

1. Suppose that  $\{Y_t\}$  is a stationary process with constant mean  $\mu_t = E(Y_t) = \mu$  and autocorrelation function  $\rho_k$ . Define

$$\bar{Y} = \frac{1}{n} \sum_{i=1}^n Y_i$$

to be the sample mean of  $Y_1, Y_2, \dots, Y_n$ .

(a) Show that  $\bar{Y}$  is an unbiased estimator of  $\mu$ ; that is, show that  $E(\bar{Y}) = \mu$ . Note that this result holds regardless of the values of  $\rho_k$ .

(b) Prove that

$$\text{var}(\bar{Y}) = \frac{\gamma_0}{n} \left[ 1 + 2 \sum_{k=1}^{n-1} \left( 1 - \frac{k}{n} \right) \rho_k \right],$$

where  $\gamma_0 = \text{var}(Y_t)$ .

(c) If  $\{Y_t\}$  is a white noise process, what does  $\text{var}(\bar{Y})$  reduce to?

(d) Suppose that  $\{Y_t\}$  is a **MA(1) process**; i.e.,

$$Y_t = e_t - \theta e_{t-1},$$

where  $\theta$  is a fixed parameter and  $\{e_t\}$  is a zero mean white noise process with  $\text{var}(e_t) = \sigma_e^2$ . Use the result in part (b) to find an expression for  $\text{var}(\bar{Y})$ .

2. Consider the process  $Y_t = \mu + X_t$ , where  $X_t = e_t - \theta e_{t-1}$  and  $e_t \sim \text{iid } \mathcal{N}(0, 1)$ . Let  $n = 100$ .

(a) Show that  $E(X_t) = 0$  and  $\text{var}(X_t) = \gamma_0 = 1 + \theta^2$ . Because linear combinations of normal random variables are normally distributed, it follows immediately that  $X_t \sim \mathcal{N}(0, 1 + \theta^2)$  and  $Y_t \sim \mathcal{N}(\mu, 1 + \theta^2)$ .

(b) Find the sampling distribution of  $\bar{Y}$ .

(c) Suppose that  $\theta = -0.5$ . Use the R code

```
> Y.t = 10 + arima.sim(list(order = c(0,0,1), ma = 0.5), n = 100)
```

to generate a realization of this process when  $\mu = 10$ . Compute a 99 percent confidence interval for  $\mu$  (which in this problem is known to be 10). Does your interval include 10? You can use the `mean(Y.t)` command to compute the sample mean of  $Y_1, Y_2, \dots, Y_{100}$ .

(d) Repeat part (c) with  $\theta = 0.5$ . You will use the same command to simulate a new realization, except use `ma = -0.5`. R uses the convention of negating the parameter  $\theta$ , as we will see in Chapter 4.

(e) Compare the confidence intervals in parts (c) and (d) in terms of interval length. Explain why one interval should be shorter.

3. The TSA library contains the data set `co2`, which lists monthly carbon dioxide levels in northern Canada from 1/1994 to 12/2004. In HW1, you examined these data and fit a straight line regression model to them. You should have concluded that this was not a very good model for these data. In this problem, we will use R to fit the model

$$\text{CO2}_t = \beta_0 + \beta_1 t + \beta_2 \cos(2\pi f t) + \beta_3 \sin(2\pi f t) + X_t,$$

where  $E(X_t) = 0$ . This deterministic part

$$\mu_t = \beta_0 + \beta_1 t + \beta_2 \cos(2\pi ft) + \beta_3 \sin(2\pi ft)$$

contains both linear and trigonometric trend components. Note that there are 12 observations per year, so we take the frequency  $f = 1$ .

(a) To fit the model, we can use the R commands:

```
> har. <- harmonic(co2,1)
> fit <- lm(co2~har. + time(co2))
> summary(fit)
```

What are the least squares estimates? Write out an equation for the fitted model. Is all of the R output relevant? Explain.

(b) Construct a graph which plots the points along with the fitted regression model superimposed. You can use the R commands:

```
> plot(ts(fitted(fit),freq=12,start=c(1994,1)),ylab="Monthly CO2 levels",
        type='l',ylim=range(c(fitted(fit),co2)))
> points(co2)
```

Is this model appear to fit the data better than the straight line model? How would you rate the fit overall?

(c) Use the remaining code to perform the model diagnostics we discussed in class (Section 3.5 in the notes).

```
> plot(rstudent(fit),ylab="Std.residuals",xlab="Year",type="o")
> abline(h=0)
> hist(rstudent(fit),main="Hist. of std.residuals",xlab="Std.residuals")
> qqnorm(rstudent(fit),main="QQ plot of std.residuals")
> shapiro.test(rstudent(fit))
> runs(rstudent(fit))
> acf(rstudent(fit),main="Sample ACF for std.residuals")
```

Interpret everything and give an overall assessment of the model we have fit in this problem. Do the residuals look to resemble a stationary white noise process? Or, is there still noticeable structure left in them?

4. Consider the ventilation data in Example 1.10 (notes, pp 11). The data are located on the course web site under the name `ventilation`. Cut and paste the data into Notepad and save the file as a `.txt` file. Then, read the data into R using the commands I give on the web site (suitably modified to the location of the `.txt` file on your computer).

(a) Fit a straight-line regression model to the data for detrending purposes.

(b) Do a thorough residual analysis of the detrended (residual) process using the techniques we discussed in class (Section 3.5 in the notes). Do the residuals from your straight-line model fit resemble a zero mean white noise process?

5. Consider the data set

5    7    9    4    5    6

(in this order).

- (a) Compute by hand  $r_1$ , the lag one autocorrelation coefficient. Show all calculations.
- (b) Compute by hand  $r_2$ , the lag two autocorrelation coefficient. Show all calculations.

6. The TSA library contains the data set `prescrip`, which lists monthly prescription costs for the months August 1986 to March 1992. These data are from the State of New Jersey Prescription Drug Program and are the cost per prescription claim during this time period.

- (a) Construct a time series plot for the data. Describe the appearance of the series.
- (b) Use the R command `diff.1 <- diff(prescrip)` to calculate the first differences  $\nabla Y_t = Y_t - Y_{t-1}$  and plot the first differences. Describe the appearance of this plot and how it compares with the plot of the original series.
- (c) Use all of the model diagnostic checks (Section 3.5 in the notes) on the difference process  $\{\nabla Y_t\}$ . Do the data differences resemble a normal zero mean white noise process?