

Exercise 3.1, page 9

Since $X_1, \dots, X_N \stackrel{iid}{\sim} \mathcal{N}(\mu, \sigma^2)$,

$$\hat{\mu} \sim \mathcal{N}(\mu, \sigma^2/N). \quad (22.18)$$

Hypothesis tests for the mean are based on

$$\frac{\text{observed} - \text{hypothesized}}{\text{standard error}} = \frac{\hat{\mu} - \mu}{\sigma/\sqrt{N}} \sim \mathcal{N}(0, 1). \quad (22.19)$$

Replacing σ^2 by the sample estimate $\hat{\sigma}^2$ leads to the t-distribution:

$$t = \frac{\hat{\mu} - \mu}{\hat{\sigma}/\sqrt{N}} \sim t_{N-1}. \quad (22.20)$$

By definition of $t_{\alpha/2, N-1}$,

$$1 - \alpha = P(-t_{\alpha/2, N-1} \leq t < t_{\alpha/2, N-1}) \quad (22.21)$$

$$= P\left(-t_{\alpha/2, N-1} \leq \frac{\hat{\mu} - \mu}{\hat{\sigma}/\sqrt{N}} < t_{\alpha/2, N-1}\right) \quad (22.22)$$

$$= P\left(-t_{\alpha/2, N-1} \leq \frac{\hat{\mu} - \mu}{\hat{\sigma}/\sqrt{N}} \text{ and } \frac{\hat{\mu} - \mu}{\hat{\sigma}/\sqrt{N}} < t_{\alpha/2, N-1}\right) \quad (22.23)$$

$$(22.24)$$

We want to solve for μ by itself:

$$1 - \alpha = P\left(\mu \leq \hat{\mu} + t_{\alpha/2, N-1} \hat{\sigma}/\sqrt{N} \text{ and } \hat{\mu} - t_{\alpha/2, N-1} \hat{\sigma}/\sqrt{N} < \mu\right) \quad (22.25)$$

$$= P\left(\hat{\mu} - t_{\alpha/2, N-1} \hat{\sigma}/\sqrt{N} < \mu \leq \hat{\mu} + t_{\alpha/2, N-1} \hat{\sigma}/\sqrt{N}\right). \quad (22.26)$$

An R code for computing the coverage of this confidence interval is

```

1 ntrials  = 100000
2 nsamp    = 20
3 mu       = 10
4 stdv    = 3
5 alpha   = 0.05
6
7 x        = array(rnorm(nsamp*ntrials,mean=mu,sd=stdv),dim=c(ntrials,nsamp))
8 x.mean  = rowMeans(x)
9 x.stdv  = sqrt(rowSums((x-x.mean)^2)/(nsamp-1)/nsamp)
10 tcrit   = qt(alpha/2,nsamp-1,lower.tail=FALSE)
11
12 mean.ci = cbind(x.mean-tcrit*x.stdv,x.mean+tcrit*x.stdv)
13 coverage = sum(mu >= mean.ci[,1] & mu <= mean.ci[,2])/ntrials
14
15 print(paste('coverage=',coverage))
16 [1] "coverage= 0.94961"

```

The coverage is very close to 95%.

Exercise 3.2, page 10

An augmented `var.equal.test` that calculates a confidence interval is the following:

```

1 var.equal.test = function(data1,data2,alpha=0.05) {
2   ### THIS FUNCTION TESTS EQUALITY OF VARIANCE OF TWO
3   ### IID NORMALLY DISTRIBUTED RANDOM VARIABLES
4   ###
5   # INPUT:
6   #   DATA1: [N1]-DIMENSIONAL VECTOR OF DATA
7   #   DATA2: [N2]-DIMENSIONAL VECTOR OF DATA
8   #   ALPHA: SIGNIFICANCE LEVEL OF THE TEST (DEFAULT = 5%)
9   # OUTPUT LIST:
10  #   F.MAX: RATIO OF VARIANCE, CONSTRUCTED TO BE GREATER THAN 1
11  #   F.CRIT: THE UPPER CRITICAL THRESHOLD OF SIGNIFICANCE
12  #   PVAL: P-VALUE OF THE F.MAX RATIO
13  #   VAR1: UNBIASED ESTIMATE OF THE VARIANCE OF DATA1
14  #   VAR2: UNBIASED ESTIMATE OF THE VARIANCE OF DATA2
15  #   RATIO: VAR1/VAR2
16  #   RATIO.LOWER: LOWER LIMIT OF ALPHA% CONFIDENCE INTERVAL FOR VARIANCE RATIO
17  #   RATIO.UPPER: UPPER LIMIT OF ALPHA% CONFIDENCE INTERVAL FOR VARIANCE RATIO
18
19  n1      = length(data1); n2      = length(data2)
20
21  mean1  = sum(data1)/n1; mean2  = sum(data2)/n2
22
23  var1   = sum((data1-mean1)^2)/(n1-1)
24  var2   = sum((data2-mean2)^2)/(n2-1)
25
26  if ( var1 > var2) {
27    f.max  = var1/var2
28    f.crit = qf(alpha/2,n1-1,n2-1,lower.tail=FALSE)
29    pval   = 2*pf(f.max,n1-1,n2-1,lower.tail=FALSE)
30  } else {
31    f.max  = var2/var1
32    f.crit = qf(alpha/2,n2-1,n1-1,lower.tail=FALSE)
33    pval   = 2*pf(f.max,n2-1,n1-1,lower.tail=FALSE)
34  }
35
36  f.upper = qf(alpha/2,n1-1,n2-1,lower.tail=FALSE)
37  f.lower = qf(alpha/2,n1-1,n2-1,lower.tail=TRUE)
38
39  ratio      = var1/var2
40  ratio.lower = ratio/f.upper
41  ratio.upper = ratio/f.lower
42
43  list(f.max=f.max,f.crit=f.crit,pval=pval,var1=var1,var2=var2,
44        ratio=ratio,ratio.lower=ratio.lower,ratio.upper=ratio.upper)
45 }
```

Note that the critical F values `f.upper` and `f.lower` were recomputed after the `if` statement so that they correspond to the ratio `var1/var2` and not its reciprocal.

The output of this function for the PDO index is the following

```
1 > nbreak = which( year == 1977)
2 > pdo1 = pdo[1:nbreak]
3 > pdo2 = pdo[(nbreak+1):nyrs]
4 >
5 > ##### TEST EQUALITY OF VARIANCE #####
6 > ## TEST EQUALITY OF VARIANCE
7 > #####
8 > source(paste('var.equal.test.R',sep=""))
9 > print(var.equal.test(pdo1,pdo2))
10 $f.max
11 [1] 1.247894
12
13 $f.crit
14 [1] 2.074389
15
16 $pval
17 [1] 0.5522268
18
19 $var1
20 [1] 0.7420675
21
22 $var2
23 [1] 0.9260218
24
25 $ratio
26 [1] 0.80135
27
28 $ratio.lower
29 [1] 0.4043894
30
31 $ratio.upper
32 [1] 1.662312
```

Or if the data is entered in the opposite direction...

```
1 > var.equal.test(pdo2,pdo1)
2 $f.max
3 [1] 1.247894
4
5 $f.crit
6 [1] 2.074389
7
8 $pval
9 [1] 0.5522268
10
11 $var1
12 [1] 0.9260218
13
14 $var2
15 [1] 0.7420675
16
17 $ratio
18 [1] 1.247894
19
20 $ratio.lower
21 [1] 0.6015718
22
23 $ratio.upper
24 [1] 2.472864
```

Accordingly, the 95% confidence interval for the ratio of variances is (0.39, 1.63), or (0.61, 2.54), depending on how the two variances are ordered in the ratio. Note that this interval includes one, implying that the ratio is not significantly different from one at the 5% significance level, consistent with the p-value.

Exercise 3.3, page 10

A modified version of `cor.equal.test` that also computes a confidence interval is the following:

```

1 cor.equal.test = function(data1,data2,alpha=0.05) {
2   ## THIS FUNCTION TESTS VANISHING CORRELATION BETWEEN
3   ## TWO BI-VARIATE, NORMALLY DISTRIBUTED RANDOM VARIABLES
4   ##
5   # INPUT:
6   #   DATA1: [N]-DIMENSIONAL VECTOR OF DATA
7   #   DATA2: [N]-DIMENSIONAL VECTOR OF DATA
8   #   ALPHA: SIGNIFICANCE LEVEL OF THE TEST (DEFAULT = 5%)
9   # OUTPUT LIST:
10  #   RHO: SAMPLE CORRELATION BETWEEN DATA1 AND DATA2
11  #   RHO.CRIT: 100*ALPHA% CRITICAL VALUE FOR THE CORRELATION
12  #   PVAL: P-VALUE OF THE STATISTIC RHO
13  #   RHO.LIMITS: [2] (1-ALPHA/2)*100% CONFIDENCE LIMITS OF THE CORRELATION
14
15  if ( length(data1) != length(data2)) stop('data sets are not the same length')
16
17  ntot      = length(data1)
18
19  mean1    = sum(data1)/ntot
20  mean2    = sum(data2)/ntot
21
22  ssr1     = sum((data1-mean1)^2)
23  ssr2     = sum((data2-mean2)^2)
24  sscx     = sum((data1-mean1)*(data2-mean2))
25
26  rho.hat  = sscx / sqrt(ssr1*ssr2)
27
28  t.crit   = qt(alpha/2,ntot-2,lower.tail=FALSE)
29  rho.crit = t.crit / sqrt( ntot - 2 + t.crit^2)
30
31  tval     = rho.hat * sqrt(ntot-2) / sqrt(1-rho.hat^2)
32  pval    = 2*pt(abs(tval),ntot-2,lower.tail=FALSE)
33
34  z        = 0.5 * log((1+rho.hat)/(1-rho.hat))
35  zc       = qnorm(alpha/2,mean=0,sd=1/sqrt(ntot-3),lower.tail=FALSE)
36  z.limits = c(z - zc, z+zc)
37
38  rho.limits = (exp(2*z.limits)-1)/(exp(2*z.limits)+1)
39
40  list(rho=rho.hat,rho.crit=rho.crit,pval=pval,tval=tval,rho.limits=rho.limits)
41 }
```

Applying this function to the PDO data gives the following

```
1 > ## test vanishing correlation for first half
2 > 10 = 1:(length(pdo1)-1)
3 > pdo1.10 = pdo1[10]; pdo1.11 = pdo1[10+1]
4 > print(cor.equal.test(pdo1.10,pdo1.11))
5 $rho
6 [1] 0.05495659
7
8 $rho.crit
9 [1] 0.3808629
10
11 $pval
12 [1] 0.7854252
13
14 $tval
15 [1] 0.2751988
16
17 $rho.limits
18 [1] -0.3319908 0.4260723
```

```
1 > ## test vanishing correlation for 2nd half
2 > 10 = 1:(length(pdo2)-1)
3 > pdo2.10 = pdo2[10]; pdo2.11 = pdo2[10+1]
4 > print(cor.equal.test(pdo2.10,pdo2.11))
5 $rho
6 [1] 0.3836199
7
8 $rho.crit
9 [1] 0.3160319
10
11 $pval
12 [1] 0.01591744
13
14 $tval
15 [1] 2.526791
16
17 $rho.limits
18 [1] 0.07748101 0.62365106
```

```
1 > ## test vanishing correlation for whole period
2 > l0 = 1:(length(pdo)-1)
3 > pdo.l0 = pdo[l0]; pdo.l1 = pdo[l0+1]
4 > print(cor.equal.test(pdo.l0,pdo.l1))
5 $rho
6 [1] 0.4915623
7
8 $rho.crit
9 [1] 0.2404471
10
11 $pval
12 [1] 2.398144e-05
13
14 $tval
15 [1] 4.550883
16
17 $rho.limits
18 [1] 0.2850066 0.6544904
```

The first confidence interval includes 0, implying that the correlation is not significantly different from zero. The second confidence interval does not include 0, but only marginally. The third confidence interval does not include 0, hence the correlation is significantly different from zero. All three conclusions are consistent with the corresponding hypothesis tests.

Exercise 3.4, page 11

A modified version of `mean.equal.test` that also computes confidence limits for the difference in means is the following:

```

1 mean.equal.test = function(data1,data2,alpha=0.05) {
2   ## THIS FUNCTION TESTS EQUALITY OF MEANS OF TWO IID
3   ## NORMALLY DISTRIBUTED RANDOM VARIABLES
4   ##
5   # INPUT:
6   #   DATA1: [N1]-DIMENSIONAL VECTOR OF DATA
7   #   DATA2: [N2]-DIMENSIONAL VECTOR OF DATA
8   #   ALPHA: SIGNIFICANCE LEVEL OF THE TEST (DEFAULT = 5%)
9   # OUTPUT LIST:
10  #   DIFF.MEAN: DIFFERENCE IN MEANS (MEAN1 - MEAN2)
11  #   DIFF.MEAN.CRIT: 100*ALPHA% LEVEL CRITICAL VALUE OF THE DIFFERENCE IN MEANS
12  #   PVAL: P-VALUE OF THE T-STATISTIC
13  #   T: T-STATISTIC FOR DIFFERENCE IN MEANS
14  #   T.CRIT: 100*ALPHA% LEVEL CRITICAL VALUE OF THE T STATISTIC
15  #   MEAN1: ESTIMATE OF THE MEAN OF DATA1
16  #   MEAN2: ESTIMATE OF THE MEAN OF DATA2
17  #   SPOOL: POOLED ESTIMATE OF THE STANDARD DEVIATION
18  #   DIFF.MEAN.LIMITS: [2] (1-ALPHA)100% CONFIDENCE LIMITS
19  #       FOR THE DIFFERENCE IN MEANS
20
21 n1      = length(data1)
22 n2      = length(data2)
23 dof     = n1 + n2 - 2
24
25 mean1   = sum(data1)/n1
26 mean2   = sum(data2)/n2
27 diff.mean = mean1 - mean2
28
29 ssr1    = sum((data1-mean1)^2)
30 ssr2    = sum((data2-mean2)^2)
31
32 spool   = sqrt((ssr1 + ssr2)/dof)
33
34 tval    = diff.mean/spool/sqrt(1/n1 + 1/n2)
35 pval    = 2* pt(abs(tval),dof,lower.tail=FALSE)
36 t.crit = qt(alpha/2,dof,lower.tail=FALSE)
37
38 diff.mean.crit = t.crit * spool * sqrt(1/n1 + 1/n2)
39
40 diff.mean.limits = diff.mean + c(-1,1)*diff.mean.crit
41
42 list(diff.mean=diff.mean,diff.mean.crit=diff.mean.crit,pval=pval,
43       tval=tval,t.crit=t.crit,mean1=mean1,mean2=mean2,spool=spool,
44       diff.mean.limits=diff.mean.limits)
45 }
```

Applying this to the PDO data gives

```
1 > print(mean.equal.test(pdo1,pdo2))
2 $diff.mean
3 [1] -1.156262
4
5 $diff.mean.crit
6 [1] 0.4537684
7
8 $pval
9 [1] 3.22733e-06
10
11 $tval
12 [1] -5.08751
13
14 $t.crit
15 [1] 1.996564
16
17 $mean1
18 [1] -0.7405952
19
20 $mean2
21 [1] 0.4156667
22
23 $spool
24 [1] 0.9223707
25
26 $diff.mean.limits
27 [1] -1.6100303 -0.7024935
```

The confidence limits do not include 0, so the difference in means is significantly different from zero. This is consistent with the hypothesis test.

Exercise 3.5, page 12

An R function for constructing a confidence interval for a correlation coefficient is the following:

```

1 cor.equal.boot = function(x,y,alpha=0.05,nboot=100000) {
2   ## ESTIMATES (1-ALPHA)100% CONFIDENCE INTERVAL
3   ## FOR THE CORRELATION COEFFICIENT USING BOOTSTRAP METHODS
4   ## INPUT:
5   ## X[NSAMP]: FIRST DATA SET
6   ## Y[NSAMP]: SECOND DATA SET
7   ## ALPHA: EQUIVALENT SIGNIFICANCE LEVEL (DEFAULT = 5%)
8   ## NBOOT: NUMBER OF BOOTSTRAP SAMPLES (DEFAULT = 100000)
9   ## OUTPUT:
10  ## CONFIDENCE LIMITS[2]
11  nsamp = length(x)
12  if (length(y) != nsamp) stop('x and y must be same length')
13
14 npic = sample(nsamp,nboot*nsamp,replace=TRUE)
15 x.boot = x[npic]
16 y.boot = y[npic]
17 dim(x.boot) = c(nboot,nsamp)
18 dim(y.boot) = c(nboot,nsamp)
19 x.boot = x.boot - rowMeans(x.boot)
20 y.boot = y.boot - rowMeans(y.boot)
21 cor.boot = rowSums(x.boot*y.boot)/sqrt(rowSums(x.boot^2)*rowSums(y.boot^2))
22 cor.boot.conf = quantile(cor.boot,probs=c(alpha/2,1-alpha/2))
23
24 cor.boot.conf
25 }
```

The result of running this function is:

```
1 nsamp    = 20
2 rho      = 0.5
3 nboot   = 100000
4 set.seed(310)
5 x         = rnorm(nsamp)
6 w         = rnorm(nsamp)
7 y         = rho * x + sqrt(1-rho^2)*w
8
9 rho.test1 = cor.equal.test(x,y)
10 rho.test2 = cor.test(x,y)
11
12 print(paste('Traditional Confidence interval:',
13   paste(signif(rho.test1$rho.limits,4),collapse=', ')))
14 print(paste('Bootstrap confidence interval  :',
15   paste(signif(cor.boot.conf,4),collapse=', ')))
16
17
18
19 [1] "Traditional Confidence interval: 0.3301, 0.8601"
20 [1] "Bootstrap confidence interval  : 0.328, 0.8761"
```

The bootstrap interval is nearly identical to the standard interval.