EFRON, B. (1982). *The Jackknife, the Bootstrap, and Other Resampling Plans*. SIAM, Philadelphia.
EFRON, B. (1987). Better bootstrap confidence intervals (with discussion). *J. Amer. Statist. Assoc.* **82** 171–200.
EFRON, B. and TIBSHIRANI, R. (1986). Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy (with discussion). *Statist. Sci.* **1** 54–77.

DEPARTMENT OF STATISTICS
SEQUOIA HALL
STANFORD UNIVERSITY
STANFORD, CALIFORNIA 94305-4065

WEI-YIN LOH[1]

University of Wisconsin-Madison

It was a pleasure to read Professor Hall's paper, which so effectively analyzes the relative performance of Efron's many recipes for bootstrap interval construction, under the assumption that the parameter is a function of vector means. In this setting, Hall shows that the percentile-t and accelerated bias-corrected methods tie for first place, the main reason being that one consults "Studentized" tables, and the other looks up "ordinary" tables, after employing an analytical correction to adjust the critical points.

I share Hall's prejudice that computer-intensive methods such as the bootstrap should not have to appeal to tedious analytic corrections and therefore agree with his preference of the percentile-t over the accelerated bias-corrected method in the present situation. Because all the bias-corrected methods look up tables "backwards," the percentile-t may also be preferred in those nonlinear and nonsmooth problems where the asymptotic distributions are asymmetric, if we know how to Studentize. One such problem is discussed in Loh (1984).

On the other hand, I believe that the idea of looking up standard tables using adjusted levels has intrinsic merit on its own, and I will present a simple way of doing this which does not look up tables backwards and does not involve difficult analytic manipulations. It turns out that, under the "smooth" model of Hall, this method yields one-sided intervals that are second-order equivalent to the STUD and ABC methods and two-sided intervals that possess coverage errors which are an order of magnitude *smaller* than those of all the methods examined in the paper. Furthermore, it requires no more bootstrap sampling than the rest.

The method I propose has its origin in the "calibrated" method introduced in Loh (1987), and the basic idea is as follows. Starting with any reasonable interval procedure, bootstrap its coverage probability $\pi(\alpha)$. (This is a distinct departure from the other bootstrap recipes because the latter all call for bootstrapping the distribution of a statistic.) After the bootstrap estimate $\hat{\pi}(\alpha)$ is obtained, a corrected α^* is computed that is then used in place of α in the original formula.

[1]Work begun while visiting the T. J. Watson Research Center for a year as an IBM Junior Faculty Research Fellow. Research also partially supported by NSF Grant DMS-85-02303.