

# Bootstrapping

Oliver d’Pug

## Bootstrap “By Hand”

The theoretical distribution of the median is often not very friendly. To get a handle on what the distribution looks like for samples from different population distributions, we can use Efron’s bootstrap.

We will make ample use of the **sample** function. This function returns either a random subset of the original data, or a permutation — depending upon the number of items that you request it returns.

```
x <- 1:10
x

## [1] 1 2 3 4 5 6 7 8 9 10

### Get a random sample of size 5 from the n=10 w/o replacement
sample(x, 5)

## [1] 5 8 10 6 7

### Get a random permutation of the n=10
sample(x)

## [1] 10 6 3 8 4 7 1 9 2 5

### Get a bootstrap sample (of size n from the n=10 w/ replacement)
bootsamp <- sample(x, replace = TRUE)
bootsamp
```

```
## [1] 5 3 3 6 1 4 2 8 8 4
```

The default is to have each item have a selection probability of  $1/n$ . We can specify other probabilities if we so choose.

```
### Create a bootstrap sample from x with probabilities that favor the first five elements
prob1 <- c(rep(.15, 5), rep(.05, 5))
prob1

## [1] 0.15 0.15 0.15 0.15 0.15 0.05 0.05 0.05 0.05 0.05

sample(10, replace = TRUE, prob=prob1)

## [1] 3 4 8 9 1 5 8 1 2 5
```

Sampling from a matrix is done similarly. We can sample from all elements of the matrix.

```
### Create a matrix filled with random values
y1 <- matrix( round(rnorm(25,5)), ncol=5)
y1

##      [,1] [,2] [,3] [,4] [,5]
## [1,]  7   5   3   6   5
## [2,]  3   5   5   6   4
```

```
## [3,] 5 5 6 5 6
## [4,] 5 4 7 4 7
## [5,] 5 5 3 5 5
```

```
### Save a sample of size 5 in the vector x1
x1 <- y1[sample(25, 5)]
x1
```

```
## [1] 4 5 5 5 5
```

Or, we can sample entire rows.

```
### Create a matrix filled with random values
y2 <- matrix( round(rnorm(40, 5)), ncol=5)
y2
```

```
##      [,1] [,2] [,3] [,4] [,5]
## [1,] 1 5 6 6 4
## [2,] 5 7 6 4 6
## [3,] 5 6 5 6 4
## [4,] 5 4 6 4 5
## [5,] 5 5 4 6 7
## [6,] 6 5 4 7 6
## [7,] 5 4 4 5 4
## [8,] 7 4 6 5 5
```

```
### Save a sample of rows from y2 in the matrix x2
x2 <- y2[sample(8, 3), ]
x2
```

```
##      [,1] [,2] [,3] [,4] [,5]
## [1,] 5 4 6 4 5
## [2,] 7 4 6 5 5
## [3,] 5 6 5 6 4
```

## Bootstrapping Using the sample Function

In the following example we find an estimate of the standard error for the estimate of the median. We will be using the **lapply** and **sapply** functions in combination with the **sample** function. (More information about the **lapply** and **sapply** functions can be found in the R help pages.)

```
### Create a data set by taking 100 random observations from a normal
### distribution with mean 5 and stdev 3. Each observation has been
### rounded to the nearest integer.
data <- round(rnorm(100, 5, 3))

### Display the first ten observations
data[1:10]
```

```
## [1] 4 5 10 6 7 7 7 3 2 9
```

The **lapply** function applies a given function to a list by passing the list elements to the function. We use **lapply** to run a function that generates a sample of the same size as the original sample by sampling with replacement.

```
### Generate 25 bootstrap samples. Note that the function has an argument i
### that is not passed to the sample function which uses the global
### data instead.
resamples <- lapply(1:25, function(i) sample(data, replace = T))
```

```
### Display the results for the third resample --- the others are similar.
resamples[3]
```

```
## [[1]]
## [1] 9 7 5 6 5 8 2 1 6 6 8 5 10 8 6 7 5 6 6 3 7 2 5 4 1
## [26] 6 5 6 9 6 9 8 6 3 6 8 12 0 6 7 5 7 1 4 6 7 8 3 7 9
## [51] 6 8 8 8 6 7 6 10 2 8 4 5 6 9 8 6 8 6 3 8 2 2 3 8 0
## [76] 6 4 5 3 6 0 4 8 9 6 5 4 9 5 5 4 9 8 7 7 6 3 9 6 7
```

We can now apply the **median** function to each of the bootstrap samples to generate a bootstrap distribution.

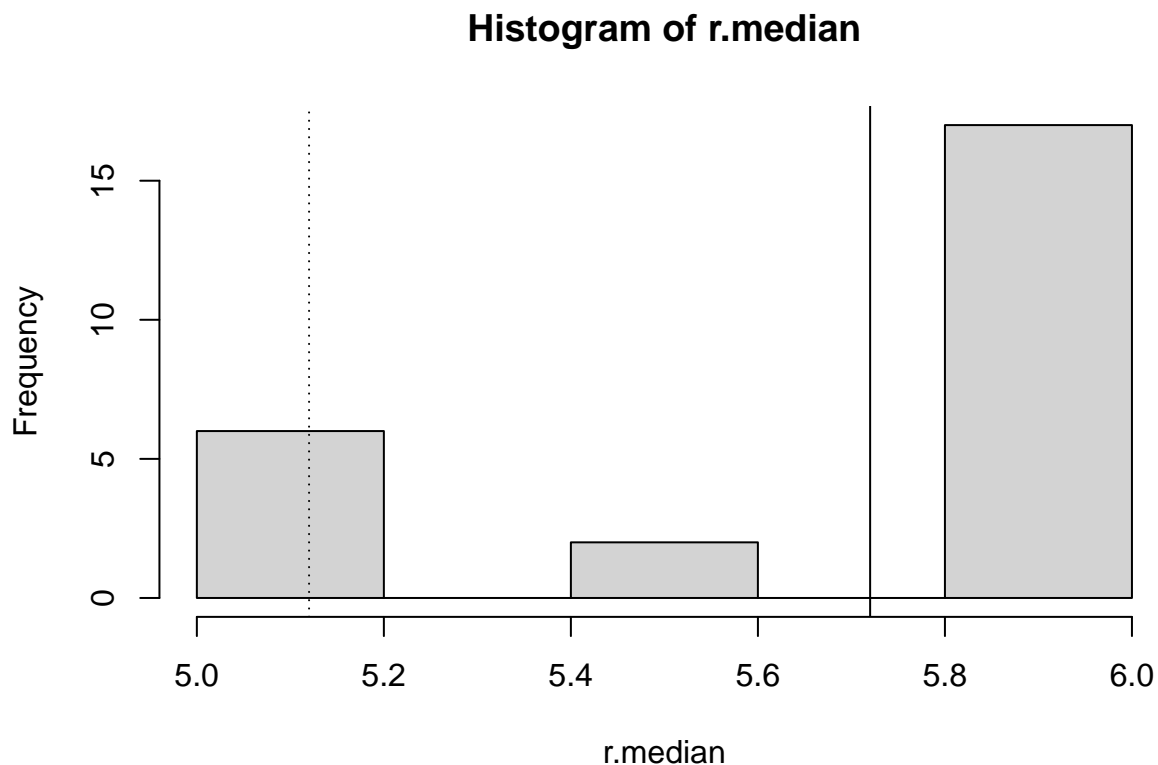
```
### Calculate the median for each bootstrap sample
r.median <- sapply(resamples, median)
r.median
```

```
## [1] 5.0 5.5 6.0 6.0 6.0 6.0 6.0 6.0 5.0 5.0 6.0 6.0 6.0 6.0 6.0 6.0 6.0 5.0 6.0
## [20] 5.0 5.5 6.0 5.0 6.0 6.0
```


```
### Compute summary statistics for the distribution of the medians
summary(r.median)
```

```
## Min. 1st Qu. Median Mean 3rd Qu. Max.
## 5.00 5.50 6.00 5.72 6.00 6.00
```

```
### Display a histogram of the distribution of the medians
hist(r.median)
abline(v=mean(data), lty=3)
abline(v=mean(r.median), lty=1)
```




```
### Display a normal probability plot
qqnorm(r.median)
qqline(r.median)
```

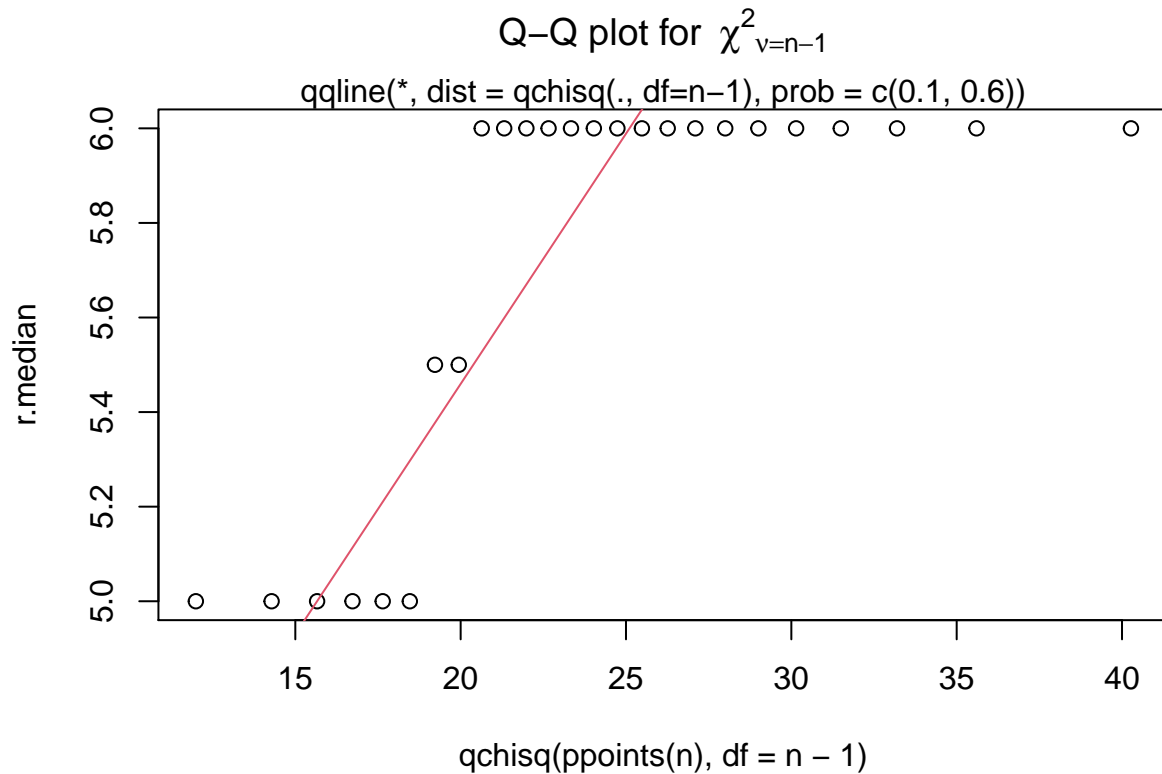


```
### Display a quantile-quantile plot using a standard normal
n <- length(r.median)
qqplot(qnorm(ppoints(n)), r.median,
       main = "Q-Q plot for Normal")
qqline(r.median, distribution = function(p){qnorm(p)},
       probs = c(0.25, 0.75), col = 2)
```

## Q-Q plot for Normal



```
### Display a quantile-quantile plot using a chiquare
n <- length(r.median)
qqplot(qchisq(ppoints(n), df = n-1), r.median,
       main = expression("Q-Q plot for" ~ {chi^2}[nu == n-1]))
qqline(r.median, qtype = 5, distribution = function(p){qchisq(p, df = n-1)},
       probs = c(0.1, 0.6), col = 2)
mtext("qqline(*, dist = qchisq(., df=n-1), prob = c(0.1, 0.6))")
```



These steps can be combined into a single function where all we would need to specify is which data set to use and how many times we want to resample in order to obtain the adjusted standard error of the median.

```

### Function to bootstrap the mean, standard error, and bias of the median
b.median <- function(data, nboot=9999) {
  resamples <- lapply(1:nboot, function(i) sample(data, replace=T))
  r.median <- sapply(resamples, median)
  d.xbar <- mean(data)
  r.xbar <- mean(r.median)
  r.bias <- r.xbar - d.xbar
  r.stderr <- sqrt(var(r.median))
  list(xbar=r.xbar, std.err=r.stderr, bias=r.bias, resamples=resamples,
       medians=r.median
      )
}

### Create some data to be used (same as in the above example)
data1 <- round(rnorm(100, 5, 3))

### Save the results of 10000 boots of the function b.median in the object b1
b1 <- b.median(data1, 10000)

### Display the third of the 10000 bootstrap samples
b1$resamples[3]

```

```

## [[1]]
## [1] 2 9 7 6 4 10 8 7 8 5 7 7 1 4 8 5 3 3 10 1 5 6 4 5 4

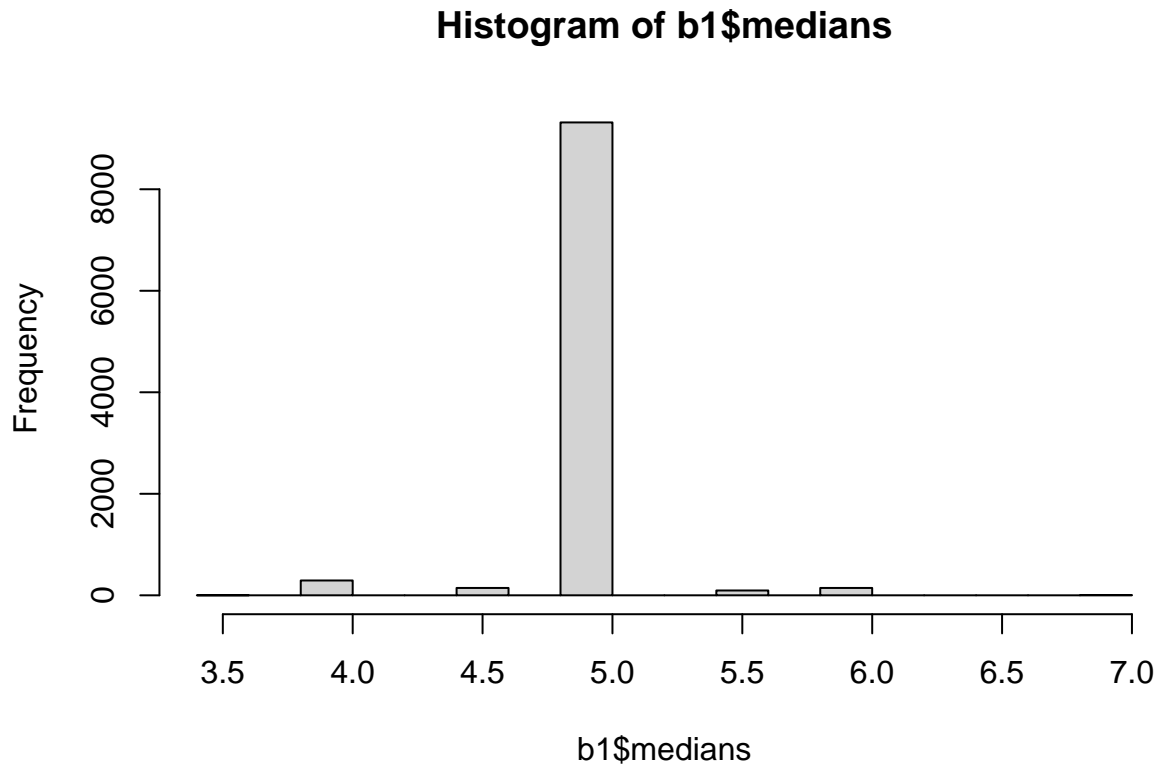
```

```
## [26] 9 8 3 4 7 1 3 3 4 -2 7 8 4 8 3 5 7 8 1 5 5 8 5 2 3
## [51] 4 3 5 7 8 9 8 0 9 4 1 7 6 3 10 2 2 6 4 5 6 4 5 8 -2
## [76] 8 0 7 6 6 5 3 4 5 2 -2 5 3 -2 5 5 5 3 5 3 2 7 1 5 3
```

```
### Display the mean, standard error, and bias of bootstrapped medians
print(b1$xbar, b1$std.err, b1$bias)
```


```
## [1] 5
```

```
### Display a histogram of the distribution of medians
hist(b1$medians)
```



```
### Make a normal probability plot
qqnorm(b1$medians)
qqline(b1$medians)
```

## Normal Q-Q Plot




```
### We can input the data directly into the function and display  
### the standard error in one line of code.  
b.median(data1)$std.err
```

```
## [1] 0.228817
```

```
### Or, we can make the histogram in one line  
hist(b.median(data1)$medians)
```



## Histogram of b.median(data1)\$medians



Note that the single line calls generate plots based upon different bootstrap samples. This approach is also not very time efficient.

It would be fairly simple to generalize the function to work for any summary statistic. We show this approach for the sample mean.

```
### Function to bootstrap the mean, standard error, and bias of a statistic
b.stat <- function(data, nboot=9999, fnct = mean, ...) {
  resamples <- lapply(1:nboot, function(i) sample(data, replace=T))
  r.boots <- sapply(resamples, fnct, ...)
  d.xbar <- mean(data)
  r.xbar <- mean(r.boots)
  r.bias <- r.xbar - d.xbar
  r.stderr <- sqrt(var(r.boots))
  list(xbar=r.xbar, std.err=r.stderr, bias=r.bias, resamples=resamples,
       boots=r.boots
      )
}

### Reuse the data created above so that we can compare mean and median
### Save the results of 10000 boots of the function b.median in the object b2
b2 <- b.stat(data1, 10000, function(x){mean(x)})

### Display the third of the 10000 bootstrap samples
b2$resamples[3]
```

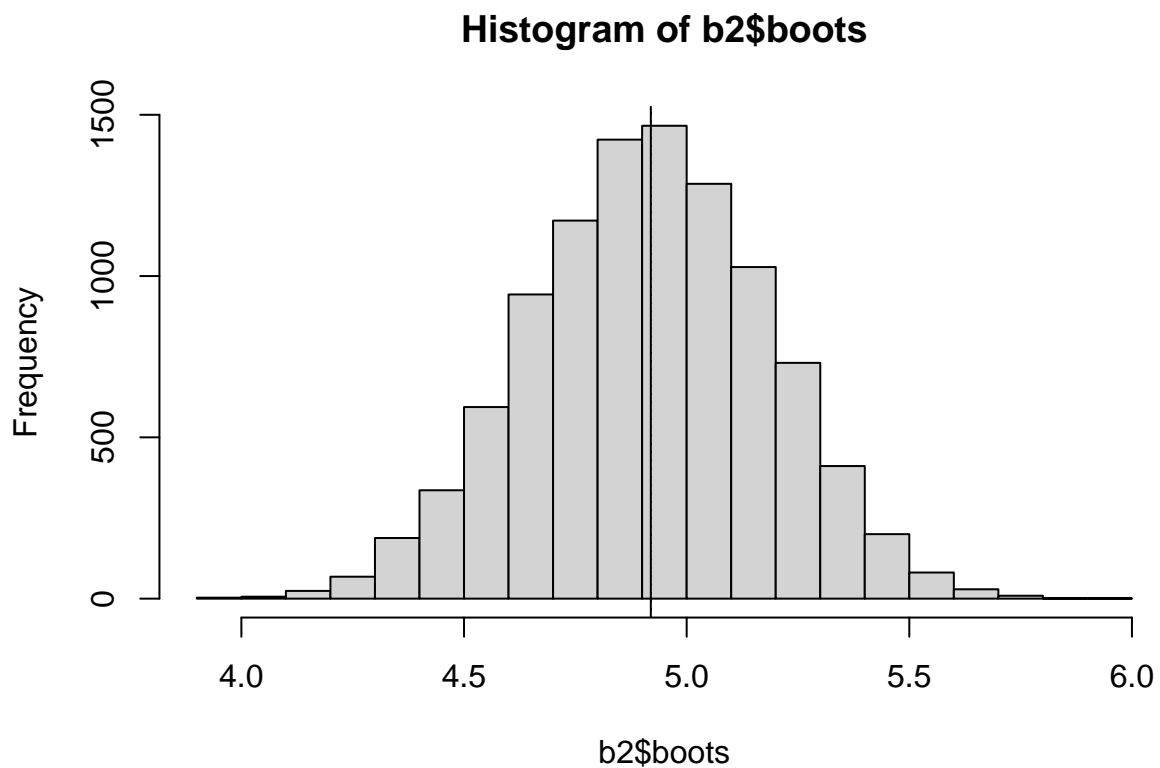
```
## [[1]]
```

```
## [1] 7 3 3 3 5 -2 6 4 5 1 4 3 6 5 3 2 10 3 7 9 6 8 4 4 5
## [26] 3 7 4 0 7 7 9 1 3 3 3 2 6 2 3 7 6 3 7 6 0 6 10 5 5
## [51] 7 8 3 -2 5 5 4 2 9 0 8 3 2 4 4 4 1 4 3 4 7 5 8 4 4
## [76] 7 4 5 6 6 5 5 6 4 5 4 7 4 4 4 5 5 -2 5 5 5 8 3 4 4
```

```
### Display the mean, standard error, and bias of bootstrapped means
print(round(c(b2$xbar, b2$std.err, b2$bias), 4))
```


```
## [1] 4.9195 0.2674 -0.0005
```

```
### Display a histogram of the distribution of means
hist(b2$boots)
abline(v = b2$xbar, lty = 1)
abline(v = mean(data1), lty = 3)
```



```
### Make a normal probability plot
qqnorm(b2$boots)
qqline(b2$boots)
```

## Normal Q-Q Plot



## The boot Package

There exists a package, **boot**, that already has a bootstrapping function that acts like the function that we created in the code chunk above. The advantage of using **boot** is that it has some additional methods of doing things.

```
### Bootstrap of the median of the data1 data from above.
### Make sure that the boot package is installed using
### install.packages("boot"), or use a package like pacman to take care
### of installation and loading.
#library(boot)
p_load(boot)


### Define the median function with data, d, and boot sample indices, i.
mystat <- function(d, i){
  median(d[i])
}

### Use the boot function to run the bootstrap
b3 <- boot(data1, mystat, R=9999)
b3
```

```
##
## ORDINARY NONPARAMETRIC BOOTSTRAP
##
##
## Call:
```

```
## boot(data = data1, statistic = mystat, R = 9999)
##
##
## Bootstrap Statistics :
##   original      bias   std. error
## t1*          5 -0.01240124  0.2256245
```

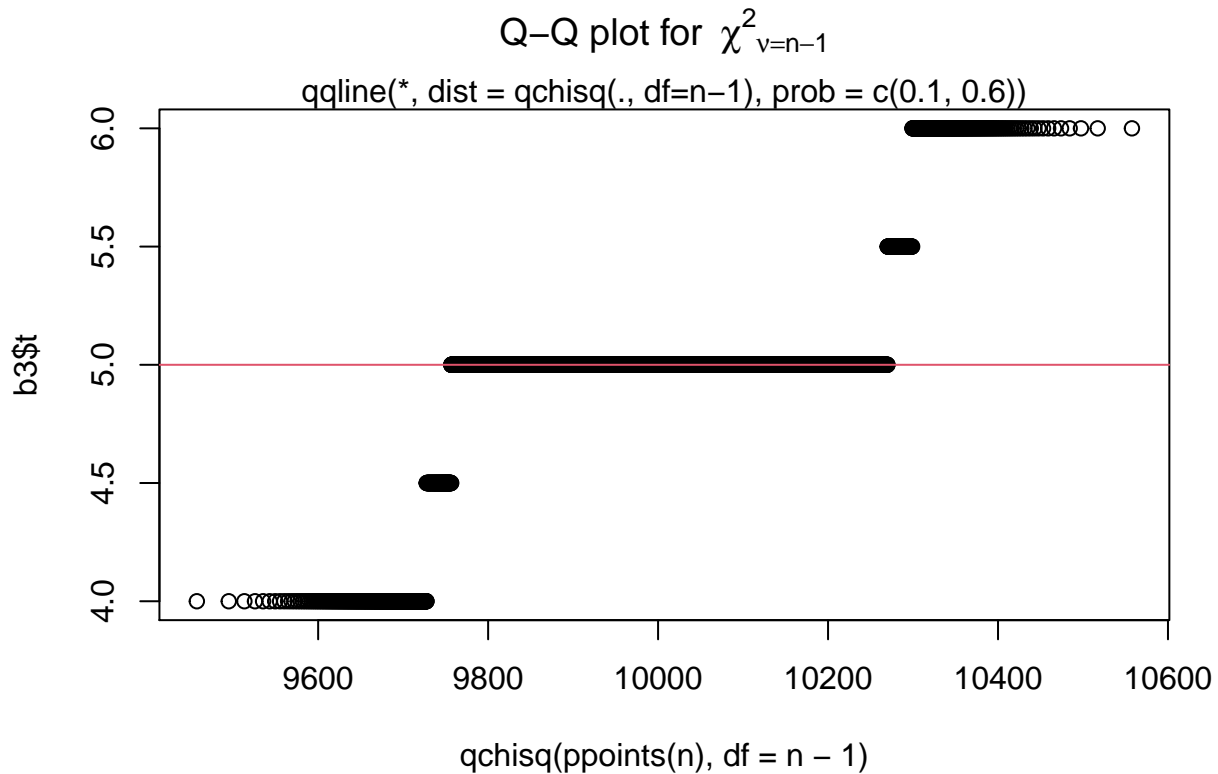
```
plot(b3)
```



```
### Display a quantile-quantile plot using a chiquare
b3$t[1:15] ### The boot statistics are stored in a variable, t, within the the boot object.
```

```
## [1] 5.0 5.0 5.0 5.0 4.5 5.0 5.0 5.0 5.0 5.0 5.0 4.0 5.0 5.0 5.0
```

```
n <- length(b3$t)
qqplot(qchisq(ppoints(n), df = n-1), b3$t,
       main = expression("Q-Q plot for" ~ ~ {chi^2}[nu == n-1]))
qqline(b3$t, qtype = 5, distribution = function(p){qchisq(p, df = n-1)},
       probs = c(0.1, 0.6), col = 2)
mtext("qqline(*, dist = qchisq(., df=n-1), prob = c(0.1, 0.6))")
```



More complex statistics can be bootstrapped. We look at the mean ratio of weight to height for college students.

```
### Get the data
htwt <- read.csv("http://facweb1.redlands.edu/fac/jim_bentley/downloads/math111/htwt.csv")
head(htwt)
```


```
##   Height Weight Group
## 1     64    159     1
## 2     63    155     2
## 3     67    157     2
## 4     60    125     1
## 5     52    103     2
## 6     58    122     2
```

```
### Usual bootstrap of the ratio of means
ratio <- function(d, i){ mean(d$Weight[i] / d$Height[i]) }
b4 <- boot(htwt, ratio, R = 9999)
b4
```

```
##
## ORDINARY NONPARAMETRIC BOOTSTRAP
##
##
## Call:
## boot(data = htwt, statistic = ratio, R = 9999)
##
##
```

```
## Bootstrap Statistics :
##   original      bias  std. error
## t1*  2.20046 -0.0003394999  0.08440419
```

```
plot(b4)
```




```
### Display a quantile-quantile plot using a gamma --- which it probably isn't
b4$t[1:15] ### The boot statistics are stored in a variable, t, within the boot object.
```

```
## [1] 2.307540 2.054809 2.099453 2.272124 2.296053 2.166432 2.083987 2.248138
## [9] 2.104033 2.196612 2.403332 2.301801 2.076991 2.117183 2.263617
```

```
n <- length(b4$t)
s <- var(b4$t)/mean(b4$t) ### Estimate the scale parameter
a <- mean(b4$t)/s ### Estimate the shape
qqplot(qgamma(ppoints(n), shape = a, scale = s), b4$t,
       main = paste0("Q-Q plot for Gamma(shape = ", round(a,2), ", scale = ", round(s,4), ")")
)
qqline(b4$t, qtype = 5, distribution = function(p){qgamma(p, shape = a, scale = s)},
       probs = c(0.1, 0.6), col = 2)
```

**Q-Q plot for Gamma(shape = 679.46, scale = 0.0032)**



Major parts of this document were taken from UCLA's stat packages web: <https://stats.idre.ucla.edu/r/faq/how-can-i-generate-bootstrap-statistics-in-r/>