

TABLE 2  
Percent coverage for nominal 90% confidence interval\*

X	Bootstrap method							
	$\sigma = 0.1$				$\sigma = 0.5$			
	T1	T1 BC	T2	T2 BC	T1	T1 BC	T2	T2 BC
0	74	94	50	50	83	88	65	63
1/8	0	7	81	50	39	76	70	67
1/4	100	100	100	94	96	89	90	90
3/8	0	3	74	48	33	76	67	66
1/2	100	100	100	95	97	90	91	90

\*T1 means type I percentile bootstrap; T1 BC is type I percentile bias-corrected; T2 is type II bootstrap; T2 BC is type II bias-corrected as described in text.

To confuse matters further, in using the bootstrap to pick a bandwidth for a kernel density estimate, the model generating bootstrap data must be an over-smoother. Failure to recognize these subtleties will result in very poor inferences.

Applying the smoother and then the bootstrap is a breeze (our simulations caused the breeze to blow 1,000 times), and we were able to commit mispractice with practically no effort. There are many other examples where hidden problems with the bootstrap will occur unless one is especially knowledgeable and careful.

Our response to Young's paper and to our example is a call to action. The statistical profession needs to communicate the good news, the bad news and the "no news yet." The bootstrap will succeed for a broad class of models and data structures. It will fail in

others; sometimes it can be rescued by modifications that attend to the structure of the problem. We need to communicate what we know about the procedure's strengths and weaknesses and to identify situations where we do not yet know the answers. This communication must reach current and potential users and thus must appear in a broad array of journals and other information sources. As we learn more, information needs to be updated. Of course, the same recommendations hold for all statistical procedures, but the attraction of the bootstrap makes the need most acute.

#### ACKNOWLEDGMENTS

Partial support was provided by Grant RO1 CA51932 from the National Cancer Institute.

## Comment

David Hinkley

#### INTRODUCTION

This is a timely article. It is likely to appear in print about the same time as first reviews of the excellent introductory book by Efron and Tibshirani (Efron and Tibshirani, 1993), a book which should allay some of the impatience and scepticism that I sense in the sophisticated user community about the bootstrap as a practical tool. We are also beginning to see the first wave of software products which claim to do bootstrap analysis: some of these are embarrass-

ingly naive. Let us hope for more good applications-oriented books and better software products.

I think that Alastair Young has done an excellent job of highlighting the key theoretical developments and has suggested some sensible steps for further research. Much of what I have to say will complement his assessment and will focus on a few practical points.

#### WHEN DOES BOOTSTRAP WORK?

This question comes up twice in the paper, in the context of nonparametric bootstrapping of a point estimator. The first time we are given a succinct mathematical characterization which is clearly useless to even the best applied statistician. The second time

*David Hinkley is Professor of Statistical Science and Head of the Department of Statistics, University of Oxford, 1 South Parks Road, Oxford OX1 3TG, England.*