

Lecture: Inference about sample mean

Key concepts:

1. Hotelling's T^2 test
2. Likelihood ratio test
3. Various confidence regions
4. Applications
5. Missing values
6. Impact of serial correlations

1 Hypothesis test of a mean vector

Let $\mathbf{x}_1, \dots, \mathbf{x}_n$ be a random sample from a p -dimensional normal population with mean $\boldsymbol{\mu}$ and positive-definite variance-covariance matrix $\boldsymbol{\Sigma}$. Consider the testing problem: $H_o : \boldsymbol{\mu} = \boldsymbol{\mu}_o$ versus $H_a : \boldsymbol{\mu} \neq \boldsymbol{\mu}_o$, where $\boldsymbol{\mu}_o$ is a known vector.

This is a generalization of the one-sample t test of the univariate case. When $p = 1$, the test statistic is

$$t = \frac{\bar{x} - \mu_o}{s/\sqrt{n}}, \quad \text{with} \quad s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2.$$

The statistic follows a Student- t distribution with $n - 1$ degrees of freedom. One rejects H_o if the p -value of t is less than the type-I error denoted by α . For generalization, we rewrite the test as

$$t^2 = (\bar{x} - \mu_o) \left(\frac{s^2}{n} \right)^{-1} (\bar{x} - \mu_o).$$

One rejects H_o if and only if $t^2 \geq t_{n-1}^2(\alpha/2)$, the upper 100($\alpha/2$) percentile of t -distribution with $n - 1$ degrees of freedom. A natural generalization of this test statistic is

$$T^2 = (\bar{\mathbf{x}} - \boldsymbol{\mu}_o)' \left(\frac{\mathbf{S}}{n} \right)^{-1} (\bar{\mathbf{x}} - \boldsymbol{\mu}_o), \tag{1}$$

where $\mathbf{S} = \frac{1}{n-1} \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \boldsymbol{\mu}_o)'$ is the sample covariance matrix.

This is the Hotelling's T^2 statistic. It is distributed as $\frac{(n-1)p}{n-p} F_{p,n-p}$, where $F_{u,v}$ denotes the F -distribution with degrees of freedom u and v . Recall that $[t_{n-1}(\alpha/2)]^2 = F_{1,n-1}(\alpha)$, where $F_{u,v}(\alpha)$ is the upper 100α percentile of the F -distribution with u and v degrees of freedom. Thus, when $p = 1$, T^2 reduces to the usual one-sample t statistic.

Example. Consider the monthly log returns of 4 Midwest companies, Boeing, Abbott Labs, Motorola, and General Motors, from 1998 to 2007. The returns are in percentages. Let $\boldsymbol{\mu}_o = \mathbf{0}$. Test the null hypothesis that the mean vector of the log returns is zero.

Answer: Except for three possible outlying observations, the chi-square QQ-plot indicates that the normal assumption is reasonable.

```
> setwd("C:/teaching/ama")
> x=read.table("m-ba4c9807.txt")

> source("hotelling.txt")
> hotelling(x,rep(0,4))
      [,1]
Hotelling-T2 2.1074254
p.value      0.7258692

%% More details
> xb=mean(x)
> S=cov(x)
> Si=solve(S)
> Tsq=120*t(xb)%*%Si%*%xb
> Tsq
      [,1]
[1,] 2.107425
> xb
      V1      V2      V3      V4
-0.26294605  0.61075048  0.67154781  0.02679363
> S
      V1      V2      V3      V4
V1 127.011719 28.776911  8.046536  50.919252
V2 28.776911  80.328617  8.385634  13.216236
V3  8.046536  8.385634  40.750353  -8.710395
V4 50.919252 13.216236 -8.710395 122.712745
```

Remark: Let $\mathbf{y} = \mathbf{C}\mathbf{x} + \mathbf{d}$, where \mathbf{C} is a $p \times p$ non-singular matrix and \mathbf{d} is a p -dimensional vector. Then, testing $H_o : \boldsymbol{\mu} = \boldsymbol{\mu}_o$ of \mathbf{x} is equivalent to testing $H_o : \boldsymbol{\mu}_y = \mathbf{C}\boldsymbol{\mu}_o + \mathbf{d}$ of \mathbf{y} . It turns out that, as expected, the Hotelling T^2 statistic is identical.

Proof: $\bar{\mathbf{y}} = \mathbf{C}\bar{\mathbf{x}} + \mathbf{d}$ and $\mathbf{S}_y = \mathbf{C}\mathbf{S}\mathbf{C}'$.

$$T^2 = n(\bar{\mathbf{y}} - \boldsymbol{\mu}_{y_o})' \mathbf{S}_y^{-1} (\bar{\mathbf{y}} - \boldsymbol{\mu}_{y_o})$$

$$\begin{aligned}
&= n(\mathbf{C}\bar{\mathbf{x}} + \mathbf{d} - \mathbf{C}\boldsymbol{\mu}_o - \mathbf{d})'(\mathbf{C}\mathbf{S}\mathbf{C}')^{-1}(\mathbf{C}\bar{\mathbf{x}} + \mathbf{d} - \mathbf{C}\boldsymbol{\mu}_o - \mathbf{d}) \\
&= n[\mathbf{C}(\bar{\mathbf{x}} - \boldsymbol{\mu}_o)]'\mathbf{C}'^{-1}\mathbf{S}^{-1}\mathbf{C}^{-1}[\mathbf{C}(\bar{\mathbf{x}} - \boldsymbol{\mu}_o)] \\
&= n(\bar{\mathbf{x}} - \boldsymbol{\mu}_o)'\mathbf{S}^{-1}(\bar{\mathbf{x}} - \boldsymbol{\mu}_o).
\end{aligned}$$

2 Likelihood ratio test

Let $\boldsymbol{\theta}$ be the parameter vector of a likelihood function $L(\boldsymbol{\theta})$ with observations $\mathbf{x}_1, \dots, \mathbf{x}_n$. Suppose that the parameter space is Θ . Consider the null hypothesis $\boldsymbol{\theta} \in \Theta_o$ versus $H_a : \boldsymbol{\theta} \ni \Theta_o$, where Θ_o is a subspace of Θ . The likelihood ratio statistic is

$$\Lambda = \frac{\max_{\boldsymbol{\theta} \in \Theta_o} L(\boldsymbol{\theta})}{\max_{\boldsymbol{\theta} \in \Theta} L(\boldsymbol{\theta})}. \quad (2)$$

One rejects H_o if $\Lambda < c$, where c is some critical value depending on the type-I error. Intuitively, one rejects H_o if the maximized likelihood function over the subspace Θ_o is much smaller than that over the parameter space Θ , indicating that it is unlikely for $\boldsymbol{\theta}$ to be in Θ_o . Specifically, under the null hypothesis H_o , the likelihood ratio statistic

$$-2 \ln(\Lambda) \sim \chi_{v-v_o}^2, \quad \text{as } n \rightarrow \infty,$$

where $v = \dim(\Theta)$ and $v_o = \dim(\Theta_o)$.

Remark: $\Theta_o \subset \Theta$ indicates that the null model is nested in the alternative model. For non-nested models, likelihood ratio test does not apply. [See some recent papers on the generalized likelihood ratio test.]

For multivariate normal distribution, the likelihood ratio statistic is relatively simple because the limiting distribution of Λ is available. Specifically, consider $H_o : \boldsymbol{\mu} = \boldsymbol{\mu}_o$ versus $H_a : \boldsymbol{\mu} \neq \boldsymbol{\mu}_o$. Here $\boldsymbol{\theta} = (\boldsymbol{\mu}', \sigma_{11}, \sigma_{21}, \sigma_{22}, \dots, \sigma_{p1}, \dots, \sigma_{pp})'$ is of dimension $p + p(p+1)/2$. Under H_o , $\boldsymbol{\mu}$ is fixed at $\boldsymbol{\mu}_o$ so that Θ_o consists of the space for the elements of $\boldsymbol{\Sigma}$. In this case,

$$\Lambda = \frac{\max_{\boldsymbol{\Sigma}} L(\boldsymbol{\mu}_o, \boldsymbol{\Sigma})}{\max_{\boldsymbol{\mu}, \boldsymbol{\Sigma}} L(\boldsymbol{\mu}, \boldsymbol{\Sigma})}.$$

Recall that, under Θ ,

$$\max_{\boldsymbol{\mu}, \boldsymbol{\Sigma}} L(\boldsymbol{\mu}, \boldsymbol{\Sigma}) = L(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}}) = \frac{1}{(2\pi)^{np/2} |\hat{\boldsymbol{\Sigma}}|^{n/2}} e^{-np/2},$$

where $\hat{\boldsymbol{\mu}} = \bar{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i$ and $\hat{\boldsymbol{\Sigma}} = \frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})'$. Under Θ_o , the likelihood function becomes

$$L(\boldsymbol{\mu}_o, \boldsymbol{\Sigma}) = \frac{1}{(2\pi)^{np/2} |\boldsymbol{\Sigma}|^{n/2}} \exp \left[-\frac{1}{2} \sum_{i=1}^n (\mathbf{x}_i - \boldsymbol{\mu}_o)' \boldsymbol{\Sigma}^{-1} (\mathbf{x}_i - \boldsymbol{\mu}_o) \right].$$

By the same method as before, we can show that

$$\max_{\Sigma} L(\boldsymbol{\mu}_o, \Sigma) = L(\boldsymbol{\mu}_o, \widehat{\Sigma}_o) = \frac{1}{(2\pi)^{np/2} |\widehat{\Sigma}_o|^{n/2}} e^{-np/2},$$

where $\widehat{\Sigma}_o = \frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i - \boldsymbol{\mu}_o)(\mathbf{x}_i - \boldsymbol{\mu}_o)'$. Consequently,

$$\Lambda = \left(\frac{|\widehat{\Sigma}|}{|\widehat{\Sigma}_o|} \right)^{n/2}.$$

The statistic $\Lambda^{2/n} = |\widehat{\Sigma}|/|\widehat{\Sigma}_o|$ is called the *Wilks' lambda*.

Result 5.1. Let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be a random sample from an $N_p(\boldsymbol{\mu}, \Sigma)$ population. Consider the hypothesis testing: $H_o : \boldsymbol{\mu} = \boldsymbol{\mu}_o$ versus $H_a : \boldsymbol{\mu} \neq \boldsymbol{\mu}_o$. Then, the likelihood ratio statistic becomes

$$\Lambda^{2/n} = \left(1 + \frac{T^2}{n-1} \right)^{-1},$$

where T^2 is the Hotelling's T^2 statistic.

Proof. The following equality of determinants is useful

$$\begin{aligned} & (-1) \left| \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})' + n(\bar{\mathbf{x}} - \boldsymbol{\mu}_o)(\bar{\mathbf{x}} - \boldsymbol{\mu}_o)' \right| \\ &= \left| \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})' \right| \left| -1 - n(\bar{\mathbf{x}} - \boldsymbol{\mu}_o)' \left(\sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})' \right)^{-1} (\bar{\mathbf{x}} - \boldsymbol{\mu}_o) \right|. \end{aligned}$$

This follows directly from the identity $|\mathbf{A}| = |\mathbf{A}_{22}| |\mathbf{A}_{11} - \mathbf{A}_{12} \mathbf{A}_{22}^{-1} \mathbf{A}_{21}| = |\mathbf{A}_{11}| |\mathbf{A}_{22} - \mathbf{A}_{21} \mathbf{A}_{11}^{-1} \mathbf{A}_{12}|$ with

$$\mathbf{A} = \begin{bmatrix} \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})' & \sqrt{n}(\bar{\mathbf{x}} - \boldsymbol{\mu}_o) \\ \sqrt{n}(\bar{\mathbf{x}} - \boldsymbol{\mu}_o)' & -1 \end{bmatrix} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix}.$$

Note also that

$$\sum_{i=1}^n (\mathbf{x}_i - \boldsymbol{\mu}_o)(\mathbf{x}_i - \boldsymbol{\mu}_o)' = \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})' + n(\bar{\mathbf{x}} - \boldsymbol{\mu}_o)(\bar{\mathbf{x}} - \boldsymbol{\mu}_o)'$$

Therefore,

$$(-1) \left| \sum_{i=1}^n (\mathbf{x}_i - \boldsymbol{\mu}_o)(\mathbf{x}_i - \boldsymbol{\mu}_o)' \right| = \left| \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})' \right| (-1) \left(1 + \frac{T^2}{n-1} \right).$$

That is,

$$|n\widehat{\Sigma}_o| = |n\widehat{\Sigma}| \left(1 + \frac{T^2}{n-1} \right).$$

Consequently,

$$\Lambda^{2/n} = \frac{|\widehat{\Sigma}|}{|\widehat{\Sigma}_o|} = \left(1 + \frac{T^2}{n-1}\right)^{-1}.$$

Remark: The prior result suggests that there is no need to calculate \mathbf{S}^{-1} in computing the Hotelling's T^2 statistic. Indeed,

$$T^2 = \frac{(n-1)|\widehat{\Sigma}_o|}{|\widehat{\Sigma}|} - (n-1).$$

3 Confidence region

Let $\boldsymbol{\theta}$ be the parameter vector and Θ be the parameter space. For a random sample $\mathbf{X}_1, \dots, \mathbf{X}_n$, let \mathbf{X} denote the data. A $100(1 - \alpha)\%$ confidence region $R(\mathbf{X})$ is defined as

$$Pr(R(\mathbf{X}) \text{ will cover the true } \boldsymbol{\theta}) = 1 - \alpha,$$

where the probability is evaluated at the unknown true parameter $\boldsymbol{\theta}$.

For the mean vector $\boldsymbol{\mu}$ of a multivariate normal distribution, the confidence region (C.R.) can be obtained by the result

$$T^2 \sim \frac{(n-1)p}{n-p} F_{p, n-p}.$$

That is,

$$Pr \left[n(\bar{\mathbf{X}} - \boldsymbol{\mu})' \mathbf{S}^{-1} (\bar{\mathbf{X}} - \boldsymbol{\mu}) \leq \frac{(n-1)p}{n-p} F_{p, n-p}(\alpha) \right] = 1 - \alpha,$$

whatever the values of $\boldsymbol{\mu}$ and Σ , where $\mathbf{S} = \frac{1}{n-1} \sum_{i=1}^n (\mathbf{X}_i - \bar{\mathbf{X}})(\mathbf{X}_i - \bar{\mathbf{X}})'$. In other words, if we measure the distance using the variance-covariance matrix $\frac{1}{n}\mathbf{S}$, then $\bar{\mathbf{X}}$ will be within $[(n-1)pF_{p, n-p}(\alpha)/(n-p)]^{1/2}$ of $\boldsymbol{\mu}$.

Remark: The quantity $\sqrt{(\bar{\mathbf{X}} - \boldsymbol{\mu})' (\mathbf{S}/n)^{-1} (\bar{\mathbf{X}} - \boldsymbol{\mu})}$ can be considered as the *Mahalanobis distance* of $\boldsymbol{\mu}$ from $\bar{\mathbf{X}}$, because the covariance matrix of $\bar{\mathbf{X}}$ is $\frac{1}{n}\Sigma$, which is consistently estimated by $\frac{1}{n}\mathbf{S}$. Compared with the Euclidean distance, Mahalanobis distance takes into consideration the covariance structure. It is a distance measure based on correlations between variables.

For normal distribution, C.R. can be viewed as ellipsoids centered at $\boldsymbol{\mu}$ and have axes determined by the eigenvalues and eigenvectors of \mathbf{S} . For $N_p(\boldsymbol{\mu}, \Sigma)$ distribution, the contours of constant density are ellipsoids defined by \mathbf{x} such that

$$(\mathbf{x} - \boldsymbol{\mu})' \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu}) = c^2.$$

These ellipsoids are centered at $\boldsymbol{\mu}$ and have axes $\pm c\sqrt{\lambda_i}\mathbf{e}_i$, where $\Sigma\mathbf{e}_i = \lambda_i\mathbf{e}_i$ for $i = 1, \dots, p$. Consequently, the C.R. are centered at $\bar{\mathbf{x}}$ and the axes of the confidence ellipsoid are

$$\pm \sqrt{\lambda_i} \sqrt{\frac{(n-1)p}{n(n-p)} F_{p, n-p}(\alpha)} \mathbf{e}_i,$$

where $\mathbf{S}\mathbf{e}_i = \lambda_i\mathbf{e}_i$ for $i = 1, \dots, p$.

Simultaneous confidence intervals: Suppose that $\mathbf{X} \sim N_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. For any non-zero p -dimensional vector \mathbf{a} , $Z = \mathbf{a}'\mathbf{X} \sim N(\mathbf{a}'\boldsymbol{\mu}, \mathbf{a}'\boldsymbol{\Sigma}\mathbf{a})$. If $\mathbf{x}_1, \dots, \mathbf{x}_n$ are a random sample from \mathbf{X} , $\{z_i = \mathbf{a}'\mathbf{x}_i\}$ is a random sample of Z . The sample mean and variance are $\bar{z} = \mathbf{a}'\bar{\mathbf{x}}$ and $s_z^2 = \mathbf{a}'\mathbf{S}\mathbf{a}$, where \mathbf{S} is the sample covariance matrix of \mathbf{x}_i .

If \mathbf{a} is fixed, the $100(1 - \alpha)\%$ confidence interval of $\mathbf{a}'\boldsymbol{\mu}$ is

$$\mathbf{a}'\bar{\mathbf{x}} - t_{n-1}(\alpha/2) \frac{\sqrt{\mathbf{a}'\mathbf{S}\mathbf{a}}}{\sqrt{n}} \leq \mathbf{a}'\boldsymbol{\mu} \leq \mathbf{a}'\bar{\mathbf{x}} + t_{n-1}(\alpha/2) \frac{\sqrt{\mathbf{a}'\mathbf{S}\mathbf{a}}}{\sqrt{n}}.$$

This interval consists of $\mathbf{a}'\boldsymbol{\mu}$ values for which

$$|t| = \left| \frac{\sqrt{n}(\mathbf{a}'\bar{\mathbf{x}} - \mathbf{a}'\boldsymbol{\mu})}{\sqrt{\mathbf{a}'\mathbf{S}\mathbf{a}}} \right| \leq t_{n-1}(\alpha/2),$$

or, equivalently,

$$t^2 = \frac{n[\mathbf{a}'(\bar{\mathbf{x}} - \boldsymbol{\mu})]^2}{\mathbf{a}'\mathbf{S}\mathbf{a}} \leq t_{n-1}^2(\alpha/2).$$

If we consider all values of \mathbf{a} for which the prior inequality holds, then we are naturally led to the determination of

$$\max_{\mathbf{a}} t^2 = \max_{\mathbf{a}} \frac{n[\mathbf{a}'(\bar{\mathbf{x}} - \boldsymbol{\mu})]^2}{\mathbf{a}'\mathbf{S}\mathbf{a}}.$$

Using the maximization lemma (see Eq. (2.50) of the text, p. 80) with $\mathbf{x} = \mathbf{a}$, $\mathbf{d} = \bar{\mathbf{x}} - \boldsymbol{\mu}$ and $\mathbf{B} = \mathbf{S}$, we obtain

$$\max_{\mathbf{a}} \frac{n[\mathbf{a}'(\bar{\mathbf{x}} - \boldsymbol{\mu})]^2}{\mathbf{a}'\mathbf{S}\mathbf{a}} = n \left[\max_{\mathbf{a}} \frac{[\mathbf{a}'(\bar{\mathbf{x}} - \boldsymbol{\mu})]^2}{\mathbf{a}'\mathbf{S}\mathbf{a}} \right] = n(\bar{\mathbf{x}} - \boldsymbol{\mu})' \mathbf{S}^{-1} (\bar{\mathbf{x}} - \boldsymbol{\mu}) = T^2,$$

with the maximum occurring for \mathbf{a} proportional to $\mathbf{S}^{-1}(\bar{\mathbf{x}} - \boldsymbol{\mu})$. Consequently, we have the following result.

Result 5.3. Let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be a random sample from an $N_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ population with $\boldsymbol{\Sigma}$ positive definite. Then, simultaneously for all \mathbf{a} , the interval

$$\left(\mathbf{a}'\bar{\mathbf{X}} - \sqrt{\frac{p(n-1)}{n(n-p)} F_{p,n-p}(\alpha) \mathbf{a}'\mathbf{S}\mathbf{a}}, \mathbf{a}'\bar{\mathbf{X}} + \sqrt{\frac{p(n-1)}{n(n-p)} F_{p,n-p}(\alpha) \mathbf{a}'\mathbf{S}\mathbf{a}} \right)$$

will contain $\mathbf{a}'\boldsymbol{\mu}$ with probability $1 - \alpha$.

For convenience, we refer to the simultaneous confidence intervals of Result 5.3 as T^2 -intervals. In particular, the choices of $\mathbf{a} = (1, 0, \dots, 0)'$, $(0, 1, 0, \dots, 0)'$, \dots , $(0, \dots, 0, 1)'$ allow us to conclude that

$$\bar{x}_i - \sqrt{\frac{p(n-1)}{p(n-p)} F_{p,n-p}(\alpha) \frac{s_{ii}}{n}} \leq \mu_i \leq \bar{x}_i + \sqrt{\frac{p(n-1)}{p(n-p)} F_{p,n-p}(\alpha) \frac{s_{ii}}{n}},$$

hold simultaneously with confidence coefficient $1 - \alpha$ for all $i = 1, \dots, p$.

Furthermore, by choosing $\mathbf{a} = (0, \dots, 0, 1, 0, \dots, 0, -1, 0, \dots, 0)'$ with 1 at the i th position and -1 at the k th position, where $i \neq k$, we have $\mathbf{a}'\boldsymbol{\mu} = \mu_i - \mu_k$ and $\mathbf{a}'\mathbf{S}\mathbf{a} = s_{ii} - 2s_{ik} + s_{kk}$, and the T^2 -intervals

$$\bar{x}_i - \bar{x}_k - \sqrt{\frac{p(n-1)}{p(n-p)} F_{p,n-p}(\alpha)} \sqrt{\frac{s_{ii} - 2s_{ik} + s_{kk}}{n}} \leq \mu_i - \mu_k \leq \bar{x}_i - \bar{x}_k + \sqrt{\frac{p(n-1)}{p(n-p)} F_{p,n-p}(\alpha)} \sqrt{\frac{s_{ii} - 2s_{ik} + s_{kk}}{n}}.$$

Comparison. An alternative approach to construct confidence intervals for the component means is to consider the component one at a time. This approach ignores the covariance structure of the variables and leads to the intervals

$$\bar{x}_i - t_{n-1}(\alpha/2) \sqrt{\frac{s_{ii}}{n}} \leq \mu_i \leq \bar{x}_i + t_{n-1}(\alpha/2) \sqrt{\frac{s_{ii}}{n}},$$

for $i = 1, \dots, p$. Here each interval has a confidence coefficient $1 - \alpha$. However, we do not know the confidence coefficient that all intervals contain their respective μ_i , except for the case that the variables X_i are independent. In the latter case, the confidence coefficient is $(1 - \alpha)^p$.

Obviously, the T^2 -interval is much wider than the individual interval for each μ_i .

Bonferroni method for multiple comparisons. In some applications, the T^2 intervals might be too wide to be of practical value and one may only concern about a finite number of linear combinations, say $\mathbf{a}_1, \dots, \mathbf{a}_m$. In this case, we may construct m confidence intervals that are shorter than the T^2 intervals. This method to multiple comparisons is called the *Bonferroni method*.

Let C_i denote a confidence statement about the value of $\mathbf{a}'_i\boldsymbol{\mu}$ with $Pr[C_i \text{ true}] = 1 - \alpha_i$, where $i = 1, \dots, m$. Then,

$$\begin{aligned} Pr[\text{all } C_i \text{ true}] &= 1 - Pr[\text{at least one } C_i \text{ false}] \\ &\geq 1 - \sum_{i=1}^m Pr[C_i \text{ false}] = 1 - \sum_{i=1}^m [1 - Pr(C_i \text{ true})] \\ &= 1 - (\alpha_1 + \dots + \alpha_m). \end{aligned}$$

This is a special case of the Bonferroni inequality and allows us to control the overall type-I error rate $\alpha_1 + \dots + \alpha_m$. The choices of α_i make the method flexible, depending on the prior knowledge on the importance of each interval. If there is no prior information, then $\alpha_i = \alpha/m$ is often used.

Let $z_i = \mathbf{a}'_i\mathbf{x}_i$ and $s_{ii} = \mathbf{a}'_i\mathbf{S}\mathbf{a}_i$. Then, each of the intervals

$$\bar{z}_i \pm t_{n-1}(\alpha/(2m)) \sqrt{\frac{s_{ii}}{n}}, \quad i = 1, \dots, m,$$

contains $\mathbf{a}'\boldsymbol{\mu}$ with probability $1 - \alpha/m$. Jointly, all $\mathbf{a}'\boldsymbol{\mu}$ are in their respective intervals with probability $p \geq 1 - m(\alpha/m) = 1 - \alpha$.

For instance, if $m = p$ and \mathbf{a}_i is the i th unit vector, then the intervals

$$\bar{x}_i - t_{n-1}(\alpha/(2p))\sqrt{\frac{s_{ii}}{n}} \leq \mu_i \leq \bar{x}_i + t_{n-1}(\alpha/(2p))\sqrt{\frac{s_{ii}}{n}}$$

hold with probability at least $1 - \alpha$. We refer to these intervals as the *Bonferroni intervals*. In general, if we use $\alpha_i = \alpha/p$, then we can compare the length of Bonferroni intervals with those of the T^2 intervals. The ratio is

$$\frac{\text{Length of Bonferroni interval}}{\text{Length of } T^2 \text{ interval}} = \frac{t_{n-1}(\alpha/(2p))}{\sqrt{\frac{p(n-1)}{n-p} F_{p,n-p}(\alpha)}}.$$

This ratio does not depend on the sample mean or covariance matrix. Table 5.4 of the text (page 234) gives some result for various n and p .

4 Large sample case

When the sample size is large, we can apply the central limit theory so that the population may not be normal. Recall that, when n is large, $n(\bar{\mathbf{X}} - \boldsymbol{\mu})'\mathbf{S}^{-1}(\bar{\mathbf{X}} - \boldsymbol{\mu}) \sim \chi_p^2$. Thus,

$$Pr[n(\bar{\mathbf{X}} - \boldsymbol{\mu})'\mathbf{S}^{-1}(\bar{\mathbf{X}} - \boldsymbol{\mu}) \leq \chi_p^2(\alpha)] \approx 1 - \alpha,$$

where $\chi_p^2(\alpha)$ is the upper 100α percentile of a chi-square distribution with p degrees of freedom. This property can be used (a) to construct asymptotic confidence intervals for the means μ_i and (b) to perform hypothesis testing about $\boldsymbol{\mu}$.

Asymptotic confidence intervals: Let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be a random sample from a population with mean $\boldsymbol{\mu}$ and positive covariance matrix $\boldsymbol{\Sigma}$. If $n - p$ is sufficiently large,

$$\mathbf{a}'\bar{\mathbf{X}} \pm \sqrt{\chi_p^2(\alpha)}\sqrt{\frac{\mathbf{a}'\mathbf{S}\mathbf{a}}{n}}$$

will contain $\mathbf{a}'\boldsymbol{\mu}$, for every \mathbf{a} , with probability approximately $1 - \alpha$.

Since the chi-square distribution does not depend on \mathbf{a} , the above confidence intervals are simultaneous confidence intervals. This result can be used to compare means of different components of \mathbf{X} .

Asymptotic testing: Let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be a random sample from a population with mean $\boldsymbol{\mu}$ and positive covariance matrix $\boldsymbol{\Sigma}$. When $n - p$ is sufficiently large, the hypothesis $H_o : \boldsymbol{\mu} = \boldsymbol{\mu}_o$ is rejected in favor of $H_a : \boldsymbol{\mu} \neq \boldsymbol{\mu}_o$, at the significance level α if

$$n(\bar{\mathbf{x}} - \boldsymbol{\mu}_o)'\mathbf{S}^{-1}(\bar{\mathbf{x}} - \boldsymbol{\mu}_o) > \chi_p^2(\alpha).$$

Remark: A simple R function is available on the course web that computes various types of confidence intervals for the means of the components of \mathbf{X} . To use the program, do the

following: (a) Download the program into your working directory, (2) use the command `source("r-cregion.txt")` to compile the function. The command `confreg` is ready to use. You can modify the function to construct confidence regions for differences between means of components of \mathbf{X} .

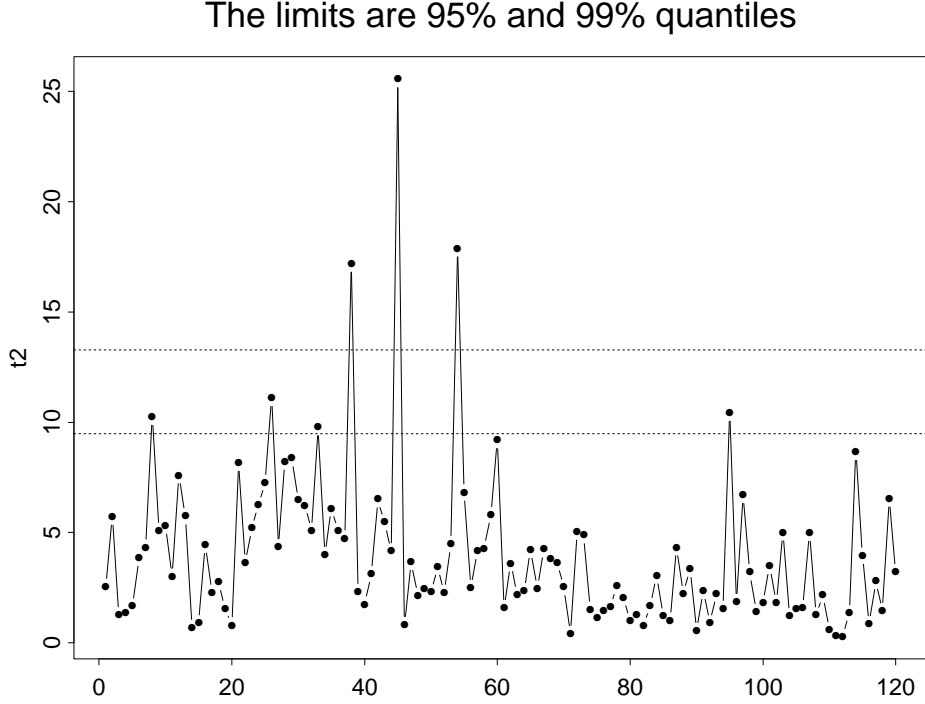
Example: Consider the monthly log returns of four Midwest companies used before. The various confidence intervals are obtained as follows:

```
> setwd("C:/teaching/ama")
> x=read.table("m-ba4c9807.txt")
> source("r-cregion.txt")
> confreg(x)
[1] "C.R. based on T^2"
      [,1]      [,2]
[1,] -3.524903  2.999011
[2,] -1.983378  3.204879
[3,] -1.176112  2.519207
[4,] -3.179484  3.233071
[1] "CR based on individual t"
      [,1]      [,2]
[1,] -2.3000743  1.774182
[2,] -1.0093115  2.230812
[3,] -0.4823362  1.825432
[4,] -1.9755624  2.029150
[1] "CR based on Bonferroni"
      [,1]      [,2]
[1,] -2.8722323  2.346340
[2,] -1.4643301  2.685831
[3,] -0.8064218  2.149517
[4,] -2.5379541  2.591541
[1] "Asymp. simu. CR"
      [,1]      [,2]
[1,] -3.431874  2.905982
[2,] -1.909395  3.130896
[3,] -1.123418  2.466514
[4,] -3.088044  3.141631
>
```

5 Multivariate control charts

We discuss two cases. In the first case, the goal is to identify unusual observations in a sample, including the possibility of drift over time. This is called the T^2 -chart, which uses the distance d_i discussed before in assessing normality assumption. For the i th observation,

Figure 1: T-sq chart for the monthly log returns of four Midwest companies: 1998-2007



the d_i^2 statistic is

$$d_i^2 = (\mathbf{x}_i - \bar{\mathbf{x}})' \mathbf{S}^{-1} (\mathbf{x}_i - \bar{\mathbf{x}}).$$

The upper control limit is then set by the upper quantiles of χ_p^2 . Typically, 95th and 99th percentiles are used to set the upper control limit and the lower control limit is zero. Once a point is found to be outside the control limit, the individual confidence intervals for the component means can be used to identify the source of the deviation.

As an illustration, consider the monthly log returns of the four Midwest companies from 1998 to 2007. The T^2 -chart is given in Figure 1. About three data points are outside of the 99% limit, indicating a volatility period.

In the second case, we consider a control chart for future observations. The theory behind this type of control chart is the following result.

Result 5.6. Let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be independently distributed as $N_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and let \mathbf{X} be a future observation from the same population. Then,

$$T^2 = \frac{n}{n+1} (\mathbf{X} - \bar{\mathbf{X}})' \mathbf{S}^{-1} (\mathbf{X} - \bar{\mathbf{X}}) \sim \frac{p(n-1)}{n-p} F_{p, n-p},$$

and a $100(1 - \alpha)\%$ p -dimensional prediction ellipsoid is given by all \mathbf{x} satisfying

$$(\mathbf{x} - \bar{\mathbf{x}})' \mathbf{S}^{-1} (\mathbf{x} - \bar{\mathbf{x}}) \leq \frac{p(n^2 - 1)}{n(n - p)} F_{p, n-p}(\alpha).$$

Proof: First, $E(\mathbf{X} - \bar{\mathbf{X}}) = \mathbf{0}$. Since \mathbf{X} and $\bar{\mathbf{X}}$ are independent,

$$\text{Cov}(\mathbf{X} - \bar{\mathbf{X}}) = \Sigma + \frac{1}{n} \Sigma = \frac{n + 1}{n} \Sigma.$$

Thus, $\sqrt{n/(n + 1)}(\mathbf{X} - \bar{\mathbf{X}}) \sim N_p(\mathbf{0}, \Sigma)$. Furthermore, by the result in Eq. (5.6) of the textbook,

$$\sqrt{\frac{n}{n + 1}}(\mathbf{X} - \bar{\mathbf{X}})' \mathbf{S}^{-1} \sqrt{\frac{n}{n + 1}}(\mathbf{X} - \bar{\mathbf{X}})$$

has a scaled F distribution as stated.

T^2 -chart for future observations. For each new observation, plot

$$T^2 = \frac{n}{n + 1} (\mathbf{x} - \bar{\mathbf{x}})' \mathbf{S}^{-1} (\mathbf{x} - \bar{\mathbf{x}}),$$

against time order. Set the lower control limit to zero and the upper control limit as

$$UCL = \frac{(n - 1)p}{n - p} F_{p, n-p}(0.01).$$

For illustration, consider the monthly log return series of the four Midwest companies. We start with initial $n = 40$ observations. The T^2 -chart for future observations is given in Figure 2.

Remark: R functions for the two control charts are available on the course web. I wrote those functions quickly; you may want to check them for accuracy.

6 Missing values in normal random sample

A key assumption: Missing at random.

Two methods are available:

1. The EM algorithm: Dempster, Laird and Rubin (1977, JRSSB)
2. Markov chain Monte Carlo (MCMC) method

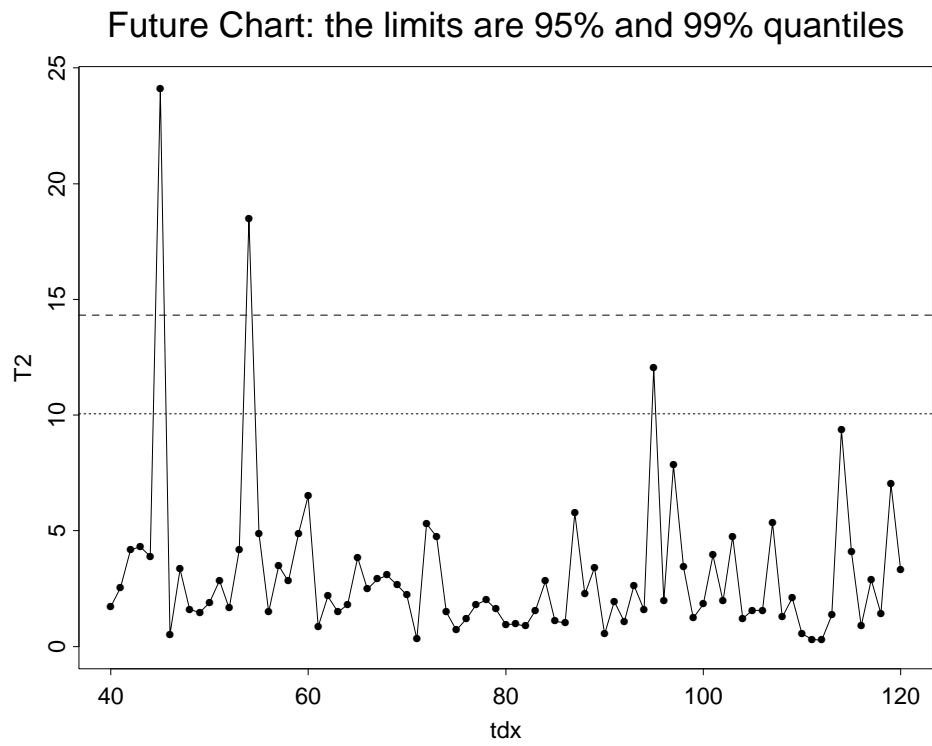
EM algorithm: Iterate between *Expectation* step and *Maximization* step.

- E-step: For each data point with missing values, use the conditional distribution

$$\mathbf{X}_1 | \mathbf{X}_2 = \mathbf{x}_2 \sim N_k(\boldsymbol{\mu}_1 + \Sigma_{12} \Sigma_{22}^{-1} (\mathbf{x}_2 - \boldsymbol{\mu}_2), \Sigma_{11} - \Sigma_{12} \Sigma_{22}^{-1} \Sigma_{21}),$$

where $k = \dim(\mathbf{X}_1)$ and the partition is based on $\mathbf{X} = (\mathbf{X}'_1, \mathbf{X}'_2)'$. Here \mathbf{X}_1 denotes the missing components and \mathbf{X}_2 the observed component.

Figure 2: T^2 Future Chart for monthly log returns of four Midwest companies: 1998-2007



- M-step: Perform MLE estimation based on the complete data.

It is helpful to make use of sufficient statistics.

MCMC method: All make use of the same conditional distribution. However, instead of using the expectation, one draws a random sample from the conditional distribution to fill the missing values.

7 Impact of serial correlations

Consider the case in which \mathbf{X}_i is serially correlated such as it follows a vector AR(1) model

$$\mathbf{X}_t - \boldsymbol{\mu} = \boldsymbol{\Phi}(\mathbf{X}_{t-1} - \boldsymbol{\mu}) + \mathbf{a}_t, \quad \mathbf{a}_t \sim_{iid} N_p(\mathbf{0}, \boldsymbol{\Sigma}),$$

where $E(\mathbf{X}_t) = \boldsymbol{\mu}$ and all eigenvalues of $\boldsymbol{\Phi}$ are less than 1 in modulus. Let $\mathbf{X}_1, \dots, \mathbf{X}_n$ be n consecutive observations of the model. Define the lag- ℓ autocovariance matrix of \mathbf{X}_t as

$$\boldsymbol{\Gamma}_\ell = E(\mathbf{X}_t - \boldsymbol{\mu})(\mathbf{X}_{t-\ell} - \boldsymbol{\mu})', \quad \ell = 0, \pm 1, \pm 2, \dots$$

It is easy to see that $\boldsymbol{\Gamma}_\ell = \boldsymbol{\Gamma}'_{-\ell}$. Also, by repeatedly substitution, we have

$$\mathbf{X}_t - \boldsymbol{\mu} = \mathbf{a}_t + \boldsymbol{\Phi}\mathbf{a}_{t-1} + \boldsymbol{\Phi}^2\mathbf{a}_{t-2} + \dots$$

Therefore,

$$\boldsymbol{\Gamma}_0 = \sum_{i=1}^{\infty} \boldsymbol{\Phi}^i \boldsymbol{\Sigma} (\boldsymbol{\Phi}^i)'$$

For $\ell > 0$, post-multiplying the model by $(\mathbf{X}_{t-\ell} - \boldsymbol{\mu})'$, taking expectation, and using the fact that $\mathbf{X}_{t-\ell}$ is uncorrelated with \mathbf{a}_t , we get

$$\boldsymbol{\Gamma}_\ell = \boldsymbol{\Phi}\boldsymbol{\Gamma}_{\ell-1}, \quad \ell = 1, 2, \dots$$

Taking the transpose of the above equation, we have

$$\boldsymbol{\Gamma}'_\ell = \boldsymbol{\Gamma}'_{\ell-1} \boldsymbol{\Phi}'.$$

Using $\boldsymbol{\Gamma}_\ell = \boldsymbol{\Gamma}'_{-\ell}$, we obtain

$$\boldsymbol{\Gamma}_{-\ell} = \boldsymbol{\Gamma}_{-(\ell-1)} \boldsymbol{\Phi}', \quad \ell > 1,$$

which is equivalent to

$$\boldsymbol{\Gamma}_\ell = \boldsymbol{\Gamma}_{\ell+1} \boldsymbol{\Phi}', \quad \ell = -1, -2, \dots$$

For the vector AR(1) model, we can show the following properties:

1. $E(\bar{\mathbf{X}}) = \boldsymbol{\mu}$.
2. $\text{Cov}(n^{-1/2} \sum_{t=1}^n \mathbf{X}_t) \rightarrow_p \boldsymbol{\Omega}$, where $\boldsymbol{\Omega} = (\mathbf{I} - \boldsymbol{\Phi})^{-1} \boldsymbol{\Gamma}_0 + \boldsymbol{\Gamma}_0 (\mathbf{I} - \boldsymbol{\Phi}')^{-1} - \boldsymbol{\Gamma}_0$, as $n \rightarrow \infty$.

3. $\mathbf{S} = \frac{1}{n-1} \sum_{t=1}^n (\mathbf{X}_t - \bar{\mathbf{X}})(\mathbf{X}_t - \bar{\mathbf{X}})' \rightarrow_p \mathbf{\Gamma}_0$ as $n \rightarrow \infty$.

Using these properties, we can show that $\sqrt{n}(\bar{\mathbf{X}} - \boldsymbol{\mu})$ is approximately normal with mean $\mathbf{0}$ and covariance $\boldsymbol{\Omega}$. This implies that $n(\bar{\mathbf{X}} - \boldsymbol{\mu})' \boldsymbol{\Omega}^{-1} (\bar{\mathbf{X}} - \boldsymbol{\mu}) \sim \chi_p^2$, not the usual statistic $n(\bar{\mathbf{X}} - \boldsymbol{\mu})' \mathbf{S}^{-1} (\bar{\mathbf{X}} - \boldsymbol{\mu})$. See Table 5.10 of the text (page 257) for the difference in coverage probability due to the effect of AR(1) serial correlations.