

Statistics 210B, Spring 1998

Class Notes

P.B. Stark

`stark@stat.berkeley.edu`

`www.stat.berkeley.edu/~stark/index.html`

Revised 25 February 2003

Eleventh Set of Notes

1 Resampling Methods

For references, see Beran, R., 1995. Stein confidence sets and the bootstrap, *Stat. Sinica*, 5, 109–127; Beran, R., 1990. Calibrating prediction regions, *J. Amer. Stat. Assoc.*, 85, 715–723; Beran, R., 1990. Refining bootstrap simultaneous confidence sets, *J. Amer. Stat. Assoc.*, 85, 417–426; Beran, R., 1987. Prepivoting to reduce level error of confidence sets, *Biometrika*, 74, 457–468; Feller, W., 1971. *An introduction to probability theory and its applications*, V. II, 2nd edition, John Wiley and Sons, Inc., New York; Lehmann, E.L., 1975. *Nonparametrics*, Holden-Day, Oakland; Efron, B., 1982. *The Jackknife, the bootstrap, and other resampling plans*, SIAM, Philadelphia; Romano, J.P., 1988. A bootstrap revival of some nonparametric distance tests, *J. Amer. Stat. Assoc.*, 83, 698–708; Romano, J.P., 1989. Bootstrap and randomization tests of some nonparametric hypotheses, *Ann. Stat.*, 17, 141–159.

1.1 Permutation Tests

Suppose we wish to test the efficacy of a purported cold remedy. We select $n = n_c + n_t$ individuals for an experiment. From these n , we randomly select n_t to receive the remedy,

and n_c to receive a placebo. The selection is done in such a way that every subset of n_t of the n subjects is equally likely to be the treatment group. There are $\binom{n}{n_t}$ such subsets. We infect everyone with the cold virus. Once they show symptoms, we give the remedy or placebo. Assume that everyone eventually recovers from a cold: we record the time interval from the onset of symptoms to the complete remission of symptoms. Let $\{X_j\}_{j=1}^{n_c}$ be the times for the control individuals, and let $\{Y_j\}_{j=1}^{n_t}$ be the times for the treated individuals. Let $\bar{X} = \frac{1}{n_c} \sum_j X_j$ and $\bar{Y} = \frac{1}{n_t} \sum_j Y_j$.

Let the null hypothesis be that the remedy does not work. If that be true, it is as if every individual's response were determined before the random assignment to treatment or control. There would be (at most)

$$\binom{n}{n_t} = \binom{n}{n_c} \tag{1}$$

possible pairs of values (\bar{X}, \bar{Y}) , and so at most that many values of the difference $\bar{X} - \bar{Y}$. We would want to reject the null hypothesis if \bar{X} is sufficiently larger than \bar{Y} (if the time-to-health with the treatment is shorter than the time-to-health with the placebo).

Because, under the null hypothesis, the assignment to treatment is random but individuals' responses are not, we can obtain a critical value for testing $\{H_0 : \text{treatment has no effect}\}$ at level α by finding a value c such that $\bar{X} - \bar{Y} > c$ for at most a fraction α of the $\binom{n}{n_t}$ partitions of the subjects into treatment and control. In principle, we could enumerate all of the $\binom{n}{n_t}$ partitions, calculate \bar{X} and \bar{Y} for each of them, and find the critical value exactly. (Incidentally, asymptotically in n , n_t and $n_c = n - n_t$ this test is equivalent to the 2-sample t -test.) In practice, finding the exact critical value is unwieldy even for fairly small values of n (but computers get faster all the time). For example, if we have 100 subjects of whom 50 are to be assigned to treatment, there are about $1.01e29$ possible assignments. Each assignment requires on the order of 100 floating point operations to compute $\bar{X} - \bar{Y}$. Using a computer that can calculate at the rate 1GFlop/s (10^9 floating point operations per second), it would take about $3.2e14$ years, to compute all the values of $\bar{X} - \bar{Y}$. The universe is about $1.5e10$ years old. (For 50 subjects with 25 assigned to treatment, the figure would be about 0.4 years.)

However, we might approximate a critical value for the test by randomly partitioning the n subjects into groups of size n_t and n_c , computing $\bar{X} - \bar{Y}$ for them, and using the $1 - \alpha$ quantile of the empirical distribution of those values among the random partitions. Such an approach *resamples* the data: we repeatedly draw samples of size n_t without replacement, compute the average response \bar{Y} of those n_t subjects, the average response \bar{X} of the $n_c = n - n_t$ not in that sample, and record $\bar{X} - \bar{Y}$.

1.2 The Bootstrap

The setting for the next few lectures is that we observe an iid sample of size n , $\{X_j\}_{j=1}^n$ iid F . Each observation is real-valued. We wish to estimate some parameter of the distribution of F that can be written as a functional of F , $T(F)$. Examples include the mean, $T(F) = \int x dF(x)$, other moments, *etc.*

The (unpenalized) nonparametric maximum likelihood estimator of F from the data $\{X_j\}$ is just the empirical distribution \hat{F}_n , which assigns mass $1/n$ to each observation:

$$\arg \max_{\text{probability distributions } G} \mathbf{P}_G\{X_j = x_j, j = 1, \dots, n\} = \hat{F}_n. \quad (2)$$

(Note, however, that the MLE of F is not generally consistent in problems with an infinite number of parameters, such as estimating a density or a distribution function.)

Using the general rule that the maximum likelihood estimator of a function of a parameter is that function of the maximum likelihood estimator of the parameter, we might be led to consider $T(\hat{F}_n)$ as an estimator of $T(F)$.

That is exactly what the sample mean does, as an estimator of the mean:

$$T(\hat{F}_n) = \int x d\hat{F}_n(x) = \sum_{j=1}^n \frac{1}{n} X_j = \frac{1}{n} \sum_j X_j. \quad (3)$$

Similarly, the maximum likelihood estimator of

$$\mathbf{Var}(X) = T(F) = \int \left(x - \int x dF \right)^2 dF \quad (4)$$

is

$$T(\hat{F}_n) = \int \left(x - \int x d\hat{F} \right)^2 d\hat{F}_n = \frac{1}{n} \sum_j \left(X_j - \frac{1}{n} \sum_k X_k \right)^2. \quad (5)$$

In these cases, we get analytically tractable expressions for $T(\hat{F}_n)$. What is often more interesting is estimating a property of the distribution of $T(\hat{F}_n)$, for example the variance of

the estimator $T(\hat{F}_n)$. The bootstrap approach is to approximate the distribution of $T(\hat{F}_n)$ by the distribution of $T(\hat{F}_n^*)$. What this notation is intended to mean is that we approximate the distribution of an estimator applied to the empirical distribution \hat{F}_n of a random sample of size n from a distribution F by the distribution of that estimator applied to a random sample \hat{F}_n^* of size n from a particular realization \hat{F}_n of the empirical distribution of a sample of size n from F . In the case T is the mean, so $T(\hat{F}_n)$ is the sample mean, we could obtain analytically the variance of the distribution of $T(\hat{F}_n^*)$: Let $\{X_j^*\}_{j=1}^n$ be an iid sample of size n from \hat{F}_n . Then

$$\mathbf{Var}_{\hat{F}_n} \frac{1}{n} \sum_{j=1}^n X_j^* = \frac{1}{n^2} \sum_{j=1}^n (X_j - \bar{X})^2, \quad (6)$$

where $\{X_j\}$ are the “original” data and \bar{X} is their mean. Generally, we would not get an analytically tractable expression for the variance of an estimator under resampling from the empirical distribution, but we could still approximate the distribution of $T(\hat{F}_n)$ by generating size- n data sets that are iid \hat{F}_n (drawing samples of size n with replacement from $\{x_j\}_{j=1}^n$), and applying T to each.

The idea of the bootstrap is to approximate the distribution (under F) of an estimator $T(\hat{F}_n)$ by the distribution of the estimator under \hat{F}_n , and to approximate *that* distribution by using a computer to take pseudo-random samples of size n from \hat{F}_n .

This basic idea is quite flexible, and can be applied to a wide variety of testing and estimation problems, including finding confidence intervals and sets. It is related to some other “resampling” schemes in which one re-weights the data to form other distributions.

Exercise. Recall the normal linear regression model of the previous section. Consider resampling as follows: Let

$$\hat{\epsilon} = A\hat{\theta}_{LS} - X, \quad (7)$$

and let \hat{F}_n be the empirical cdf of the components of $\hat{\epsilon}$. Consider “pseudo-data” distributed as

$$X^* = A\hat{\theta}_{LS} + \epsilon^*, \quad (8)$$

where the components of ϵ^* are iid \hat{F}_n . Consider applying least squares to the pseudo-data, and calculating the variance of the components of θ . (If we were using Monte-Carlo, we would generate many independent sets of pseudo-data, apply least-squares to each, and get an empirical joint distribution for $\hat{\theta}_{LS}$ under $\epsilon^* \sim \hat{F}_n$.) Give a formula for the resulting estimate of the variance of θ_j .

Before doing more theory with the bootstrap, let's examine the jackknife.

1.3 The Jackknife

The idea behind the jackknife, which is originally due to Tukey and Quenouille, is to form from the data $\{X_j\}_{j=1}^n$, n sets of $n - 1$ data, leaving each datum out in turn. The “distribution” of T applied to these n sets is used to approximate the distribution of $T\hat{F}_n$. Let $\hat{F}_{(i)}$ denote the empirical distribution of the data set with the i th value deleted; $T_{(i)} = T(\hat{F}_{(i)})$ is the corresponding estimate of $T(F)$. An estimate of the expected value of $T(\hat{F}_n)$ is

$$\hat{T}_{(\cdot)} = \frac{1}{n} \sum_{i=1}^n T(\hat{F}_{(i)}). \quad (9)$$

Consider the bias of $T(\hat{F}_n)$:

$$E_F T(\hat{F}_n) - T(F). \quad (10)$$

Quenouille's jackknife estimate of the bias is

$$\widehat{\text{BIAS}} = (n - 1)(\hat{T}_{(\cdot)} - T(\hat{F}_n)). \quad (11)$$

It can be shown that if the bias of T has a homogeneous polynomial expansion in n^{-1} whose coefficients do not depend on n , then the bias of the bias-corrected estimate

$$\tilde{T} = nT(\hat{F}_n) - (n - 1)\hat{T}_{(\cdot)} \quad (12)$$

is $O(n^{-2})$ instead of $O(n^{-1})$.

Applying the jackknife estimate of bias to correct the plug-in estimate of variance reproduces the formula for the sample variance (with $1/(n - 1)$) from the formula with $1/n$:

Define

$$\bar{X} = \frac{1}{n} \sum_{j=1}^n X_j, \quad (13)$$

$$\bar{X}_{(i)} = \frac{1}{n - 1} \sum_{j \neq i} X_j, \quad (14)$$

$$T(\hat{F}_n) = \hat{\sigma}^2 = \frac{1}{n} \sum_{j=1}^n (X_j - \bar{X})^2, \quad (15)$$

$$T(\hat{F}_{(i)}) = \frac{1}{n - 1} \sum_{j \neq i} (X_j - \bar{X}_{(i)})^2, \quad (16)$$

$$T(\hat{F}_{(\cdot)}) = \frac{1}{n} \sum_{i=1}^n T(\hat{F}_{(i)}). \quad (17)$$

Now

$$\bar{X}_{(i)} = \frac{n\bar{X} - X_i}{n-1} = \bar{X} + \frac{1}{n-1}(\bar{X} - X_i), \quad (18)$$

so

$$\begin{aligned} (X_j - \bar{X}_{(i)})^2 &= \left(X_j - \bar{X} + \frac{1}{n-1}(\bar{X} - X_i) \right)^2 \\ &= (X_j - \bar{X})^2 + \frac{2}{n-1}(X_j - \bar{X})(X_i - \bar{X}) + \frac{1}{(n-1)^2}(X_i - \bar{X})^2. \end{aligned} \quad (19)$$

Note also that

$$\sum_{j \neq i} (X_j - \bar{X}_{(i)})^2 = \sum_{j=1}^n (X_j - \bar{X}_{(i)})^2 - (X_i - \bar{X}_{(i)})^2. \quad (20)$$

Thus

$$\begin{aligned} \sum_{i=1}^n \sum_{j \neq i} (X_j - \bar{X}_{(i)})^2 &= \frac{1}{n-1} \sum_{i=1}^n \left[\sum_{j=1}^n \left[(X_j - \bar{X})^2 + \frac{2}{n-1}(X_j - \bar{X})(X_i - \bar{X}) + \right. \right. \\ &\quad \left. \left. + \frac{1}{(n-1)^2}(X_i - \bar{X})^2 \right] - (X_i - \bar{X})^2 - \right. \\ &\quad \left. - \frac{2}{n-1}(X_i - \bar{X})^2 - \frac{1}{(n-1)^2}(X_i - \bar{X})^2 \right]. \end{aligned} \quad (21)$$

The last three terms all are multiples of $(X_i - \bar{X})^2$; the sum of the coefficients is

$$1 + 2/(n-1) + 1/(n-1)^2 = n^2/(n-1)^2. \quad (22)$$

The middle term of the inner sum is a constant times $(X_j - \bar{X})$, which sums to zero over j .

Simplifying the previous displayed equation yields

$$\begin{aligned} \sum_{i=1}^n \sum_{j \neq i} (X_j - \bar{X}_{(i)})^2 &= \frac{1}{n-1} \sum_{i=1}^n \left(n\hat{\sigma}^2 + \frac{n}{(n-1)^2}(X_i - \bar{X})^2 - \frac{n^2}{(n-1)^2}(X_i - \bar{X})^2 \right) \\ &= \frac{1}{n-1} \sum_{i=1}^n \left(n\hat{\sigma}^2 - \frac{n}{n-1}(X_i - \bar{X})^2 \right) \\ &= \frac{1}{n-1} \left[n^2\hat{\sigma}^2 - \frac{n^2}{n-1}\hat{\sigma}^2 \right] \\ &= \frac{n(n-2)}{(n-1)^2}\hat{\sigma}^2. \end{aligned} \quad (23)$$

The jackknife bias estimate is thus

$$\widehat{\text{BIAS}} = (n-1) \left(T(\hat{F}_{(\cdot)}) - T(\hat{F}_n) \right) = \hat{\sigma}^2 \frac{n(n-2) - (n-1)^2}{n-1} = \frac{-\hat{\sigma}^2}{n-1}. \quad (24)$$

The bias-corrected MLE variance estimate is therefore

$$\hat{\sigma}^2 \left(1 - \frac{1}{n-1} \right) = \hat{\sigma}^2 \frac{n}{n-1} = \frac{1}{n-1} \sum_{j=1}^n (X_j - \bar{X})^2 = S^2, \quad (25)$$

the usual sample variance.

The jackknife also can be used to estimate other properties of an estimator, such as its variance. The jackknife estimate of the variance of $T(\hat{F}_n)$ is

$$\widehat{\mathbf{Var}}(T) = \frac{n-1}{n} \sum_{j=1}^n (T_{(j)} - T_{(\cdot)})^2. \quad (26)$$

It is convenient to think of distributions on data sets to compare the jackknife and the bootstrap. We shall follow the notation in Efron (1982). We condition on $(X_i = x_i)$ and treat the data as fixed in what follows. Let \mathcal{S}_n be the n -dimensional simplex

$$\mathcal{S}_n \equiv \{\mathbf{P}^* = (P_i^*)_{i=1}^n \in \mathbf{R}^n : P_i^* \geq 0 \text{ and } \sum_{i=1}^n P_i^* = 1\}. \quad (27)$$

A *resampling vector* $\mathbf{P}^* = (P_k^*)_{k=1}^n$ is any element of \mathcal{S}_n ; *i.e.*, an n -dimensional discrete probability vector. To each $\mathbf{P}^* = (P_k^*) \in \mathcal{S}_n$ there corresponds a re-weighted empirical measure $\hat{F}^*(\mathbf{P}^*)$ which puts mass P_k^* on x_k , and a value of the estimator $T^* = T(\hat{F}(\mathbf{P}^*)) = T(\mathbf{P}^*)$. The resampling vector $\mathbf{P}^0 = (1/n)_{j=1}^n$ corresponds to the empirical distribution \hat{F}_n (each datum x_j has the same mass). The resampling vector

$$\mathbf{P}_{(i)} = \frac{1}{n-1} (1, 1, \dots, 0, 1, \dots, 1), \quad (28)$$

which has the zero in the i th place, is one of the n resampling vectors the jackknife visits; denote the corresponding value of the estimator T by $T_{(i)}$. The bootstrap visits all resampling vectors whose components are multiples of $1/n$.

The bootstrap estimate of variance tends to be better than the jackknife estimate of variance for nonlinear estimators because of the distance between the empirical measure and the resampled measures:

$$\|\mathbf{P}^* - \mathbf{P}^0\| = O_P(n^{-1/2}), \quad (29)$$

while

$$\|\mathbf{P}_{(k)} - \mathbf{P}^0\| = O(n^{-1}). \quad (30)$$

To see the former, recall that the difference between the empirical distribution and the true distribution is $O_P(n^{-1/2})$: For any two probability distributions $\mathbf{P}_1, \mathbf{P}_2$, on \mathbf{R} , define the Kolmogorov-Smirnov distance

$$d_{KS}(\mathbf{P}_1, \mathbf{P}_2) \equiv \|\mathbf{P}_1 - \mathbf{P}_2\|_{KS} \equiv \sup_{x \in \mathbf{R}} |\mathbf{P}_1\{(-\infty, x]\} - \mathbf{P}_2\{(-\infty, x]\}|. \quad (31)$$

There exist universal constants χ_α so that for every continuous (w.r.t. Lebesgue measure) distribution F ,

$$\mathbf{P}_F \left\{ \|F - \hat{F}_n\|_{KS} \geq \chi_n(\alpha) \right\} = \alpha. \quad (32)$$

This is the Dvoretzky-Kiefer-Wolfowitz inequality. Massart (*Ann. Prob.*, 18, 1269–1283, 1990; stable URL <http://links.jstor.org/sici?sici=0091-1798%28199007%2918%3A3%3C1269%3ATTCTD%3E> Q) showed that the constant

$$\chi_n(\alpha) \leq \sqrt{\frac{\ln \frac{2}{\alpha}}{2n}} \quad (33)$$

is *tight*. Thinking of the bootstrap distribution (the empirical distribution \hat{F}_n) as the true cdf and the resamples from it as the data gives the result that the distance between the cdf of the bootstrap resample and the empirical cdf of the original data is $O_P(n^{-1/2})$.

To see that the cdfs of the jackknife samples are $O(n^{-1})$ from the empirical cdf \hat{F}_n , note that for univariate real-valued data, the difference between \hat{F}_n and the cdf of the jackknife data set that leaves out the j th ranked observation $X_{(j)}$ is largest either at $X_{(j-1)}$ or at $X_{(j)}$. For $j = 1$ or $j = n$, the jackknife samples that omit the smallest or largest observation, the L_1 distance between the jackknife measure and the empirical distribution is exactly $1/n$. Consider the jackknife cdf $\hat{F}_{n,(j)}$, the cdf of the sample without $X_{(j)}$, $1 < j < n$.

$$\hat{F}_{n,(j)}(X_{(j)}) = (j-1)/(n-1), \quad (34)$$

while $\hat{F}_n(X_{(j)}) = j/n$; the difference is

$$\frac{j}{n} - \frac{j-1}{n-1} = \frac{j(n-1) - n(j-1)}{n(n-1)} = \frac{n-j}{n(n-1)} = \frac{1}{n-1} - \frac{j}{n(n-1)}. \quad (35)$$

On the other hand,

$$\hat{F}_{n,(j-1)}(X_{(j-1)}) = (j-1)/(n-1), \quad (36)$$

while $\hat{F}_n(X_{(j-1)}) = (j-1)/n$; the difference is

$$\frac{j-1}{n-1} - \frac{j-1}{n} = \frac{n(j-1) - (n-1)(j-1)}{n(n-1)} = \frac{j-1}{n(n-1)}. \quad (37)$$

Thus

$$\|\hat{F}_{n,(j)} - \hat{F}_n\| = \frac{1}{n(n-1)} \max\{n-j, j-1\}. \quad (38)$$

But $n/2 \leq \max\{n-j, j-1\} \leq n-1$, so

$$\|\hat{F}_{n,(j)} - \hat{F}_n\| = O(n^{-1}). \quad (39)$$

The neighborhood that the bootstrap samples is larger, and is probabilistically of the right size to correspond to the uncertainty of the empirical distribution function as an estimator of the underlying distribution function F (recall the Kiefer-Dvoretzky-Wolfowitz inequality—a K-S ball of radius $O(n^{-1/2})$ has fixed coverage probability). For linear functionals, this does not matter, but for strongly nonlinear functionals, the bootstrap estimate of the variability tends to be more accurate than the jackknife estimate of the variability.

Let us have a quick look at the distribution of the K-S distance between a continuous distribution and the empirical distribution of a sample $\{X_j\}_{j=1}^n$ iid F . The discussion follows *Feller* (1971, pp. 36ff). First we show that for continuous distributions F , the distribution of $\|\hat{F}_n - F\|_{KS}$ does not depend on F . To see this, note that $F(X_j) \sim U[0, 1]$: Let $x_t \equiv \inf\{x \in \mathbf{R} : F(x) = t\}$. Continuity of F ensures that x_t exists for all $t \in [0, 1]$. Now the event $\{X_j \leq x_t\}$ is equivalent to the event $\{F(X_j) \leq F(x_t)\}$ up to a set of F -measure zero. Thus

$$t = \mathbf{P}_F\{X_j \leq x_t\} = \mathbf{P}_F\{F(X_j) \leq F(x_t)\} = \mathbf{P}_F\{F(X_j) \leq t\}, \quad t \in [0, 1]. \quad (40)$$

Thus $\{F(X_j)\}_{j=1}^n$ are iid $U[0, 1]$. Let

$$\hat{G}_n(t) = \#\{F(X_j) \leq t\}/n = \#\{X_j \leq x_t\}/n = \hat{F}_n(x_t) \quad (41)$$

be the empirical cdf of the $F(X_j)$ s. Note that

$$\sup_{x \in \mathbf{R}} |\hat{F}_n(x) - F(x)| = \sup_{t \in [0, 1]} |\hat{F}_n(x_t) - F(x_t)| = \sup_{t \in [0, 1]} |\hat{G}_n(t) - t|. \quad (42)$$

The probability distribution of \hat{G}_n is that of the cdf of n iid $U[0, 1]$ random variables (it does not depend on F), so the distribution of the K-S distance between the empirical cdf and the true cdf is the same for every continuous distribution. It turns out that for distributions with atoms, the K-S distance between the empirical and the true distribution functions is stochastically smaller than it is for continuous distributions.

1.4 Bootstrap and Randomization Tests

This section is about Romano's papers. The set-up is as follows: We observe $\{X_j\}_{j=1}^n$ iid F , where F is a distribution on an abstract sample space \mathcal{S} . The distribution $F \in \Omega$, a collection of distributions on \mathcal{S} . The null hypothesis is that $F \in \Omega_0 \subset \Omega$. We assume that Ω_0 can be characterized as a set of distributions that are invariant under a transformation on Ω : let $\tau : \Omega \rightarrow \Omega_0$; we assume that $\tau(G) = G$ for all $G \in \Omega_0$.

Let \mathcal{V} be a collection of subsets of a set \mathcal{S} . For a finite set $D \subset \mathcal{S}$, let $\Delta^{\mathcal{V}}(D)$ be the number of distinct sets $\{V \cap D : V \in \mathcal{V}\}$. For positive integers n , let

$$m^{\mathcal{V}}(n) = \max_{D \subset \mathcal{S}: \#D=n} \Delta^{\mathcal{V}}(D). \quad (43)$$

Let

$$c(\mathcal{V}) \equiv \inf\{n : m^{\mathcal{V}}(n) < 2^n\}. \quad (44)$$

If $c(\mathcal{V}) < \infty$, \mathcal{V} is a Vapnik-Cervonenkis (V-C) class. That is, \mathcal{V} is a V-C class if the maximum number of distinct intersections of sets in \mathcal{V} with sets containing n points grows sub-exponentially with n . Intersections, finite unions, and Cartesian products of V-C classes are V-C classes. In \mathbf{R}^n , the set of all ellipsoids, the set of all half-spaces, the set of all lower-left quadrants, and the set of all convex sets with at most p extreme points are all V-C classes.

An alternative, equivalent definition of a V-C class is based on the following

Definition 1 *Suppose \mathcal{V} is a collection of subsets of a set \mathcal{X} , and that D is a finite subset of \mathcal{X} . We say D is shattered by \mathcal{V} if every subset $d \subset D$ can be written $d = V \cap D$ for some $V \in \mathcal{V}$.*

Because there are 2^n subsets of a set with n elements, this is equivalent to saying that there are 2^n different subsets of the form $D \cap V$ as V ranges over \mathcal{V} .

A collection \mathcal{V} is a V-C class if for some finite integer n , there exists a set $D \subset \mathcal{X}$ with n elements that is not shattered by \mathcal{V} .

Exercise. Show that intersections and finite unions of V-C classes are V-C classes. Give a (counter) example showing that a countable union of V-C classes need not be a V-C class.

Let \mathcal{V} be a VC class of subsets of \mathcal{S} . Define the pseudo-metric

$$\begin{aligned} \delta : \Omega \times \Omega &\rightarrow \mathbf{R}^+ \\ (G, H) &\rightarrow \sup_{V \in \mathcal{V}} |G(V) - H(V)|. \end{aligned} \quad (45)$$

This is a generalization of the Kolmogorov-Smirnov distance for distributions on the line. In that case, the sets in \mathcal{V} are the half-lines $(-\infty, y]$ (which comprise a V-C class). Consider a set $D = \{x_j\}_{j=1}^n$ of points on the real line. How many sets are there of the form $V \cap D$, for V a half-line $(-\infty, y]$? Just $n + 1$. Suppose the points are in increasing order. Then the possibilities for $V \cap D$ are $\{\}, \{x_1\}, \{x_1, x_2\}, \dots, \{x_j\}_{j=1}^n$. Thus $m^{\mathcal{V}}(n) = n + 1$, and

$c(\mathcal{V}) \equiv \inf\{n : m^{\mathcal{V}}(n) < 2^n\} = 2$ (for $n = 0$, we have $0 + 1 = 2^0$, and for $n = 1$, we have $1 + 1 = 2^1$, but for $n = 2$, we have $2 + 1 < 2^2$).

Assume that \mathcal{V} has been selected such that $\delta(G, \tau G) = 0$ iff $G \in \Omega_0$. Romano proposes using the test statistic

$$T_n = n^{1/2} \delta(\hat{P}_n, \tau \hat{P}_n), \quad (46)$$

where \hat{P}_n is the empirical measure of $\{X_j\}_{j=1}^n$. One rejects the hypothesis when $\tau \hat{P}_n$ is far from \hat{P}_n .

One way to obtain a critical value from the test is by the bootstrap; one resamples from \hat{P}_n , and looks at the distribution of the distance between the empirical distribution of the bootstrap samples and τ applied to them.

Suppose that there is a (known) group \mathcal{G}_n of transformations of the sample space \mathcal{S}_n such that under the null hypothesis, F is invariant under \mathcal{G}_n . Then we can also construct a *randomization test* of the hypothesis H_0 . For simplicity, suppose that \mathcal{G}_n is finite, with M_n elements $\{g_{nj}\}_{j=1}^{M_n}$. Under the null hypothesis, conditional on $X = x$, the values $\{g_{nj}x\}_{j=1}^{M_n}$ are equally likely. The *orbit* of an point x in a space S acted on by a group \mathcal{G} is the set of all elements of S that can be obtained by applying elements of \mathcal{G} to x . That is, it is the set $\{g(x) : g \in \mathcal{G}\}$. For example, consider points in the plane and the group of rotations about the origin. Then the orbit of a point x is the circle with radius $\|x\|$. Compute the test statistic for each $g_{nj}x$ in the orbit of x . Reject the null hypothesis if the statistic for x exceeds the α quantile of the test statistic for the set of values obtained from the orbit; do not reject if it is less; reject with a given probability if the statistic equals the α quantile in such a way as to get a level α test. This is a randomization test. Because the level of the randomization test is α , conditional on the data, integrating over the distribution of the data shows that it is α unconditionally.

1.4.1 Examples of hypotheses and functions τ

. Examples Romano gives include testing for independence of the components of each X_j , testing for exchangeability of the components of each X_j , testing for spherical symmetry of the distribution of X_j , testing for homogeneity among the X_j , and testing for a change point.

In the example of testing for independence, the mapping τ takes the marginal distributions of the joint distribution, then constructs a joint distribution that is the product of

the marginals. For distributions with independent components, this is the identity; otherwise, it maps a distribution into one with the same marginals, but whose components are independent. For testing for spherical symmetry, τ maps a distribution into one with the same mass at every distance from the origin, but that is uniform on spherical shells. For testing for exchangeability, Romano proposes looking at the largest difference between F and a permutation of the coordinates of F , over all permutations of the coordinates. See his excellent and lucid paper for more details.

Romano shows that these tests are consistent against all alternatives, and that the critical values given by the bootstrap and by randomization are asymptotically equal with probability one. Because the randomization tests are exact level α tests, they might be preferred. Romano also discusses how to implement the tests computationally.

1.5 Bootstrap Confidence Sets

Let \mathcal{U} be an index set (not necessarily countable). Recall that a collection $\{\mathcal{I}_u\}_{u \in \mathcal{U}}$ of confidence intervals for parameters $\{\theta_u\}_{u \in \mathcal{U}}$ has simultaneous $1 - \alpha$ coverage probability if

$$\mathbf{P}_\theta \{ \bigcap_{u \in \mathcal{U}} \{ \mathcal{I}_u \ni \theta_u \} \} \geq 1 - \alpha. \quad (47)$$

If $\mathbf{P}\{\mathcal{I}_u \ni \theta_u\}$ does not depend on u , the confidence intervals are said to be *balanced*.

Many of the procedures for forming joint confidence sets we have seen depend on *pivots*, which are functions of the data and the parameter(s) whose distribution is known (even though the parameter and the parent distribution are not). For example, the Scheffé method relies on the fact that the sum of squared differences between the data and the corresponding parameters, divided by the variance estimate, has an F distribution, regardless of the parameter values. Similarly, Tukey's maximum modulus method relies on the fact that the distribution of the maximum of the studentized absolute differences between the data and the corresponding parameters does not depend on the parameters. Both of those examples are parametric, but the idea is more general: the procedure we looked at for finding bounds on the density function subject to shape restrictions just relied on the fact that there are uniform bounds on the probability that the K-S distance between the empirical distribution and the true distribution exceeds some threshold.

Even in cases where there is no known exact pivot, one can sometimes show that some function of the data and parameters is asymptotically a pivot. Working out the distribu-

tions of the functions involved is not typically straightforward, and a general method of constructing (possibly simultaneous) confidence sets would be nice.

Efron gives several methods of basing confidence sets on the bootstrap. Those methods are substantially improved (in theory, and in my experience) by Beran's pre-pivoting approach, which leads to iterating the bootstrap.

Let X_n denote a sample of size n from F . Let $R_n(\theta) = R_n(X_n, \theta)$ have cdf H_n , and let $H_n^{-1}(\alpha)$ be the largest α quantile of the distribution of R_n . Then

$$\{\gamma \in \Theta : R_n(\gamma) \leq H_n^{-1}(1 - \alpha)\} \tag{48}$$

is a $1 - \alpha$ confidence set for θ .

1.5.1 The Percentile Method

The idea of the percentile method is to use the empirical bootstrap percentiles of some quantity to approximate the true percentiles. Consider constructing a confidence interval for a single real parameter $\theta = T(F)$. We will estimate θ by $\hat{\theta} = T(\hat{F}_n)$. We would like to know the distribution function $H_n = H_n(\cdot, F)$ of $D_n(\theta) = T(\hat{F}_n) - \theta$. Suppose we did. Let $H_n^{-1}(\cdot) = H_n^{-1}(\cdot, F)$ be the inverse cdf of D_n . Then

$$\mathbf{P}_F\{H_n^{-1}(\alpha/2) \leq T(\hat{F}_n) - \theta \leq H_n^{-1}(1 - \alpha/2)\} = 1 - \alpha, \tag{49}$$

so

$$\mathbf{P}_F\{\theta \leq T(\hat{F}_n) - H_n^{-1}(\alpha/2) \text{ and } \theta \geq T(\hat{F}_n) - H_n^{-1}(1 - \alpha/2)\} = 1 - \alpha, \tag{50}$$

or, equivalently,

$$\mathbf{P}_F\{[T(\hat{F}_n) - H_n^{-1}(1 - \alpha/2), T(\hat{F}_n) - H_n^{-1}(\alpha/2)] \ni \theta\} = 1 - \alpha, \tag{51}$$

so the interval $[T(\hat{F}_n) - H_n^{-1}(1 - \alpha/2), T(\hat{F}_n) - H_n^{-1}(\alpha/2)]$ would be a $1 - \alpha$ confidence interval for θ .

The idea behind the percentile method is to approximate $H_n(\cdot, F)$ by $\hat{H}_n = H_n(\cdot, \hat{F}_n)$, the distribution of D_n under resampling from \hat{F}_n rather than F . An alternative approach would be to take $D_n(\theta) = |T(\hat{F}_n) - \theta|$; then

$$\mathbf{P}_F\{|T(\hat{F}_n) - \theta| \leq H_n^{-1}(1 - \alpha)\} = 1 - \alpha, \tag{52}$$

so

$$\mathbf{P}_F\{[T(\hat{F}_n - H_n^{-1}(1 - \alpha)), T(\hat{F}_n + H_n^{-1}(1 - \alpha))] \ni \theta\} = 1 - \alpha. \quad (53)$$

In either case, the “raw” bootstrap approach is to approximate H_n by resampling under \hat{F}_n .

Beran proves a variety of results under the following condition:

Condition 1. (Beran, 1987) For any sequence $\{F_n\}$ that converges to F in a metric d on cdfs, $H_n(\cdot, F_n)$ converges weakly to a continuous cdf $H = H(\cdot, F)$ that depends only on F , and not the sequence $\{F_n\}$.

Suppose Condition 1 holds. Then because \hat{F}_n is consistent for F , the estimate \hat{H}_n converges in probability to H in sup norm; moreover, the distribution of $\hat{H}_n(R_n(\theta))$ converges to $U[0, 1]$.

Instead of D_n , consider $R_n(\theta) = |T(\hat{F}_n) - \theta|$ or some other (approximate) pivot. Let $\hat{H}_n(\cdot, \hat{F}_n)$ be the bootstrap estimate of the cdf of R_n ; The set

$$\begin{aligned} B_n &= \{\gamma \in \Theta : \hat{H}_n(R_n(\gamma)) \leq 1 - \alpha\} \\ &= \{\gamma \in \Theta : R_n(\gamma) \leq \hat{H}_n^{-1}(1 - \alpha)\} \end{aligned} \quad (54)$$

is (asymptotically) a $1 - \alpha$ confidence set for θ .

The level of this set for finite samples tends to be inaccurate. It can be improved in the following way, due to Beran.

The original root, $R_n(\theta)$, whose limiting distribution depends on F , was transformed into a new root $R_{n,1}(\theta) = \hat{H}_n(R_n(\theta))$, whose limiting distribution is $U[0, 1]$. The distribution of $R_{n,1}$ depends less strongly on F than does that of R_n ; Beran calls mapping R_n into $R_{n,1}$ *prepivoting*. The confidence set 54 acts as if the distribution of $R_{n,1}$ really is uniform, which is not generally true. One could instead treat $R_{n,1}$ itself as a root, and pivot to reduce the dependence on F .

Let $H_{n,1} = H_{n,1}(\cdot, F)$ be the cdf of the new root $R_{n,1}(\theta)$, estimate $H_{n,1}$ by $\hat{H}_{n,1} = H_{n,1}(\cdot, \hat{F}_n)$, and define

$$\begin{aligned} B_{n,1} &= \{\gamma \in \Theta : \hat{H}_{n,1}(R_{n,1}(\gamma)) \leq 1 - \alpha\} \\ &= \{\gamma \in \Theta : \hat{H}_{n,1}(\hat{H}_n(R_n(\gamma))) \leq 1 - \alpha\} \\ &= \{\gamma \in \Theta : R_n(\gamma) \leq \hat{H}_n^{-1}(\hat{H}_{n,1}^{-1}(1 - \alpha))\}. \end{aligned} \quad (55)$$

Beran shows that this confidence set tends to have smaller error in its level than does B_n .

The transformation can be iterated further, typically resulting in additional reductions in the level error.

1.6 Approximating $B_{n,1}$ by Monte Carlo

I'll follow Beran's (1987) notation (mostly).

Let x_n denote the “real” sample of size n . Let x_n^* be a bootstrap sample of size n drawn from the empirical cdf \hat{F}_n . The components of x_n^* are conditionally iid given x_n . Let \hat{F}_n^* denote the “empirical” cdf of the bootstrap sample x_n^* . Let x_n^{**} denote a sample of size n drawn from \hat{F}_n^* ; the components of x_n^{**} are conditionally iid given x_n and x_n^* . Let $\hat{\theta}_n = T(\hat{F}_n)$, and $\hat{\theta}_n^* = T(\hat{F}_n^*)$. Then

$$H_n(s, F) = \mathbf{P}_F\{R_n(x_n, \theta) < s\}, \quad (56)$$

and

$$H_{n,1}(s, F) = \mathbf{P}_F\left\{\mathbf{P}_{\hat{F}_n}\{R_n(x_n^*, \hat{\theta}_n) < R_n(x_n, \theta)\} < s\right\}. \quad (57)$$

The bootstrap estimates of these cdfs are

$$\hat{H}_n(s) = H_n(s, \hat{F}_n) = \mathbf{P}_{\hat{F}_n}\{R_n(x_n^*, \hat{\theta}_n) < s\}, \quad (58)$$

and

$$\hat{H}_{n,1}(s) = H_{n,1}(s, \hat{F}_n) = \mathbf{P}_{\hat{F}_n}\left\{\mathbf{P}_{\hat{F}_n^*}\{R_n(x_n^{**}, \hat{\theta}_n^*) < R_n(x_n^*, \hat{\theta}_n)\}\right\}. \quad (59)$$

The Monte Carlo approach is as follows:

1. Draw $\{y_k^*\}_{k=1}^M$ bootstrap samples of size n from \hat{F}_n . The ecdf of $\{R_n(y_k^*, \hat{\theta}_n)\}_{k=1}^M$ is an approximation to \hat{H}_n .
2. For $k = 1, \dots, M$, let $\{y_{k\ell}^{**}\}_{\ell=1}^N$ be N size n bootstrap samples from the ecdf of y_k^* . Let $\hat{\theta}_{n,k}^* = T(\hat{F}_{n,k}^*)$. Let Z_k be the fraction of the values

$$\{R_n(y_{k,\ell}^{**}, \hat{\theta}_{n,k}^*)\}_{\ell=1}^N \quad (60)$$

that are less than $R_n(y_k^*, \hat{\theta}_n)$. The ecdf of $\{Z_k\}$ is an approximation to $\hat{H}_{n,1}$ that improves (in probability) as M and N grow.

Note that this approach is extremely general. Beran gives examples for confidence sets for directions, *etc.* The pivot can in principle be a function of any number of parameters, which can yield simultaneous confidence sets for parameters of any dimension.

1.7 Other approaches to improving coverage probability

There are other ways of iterating the bootstrap to improve the level accuracy of bootstrap confidence sets. Efron suggests trying to attain a different coverage probability so that the coverage attained in the second generation samples is the nominal coverage probability. That is, if one wants a 95% confidence set, one tries different percentiles so that in resampling from the sample, the attained coverage probability is 95%. Typically, the percentile one uses in the second generation will be higher than 95%. Here is a sketch of the Monte-Carlo approach:

- Set a value of α^* (initially taking $\alpha^* = \alpha$ is reasonable)
- From the sample, draw M size- n samples that are each iid \hat{F}_n . Denote the ecdfs of the samples by $\{\hat{F}_{n,j}^*\}$.
- For each $j = 1, \dots, M$, apply the percentile method to make a (nominal) level $1 - \alpha^*$ confidence interval for $T(\hat{F}_n)$. This gives M confidence intervals; a fraction $1 - \alpha'$ will cover $T(\hat{F}_n)$. Typically, $1 - \alpha' \neq 1 - \alpha$.
- If $1 - \alpha' < 1 - \alpha$, decrease α^* and return to the previous step. If $1 - \alpha' > 1 - \alpha$, increase α^* and return to the previous step. If $1 - \alpha' \approx 1 - \alpha$ to the desired level of precision, go to the next step.
- Report as a $1 - \alpha$ confidence interval for $T(F)$ the (first generation) bootstrap quantile confidence interval that has nominal $1 - \alpha^*$ coverage probability.

An alternative approach to increasing coverage probability by iterating the bootstrap is to use the same root, but to use a quantile (among 2nd generation bootstrap samples) of its $1 - \alpha$ quantile rather than the quantile observed in the first generation. The heuristic justification is that we would ideally like to know the $1 - \alpha$ quantile of the pivot under sampling from the true distribution F . We don't. The percentile method estimates the $1 - \alpha$ quantile of the pivot under F by the $1 - \alpha$ quantile of the pivot under \hat{F}_n , but this is subject to sampling variability. To try to be conservative, we could use the bootstrap a second time find an (approximate) upper $1 - \alpha^*$ confidence interval for the $1 - \alpha$ quantile of the pivot.

Here is a sketch of the Monte-Carlo approach:

- Pick a value $\alpha^* \in (0, 1/2)$ (e.g., $\alpha^* = \alpha$). This is a tuning parameter.
- From the sample, draw M size- n samples that are each iid \hat{F}_n . Denote the ecdfs of the samples by $\{\hat{F}_{n,j}^*\}$.
- For each $j = 1, \dots, M$, draw N size- n samples, each iid $\hat{F}_{n,j}$. Find the $1 - \alpha$ quantile of the pivot. This gives M values of the $1 - \alpha$ quantile. Let c be the $1 - \alpha^*$ quantile of the $1 - \alpha$ quantiles.
- Report as a $1 - \alpha$ confidence interval for $T(F)$ the interval one gets by taking c to be the estimate of the $1 - \alpha$ quantile of the pivot.

In a variety of simulations, this tends to be more conservative than Beran’s method, and more often attains at least the nominal coverage probability.

Exercise. Consider forming a two-sided 95% confidence interval for the mean θ of a distribution F based on the sample mean, using $|\bar{X} - \theta|$ as a pivot. Implement the three “double-bootstrap” approaches to finding a confidence interval (Beran’s pre-pivoting, Efron’s calibrated target percentile, and the percentile-of-percentile). Generate 100 synthetic samples of size 100 from the following distributions: normal, lognormal, Cauchy, mixtures of normals with the same mean but quite different variances (try different mixture coefficients), and mixtures of normals with different means and different variances (the means should differ enough that the result is bimodal). Apply the three double bootstrap methods to each, resampling 1000 times from each of 1000 first-generation bootstrap samples. Which method on the average has the lowest level error? Which method tends to be most conservative? Try to provide some intuition about the circumstances under which each method fails, and the circumstances under which each method would be expected to perform well. How do you interpret coverage for the Cauchy? **Warning:** You might need to be clever in how you implement this to make it a feasible calculation in S or Matlab. If you try to store all the intermediate results, the memory requirement is huge. On the other hand, if you use too many loops, the execution time will be long.

1.8 Bootstrap confidence sets based on Stein (shrinkage) estimates

Beran (1995) discusses finding a confidence region for the mean vector $\theta \in \mathbf{R}^q$, $q \geq 3$, from data $X \sim N(\theta, I)$. This is an example illustrating that *what* one bootstraps is important, and that naive plug-in bootstrapping doesn't always work.

The sets are spheres centered at the shrinkage estimate

$$\hat{\theta}_S = \left(1 - \frac{q-2}{\|X\|^2}\right) X, \quad (61)$$

with random diameter \hat{d} . That is, the confidence sets C are of the form

$$C(\hat{\theta}_S, \hat{d}) = \{\gamma \in \mathbf{R}^q : \|\hat{\theta}_S - \gamma\| \leq \hat{d}\}. \quad (62)$$

The problem is how to find $\hat{d} = \hat{d}(X; \alpha)$ such that

$$\mathbf{P}_\gamma\{C(\hat{\theta}_S, \hat{d}) \ni \gamma\} \geq 1 - \alpha \quad (63)$$

whatever be $\gamma \in \mathbf{R}^q$.

This problem is parametric: F is known up to the q -dimensional mean vector θ . We can thus use a “parametric bootstrap” to generate data that are approximately from F , instead of drawing directly from \hat{F}_n : if we have an estimate $\hat{\theta}$ of θ , we can generate artificial data distributed as $N(\hat{\theta}, I)$. If $\hat{\theta}$ is a good estimator, the artificial data will be distributed nearly as F . The issue is in what sense $\hat{\theta}$ needs to be good.

Beran shows (somewhat surprisingly) that resampling from $N(\hat{\theta}_S, I)$ or from $N(X, I)$ do not tend to work well in calibrating \hat{d} . The crucial thing in using the bootstrap to calibrate the radius of the confidence sphere seems to be to estimate $\|\theta\|$ well.

Definition 2 *The geometrical risk of a confidence set C for the parameter $\theta \in \mathbf{R}^q$ is*

$$G_q(C, \theta) \equiv q^{-1/2} E_\theta \sup_{\gamma \in C} \|\gamma - \theta\|. \quad (64)$$

That is, the geometrical risk is the expected distance to the parameter from the most distant point in the confidence set.

For confidence spheres

$$C = C(\hat{\theta}, \hat{d}) = \{\gamma \in \mathbf{R}^q : \|\gamma - \hat{\theta}\| \leq \hat{r}\}, \quad (65)$$

the geometrical risk can be decomposed further: the distance from θ to the most distant point in the confidence set is the distance from θ to the center of the sphere, plus the radius of the sphere, so

$$\begin{aligned} G_q(C(\hat{\theta}, \hat{r}), \theta) &= q^{-1/2} E_\theta (\|\hat{\theta} - \theta\| + \hat{d}) \\ &= q^{-1/2} E_\theta \|\hat{\theta} - \theta\| + q^{-1/2} E_\theta \hat{d}. \end{aligned} \tag{66}$$

Lemma 1 (*Beran, 1995, Lemma 4.1*). *Define*

$$W_q(X, \gamma) \equiv (q^{-1/2}(\|X - \gamma\|^2 - q), q^{-1/2}\gamma'(X - \gamma)). \tag{67}$$

Suppose $\{\gamma_q \in \mathbf{R}^q\}$ is any sequence such that

$$\frac{\|\gamma_q\|^2}{q} \rightarrow a < \infty \text{ as } q \rightarrow \infty. \tag{68}$$

Then

$$\begin{aligned} W_q(X, \gamma_q) &\rightarrow (\sqrt{2}Z_1, \sqrt{a}Z_2) \\ &W \end{aligned} \tag{69}$$

under \mathbf{P}_{γ_q} , where Z_1 and Z_2 are iid standard normal random variables. (The symbol $\rightarrow W$ denotes weak convergence of distributions.)

Proof. Under \mathbf{P}_{γ_q} , the distribution of $X - \gamma$ is rotationally invariant, so the distribution of the components of W_q depend on γ only through $\|\gamma\|$. Wlog, we may take each component of γ_q to be $q^{-1/2}\|\gamma_q\|$. The distribution of the first component of W_q is then that of the sum of squares of q iid standard normals (a chi-square rv with q df), minus the expected value of that sum, times $q^{-1/2}$. The standard deviation of a chi-square random variable with q df is $\sqrt{2q}$, so the first component of W_q is $\sqrt{2}$ times a standardized variable whose distribution is asymptotically (in q) normal. The second component of W_q is a linear combination of iid standard normals; by symmetry (as argued above), its distribution is that of

$$\begin{aligned} q^{-1/2} \sum_{j=1}^q q^{-1/2} \|\gamma_q\| Z_j &= \|\gamma_q\| \sum_{j=1}^q Z_j \\ &\rightarrow a^{1/2} Z_2. \end{aligned} \tag{70}$$

Recall that the squared-error risk (normalized by $q^{-1/2}$) of the James-Stein estimator is $1 - q^{-1}E_\theta\{(q-2)^2/\|X\|^2\} < 1$. The difference between the loss of $\hat{\theta}_S$ and an unbiased estimate of its risk is

$$D_q(X, \theta) = q^{-1/2}\{\|\hat{\theta}_S - \theta\|^2 - [q - (q-2)^2/\|X\|^2]\}. \quad (71)$$

By rotational invariance, the distribution of this quantity depends on θ only through $\|\theta\|$; Beran writes the distribution as $H_q(\|\theta\|^2/q)$. Beran shows that if $\{\gamma_q \in \mathbf{R}^q\}$ satisfies 68, then

$$H_q(\|\gamma_q\|^2/q) \xrightarrow{W} N(0, \sigma^2(a)), \quad (72)$$

where

$$\sigma^2(t) \equiv 2 - 4t/(1+t)^2 \geq 1. \quad (73)$$

Define

$$\hat{\theta}_{\text{CL}} = [1 - (q-2)/\|X\|^2]_+^{1/2} X. \quad (74)$$

Theorem 1 (*Beran, 1995, Theorem 3.1*) *Suppose $\{\gamma_q \in \mathbf{R}^q\}$ satisfies 68. Then*

$$H_q(\|\hat{\theta}_{\text{CL}}\|^2/q) \xrightarrow{W} N(0, \sigma^2(a)), \quad (75)$$

$$H_q(\|X\|^2/q) \xrightarrow{W} N(0, \sigma^2(1+a)), \quad (76)$$

and

$$H_q(\|\hat{\theta}_S\|^2/q) \rightarrow N(0, \sigma^2(a^2/(1+a))), \quad (77)$$

all in P_{γ_q} probability.

It follows that to estimate H_q by the bootstrap consistently, one should use

$$\hat{H}_B = H_q(\|\hat{\theta}_{\text{CL}}\|^2/q) \quad (78)$$

rather than estimating using either the norm of X or the norm of the James-Stein estimate $\hat{\theta}_S$ of θ .

Proof. Lemma 1 implies that under the conditions of the theorem, $\|\hat{\theta}_{\text{CL}}\|^2/q \rightarrow a$, $\|X\|^2/q \rightarrow 1+a$, and $\|\hat{\theta}_S\|^2/q \rightarrow a^2/(1+a)$.