Statistica Sinica **5**(1995), 161-171

FIELLER'S PROBLEMS AND RESAMPLING TECHNIQUES

J. T. Gene Hwang
(formerly Jiunn T. Hwang)

Cornell University

Abstract. Fieller's problems occur in many areas such as Bioassy and Calibration. The classical solutions based on normality assumptions were proposed in Bliss (1935) and Fieller (1954) and are compared in this paper to the more modern solutions. Based on resampling techniques, confidence intervals have been constructed without the normality assumption. These confidence intervals, however, have low coverage probabilities. An alternative bootstrapping technique is proposed here for Fieller's problems. This produces parametric and nonparametric confidence intervals that closely mimic Fieller's intervals and have good coverage probabilities for the normal model and many other parametric models. Also the nonparametric confidence intervals are demonstrated to be second order correct whereas the Fieller's intervals are only first order correct for the nonnormal observations.

Key words and phrases: Bootstrap, nonparametric confidence interval, coverage probability.

1. Introduction

Fieller's problems in this paper refer to the problems of constructing confidence sets for a ratio of parameters. These problems occur frequently at least in three areas: bioassay, bioequivalence and calibration.

In a bioassay problem (see Finney (1978)) and a bioequivalence problem (see Chow and Liu (1992)), often the relative potency of a new drug to that of a standard drug is expressed in terms of a ratio. Consider a simple example below.

Example 1.1. (bioassay) Let $X = (X_1, \ldots, X_m)$ and $Y = (Y_1, \ldots, Y_n)$ be independent observations of the potency of a new drug and a standard drug respectively. Assume that X_i are independently identically distributed (i.i.d.) with $N(\mu_1, \sigma^2)$ and Y_i are i.i.d. with $N(\mu_2, \sigma^2)$, where μ_1 and μ_2 are the true potencies. The problem is to construct a confidence interval for the relative potency $\theta = \mu_1/\mu_2$ of the new drug to the standard one.

For this example, Bliss (1935) and Fieller (1954) constructed a confidence interval based on the statistic

$$T_0 = (\bar{X} - \theta\bar{Y})/[(1/n + \theta^2/m)^{1/2}\hat{\sigma}], \tag{1.1}$$

where $\bar{X} = \frac{1}{m}\sum X_i$, $\bar{Y} = \frac{1}{n}\sum Y_j$, and

$$\hat{\sigma}^2 = \left[\sum(X_i - \bar{X})^2 + \sum(Y_i - \bar{Y})^2\right] \Big/ (m+n). \tag{1.2}$$

It is obvious $[(m+n-2)/(m+n)]^{1/2}T_0$ has a t distribution with $(m+n-2)$ degrees of freedom. Therefore if q_α is the α upper quantile of the distribution, the probability that

$$[(m+n-2)/(m+n)]^{1/2}|T_0| < q_\alpha \tag{1.3}$$

is $1-2\alpha$. Solving θ in the inequality (1.3) will then give a confidence interval with coverage probability exactly $1-2\alpha$. In doing so, however, quadratic inequalities are involved. Consequently, the interval for θ can be a bounded interval or the complement of a bounded interval or even $(-\infty, \infty)$.

Although it seems to be unfortunate for a confidence interval to be unbounded, by Gleser and Hwang (1987) this is inevitable for a confidence interval with a positive confidence level. To the author, when one ends up with an unbounded interval, it actually serves as a proper warning that perhaps μ_2 is too close to zero for the statistician to draw any conclusion from the study.

The two sample model in Example 1.1, can be generalized to a linear regression model. Fieller's interval can be similarly constructed for the ratio of regression parameters based on a t distributed pivot (see Hwang (1989)). Calibration or inverse regression problems are all special cases.

Fieller's interval is based on normality. By using resampling techniques, alternative confidence intervals were previously constructed without assuming normality. However, the coverage probabilities of these intervals could be low, no matter what the true distribution is.

In this paper, we propose a different approach by bootstrapping the Fieller's statistic such as T_0 in (1.1). Some analytic and numerical evidence presented here show that our proposed intervals have good coverage probabilities.

2. A Brief Explanation

In this section, we first explain why the resampling approaches in the literature produce confidence intervals with low coverage probability when the true parameter is in some parameter space. Taking the model in Example 1.1 as an example, one can establish the following theorem.

Theorem 2.1. *Any confidence interval which has finite length for almost every observation has minimum zero coverage probability while minimizing over* (μ_1, μ_2).

Since all the existing resampling intervals (the BC interval in Efron (1985), and other intervals in Hinkley and Wei (1984), Wu (1986), Simonoff and Tsai (1986) and Chan and Srivastava (1988)) have finite length for every observation, application of this theorem leads one to conclude that their coverage probability can be arbitrarily small.

The argument for Theorem 2.1 is basically as follows. Assume that $C(X,Y)$ denotes a confidence set for θ and the length of $C(X,Y)$ is finite for almost every (X,Y). It then follows that

$$\inf_{\mu_1,\mu_2} P(\theta \in C(X,Y)) \leq \lim_{\substack{\mu_2 \to 0 \\ \mu_1 = 1}} P(\theta \in C(X,Y)) = P_{\substack{\mu_1=1 \\ \mu_2=0}}(\infty \in C(X,Y)),$$

which equals zero since $C(X,Y)$ has finite length. Passing the limit inside the integration can be rigorously justified by using the Bounded Convergence Theorem. Note that the normality assumption is not needed. The theorem holds as long as the probability density function of the data is continuous with respect to the parameters.

A general theorem of this type has been given in Gleser and Hwang (1987), although the argument therein may not be as clear as what is depicted here for our special case.

The theorem of Gleser and Hwang (1987) applies to the estimation of ratios of regression parameters of a linear model (i.e., Fieller's problem), and to the estimation of regression parameters of any linear (and most nonlinear) errors-in-variables models. The contrapositive statement of the Gleser and Hwang (1987) theorem asserts that confidence intervals whose minimum coverage probability is nonzero would inevitably give an unbounded interval with positive probability similar to what Fieller's intervals do.

3. Bootstrapping Fieller's Problems

To construct confidence intervals with good coverage probabilities for the problem in Example 1.1, it appears that one should bootstrap the $T_0(X,Y,\theta) \equiv T_0$ statistic in (1.1).

Now let $X = (X_1, \ldots, X_m)$ and $Y = (Y_1, \ldots, Y_n)$ be the observation vectors. Specifically, one generates bootstrap samples $X^* = (X_1^*, \ldots, X_m^*)$ and $Y^* = (Y_1^*, \ldots, Y_n^*)$ according to various ways to be described after the next paragraph.

To construct a $1 - 2\alpha$ confidence intervals, we then find (by simulation) the α and $1-\alpha$ quantiles q_α and $q_{1-\alpha}$, respectively of the distribution of $T_0^* =$

$T_0(X^*, Y^*, \hat{\theta})$, where $\hat{\theta} = \bar{X}/\bar{Y}$. The $(1-2\alpha)$ equal tailed confidence set consists of θ such that

$$q_\alpha < T_0(X, Y, \theta) < q_{1-\alpha}.$$

Since the denominator of T_0 involves θ, the resultant interval may sometimes have infinite length. Therefore Gleser and Hwang's theorem can not be used to conclude that these confidence intervals have zero minimum coverage probability. In fact, we show in Theorem 3.1 that these types of confidence intervals have minimum coverage probability above a positive nominal level. Hinkley and Wei (1984) proposed bootstrapping a pivot similar to (1.1) except that the denominator is based on some jackknife estimate of the variance. Since their denominator does not involve θ, the resultant interval has finite length. Therefore Theorem 3.1 applies.

There are several ways of generating bootstrap samples, including the four approaches below: (1) parametric bootstrap, (2) independently, similarly distributed nonparametric bootstrap (ISD-bootstrap), (3) independently nonparametric bootstrap (I-bootstrap), and (4) paired nonparametric bootstrap. The corresponding interval will be called, respectively, a parametric bootstrap T_0 interval, an ISD-bootstrap T_0 interval, an I-bootstrap T_0 interval, and an NP-bootstrap T_0 interval.

In parametric bootstrapping, one assumes a parametric model $f_\mu(x, y)$ and generates

$$(X^*, Y^*) \sim f_{\hat{\mu}}(\cdot, \cdot),$$

where $\hat{\mu} = \hat{\mu}(X, Y)$ is an estimator of μ (typically the maximum likelihood estimator).

In ISD-bootstrapping, one assumes independence of X and Y and also assumes that X and Y have the same distribution except that they are shifted by possibly different locations. In this case, let $\epsilon_i = X_i - \bar{X}, 1 \le i \le m$, and $\epsilon_{j+m} = Y_j - \bar{Y}, 1 \le j \le n$. Also let $\epsilon_i^*, 1 \le i \le m+n$, be the number randomly drawn with replacement from $\{\epsilon_1, \ldots, \epsilon_{m+n}\}$. The bootstrap samples then are $X_i^* = \bar{X} + \epsilon_i^*$ and $Y_j^* = \bar{Y} + \epsilon_{j+m}^*$.

In I-bootstrapping, one assumes only the independence of X and Y. Naturally one defines X_i^* to be the quantity randomly drawn with replacement from $\{X_1, \ldots, X_m\}$. Similarly, Y_i^* is randomly chosen from $\{Y_1, \ldots, Y_n\}$.

For the paired nonparametric situation one does not even assume independence of X and Y. One naturally assumes that we have matched data $(X_1, Y_1), \ldots, (X_m, Y_m)$. (Therefore $m = n$.) In such a case one can draw the pair (X_i^*, Y_i^*) randomly with replacement from the n pairs (X_i, Y_i).

For the normal setting in Example 1.1, the normal (parametric) bootstrapping T_0 is appropriate. Namely X_i^* and Y_i^* are independently generated from

$N(\bar{X}, \hat{\sigma}^2)$ and $N(\bar{Y}, \hat{\sigma}^2)$ where $\hat{\sigma}^2$ are as in (1.2). Since T_0 is a pivot (i.e., T_0 has a distribution independent of all the parameters), the distribution of T_0^* is the same as T_0. In fact T_0 and T_0^* after multiplication by a constant, will have a t distribution. Hence the $(1-2\alpha)$ normal bootstrapping T_0 is exactly the $(1-2\alpha)$ Fieller's interval.

Could the zero confidence level phenomenon still happen to the parametric bootstrapping T_0 interval or nonparametric bootstrapping T_0? The answer is no, at least for the case where the variances of X and Y are known to be σ_x^2 and σ_y^2, and when the sample sizes m and n are identical. For this case, we naturally consider

$$T_0 = \sqrt{n}(\bar{X} - \theta\bar{Y})/(\sigma_x^2 + \theta^2\sigma_y^2)^{1/2} \tag{3.1}$$

and

$$T_0^* = \sqrt{n}(\bar{X}^* - \hat{\theta}\bar{Y}^*)/(\hat{\sigma}_x^2 + \hat{\theta}^2\hat{\sigma}_y^2)^{1/2}.$$

Note that σ_x^2 and σ_y^2 in T_0, even though known, are also replaced by their normal M.L.E. estimates $\hat{\sigma}_x^2$ and $\hat{\sigma}_y^2$ of σ_x^2 and σ_y^2 in T_0^*. The replacement leads to a more accurate interval, being second order correct instead of first order correct. This is due to the fact that the variance of T_0^* is 1, identical to that of T_0. It suffices to consider one-sided intervals. We focus on I-bootstrapping below. A similar theorem can be established for ISD-bootstrapping.

Theorem 3.1. (I-Bootstrapping) *Let X_i, $1 \le i \le n$, be i.i.d. random variables with mean μ_1 and $Y_i, 1 \le i \le n$, be i.i.d. random variables with mean μ_2. Let $U(X,Y)$ be the $(1-\alpha)$ quantile of T_0^*. Namely,*

$$P(T_0^* < U(X,Y)) = 1 - \alpha.$$

Then as $n \to \infty$

$$\sup_{\mu_1,\mu_2,\sigma_x,\sigma_y} |P(T_0 < U(X,Y)) - (1-\alpha)| \to 0,$$

provided that for some finite numbers c_1 and c_2,

$$\sup_{\mu_1,\mu_2,\sigma_x,\sigma_y} E|\frac{X_1-\mu_1}{\sigma_x}|^3 + E|\frac{Y_1-\mu_2}{\sigma_y}|^3 < c_1, \tag{3.2}$$

and as $n \to \infty$

$$\sup_{\mu_1,\mu_2,\sigma_x^2\sigma_y^2} E\hat{m}_3 \to c_2, \tag{3.3}$$

where

$$\hat{m}_3 = \frac{\frac{1}{n}\sum|X_i - \bar{X}|^3}{\hat{\sigma}_x^3} + \frac{\frac{1}{n}\sum|Y_i - \bar{Y}|^3}{\hat{\sigma}_y^3}.$$

Proof. Two applications of Berry-Esseen's Theorem. For details, see Hwang (1989).

We remark here that (3.2) is satisfied for any location-scale family with finite third moment. In such a case, the expectation of $\hat{m}_3$ is finite due to the inequality

$$\sum |X_i - \bar{X}|^3/(n\hat{\sigma}_x^3) < \sqrt{n}.$$

Also for any location-scale family (3.3) is equivalent to

$$E_{\sigma_x^2=\sigma_y^2=1,\mu_1=\mu_2=0}\Big[\frac{\frac{1}{n}\sum|X_i-\bar{X}|^3}{\hat{\sigma}_x^3}+\frac{\frac{1}{n}\sum|Y_i-\bar{Y}|^3}{\hat{\sigma}_y^3}\Big] \to c_2 < \infty.$$

This is satisfied for random variables X and Y such that the random variable inside the expectation of the last expression is uniformly integrable.

The last theorem provides asymptotic justification for Bootstrapping T_0 intervals. Now we turn to finite-sample numerical results.

Our simulation studies show that the ISD-bootstrap interval is quite similar to the Fieller's solution for the normal case. I-bootstrap and NP-bootstrap intervals are qualitatively similar to Fieller's intervals (or equivalently normal bootstrap intervals). Namely, they are respectively $(-\infty,\infty)$, complements of bounded intervals, and bounded intervals with similar probabilities (see Table 4). Quantitatively, the differences (in terms of absolute differences) are generally small. For more information see Hwang (1989).

We calculate the coverage probabilities in Tables 1–3 based on simulating 3000 replicates, each consisting of two independent 15 dimensional normal random vectors X and Y as described in Example 1.1. Table 1 examines the known variance case. The coverage probabilities of the 90% BC interval can be low (69.4%) for $(\mu_1/\text{S.D.}, \mu_2/\text{S.D.}) = (1,1)$ and near zero for $\mu_2/\text{S.D.} = 10^{-4}$ which agrees with Theorem 2.1. Here S.D. denotes the common standard deviation of $\bar{X}$ and $\bar{Y}$. We consider T_0 which is the same as (1.1) except that $\hat{\sigma}$ is replaced by the known standard deviation and

$$T_0^* = (\bar{X}^* - \hat{\theta}\bar{Y}^*)/(1/n + \hat{\theta}^2/n)^{1/2}\hat{\sigma}$$

with $\hat{\sigma}$ being defined in (1.2). As shown in Table 1, the coverage probabilities of the I-bootstrap T_0 interval are very close to the target probability .9. Similar conclusion extends to the unknown variance case in Table 2 with T_0 and T_0^* being defined in and after (3.5) and the unknown unequal variance case in Table 3 with T_0 being defined as below:

$$(\bar{X} - \theta\bar{Y})/[\hat{\sigma}_x^2/m + \theta^2\hat{\sigma}_y^2/n]^{1/2}. \tag{3.4}$$

We did not calculate the coverage probabilities of the ISD-bootstrap T_0 interval but do expect the coverage probabilities to be even closer to the target level, since the ISD-bootstrap T_0 interval uses the correct information that X and Y are similarly distributed.

Although Fieller-type procedures have good coverage probabilities, an issue of concern is that they may be unbounded or even cover the whole space. How often will this happen? In Table 4, we report the probabilities. For Fieller's intervals, probabilities of being unbounded or $(-\infty, \infty)$ are calculated based on 3000 simulations while those related to I-bootstrap T_0 1000 simulations. The only exception is that the entries in the first row are based on exact numerical integration. These probabilities show that the bootstrap T_0 intervals behave very similarly to the Fieller's intervals.

Table 1. Coverage probabilities based on 3000 replicate and 1000 bootstrap samples each. Data generated as in Example 1.1 with $n = m = 15$.

$\left[\frac{\mu_1}{\text{S.D.}}, \frac{\mu_2}{\text{S.D.}}\right]$	$(4, 8)$	$(8, 4)$	$(1, 1)$	$(1, 10^{-4})$
Fieller's interval and normal bootstrap T_0	.9	.9	.9	.9
BC interval	.9	.908	.694	.03
I-Bootstrap T_0	.894	.900	.904	.909

Table 2. Coverage probabilities of the intervals constructed by bootstrapping (3.5) and by assuming no knowledge of σ^2. Data were generated as in Table 1.

$\left[\frac{\mu_1}{\text{S.D.}}, \frac{\mu_2}{\text{S.D.}}\right]$	$(4, 8)$	$(8, 4)$	$(1, 1)$	$(1, 10^{-4})$
Fieller's interval and normal bootstrap T_0	.9	.9	.9	.9
I-Bootstrap T_0	.897	.905	.904	.903

Table 3. Coverage probabilities of the intervals constructed by bootstrapping (3.4). The variances of X and Y are unknown and unequal. Data generated similar to Table 1.

$\left[\frac{\mu_1}{\text{S.D.}}, \frac{\mu_2}{\text{S.D.}}\right]$	$(4, 8)$	$(8, 4)$	$(1, 1)$	$(1, 10^{-4})$
Fieller's interval	.917	.905	.908	.889
I-Boot T_0	.915	.894	.902	.882

Table 4. Probabilities of being the whole space are calculated based on simulation. For each type of intervals, entries in the first row correspond to the scenario in Table 1, whereas those in the second row correspond to Table 2. The entries in brackets are the simulated probabilities of obtaining unbounded intervals.

$\left[\frac{\mu_1}{\text{S.D.}}, \frac{\mu_2}{\text{S.D.}}\right]$	$(4,8)$	$(8,4)$	$(1,1)$	$(1,10^{-4})$
Fieller's	.0[.0]	.0[.0093]	.449[.736]	.581[.9]
interval and normal bootstrap T_0	.0[.0]	.0[.014]	.467[.731]	.586[.9]
I-Boot T_0	.0[.0]	.0[.011]	.432[.748]	.562[.904]
	.0[.0]	.0[.015]	.459[.755]	.629[.918]

The probabilities could be high if the model is close to the singularity point, i.e., μ_2 is close to zero. As $\mu_2 \to 0$, the theoretical value of the probability of unboundedness of Fieller's intervals is $1 - \alpha = .9$. In fact, for any confidence interval with minimum coverage probability $1 - \alpha$, the probability of unboundedness is at least $1 - \alpha$ when $\mu_2 \to 0$. This can be proved similar to Theorem 2.1.

Even if $\mu_2/\text{S.D.}$ is only as large as 1, the probability of unboundedness is surpringly high around .7. However, the probability of being $(-\infty, \infty)$, the truly uninformative case, drops to about .46. For the other two less extreme cases, the probabilities are either zero or tiny and should cause no concern.

Due to the problem of unboundedness, one may argue that there is no advantage of the Fieller-type interval over the BC interval. However, the latter provides misleading answers for the cases where probabilities are low, since it would miss the true answer too often. A misleading answer seems to be worse than an uninformative one. Moreover, when Fieller's interval is $(-\infty, \infty)$ or unbounded, it is giving a proper warning that perhaps the model is too close to the singularity point to draw a precise conclusion. Therefore, to the author, intervals with good coverage probabilities are preferred.

From the above discussion, the Bootstrapping T_0 interval seems to have a reasonable confidence level. However, a more basic question is what is the advantage of the bootstrapping T_0 interval over the usual Fieller's interval? The following theorem gives a theoretical justification. It shows that bootstrapping T_0 produce intervals that are second order correct whereas Fieller's intervals are only first order correct when the distributions are not normal (or more precisely when the skewness is not zero). In the theorem below we assume that the numbers of X_i's and Y_i's are the same and $X_i - \theta$ and $Y_i - \theta$ are i.i.d. with unknown variance σ^2. The results could possibly be extended to the unequal sample size case, and also the case where $X_i - \theta$, $Y_i - \theta$ are not identically distributed by using an appropriate pivot. However, we focus on the simpler case $n = m$ for ease of presentation. Here the pivot (1.1) reduces to

$$T_0 = \sqrt{n}(\bar{X} - \theta\bar{Y})/[\hat{\sigma}(1+\theta^2)^{1/2}], \tag{3.5}$$

where

$$\hat{\sigma}^2 = [\sum(X_i - \bar{X})^2 + \sum(Y_i - \bar{Y})^2]/2n,$$

and

$$T_0^* = \sqrt{n}(\bar{X}^* - \hat{\theta}\bar{Y}^*)/[\hat{\sigma}^*(1+\hat{\theta}^2)^{1/2}].$$

In the following theorem, the resampling scheme used is the ISD-bootstrapping.

Theorem 3.2. *Let q_α^* be such that*

$$P_*(T_0^* \le q_\alpha^*) = 1 - \alpha.$$

Then q_α^* is second order correct for estimating q_α.

Proof. Straightforward calculations.

Theorem 3.2 is not surprising, given the general results of Hall (1986 and 1988), which also provide the definition of second order correctness. His theorems do not apply directly to our problem since in his case the denominator of the statistic does not depend on θ.

Note for the Fieller interval, the cutoff point is $z^{(\alpha)}$, the α-upper quantile of a standard normal distribution, and hence is only correct up to the first order when the skewness of the underlining distribution is nonzero. Therefore, in such a situation, bootstrapping T_0 is more accurate asymptotically.

Hwang (1989) has also performed a numerical study for the moderate sample size $n = 15$ which is not reported here. In the studies, we assume that the variances are unknown and not necessarily equal. Consequently, I-bootstrapping, not ISD-bootstrapping, is considered. In the simulation, the data X_i and Y_i are independently generated from exponential distribution with means μ_1 and μ_2, and common variances σ^2; however, there the statistic (3.4) was used and $\hat{\sigma}_x^2$ and $\hat{\sigma}_y^2$ are still the (normal) m.l.e. estimate of σ_x^2 and σ_y^2 respectively. Even though the normal theory cutoff point is completely off, the bootstrapping T_0 intervals work well. See Hwang (1989).

4. Generalization

The bootstrap T_0 technique presented in this paper can be easily generalized to more complicated situations.

For the situation of correlated matched data $(X_1, Y_1), (X_2, Y_2), \ldots, (X_n, Y_n)$, one may consider

$$T_0 = \frac{\bar{X} - \theta\bar{Y}}{\sqrt{D}},$$

where

$$D = \frac{1}{n-1}\sum_{i=1}^{n}(X_i - \theta Y_i - (\bar{X} - \theta\bar{Y}))^2 = \hat{\sigma}_X^2 - \theta\hat{\sigma}_{XY} + \theta^2\hat{\sigma}_Y^2;$$

and $\hat{\sigma}_X^2$ and $\hat{\sigma}_Y^2$ are respectively the sample variances of X and Y and $\hat{\sigma}_{XY}$ is the sample covariance of X and Y. Under the normality assumption even when X_i and Y_i are correlated, T_0 has an exact t-distribution. Therefore all the analytic results and simulation results concerning bootstrapping T_0 should remain valid.

In another scenario, Fieller (1954) discussed the problem of setting a confidence interval for the root θ of the equation

$$\beta_1 F_1(\theta) + \beta_2 F_2(\theta) + \cdots + \beta_p F_p(\theta) = 0.$$

Let $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_p)'$ and assume that it is the regression parameter of a linear model with homoscedastic errors. The least squares estimator is denoted by $\hat{\beta}$. Therefore $\text{cov}\hat{\beta} = \sigma^2\Sigma$, where Σ depends only on the design matrix and is known. The present work seems to indicate that one should bootstrap

$$T_0 = F(\theta)'\hat{\beta}/[\hat{\sigma}(F(\theta)'\Sigma F(\theta))^{1/2}],$$

where $F(\theta)' = (F_1(\theta), \ldots, F_p(\theta))$. The above expression is exactly a pivot under normality and hence normal bootstrapping T_0 leads to an interval of exactly the right coverage probability. The recommended T_0^* is $F(\hat{\theta})'\hat{\beta}^*/[\hat{\sigma}^*(F(\hat{\theta})'\Sigma F(\hat{\theta}))^{1/2}]$. For nonparametric bootstrapping one can use the simulated residual approach as in (5.17) of Efron (1982). This is a generalization of ISD nonparametric bootstrap which has been shown working well numerically for this type of problem in Section 3 in the simpler case.

Acknowledgement

This work is partially supported by NSF Grant DMS-9107842 and DMS-9307499. The author thanks Professors Bill Brown, Brad Efron, Peter Hall, Tom DiCiccio, and Joe Romano for valuable discussions.

References

Bliss, C. I. (1935). The comparison of dosage-mortality curve. Ann. Appl. Biology **22**, 134-167.

Chan, Y. M. and Srivastava, M. S. (1988). Roubusness of Fieller's theorem and comparison with bootstrap method. Technical Report No. 7, University of Toronto.

Chow, S. C. and Liu, J. P. (1992). Design and Analysis of Bioavailability and Bioequivalence Studies. Marcel Dekker, Inc., New York.

Efron, B. (1982). The Jackknife, the Bootstrap, and other Resampling Plans. BMS 38, SIAM-NSF.

Efron, B. (1985). Bootstrap confidence intervals for a class of parametric problems. Biometrika **72**, 45-58.

Fieller, E. C. (1954). Some problems in interval estimation. J. Roy. Statist. Soc. Ser.B **16**, 175-185.

Finney, D. J. (1978). Statistical Method in Biological Assay, 3rd edition. Macmillan, New York.

Gleser, L. J. and Hwang, J. T. (1987). The nonexistence of $100(1-\alpha)\%$ confidence sets of finite expected diameter in errors-in-variables and related models. Ann. Statist. **15**, 1351-1362.

Hall, P. (1986). On the bootstrap and confidence intervals. Ann. Statist. **14**, 1431-1452.

Hall, P. (1988). Theoretical comparison of bootstrap confidence intervals. Ann. Statist. **16**, 927-953.

Hwang, J. T. (1989). Fieller's problems and resampling techniques. Cornell Technical Report.

Hinkley, D. and Wei, B.-C. (1984). Improvements of jackknife confidence limit methods. Biometrika **71**, 331-339.

Simonoff, J. S. and Tsai, C.-L. (1986). Discussion of C. F. J. Wu's (1986) paper. Ann. Statist. **14**, 1326-1328.

Wu, C. F. J. (1986). Jackknife, bootstrap and other resampling methods in regression analysis. Ann. Statist. **14**, 1261-1295.

Department of Mathematics, Cornell University, Ithaca, NY 14853, U.S.A.

(Received September 1991; accepted August 1994)