statistics to have the asymptotic uniform distribution using (1). The scatter plots of those transformed statistics are presented in Figure 2 based on 500 simulations. They are approximately uniformly distributed in the unit square. This demonstrates that the test statistics are indeed asymptotically independent and that the marginal distributions are accurate for the given setting.

The simulation results presented in Figures 1 and 2 suggest that (1) holds for more general designs, not just for orthogonal designs. This corresponds to suggesting that Theorem 1 of the main paper holds more generally.

For a more general case in Theorem 3, the authors give a nice upper bound. It requires a sure screening property and other conditions. A large number of false positives in the set A_0 of the lasso path at step k_0 should make the upper bound very crude and the upper bound is tight when A_0 is model selection consistent. This can easily be seen from the orthogonal design case with the global null true model. In this case, from (1),

$$(2) \quad T_{k_0+1} \stackrel{a}{\sim} \text{Exp}(1)/k_0,$$

which is of course stochastically bounded by $\text{Exp}(1)$ but this bound can be very crude when k_0 is large.

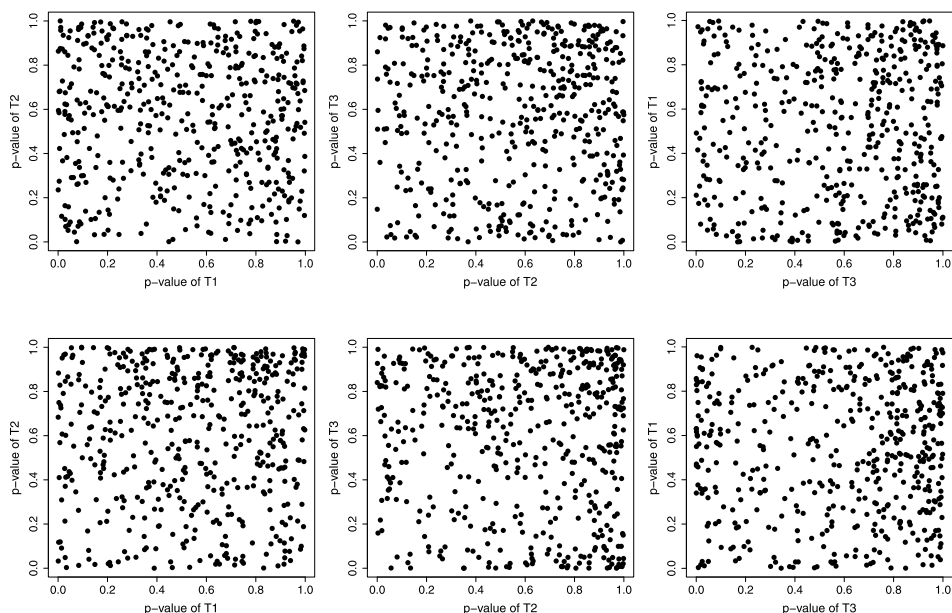


FIG. 2. p -values of T_{k_0} versus those of T_{k_0+1} for $n = 100$ and $p = 10$ with the orthogonal design matrices (top panel) and $n = 100$ and $p = 50$ with the equal-correlation design matrices (bottom panel) based on 500 simulations.

Getting the sure screening property is difficult for lasso when the irresrepresentable condition [Zhao and Yu (2006)] does not hold. This was demonstrated in Fan and Song (2010) in which the design matrix is generated such that $\{X_j\}_{j=1}^{p-50}$ are i.i.d. standard normal variables and the last 50 predictor variables are

$$X_k = \sum_{j=1}^s \frac{(-1)^{j+1}}{5} X_j + \frac{\sqrt{25-s}}{5} \varepsilon_k, \quad k = p-49, \dots, p,$$

where $\{\varepsilon_k\}_{k=p-49}^p$ are i.i.d. standard normal variables and $\{X_j\}_{j=1}^s$ are important variables. They also noted that the larger the intrinsic model size s , the harder the irresrepresentable condition to hold; the larger the dimensionality, the harder the condition. These follow from the definition of irresrepresentable condition. The question then arises what the null distribution is when there are many false positives or even some false negatives.

To provide the insights, we fix $n = 600$, $p = 2000$, and $s = 6$, take the regression coefficient vector β with $\beta_1 = \dots = \beta_s = 5$ and $\beta_{s+1} = \dots = \beta_p = 0$, and simulated 500 data sets. We computed the test statistic T_{k_0+1} at $k_0 = 6$ and $k_0 = 15$. The results are shown in Figure 3. As expected, Theorem 3 continues to hold, but the bound is uselessly crude. For $k_0 = 6$, there are only 43.4% of the lasso paths that have the sure screening or equivalently the model selection consistency; others have both false positives and false negatives. As a result, while

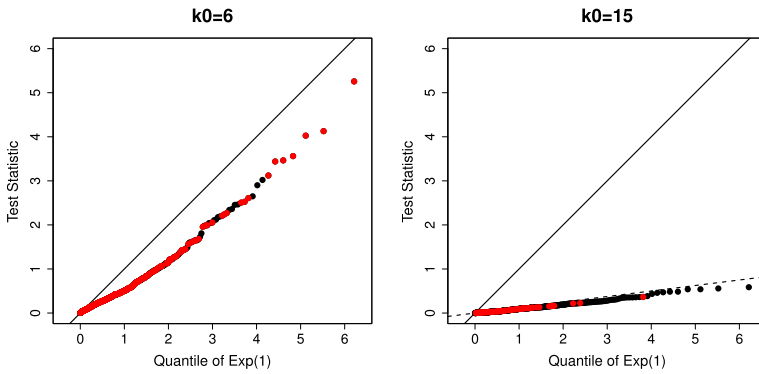


FIG. 3. *Quantile–quantile plots of the covariance test statistics versus Exp(1). By taking $k_0 = 6$ and $k_0 = 15$, the percentiles of sure screening are 43.4% and 87.8%, respectively. The black/red dots correspond to simulations with/without sure screening at step k_0 . The dash line on the right panel has slope $1/9$, matching the distribution given by (1) with $k_0 = 6$ and $d = 9$.*

Theorem 3 continues to hold, the bound is too crude. We have also taken $k_0 = 15$, which makes 87.8% of lasso paths to have sure screening. In this case, there are many (at least 9) false positives. Not knowing the true model size is 6, we compare it with Exp(1) distribution, which shows again that Theorem 3 is correct, but the bound is too crude to be useful. Interestingly, although this is not supported by Theorems 1–3, the test statistic T_{k_0+1} with $k_0 = 15$ is very close to Exp(1/9), even though there are many false positives or even false negatives. Is there any deeper theory underpinning the plot or is it just a coincidence?

Another important condition is the normality assumption. This does not seem as critical, thanks to the central limit theorem. For the orthogonal design case, the variable $X_j^T y$ is approximately normal under some mild conditions. For the logistic regression and Cox’s proportional hazards models, Figures 8 and 9 of the main paper show that the covariance test statistic has approximately Exp(1) distribution. Formal verifications of these results pose some technical challenges, but are interesting research problems.

2. Choice of the model size k_0 . The choice of model size k_0 is critically important. First, it should be large enough to ensure the sure screening. Second, it should not be too large to make overconservative inferences. For the current paper, k_0 directly relates to the null distribution that is used for computing p -values.

Let T_k be the covariance statistic, defined by (7) and simplified in (9) in the main paper. For a given k_0 , define

$$\tilde{T}_{k_0,j} = jT_{k_0+j} \quad \text{for } j = 1, \dots, d.$$

When k_0 is the correct model size so that the model selection consistency holds, from (1), $\{\tilde{T}_{k_0,j}\}_{j=1}^d$ is a sequence of i.i.d. Exp(1) random variables. Therefore, the

average

$$(3) \quad Q_{k_0} = \frac{1}{d} \sum_{j=1}^d \tilde{T}_{k_0, j} \approx 1.$$

A natural choice of k_0 is the one that makes Q_{k_0} closest to its expected value 1, namely

$$(4) \quad \hat{k}_0 = \underset{k \text{ in a range}}{\operatorname{arg\,min}} |Q_k - 1|.$$

The rationale is that when k_0 is the true model size, for example,

$$E(Q_{k_0+1}) = \frac{1}{d} \sum_{j=1}^d \frac{j}{j+1} = 1 - \frac{1}{d} \sum_{j=1}^d (j+1)^{-1},$$

which is less than 1 and when $k < k_0$, $E Q_k$ is expected to be much bigger than 1 (see Table 1).

To see the accuracy of this method, we note that it is typically the hardest to differentiate the choice of k_0 and $k_0 + 1$ when the true model size is k_0 . The variance of the difference is

$$\begin{aligned} \operatorname{var}(Q_{k_0} - Q_{k_0+1}) &= d^{-2} \operatorname{var}(T_{k_0+1} + \cdots + T_{k_0+d} - dT_{k_0+d+1}) \\ &= d^{-2}(1 + 2^{-2} + \cdots + d^{-2} + d^2/(d+1)^2). \end{aligned}$$

It follows that

$$\begin{aligned} \frac{E(Q_{k_0} - Q_{k_0+1})}{\operatorname{var}(Q_{k_0} - Q_{k_0+1})^{1/2}} &= \frac{\sum_{j=1}^d (j+1)^{-1}}{(1 + 2^{-2} + \cdots + d^{-2} + d^2/(d+1)^2)^{1/2}} \\ &\asymp \frac{\log(d)}{\sqrt{1 + \pi^2/6}} \quad \text{as } d \rightarrow \infty. \end{aligned}$$

TABLE 1

Selection of the model size k_0 . $n = 500$ and the true $k_0 = 2$. Based on 1000 simulations, the mean of Q_k (with standard deviation in the parenthesis) and the distribution of selected \hat{k}_0 are displayed

p	d	k	0	1	2	3	4
10	6	mean Q_k	9.30 (2.3)	4.40 (1.3)	0.76 (0.43)	0.48 (0.28)	0.33 (0.22)
		prob($\hat{k}_0 = k$)	0.0%	0.5%	79.9%	15.2%	4.4%
1000	6	mean Q_k	6.31 (2.0)	3.00 (1.1)	0.93 (0.39)	0.66 (0.29)	0.53 (0.24)
		prob($\hat{k}_0 = k$)	0.0%	0.04%	64.5%	20.5%	10.8%
1000	20	mean Q_k	2.58 (0.62)	1.53 (0.36)	0.85 (0.20)	0.72 (0.17)	0.64 (0.16)
		prob($\hat{k}_0 = k$)	0.0%	22.3%	61.6%	11.6%	4.5%

Thus, the signal to noise ratio is large when d is large, but increases slowly with d . Therefore, in practice, we do not wish to take a too large d due to the accuracy of approximation (1).

We conducted a numerical experiment where $n = 500$ and $p = 10$ and 1000. The predictors $\{X_j\}_{j=1}^p$ are i.i.d. standard normal variables. Let $\beta = (6, 6, 0, \dots, 0)^T$, so the true $k_0 = 2$. For fixed $d = 6$ and $d = 20$ (only when $p = 1000$), we selected \hat{k}_0 from $\{0, \dots, 4\}$ to minimize $|Q_k - 1|$. Table 1 summarizes the results based on 1000 simulations. When $p = 10$, the percentiles of $\hat{k}_0 = 2$ (accurate) and $\hat{k}_0 = 3$ (overshooting by 1) are about 80% and 15%, and there are almost no undershootings ($\hat{k}_0 < 2$). When $p = 1000$, the accuracy decreases to about 65%, but there are still almost no undershootings. Interestingly, when we increase d to 20, the results become inferior, with about 22% of undershootings. This suggests that d should not be chosen too large that smooths out the signals in Q_k for $k < k_0$ and makes (1) inaccurate.

3. Power of the tests. When we test the first few variables that enter the active set of lasso, it is very often that there remain true active variables not yet selected. The covariance test statistics are just one of many possibilities, constructed carefully and intriguingly and supported by the nice asymptotic null distribution. However, they are not necessarily the most powerful tests.

To understand the possible loss of the power of the covariance test, we consider again the simplest setting where the design matrix is orthogonal, $k_0 = 0$ (so the null hypothesis is the global null) and $\sigma = 1$. It follows that

$$T_1 = V_1(V_1 - V_2),$$

where V_1 and V_2 are the first and second largest elements of $\{|X_j^T y| : 1 \leq j \leq p\}$. The factor $V_1 - V_2$ makes the null distribution very beautiful, but this can also reduce the power of the statistic V_1 , which is equal to square root of the maximum drop in RSS.

To demonstrate this, consider the specific alternative

$$\beta_1 = \beta_2 = \theta, \quad \beta_3 = \dots = \beta_p = 0,$$

where $\theta \gg \sqrt{\log(p)}$. With probability tending to 1, $|X_1^T y|$ and $|X_2^T y|$ are the largest two elements. As a result, V_1 is stochastically equivalent to that of $(\theta + \max\{\epsilon_1, \epsilon_2\}) + o_p(1)$ and $V_1 - V_2 = |\epsilon_1 - \epsilon_2| + o_p(1)$ with ϵ_1, ϵ_2 being independent standard normal variables. It follows that

$$(5) \quad T_1/\theta \xrightarrow{d} |N(0, \sqrt{2})| \quad \text{and} \quad V_1/\theta \xrightarrow{d} 1.$$

The statistic T_1 and the maximum drop of RSS V_1^2 indeed have asymptotic power one. On the other hand, (5) shows clearly that T_1 is corrupted by an extra noise $|N(0, \sqrt{2})|$ and is therefore less powerful.




FIG. 4. Power curves based on 1000 simulations. $n = 100$, $p = 10$ and the predictors are drawn *i.i.d.* from $N(0, 1)$. “Cov/Theory” and “Cov/Simulation” refer to the covariance test statistic T_1 , with critical value being the 95% quantile of $\text{Exp}(1)$ and the sample 95% quantile. “MaxCov/Simulation” refers to the maximum of T_1 and T_2 , and “RSSdrop/Simulation” refers to the maximum drop in RSS, with critical values being the sample 95% quantile.

We illustrate this point using a small-scale numerical study. We use similar settings as the left panel of Figure 4 in the main paper ($n = 100$, $p = 10$ and “orthogonal design”). Instead of having only one truly nonzero regression coefficient, we set two equal nonzero regression coefficients. Figure 4 shows the estimated power curves. When there is only one true nonzero coefficient, the covariance test statistic and the maximum drop in RSS have similar powers as shown in Figure 4 of the main paper. On the other hand, when there are two equal nonzero coefficients, the statistic of maximum drop in RSS has a larger power, especially when the signal strength θ is large. Interestingly, when we compute the covariance test statistics in this case, it is quite often that the first entering variable is not very significant but the second one is. We also compute the power when looking at the maximum of T_1 and T_2 . It turns out that this test is more powerful than using T_1 only. See Figure 4.

4. Validity of the results to other penalties. A natural question is whether or not the results in the paper are tied to the lasso path. Given many nice bias properties of folded concave penalty [Fan and Li (2001)] and weighted lasso penalty [Zou (2006)] functions, it is natural to examine the solution paths created by those penalty functions.

For a general penalty function $p_\lambda(\cdot)$, we define the covariance test statistic at the knot λ_k the same as (5) in the main paper, except that $\hat{\beta}(\lambda_{k+1})$ and $\tilde{\beta}_A(\lambda_{k+1})$ are computed with $\|\beta\|_1$ replaced by $\sum_j p_\lambda(\beta_j)$ in the expressions. Although there are the issues on the uniqueness of the folded concave penalized least-squares, Fan and Lv (2011) show that folded concave penalized least-squares estimator is indeed unique in the sense of restricted global optimality.

As in the main paper, we examine the showcase example in which the design matrix is orthogonal. In this case, the penalized least-squares with folded concave penalty is unique [Fan and Li (2001)]. By direct calculation,

$$(6) \quad T_k = V_k \cdot h_{V_{k+1}}(V_k)/\sigma^2,$$

where $h_\lambda(\cdot)$ is a thresholding function defined by $h_\lambda(x) = \arg \min_u \{\frac{1}{2}(u - x)^2 + p_\lambda(u)\}$. For the SCAD penalty [Fan and Li (2001)] with a parameter $a > 2$,

$$T_k^{\text{scad}} = \begin{cases} V_k(V_k - V_{k+1})/\sigma^2, & V_k \leq 2V_{k+1}, \\ \frac{a-1}{a-2}V_k\left(V_k - \frac{a}{a-2}V_{k+1}\right) / \sigma^2, & 2V_{k+1} < V_k < aV_{k+1}, \\ V_k^2/\sigma^2, & V_k \geq aV_{k+1}. \end{cases}$$

We can similarly show that for any fixed $k \geq 1$,

$$(7) \quad (T_1^{\text{scad}}, T_2^{\text{scad}}, \dots, T_k^{\text{scad}}) \xrightarrow{d} (\text{Exp}(1), \text{Exp}(1/2), \dots, \text{Exp}(1/k)),$$

under the global null true model.

PROOF. Let $F(x) = (2\Phi(x) - 1)I\{x > 0\}$. From the proof of Lemma 3 in the main paper, for $a_p = F^{-1}(1 - 1/p)$ and $b_p = pF'(a_p)$, the random variables $W_0 = b_p(V_{k+1} - a_p)$ and $W_i = b_p(V_i - V_{i+1}), i = 1, \dots, k$, converge jointly:

$$(8) \quad (W_0, W_1, W_2, \dots, W_k) \xrightarrow{d} (-\log G_0, E_1, E_2/2, \dots, E_k/k),$$

where G_0, E_1, \dots, E_k are independent, G_0 is Gamma distributed with scale parameter 1 and shape parameter k , and E_1, \dots, E_k are standard exponentials. In addition, $a_p, b_p \rightarrow \infty$ and $a_p/b_p \rightarrow 1$ as $p \rightarrow \infty$.

Note that $T_i^{\text{scad}} = T_i^{\text{lasso}}, i = 1, \dots, k$, on the event $B = \{V_i \leq 2V_{i+1}, 1 \leq i \leq k\}$. By (8) and the fact that $a_p b_p \rightarrow \infty$,

$$P(V_i > 2V_{i+1}) = P\left(iE_i + \log G_0 - \sum_{j=i+1}^k jE_j > a_p b_p\right) = o(1).$$

Then $P(B^c) \leq \sum_{i=1}^k P(V_i > 2V_{i+1}) = o(1)$. Therefore, (7) follows immediately from Lemma 3 and the Slutsky's lemma. \square

For the MCP penalty [Zhang (2010)] with a parameter $\gamma > 1$, it can be shown similarly that

$$(T_1^{\text{mcp}}, T_2^{\text{mcp}}, \dots, T_k^{\text{mcp}}) \xrightarrow{d} \frac{\gamma}{\gamma - 1}(\text{Exp}(1), \text{Exp}(1/2), \dots, \text{Exp}(1/k)).$$

For the weighted lasso penalty [Zou (2006)], the solution path depends on order statistics of variables $\{w_j^{-1}|X_j^T y|\}_{j=1}^p$, where w_j is the weight for variable j . These variables are not identically distributed. It remains an interesting question to what extent the current results can be generalized.

5. Further comments. The mathematical results are derived when d and k_0 are finite. A more interesting asymptotic framework is to let both d and k_0 diverge with n . See, for example, [Fan and Lv \(2011\)](#) for the joint asymptotic distribution when the dimensionality grows with sample size.

The beautiful results in the paper are derived under the assumptions that the signals are very strong and the designs are so nice that sure screening is possible. These assumptions are difficult to meet in practice. Even when they are met, we need to specify k_0 which is hoped to be small and contains all important variables (sure screening). Sure screening assumption implies that the null hypothesis is true. What are we testing: sure screening hypothesis or significance of the newly entered variable? Under the sure screening assumption, why not run the least-squares based on the screened predictors and use splitted data (when needed), as suggested in [Fan and Lv \(2008\)](#) and [Wasserman and Roeder \(2009\)](#)? The statistical inference can be based upon the low-dimensional least-squares theory. To utilize the asymptotic null distribution without conservatism, we need to have the model selection consistency assumption: the first k_0 variables contain all important variables. If so, why do we need the significance tests of the newly entered variables? Relaxing the model selection consistency to sure screening does not help the matter very much. Using the standard exponential distribution as the upper bound of the p -values, we can mislabel many “important variables” as “unimportant ones,” a missed discovery that we strive to avoid in high-dimensional inference.

The authors mentioned in the paper that they plan to construct confidence regions for the lasso $\hat{\beta}(\lambda)$ at specific λ . The challenge here is that there are biases involved in the lasso fit. Another challenge is to give a formal confidence assessment that a group of “unimportant variables” are really unimportant. The efforts are certainly welcome [see, e.g., [Meinshausen, Meier and Bühlmann \(2009\)](#) and [Zhang and Zhang \(2014\)](#)]. We would like to note that for the folded-concave penalized least-squares or likelihood, the resulting estimator is the oracle estimator with probability tending to one [[Fan and Lv \(2011\)](#)]. Therefore, the confidence intervals can easily be constructed based on the low-dimensional likelihood inference. However, it also remains to give confidence assessment that a group of “unimportant variables” are really unimportant.

The authors have mentioned a couple of times the null distribution of the largest RSS drop. This is equivalent to $(1 - \gamma_n^2)$, where γ_n is the maximum correlation coefficient between the residuals at the current step of the forward regression and the covariates. Under the global null true model, this is the maximum spurious correlation between the response and each variable. The asymptotic distribution for the maximum spurious correlation in case where all predictors are independent has been derived in [Cai and Jiang \(2011\)](#). However, we do not expect that the asymptotic null distribution is accurate enough for many applications.

In conclusion, the idea and results in the main paper are insightful and amazing. The technical arguments are delicate and ingenious. The authors should be

congratulated again for successful adaptive inference based on the lasso solution path. We hope that our comments contribute positively to the understanding of this seminal article.

REFERENCES

- CAI, T. T. and JIANG, T. (2011). Limiting laws of coherence of random matrices with applications to testing covariance structure and construction of compressed sensing matrices. *Ann. Statist.* **39** 1496–1525. [MR2850210](#)
- FAN, J. and LI, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.* **96** 1348–1360. [MR1946581](#)
- FAN, J. and LV, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **70** 849–911. [MR2530322](#)
- FAN, J. and LV, J. (2011). Nonconcave penalized likelihood with NP-dimensionality. *IEEE Trans. Inform. Theory* **57** 5467–5484. [MR2849368](#)
- FAN, J. and SONG, R. (2010). Sure independence screening in generalized linear models with NP-dimensionality. *Ann. Statist.* **38** 3567–3604. [MR2766861](#)
- MEINSHAUSEN, N., MEIER, L. and BÜHLMANN, P. (2009). *P*-values for high-dimensional regression. *J. Amer. Statist. Assoc.* **104** 1671–1681. [MR2750584](#)
- WASSERMAN, L. and ROEDER, K. (2009). High-dimensional variable selection. *Ann. Statist.* **37** 2178–2201. [MR2543689](#)
- ZHANG, C.-H. (2010). Nearly unbiased variable selection under minimax concave penalty. *Ann. Statist.* **38** 894–942. [MR2604701](#)
- ZHANG, C.-H. and ZHANG, S. S. (2014). Confidence intervals for low-dimensional parameters in high-dimensional linear models. *J. R. Stat. Soc. Ser. B.* **76** 217–242. [MR3153940](#)
- ZHAO, P. and YU, B. (2006). On model selection consistency of Lasso. *J. Mach. Learn. Res.* **7** 2541–2563. [MR2274449](#)
- ZOU, H. (2006). The adaptive lasso and its oracle properties. *J. Amer. Statist. Assoc.* **101** 1418–1429. [MR2279469](#)

DEPARTMENT OF OPERATIONS RESEARCH
AND FINANCIAL ENGINEERING
PRINCETON UNIVERSITY
PRINCETON, NEW JERSEY 08544
USA
E-MAIL: jqfan@princeton.edu
zke@princeton.edu