The Annals of Statistics
2014, Vol. 42, No. 2, 518–531
DOI: 10.1214/14-AOS1175REJ
Main article DOI: 10.1214/13-AOS1175

REJOINDER: "A SIGNIFICANCE TEST FOR THE LASSO"

BY RICHARD LOCKHART, JONATHAN TAYLOR,
RYAN J. TIBSHIRANI AND ROBERT TIBSHIRANI

Simon Fraser University, Stanford University, Carnegie Mellon University and Stanford University

We would like to thank the Editors and referees for their considerable efforts that improved our paper, and all of the discussants for their feedback, and their thoughtful and stimulating comments. Linear models are central in applied statistics, and inference for adaptive linear modeling is an important active area of research. Our paper is clearly not the last word on the subject. Several of the discussants introduce novel proposals for this problem; in fact, many of the discussions are interesting "mini-papers" on their own, and we will not attempt to reply to all of the points that they raise. Our hope is that our paper and the excellent accompanying discussions will serve as a helpful resource for researchers interested in this topic.

Since the writing of our original paper, we have (with many our of graduate students) extended the work considerably. Before responding to the discussants, we will first summarize this new work because it will be relevant to our responses.

- As mentioned in the last section of the paper, we have derived a "spacing" test of the global null hypothesis, $\beta^* = 0$, which takes the form

$$\frac{1 - \Phi(\lambda_1/\sigma)}{1 - \Phi(\lambda_2/\sigma)} \sim \mathrm{Unif}(0, 1) \tag{1}$$

  for unit normed predictors, $\|X_i\|_2 = 1$, $i = 1, \dots, p$. As opposed to the covariance test theory, this result is exact in finite samples, that is, it is valid for any n and p (and so nonasymptotic). It requires (essentially) only normality of the errors, and no truly stringent assumptions about the predictor matrix X. In many cases, the agreement between this test and the covariance test is very high; details are in Taylor, Loftus and Tibshirani (2013) and Taylor et al. (2014).
- The spacing test (1) is designed for the first step of the lasso path. In Taylor et al. (2014), we generalize it to subsequent steps (this work is most clearly explained when we assume no variable deletions occur along the path, i.e., when we assume the least angle regression path, but can also be extended to the lasso path). In addition, we study a more general pivot that can be inverted to yield "selection intervals" for coefficients of active variables at any step.

Received February 2014.