

~~Most simulations would have $M > 1000$, but M between 10,000 and 25,000 is not uncommon. One important guideline for determining the number of trials, is the purpose of the simulation. If the tail of the distribution is of interest (e.g., estimating Type I error, getting p values, etc.), then more trials are needed to ensure that there will be a good estimate of that area.~~

6.4 Bootstrap Methods

The treatment of the bootstrap methods described here comes from Efron and Tibshirani [1993]. The interested reader is referred to that text for more information on the underlying theory behind the bootstrap. There does not seem to be a consistent terminology in the literature for what techniques are considered bootstrap methods. Some refer to the resampling techniques of the previous section as bootstrap methods. Here, we use *bootstrap* to refer to Monte Carlo simulations that treat the original sample as the pseudo-population or as an estimate of the population. Thus, in the steps where we randomly sample from the pseudo-population, we now resample from the original sample.

In this section, we discuss the general bootstrap methodology, followed by some applications of the bootstrap. These include bootstrap estimates of the standard error, bootstrap estimates of bias, and bootstrap confidence intervals.

General Bootstrap Methodology

The bootstrap is a method of Monte Carlo simulation where no parametric assumptions are made about the underlying population that generated the random sample. Instead, we use the sample as an estimate of the population. This estimate is called the empirical distribution \hat{F} where each x_i has probability mass $1/n$. Thus, each x_i has the same likelihood of being selected in a new sample taken from \hat{F} .

When we use \hat{F} as our pseudo-population, then we resample *with replacement* from the original sample $\mathbf{x} = (x_1, \dots, x_n)$. We denote the new sample obtained in this manner by $\mathbf{x}^* = (x_1^*, \dots, x_n^*)$. Since we are sampling with replacement from the original sample, there is a possibility that some points x_i will appear more than once in \mathbf{x}^* or maybe not at all. We are looking at the univariate situation, but the bootstrap concepts can also be applied in the d -dimensional case.

A small example serves to illustrate these ideas. Let's say that our random sample consists of the four numbers $\mathbf{x} = (5, 8, 3, 2)$. The following are possible samples \mathbf{x}^* , when we sample with replacement from \mathbf{x} :

$$\mathbf{x}^{*1} = (x_4, x_4, x_2, x_1) = (2, 2, 8, 5)$$

$$\mathbf{x}^{*2} = (x_4, x_2, x_3, x_4) = (2, 8, 3, 2).$$

We use the notation \mathbf{x}^{*b} , $b = 1, \dots, B$ for the b -th bootstrap data set.

In many situations, the analyst is interested in estimating some parameter θ by calculating a statistic from the random sample. We denote this estimate by

$$\hat{\theta} = T = t(x_1, \dots, x_n). \quad (6.11)$$

We might also like to determine the standard error in the estimate $\hat{\theta}$ and the bias. The bootstrap method can provide an estimate of this when analytical methods fail. The method is also suitable for situations when the estimator $\hat{\theta} = t(\mathbf{x})$ is complicated.

To get estimates of bias or standard error of a statistic, we obtain B bootstrap samples by sampling with replacement from the original sample. For every bootstrap sample, we calculate the same statistic to obtain the *bootstrap replications* of θ , as follows

$$\hat{\theta}^{*b} = t(\mathbf{x}^{*b}); \quad b = 1, \dots, B. \quad (6.12)$$

These B bootstrap replicates provide us with an estimate of the distribution of $\hat{\theta}$. This is similar to what we did in the previous section, except that we are not making any assumptions about the distribution for the original sample. Once we have the bootstrap replicates in Equation 6.12, we can use them to understand the distribution of the estimate.

The steps for the basic bootstrap methodology are given here, with detailed procedures for finding specific characteristics of $\hat{\theta}$ provided later. The issue of how large to make B is addressed with each application of the bootstrap.

PROCEDURE - BASIC BOOTSTRAP

1. Given a random sample, $\mathbf{x} = (x_1, \dots, x_n)$, calculate $\hat{\theta}$.
2. Sample with replacement from the original sample to get $\mathbf{x}^{*b} = (x_1^{*b}, \dots, x_n^{*b})$.
3. Calculate the same statistic using the bootstrap sample in step 2 to get, $\hat{\theta}^{*b}$.
4. Repeat steps 2 through 3, B times.
5. Use this estimate of the distribution of $\hat{\theta}$ (i.e., the bootstrap replicates) to obtain the desired characteristic (e.g., standard error, bias or confidence interval).

Efron and Tibshirani [1993] discuss a method called the *parametric bootstrap*. In this case, the data analyst makes an assumption about the distribution that generated the original sample. Parameters for that distribution are estimated from the sample, and resampling (in step 2) is done using the assumed distribution and the estimated parameters. The parametric bootstrap is closer to the Monte Carlo methods described in the previous section.

For instance, say we have reason to believe that the data come from an exponential distribution with parameter λ . We need to estimate the variance and use

$$\hat{\theta} = \frac{1}{n} \sum_{i=1}^n (x_i - \bar{x})^2 \quad (6.13)$$

as the estimator. We can use the parametric bootstrap as outlined above to understand the behavior of $\hat{\theta}$. Since we assume an exponential distribution for the data, we estimate the parameter λ from the sample to get $\hat{\lambda}$. We then resample from an exponential distribution with parameter $\hat{\lambda}$ to get the bootstrap samples. The reader is asked to implement the parametric bootstrap in the exercises.

Bootstrap Estimate of Standard Error

When our goal is to estimate the standard error of $\hat{\theta}$ using the bootstrap method, we proceed as outlined in the previous procedure. Once we have the estimated distribution for $\hat{\theta}$, we use it to estimate the standard error for $\hat{\theta}$. This estimate is given by

$$\hat{SE}_B(\hat{\theta}) = \left\{ \frac{1}{B-1} \sum_{b=1}^B (\hat{\theta}^{*b} - \bar{\hat{\theta}}^*)^2 \right\}^{\frac{1}{2}}, \quad (6.14)$$

where

$$\bar{\hat{\theta}}^* = \frac{1}{B} \sum_{b=1}^B \hat{\theta}^{*b}. \quad (6.15)$$

Note that Equation 6.14 is just the sample standard deviation of the bootstrap replicates, and Equation 6.15 is the sample mean of the bootstrap replicates.

Efron and Tibshirani [1993] show that the number of bootstrap replicates B should be between 50 and 200 when estimating the standard error of a statistic. Often the choice of B is dictated by the computational complexity of $\hat{\theta}$, the sample size n , and the computer resources that are available. Even using

a small value of B , say $B = 25$, the analyst will gain information about the variability of $\hat{\theta}$. In most cases, taking more than 200 bootstrap replicates to estimate the standard error is unnecessary.

The procedure for finding the bootstrap estimate of the standard error is given here and is illustrated in Example 6.9

PROCEDURE - BOOTSTRAP ESTIMATE OF THE STANDARD ERROR

1. Given a random sample, $\mathbf{x} = (x_1, \dots, x_n)$, calculate the statistic $\hat{\theta}$.
2. Sample with replacement from the original sample to get $\mathbf{x}^{*b} = (x_1^{*b}, \dots, x_n^{*b})$.
3. Calculate the same statistic using the sample in step 2 to get the bootstrap replicates, $\hat{\theta}^{*b}$.
4. Repeat steps 2 through 3, B times.
5. Estimate the standard error of $\hat{\theta}$ using Equations 6.14 and 6.15.

Example 6.9

The lengths of the forearm (in inches) of 140 adult males are contained in the file `forearm` [Hand, et al., 1994]. We use these data to estimate the skewness of the population. We then estimate the standard error of this statistic using the bootstrap method. First we load the data and calculate the skewness.

```
load forearm
% Sample with replacement from this.
% First get the sample size.
n = length( forearm );
B = 100; % number of bootstrap replicates
% Get the value of the statistic of interest.
theta = skewness( forearm );
```

The estimated skewness in the `forearm` data is -0.11. To implement the bootstrap, we use the MATLAB Statistics Toolbox function `unidrnd` to sample with replacement from the original sample. The corresponding function from the Computational Statistics Toolbox can also be used. The output from this function will be indices from 1 to n that point to what observations have been selected for the bootstrap sample.

```
% Use unidrnd to get the indices to the resamples.
% Note that each column corresponds to indices
% for a bootstrap resample.
inds = unidrnd( n, n, B );
% Extract these from the data.
xboot = forearm( inds );
% We can get the skewness for each column using the
% MATLAB Statistics Toolbox function skewness.
```

```

thetab = skewness(xboot);
seb = std(thetab);

```

From this we get an estimated standard error in the skewness of 0.14. Efron and Tibshirani [1993] recommend that one look at histograms of the bootstrap replicates as a useful tool for understanding the distribution of $\hat{\theta}$. We show the histogram in [Figure 6.6](#).

The MATLAB Statistics Toolbox has a function called `bootstrap` that returns the bootstrap replicates. We now show how to get the bootstrap estimate of standard error using this function.

```

% Now show how to do it with MATLAB Statistics Toolbox
% function: bootstrap.
Bmat = bootstrap(B, 'skewness', forearm);
% What we get back are the bootstrap replicates.
% Get an estimate of the standard error.
sebmat = std(Bmat);

```

Note that one of the arguments to `bootstrap` is a string representing the function that calculates the statistics. From this, we get an estimated standard error of 0.12.

□

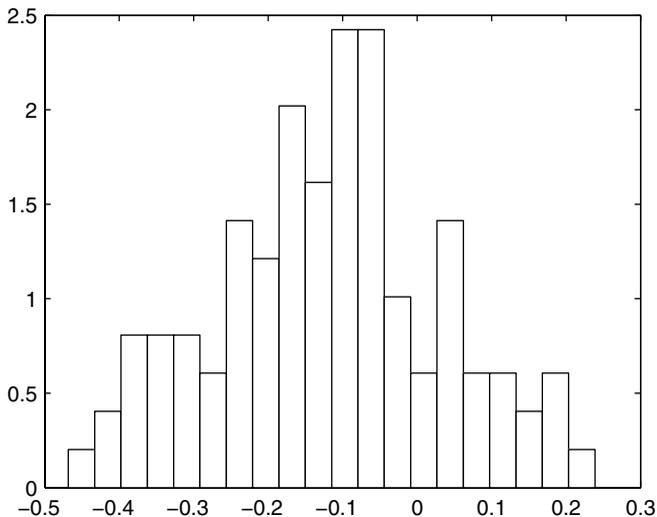


FIGURE 6.6 This is a histogram for the bootstrap replicates in Example 6.9. This shows the estimated distribution of the sample skewness of the `forearm` data.

Bootstrap Estimate of Bias

The standard error of an estimate is one measure of its performance. Bias is another quantity that measures the statistical accuracy of an estimate. From [Chapter 3](#), the bias is defined as the difference between the expected value of the statistic and the parameter,

$$\text{bias}(T) = E[T] - \theta. \quad (6.16)$$

The expectation in Equation 6.16 is taken with respect to the true distribution F . To get the bootstrap estimate of bias, we use the empirical distribution \hat{F} as before. We resample from the empirical distribution and calculate the statistic using each bootstrap resample, yielding the bootstrap replicates $\hat{\theta}^{*b}$. We use these to estimate the bias from the following:

$$\hat{\text{bias}}_B = \bar{\hat{\theta}}^* - \hat{\theta}, \quad (6.17)$$

where $\bar{\hat{\theta}}^*$ is given by the mean of the bootstrap replicates (Equation 6.15).

Presumably, one is interested in the bias in order to correct for it. The bias-corrected estimator is given by

$$\hat{\theta} = \hat{\theta} - \hat{\text{bias}}_B. \quad (6.18)$$

Using Equation 6.17 in Equation 6.18, we have

$$\hat{\theta} = 2\hat{\theta} - \bar{\hat{\theta}}^*. \quad (6.19)$$

More bootstrap samples are needed to estimate the bias, than are required to estimate the standard error. Efron and Tibshirani [1993] recommend that $B \geq 400$.

It is useful to have an estimate of the bias for $\hat{\theta}$, but caution should be used when correcting for the bias. Equation 6.19 will hopefully yield a less biased estimate, but it could turn out that $\hat{\theta}$ will have a larger variation or standard error. It is recommended that if the estimated bias is small relative to the estimate of standard error (both of which can be estimated using the bootstrap method), then the analyst should not correct for the bias [Efron and Tibshirani, 1993]. However, if this is not the case, then perhaps some other, less biased, estimator should be used to estimate the parameter θ .

PROCEDURE - BOOTSTRAP ESTIMATE OF THE BIAS

1. Given a random sample, $\mathbf{x} = (x_1, \dots, x_n)$, calculate the statistic $\hat{\theta}$.

2. Sample with replacement from the original sample to get $\mathbf{x}^{*b} = (x_1^{*b}, \dots, x_n^{*b})$.
3. Calculate the same statistic using the sample in step 2 to get the bootstrap replicates, $\hat{\theta}^{*b}$.
4. Repeat steps 2 through 3, B times.
5. Using the bootstrap replicates, calculate $\overline{\hat{\theta}^*}$.
6. Estimate the bias of $\hat{\theta}$ using Equation 6.17.

Example 6.10

We return to the `forearm` data of Example 6.9, where now we want to estimate the bias in the sample skewness. We use the same bootstrap replicates as before, so all we have to do is to calculate the bias using Equation 6.17.

```
% Use the same replicates from before.
% Evaluate the mean using Equation 6.15.
meanb = mean(thetab);
% Now estimate the bias using Equation 6.17.
biasb = meanb - theta;
```

We have an estimated bias of -0.011. Note that this is small relative to the standard error.

□

In the next chapter, we discuss another method for estimating the bias and the standard error of a statistic called the jackknife. The jackknife method is related to the bootstrap. However, since it is based on the reuse or partitioning of the original sample rather than resampling, we do not include it here.

Bootstrap Confidence Intervals

There are several ways of constructing confidence intervals using the bootstrap. We discuss three of them here: the standard interval, the bootstrap- t interval and the percentile method. Because it uses the jackknife procedure, an improved bootstrap confidence interval called the BC_a will be presented in the next chapter.

Bootstrap Standard Confidence Interval

The *bootstrap standard confidence interval* is based on the parametric form of the confidence interval that was discussed in Section 6.2. We showed that the $(1 - \alpha) \cdot 100\%$ confidence interval for the mean can be found using

$$P\left(\bar{X} - z^{(1-\alpha/2)} \frac{\sigma}{\sqrt{n}} < \mu < \bar{X} - z^{(\alpha/2)} \frac{\sigma}{\sqrt{n}}\right) = 1 - \alpha. \quad (6.20)$$

Similar to this, the bootstrap standard confidence interval is given by

$$(\hat{\theta} - z^{(1-\alpha/2)} SE_{\hat{\theta}}, \hat{\theta} - z^{(\alpha/2)} SE_{\hat{\theta}}), \quad (6.21)$$

where $SE_{\hat{\theta}}$ is the standard error for the statistic $\hat{\theta}$ obtained using the bootstrap [Mooney and Duval, 1993]. The confidence interval in Equation 6.21 can be used when the distribution for θ is normally distributed or the normality assumption is plausible. This is easily coded in MATLAB using previous results and is left as an exercise for the reader.

Bootstrap-t Confidence Interval

The second type of confidence interval using the bootstrap is called the *bootstrap-t*. We first generate B bootstrap samples, and for each bootstrap sample the following quantity is computed:

$$z^{*b} = \frac{\hat{\theta}^{*b} - \hat{\theta}}{\hat{SE}^{*b}}. \quad (6.22)$$

As before, $\hat{\theta}^{*b}$ is the bootstrap replicate of $\hat{\theta}$, but \hat{SE}^{*b} is the estimated standard error of $\hat{\theta}^{*b}$ for that bootstrap sample. If a formula exists for the standard error of $\hat{\theta}^{*b}$, then we can use that to determine the denominator of Equation 6.22. For instance, if $\hat{\theta}$ is the mean, then we can calculate the standard error as explained in [Chapter 3](#). However, in most situations where we have to resort to using the bootstrap, these formulas are not available. One option is to use the bootstrap method of finding the standard error, keeping in mind that you are estimating the standard error of $\hat{\theta}^{*b}$ using the bootstrap sample \mathbf{x}^{*b} . In other words, one resamples with replacement from the bootstrap sample \mathbf{x}^{*b} to get an estimate of SE^{*b} .

Once we have the B bootstrapped z^{*b} values from Equation 6.22, the next step is to estimate the quantiles needed for the endpoints of the interval. The $\alpha/2$ -th quantile, denoted by $\hat{t}^{(\alpha/2)}$ of the z^{*b} , is estimated by

$$\alpha/2 = \frac{\#(z^{*b} \leq \hat{t}^{(\alpha/2)})}{B}. \quad (6.23)$$

This says that the estimated quantile is the $\hat{t}^{(\alpha/2)}$ such that $100 \cdot \alpha/2$ % of the points z^{*b} are less than this number. For example, if $B = 100$ and $\alpha/2 = 0.05$, then $\hat{t}^{(0.05)}$ could be estimated as the fifth largest value of the z^{*b} ($B \cdot \alpha/2 = 100 \cdot 0.05 = 5$). One could also use the quantile estimates discussed previously in [Chapter 3](#) or some other suitable estimate.

We are now ready to calculate the bootstrap-t confidence interval. This is given by

$$(\hat{\theta} - \hat{t}^{(1-\alpha/2)} \cdot \hat{SE}_{\hat{\theta}}, \hat{\theta} - \hat{t}^{(\alpha/2)} \cdot \hat{SE}_{\hat{\theta}}), \quad (6.24)$$

where \hat{SE} is an estimate of the standard error of $\hat{\theta}$. The bootstrap- t interval is suitable for location statistics such as the mean or quantiles. However, its accuracy for more general situations is questionable [Efron and Tibshirani, 1993]. The next method based on the bootstrap percentiles is more reliable.

PROCEDURE - BOOTSTRAP-T CONFIDENCE INTERVAL

1. Given a random sample, $\mathbf{x} = (x_1, \dots, x_n)$, calculate $\hat{\theta}$.
2. Sample with replacement from the original sample to get $\mathbf{x}^{*b} = (x_1^{*b}, \dots, x_n^{*b})$.
3. Calculate the same statistic using the sample in step 2 to get $\hat{\theta}^{*b}$.
4. Use the bootstrap sample \mathbf{x}^{*b} to get the standard error of $\hat{\theta}^{*b}$. This can be calculated using a formula or estimated by the bootstrap.
5. Calculate z^{*b} using the information found in steps 3 and 4.
6. Repeat steps 2 through 5, B times, where $B \geq 1000$.
7. Order the z^{*b} from smallest to largest. Find the quantiles $\hat{t}^{(1-\alpha/2)}$ and $\hat{t}^{(\alpha/2)}$.
8. Estimate the standard error $\hat{SE}_{\hat{\theta}}$ of $\hat{\theta}$ using the B bootstrap replicates of $\hat{\theta}^{*b}$ (from step 3).
9. Use Equation 6.24 to get the confidence interval.

The number of bootstrap replicates that are needed is quite large for confidence intervals. It is recommended that B should be 1000 or more. If no formula exists for calculating the standard error of $\hat{\theta}^{*b}$, then the bootstrap method can be used. This means that there are two levels of bootstrapping: one for finding the \hat{SE}^{*b} and one for finding the z^{*b} , which can greatly increase the computational burden. For example, say that $B = 1000$ and we use 50 bootstrap replicates to find \hat{SE}^{*b} , then this results in a total of 50,000 resamples.

Example 6.11

Say we are interested in estimating the variance of the **forearm** data, and we decide to use the following statistic,

$$\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n (X_i - \bar{X})^2,$$

which is the sample second central moment. We write our own simple function called `mom` (included in the Computational Statistics Toolbox) to estimate this.

```
% This function will calculate the sample 2nd
% central moment for a given sample vector x.
function mr = mom(x)
n = length(x);
mu = mean(x);
mr = (1/n)*sum((x-mu).^2);
```

We use this function as an input argument to `bootstrap` to get the bootstrap confidence interval. The MATLAB code given below also shows how to get the bootstrap estimate of standard error for each bootstrap sample. First we load the data and get the observed value of the statistic.

```
load forearm
n = length(forearm);
alpha = 0.1;
B = 1000;
thetahat = mom(forearm);
```

Now we get the bootstrap replicates using the function `bootstrap`. One of the optional output arguments from this function is a matrix of indices for the resamples. As shown below, each column of the output `bootsam` contains the indices to a bootstrap sample. We loop through all of the bootstrap samples to estimate the standard error of the bootstrap replicate using that resample.

```
% Get the bootstrap replicates and samples.
[bootreps, bootsam] = bootstrap(B, 'mom', forearm);
% Set up some storage space for the SE's.
sehats = zeros(size(bootreps));
% Each column of bootsam contains indices
% to a bootstrap sample.
for i = 1:B
    % Extract the sample from the data.
    xstar = forearm(bootsam(:,i));
    bvals(i) = mom(xstar);
    % Do bootstrap using that sample to estimate SE.
    sehats(i) = std(bootstrap(25, 'mom', xstar));
end
zvals = (bootreps - thetahat)./sehats;
```

Then we get the estimate of the standard error that we need for the endpoints of the interval.

```
% Estimate the SE using the bootstrap.
SE = std(bootreps);
```

Now we get the quantiles that we need for the interval given in Equation 6.24 and calculate the interval.

```
% Get the quantiles.
k = B*alpha/2;
szval = sort(zvals);
tlo = szval(k);
thi = szval(B-k);
% Get the endpoints of the interval.
blo = thetahat - thi*SE;
bhi = thetahat - tlo*SE;
```

The bootstrap- t interval for the variance of the `forearm` data is (1.00, 1.57).
□

Bootstrap Percentile Interval

An improved bootstrap confidence interval is based on the quantiles of the distribution of the bootstrap replicates. This technique has the benefit of being more stable than the bootstrap- t , and it also enjoys better theoretical coverage properties [Efron and Tibshirani, 1993]. The *bootstrap percentile confidence interval* is

$$(\hat{\theta}_B^{*(\alpha/2)}, \hat{\theta}_B^{*(1-\alpha/2)}), \quad (6.25)$$

where $\hat{\theta}_B^{*(\alpha/2)}$ is the $\alpha/2$ quantile in the bootstrap distribution of $\hat{\theta}^*$. For example, if $\alpha/2 = 0.025$ and $B = 1000$, then $\hat{\theta}_B^{*(0.025)}$ is the $\hat{\theta}^{*b}$ in the 25th position of the ordered bootstrap replicates. Similarly, $\hat{\theta}_B^{*(0.975)}$ is the replicate in position 975. As discussed previously, some other suitable estimate for the quantile can be used.

The procedure is the same as the general bootstrap method, making it easy to understand and to implement. We outline the steps below.

PROCEDURE - BOOTSTRAP PERCENTILE INTERVAL

1. Given a random sample, $\mathbf{x} = (x_1, \dots, x_n)$, calculate $\hat{\theta}$.
2. Sample with replacement from the original sample to get $\mathbf{x}^{*b} = (x_1^{*b}, \dots, x_n^{*b})$.
3. Calculate the same statistic using the sample in step 2 to get the bootstrap replicates, $\hat{\theta}^{*b}$.
4. Repeat steps 2 through 3, B times, where $B \geq 1000$.
5. Order the $\hat{\theta}^{*b}$ from smallest to largest.
6. Calculate $B \cdot \alpha/2$ and $B \cdot (1 - \alpha/2)$.

7. The lower endpoint of the interval is given by the bootstrap replicate that is in the $B \cdot \alpha/2$ -th position of the ordered $\hat{\theta}^{*b}$, and the upper endpoint is given by the bootstrap replicate that is in the $B \cdot (1 - \alpha/2)$ -th position of the same ordered list. Alternatively, using quantile notation, the lower endpoint is the estimated quantile $\hat{q}_{\alpha/2}$ and the upper endpoint is the estimated quantile $\hat{q}_{1-\alpha/2}$, where the estimates are taken from the bootstrap replicates.

Example 6.12

Let's find the bootstrap percentile interval for the same `forearm` data. The confidence interval is easily found from the bootstrap replicates, as shown below.

```
% Use Statistics Toolbox function
% to get the bootstrap replicates.
bvals = bootstrp(B, 'mom', forearm);
% Find the upper and lower endpoints
k = B*alpha/2;
sbval = sort(bvals);
blo = sbval(k);
bhi = sbval(B-k);
```

This interval is given by (1.03, 1.45), which is slightly narrower than the bootstrap- t interval from Example 6.11.

□

So far, we discussed three types of bootstrap confidence intervals. The standard interval is the easiest and assumes that $\hat{\theta}$ is normally distributed. The bootstrap- t interval estimates the standardized version of $\hat{\theta}$ from the data, avoiding the normality assumptions used in the standard interval. The percentile interval is simple to calculate and obtains the endpoints directly from the bootstrap estimate of the distribution for $\hat{\theta}$. It has another advantage in that it is range-preserving. This means that if the parameter θ can take on values in a certain range, then the confidence interval will reflect that. This is not always the case with the other intervals.

According to Efron and Tibshirani [1993], the bootstrap- t interval has good coverage probabilities, but does not perform well in practice. The bootstrap percentile interval is more dependable in most situations, but does not enjoy the good coverage property of the bootstrap- t interval. There is another bootstrap confidence interval, called the BC_a interval, that has both good coverage and is dependable. This interval is described in the next chapter.

The bootstrap estimates of bias and standard error are also random variables, and they have their own error associated with them. So, how accurate are they? In the next chapter, we discuss how one can use the jackknife method to evaluate the error in the bootstrap estimates.

As with any method, the bootstrap is not appropriate in every situation. When analytical methods are available to understand the uncertainty associ-

ated with an estimate, then those are more efficient than the bootstrap. In what situations should the analyst use caution in applying the bootstrap? One important assumption that underlies the theory of the bootstrap is the notion that the empirical distribution function is representative of the true population distribution. If this is not the case, then the bootstrap will not yield reliable results. For example, this can happen when the sample size is small or the sample was not gathered using appropriate random sampling techniques. Chernick [1999] describes other examples from the literature where the bootstrap should not be used. We also address a situation in [Chapter 7](#) where the bootstrap fails. This can happen when the statistic is non-smooth, such as the median.

6.5 MATLAB Code

We include several functions with the Computational Statistics Toolbox that implement some of the bootstrap techniques discussed in this chapter. These are listed in [Table 6.2](#). Like `bootstrap`, these functions have an input argument that specifies a MATLAB function that calculates the statistic.

TABLE 6.2
List of MATLAB Functions for [Chapter 6](#)

Purpose	MATLAB Function
General bootstrap: resampling, estimates of standard error and bias	<code>csboot</code> <code>bootstrap</code>
Constructing bootstrap confidence Intervals	<code>csbootint</code> <code>csbootperint</code> <code>csbootbca</code>

As we saw in the examples, the MATLAB Statistics Toolbox has a function called `bootstrap` that will return the bootstrap replicates from the input argument `bootfun` (e.g., `mean`, `std`, `var`, etc.). It takes an input data set, finds the bootstrap resamples, applies the `bootfun` to the resamples, and stores the replicate in the first row of the output argument. The user can get two outputs from the function: the bootstrap replicates and the indices that correspond to the points selected in the resample.

There is a Bootstrap MATLAB Toolbox written by Zoubir and Iskander at the Curtin University of Technology. It is available for download at

www.atri.curtin.edu.au/csp. It requires the MATLAB Statistics Toolbox and has a postscript version of the reference manual.

Other software exists for Monte Carlo simulation as applied to statistics. The Efron and Tibshirani [1993] book has a description of S code for implementing the bootstrap. This code, written by the authors, can be downloaded from the statistics archive at Carnegie-Mellon University that was mentioned in [Chapter 1](#). Another software package that has some of these capabilities is called Resampling Stats® [Simon, 1999], and information on this can be found at www.resample.com. Routines are available from Resampling Stats for MATLAB [Kaplan, 1999] and Excel.

6.6 Further Reading

Mooney [1997] describes Monte Carlo simulation for inferential statistics that is written in a way that is accessible to most data analysts. It has some excellent examples of using Monte Carlo simulation for hypothesis testing using multiple experiments, assessing the behavior of an estimator, and exploring the distribution of a statistic using graphical techniques. The text by Gentle [1998] has a chapter on performing Monte Carlo studies in statistics. He discusses how simulation can be considered as a scientific experiment and should be held to the same high standards. Hoaglin and Andrews [1975] provide guidelines and standards for reporting the results from computations. Efron and Tibshirani [1991] explain several computational techniques, written at a level accessible to most readers. Other articles describing Monte Carlo inferential methods can be found in Joeckel [1991], Hope [1968], Besag and Diggle [1977], Diggle and Gratton [1984], Efron [1979], Efron and Gong [1983], and Teichroew [1965].

There has been a lot of work in the literature on bootstrap methods. Perhaps the most comprehensive and easy to understand treatment of the topic can be found in Efron and Tibshirani [1993]. Efron's [1982] earlier monogram on resampling techniques describes the jackknife, the bootstrap and cross-validation. A more recent book by Chernick [1999] gives an updated description of results in this area, and it also has an extensive bibliography (over 1,600 references!) on the bootstrap. Hall [1992] describes the connection between Edgeworth expansions and the bootstrap. A volume of papers on the bootstrap was edited by LePage and Billard [1992], where many applications of the bootstrap are explored. Politis, Romano, and Wolf [1999] present subsampling as an alternative to the bootstrap. A subset of articles that present the theoretical justification for the bootstrap are Efron [1981, 1985, 1987]. The paper by Boos and Zhang [2000] looks at a way to ease the computational burden of Monte Carlo estimation of the power of tests that uses resampling methods. For a nice discussion on the coverage of the bootstrap percentile confidence interval, see Polansky [1999].