CHAPTER 14

ESTIMATING THE SPECTRUM

Several alternative methods for constructing reasonable estimators of the spectral density have been proposed and investigated over the years. We will highlight just a few of them that have gained the most acceptance in light of present-day computing power. So-called nonparametric estimation of the spectral density (that is, smoothing of the sample spectral density) assumes very little about the shape of the "true" spectral density. Parametric estimation assumes that an autoregressive model—perhaps of high order—provides an adequate fit to the time series. The estimated spectral density is then based on the theoretical spectral density of the fitted AR model. Some other methods are touched on briefly.

14.1 Smoothing the Spectral Density

The basic idea here is that most spectral densities will change very little over small intervals of frequencies. As such, we should be able to average the values of the sample spectral density over small intervals of frequencies to gain reduced variability. In doing so, we must keep in mind that we may introduce bias into the estimates if, in fact, the theoretical spectral density does change substantially over that interval. *There will always be a trade-off between reducing variability and introducing bias.* We will be required to use judgment to decide how much averaging to perform in a particular case.

Let f be a Fourier frequency. Consider taking a simple average of the neighboring sample spectral density values centered on frequency f and extending m Fourier frequencies on either side of f. We are averaging $2m + 1$ values of the sample spectrum, and the *smoothed sample spectral density* is given by

$$\bar{S}(f) = \frac{1}{2m+1} \sum_{j=-m}^{m} \hat{S}\left(f + \frac{j}{n}\right) \tag{14.1.1}$$

(When averaging for frequencies near the end points of 0 and ½, we treat the periodogram as symmetric about 0 and ½.)

More generally, we may smooth the sample spectrum with a weight function or *spectral window* $W_m(f)$ with the properties

$$\left.\begin{aligned} W_m(k) &\geq 0 \\ W_m(k) &= W_m(-k) \\ \sum_{k=-m}^{m} W_m(k) &= 1 \end{aligned}\right\} \tag{14.1.2}$$

and obtain a smoothed estimator of the spectral density as

$$\bar{S}(f) = \sum_{k=-m}^{m} W_m(k)\hat{S}\left(f + \frac{k}{n}\right) \tag{14.1.3}$$

The simple averaging shown in Equation (14.1.1) corresponds to the rectangular spectral window

$$W_m(k) = \frac{1}{2m+1} \quad \text{for } -m \leq k \leq m \tag{14.1.4}$$

For historical reasons, this spectral window is usually called the *Daniell spectral window* after P. J. Daniell, who first used it in the 1940s.

As an example, consider the simulated AR(1) series whose sample spectral density was shown in Exhibit 13.20 on page 341. Exhibit 14.1 displays the smoothed sample spectrum using the Daniell window with $m = 5$. The true spectrum is again shown as a dotted line. The smoothing did reduce some of the variability that we saw in the sample spectrum.

**Exhibit 14.1   Smoothed Spectrum Using the Daniell Window With $m = 5$**
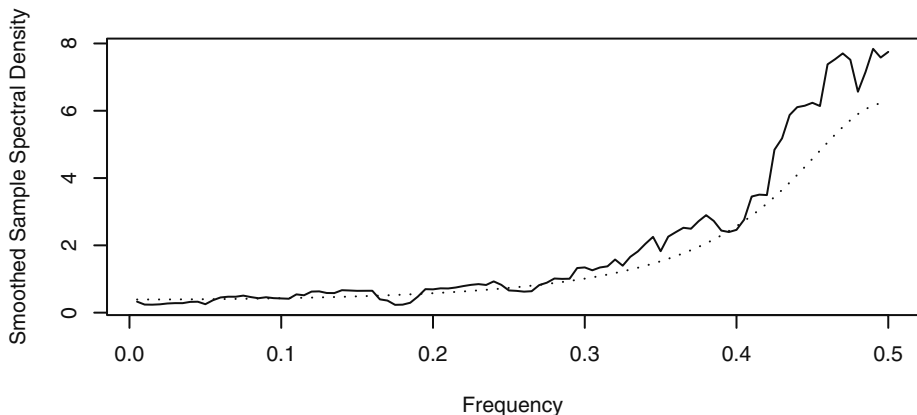

```
> win.graph(width=4.875,height=2.5,pointsize=8)
> set.seed(271435); n=200; phi=-0.6
> y=arima.sim(model=list(ar=phi),n=n)
> k=kernel('daniell',m=5)
```

```
> sp=spec(y,kernel=k,log='no',sub='',xlab='Frequency',
    ylab='Smoothed Sample Spectral Density')
> lines(sp$freq,ARMAspec(model=list(ar=phi),freq=sp$freq,
    plot=F)$spec,lty='dotted')
```

If we make the smoothing window wider (that is, increase m) we will reduce the variability even further. Exhibit 14.2 shows the smoothed spectrum with a choice of $m = 15$. The danger with more and more smoothing is that we may lose important details in the spectrum and introduce bias. The amount of smoothing needed will always be a matter of judgmental trial and error, recognizing the trade-off between reducing variability at the expense of introducing bias.

Exhibit 14.2   Smoothed Spectrum Using the Daniell Window With $m = 15$
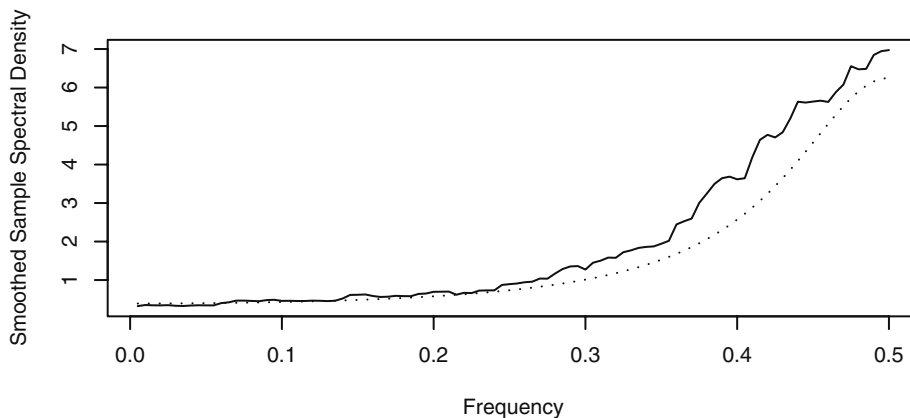

```
> k=kernel('daniell',m=15)
> sp=spec(y,kernel=k,log='no',sub='',xlab='Frequency',
    ylab='Smoothed Sample Spectral Density')
> lines(sp$freq,ARMAspec(model=list(ar=phi),freq=sp$freq,
    plot=F)$spec,lty='dotted')
```
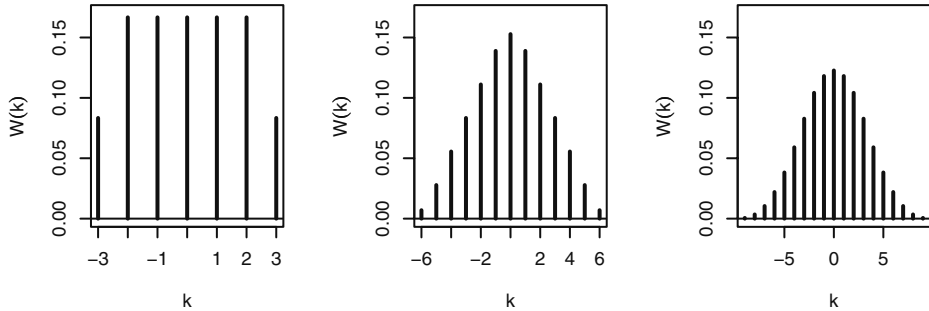
Other Spectral Windows

Many other spectral windows have been suggested over the years. In particular, the abrupt change at the end points of the Daniell window could be softened by making the weights decrease at the extremes. The so-called modified Daniell spectral window simply defines the two extreme weights as half of the other weights still retaining the property that the weights sum to 1. The leftmost graph in Exhibit 14.3 shows the modified Daniell spectral window for $m = 3$.

Exhibit 14.3   The Modified Daniell Spectral Window and Its Convolutions

Another common way to modify spectral windows is to use them to smooth the periodogram more than once. Mathematically, this amounts to using the convolution of the spectral windows. If the modified Daniell spectral window with $m = 3$ is used twice (convolved with itself), we in fact are using the (almost) triangular-shaped window shown in the middle display of Exhibit 14.3. A third smoothing (with $m = 3$) is equivalent to using the spectral window shown in the rightmost panel. This spectral window appears much like a normal curve. We could also use different values of m in the various components of the convolutions.

Most researchers agree that the shape of the spectral window is not nearly as important as the choice of m (or the bandwidth—see below). We will use the modified Daniell spectral window—possibly with one or two convolutions—in our examples.[†]

14.2 Bias and Variance

If the theoretical spectral density does not change much over the range of frequencies that the smoothing window covers, we expect the smoothed estimator to be approximately unbiased. A calculation using this approximation, the spectral window properties in Equations (14.1.2), and a short Taylor expansion produces

$$E[\bar{S}(f)] \approx \sum_{k=-m}^{m} W_m(k) S\left(f + \frac{k}{n}\right)$$

$$\approx \sum_{k=-m}^{m} W_m(k)\left[S(f) + \frac{k}{n}S'(f) + \frac{1}{2}\left(\frac{k}{n}\right)^2 S''(f) \right]$$

or

$$E[\bar{S}(f)] \approx S(f) + \frac{1}{n^2}\frac{S''(f)}{2} \sum_{k=-m}^{m} k^2 W_m(k) \tag{14.2.1}$$

[†] In R, the modified Daniell kernel is the default kernel for smoothing sample spectra, and m may be specified by simply specifying span $= 2m + 1$ in the spec function where span is an abbreviation of the spans argument.

So an approximate value for the bias in the smoothed spectral density is given by

$$\text{bias} \approx \frac{1}{n^2}\frac{S''(f)}{2}\sum_{k=-m}^{m} k^2 W_m(k) \tag{14.2.2}$$

For the Daniell rectangular spectral window, we have

$$\frac{1}{n^2}\sum_{k=-m}^{m} k^2 W_m(k) = \frac{2}{n^2(2m+1)}\left(\frac{m^3}{3}+\frac{m^2}{2}+\frac{m}{6}\right) \tag{14.2.3}$$

and thus the bias tends to zero as $n \to \infty$ as long as $m/n \to 0$.

   Using the fact that the sample spectral density values at the Fourier frequencies are approximately uncorrelated and Equation (13.6.5) on page 341, we may also obtain a useful approximation for the variance of the smoothed spectral density as

$$Var[\bar{S}(f)] \approx \sum_{k=-m}^{m} W_m^2(k) Var\left[\hat{S}\left(f+\frac{k}{n}\right)\right] \approx \sum_{k=-m}^{m} W_m^2(k)S^2(f)$$

so that

$$Var[\bar{S}(f)] \approx S^2(f)\sum_{k=-m}^{m} W_m^2(k) \tag{14.2.4}$$

Note that for the Daniell or rectangular spectral window $\sum_{k=-m}^{m} W_m^2(k) = \frac{1}{2m+1}$, so that as long as $m \to \infty$ (as $n \to \infty$) we have consistency.

   In general, we require that as $n \to \infty$ we have $m/n \to 0$ to reduce bias and $m \to \infty$ to reduce variance. As a practical matter, the sample size n is usually fixed and we must choose m to balance bias and variance considerations.

   Jenkins and Watts (1968) suggest trying three different values of m. A small value will give an idea where the large peaks in $S(f)$ are but may show a large number of peaks, many of which are spurious. A large value of m may produce a curve that is likely to be too smooth. A compromise may then be achieved with the third value of m. Chatfield (2004, p. 135) suggests using $m = \sqrt{n}$. Often trying values for m of $2\sqrt{n}$, $\sqrt{n}$, and $\frac{1}{2}\sqrt{n}$ will give you some insight into the shape of the true spectrum. Since the width of the window decreases as m decreases, this is sometimes called *window closing*. As Hannan (1973, p. 311) says, "Experience is the real teacher and cannot be got from a book."

14.3 Bandwidth

In the approximate bias given by Equation (14.2.2), notice that the factor $S''(f)$ depends on the curvature of the true spectral density and will be large in magnitude if there is a sharp peak in $S(f)$ near f but will be small when $S(f)$ is relatively flat near f. This makes intuitive sense, as the motivation for the smoothing of the sample spectral density assumed that the true density changed very little over the range of frequencies used in the spectral window. The square root of the other factor in the approximate bias from

Equation (14.2.2) is sometimes called the *bandwidth*, *BW*, of the spectral window, namely

$$BW = \frac{1}{n}\sqrt{\sum_{k=-m}^{m} k^2 W_m(k)} \tag{14.3.1}$$

As we noted in Equation (14.2.3), for the Daniell window this *BW* will tend to zero as $n \to \infty$ as long as $m/n \to 0$. From Equations (14.1.2) on page 352 a spectral window has the mathematical properties of a discrete zero-mean probability density function, so the *BW* defined here may be viewed as proportional to the standard deviation of the spectral window. As such, it is one way to measure the width of the spectral window. It is interpreted as a measure of width of the band of frequencies used in smoothing the sample spectral density. If the true spectrum contains two peaks that are close relative to the bandwidth of the spectral window, those peaks will be smoothed together when we calculate $\bar{S}(f)$ and they will not be seen as separate peaks. It should be noted that there are many alternative definitions of bandwidth given in the time series literature. Priestley (1981, pp. 513–528) spends considerable time discussing the advantages and disadvantages of the various definitions.

14.4 Confidence Intervals for the Spectrum

The approximate distributional properties of the smoothed spectral density may be easily used to obtain confidence intervals for the spectrum. The smoothed sample spectral density is a linear combination of quantities that have approximate chi-square distributions. A common approximation in such a case is to use some multiple of another chi-square distribution with degrees of freedom obtained by matching means and variances. Assuming $\bar{S}(f)$ to be roughly unbiased with variance given by Equation (14.2.4), matching means and variances leads to approximating the distribution of

$$\frac{\nu\bar{S}(f)}{S(f)} \tag{14.4.1}$$

by a chi-square distribution with degrees of freedom given by

$$\nu = \frac{2}{\sum_{k=-m}^{m} W_m^2(k)} \tag{14.4.2}$$

Letting $\chi^2_{\nu,\,\alpha/2}$ be the $100(\alpha/2)$th percentile of a chi-square distribution with ν degrees of freedom, the inequality

$$\chi^2_{\nu,\,\alpha/2} < \frac{\nu\bar{S}(f)}{S(f)} < \chi^2_{\nu,\,1-\alpha/2}$$

can be converted into a $100(1-\alpha)\%$ confidence statement for $S(f)$ as

$$\frac{v\overline{S}(f)}{\chi^2_{v,\,1-\alpha/2}} < S(f) < \frac{v\overline{S}(f)}{\chi^2_{v,\,\alpha/2}} \tag{14.4.3}$$

In this formulation, the width of the confidence interval will vary with frequency. A review of Equation (14.2.4) on page 355 shows that the variance of $\overline{S}(f)$ is roughly proportional to the square of its mean. As we saw earlier in Equations (5.4.1) and (5.4.2) on page 98, this suggests that we take the logarithm of the smoothed sample spectral density to stabilize the variance and obtain confidence intervals with width independent of frequency as follows:

$$\log[\overline{S}(f)] + \log\left[\frac{v}{\chi^2_{v,\,1-\alpha/2}}\right] \le \log[S(f)] \le \log[\overline{S}(f)] + \log\left[\frac{v}{\chi^2_{v,\,\alpha/2}}\right] \tag{14.4.4}$$

For these reasons it is common practice to plot the logarithms of estimated spectra. If we redo Exhibit 14.2 on page 353 in logarithm terms, we obtain the display shown in Exhibit 14.4, where we have also drawn in the 95% confidence limits (dotted) and the true spectral density (dashed) from the AR(1) model. With a few exceptions, the confidence limits capture the true spectral density.

Exhibit 14.4   Confidence Limits from the Smoothed Spectral Density
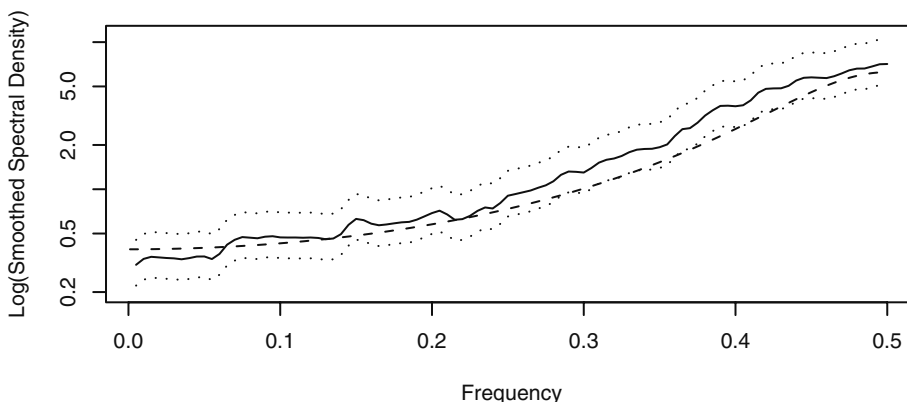

```
> set.seed(271435); n=200; phi=-0.6
> y=arima.sim(model=list(ar=phi),n=n)
> k=kernel('daniell',m=15)
> sp=spec(y,kernel=k,sub='',xlab='Frequency',
    ylab='Log(Smoothed Spectral Density)', ci.plot=T,ci.col=NULL)
> lines(sp$freq,ARMAspec(model=list(ar=phi),sp$freq,plot=F)$spec,
    lty='dashed')
```

Exhibit 14.5 shows a less cluttered display of confidence limits. Here a 95% confidence interval and bandwidth guide is displayed in the upper right-hand corner—the "crosshairs." The vertical length gives the length (width) of a confidence interval, while

the horizontal line segment indicates the central point[†] of the confidence interval, and its width (length) matches the bandwidth of the spectral window. If you visualize the guide repositioned with the crosshairs centered on the smoothed spectrum above any frequency, you have a visual display of a vertical confidence interval for the "true" spectral density at that frequency and a rough guide of the extent of the smoothing. In this simulated example, we also show the true spectrum as a dotted line.

Exhibit 14.5   Logarithm of Smoothed Spectrum from Exhibit 14.2
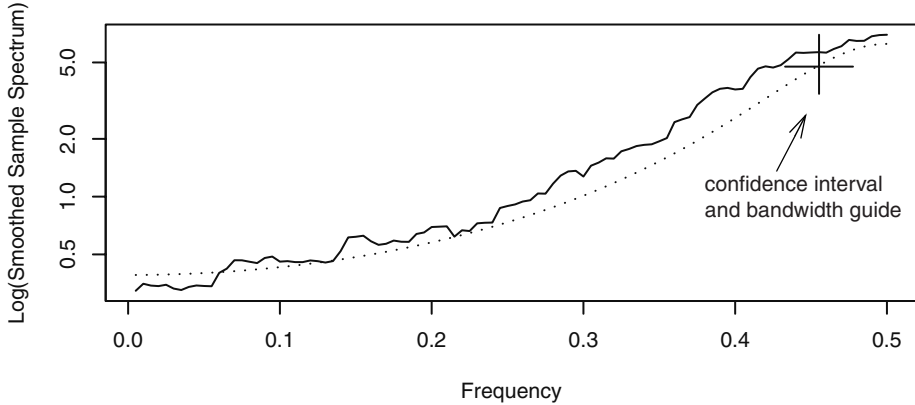

```
> sp=spec(y,span=31,sub='',xlab='Frequency',
    ylab='Log(Smoothed Sample Spectrum)')
> lines(sp$freq,ARMAspec(model=list(ar=phi),sp$freq,
    plot=F)$spec,lty='dotted')
```

14.5 Leakage and Tapering

Much of the previous discussion has assumed that the frequencies of interest are the Fourier frequencies. What happens if that is not the case? Exhibit 14.6 displays the periodogram of a series of length $n = 96$ with two pure cosine-sine components at frequencies $f = 0.088$ and $f = 14/96$. The model is simply

$$Y_t = 3\cos[2\pi(0.088)t] + \sin\left[2\pi\left(\frac{14}{96}\right)t\right] \tag{14.5.1}$$

Note that with $n = 96$, $f = 0.088$ is *not* a Fourier frequency. The peak with lower power at the Fourier frequency $f = 14/96$ is clearly indicated. However, the peak at $f = 0.088$ is not

---

[†] The central point is not, in general, halfway between the endpoints, as Equation (14.4.4) determines *asymmetric* confidence intervals. In this example, using the modified Daniell window with $m = 15$, we have $\nu = 61$ degrees of freedom, so the chi-square distribution used is effectively a normal distribution, and the confidence intervals are nearly symmetric.
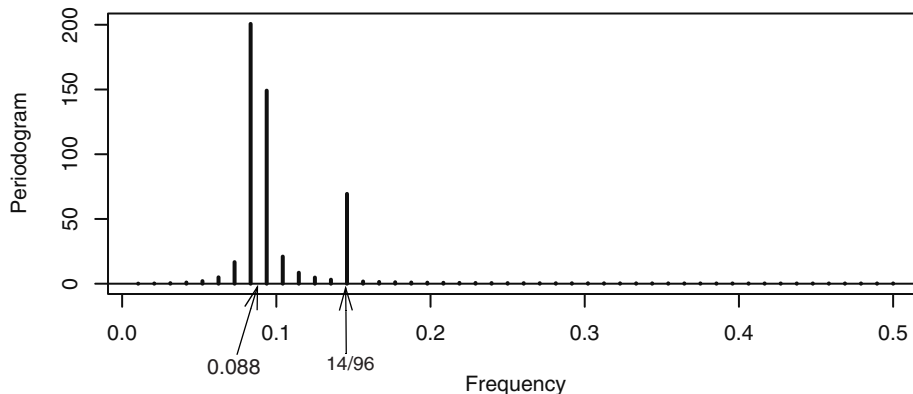
there. Rather, the power at this frequency is blurred across several nearby frequencies, giving the appearance of a much wider peak.

Exhibit 14.6   Periodogram of Series with Peaks at $f = 0.088$ and $f = 14/96$


```
> win.graph(width=4.875,height=2.5,pointsize=8)
> t=1:96; f1=0.088; f2=14/96
> y=3*cos(f1*2*pi*t)+sin(f2*2*pi*t)
> periodogram(y); abline(h=0)
```

An algebraic analysis[†] shows that we may view the periodogram as a "smoothed" spectral density formed with the Dirichlet kernel spectral window given by

$$D(f) = \frac{1}{n}\frac{\sin(n\pi f)}{\sin(\pi f)} \qquad (14.5.2)$$

Note that for all Fourier frequencies $f = j/n$, $D(f) = 0$, so this window has no effect whatsoever at those frequencies. However, the plot of $D(f)$ given on the left-hand side of Exhibit 14.7 shows significant "side lobes" on either side of the main peak. This will cause power at non-Fourier frequencies to leak into the supposed power at the nearby Fourier frequencies, as we see in Exhibit 14.6.

Tapering is one method used to improve the issue with the side lobes. Tapering involves decreasing the data magnitudes at both ends of the series so that the values move gradually toward the data mean of zero. The basic idea is to reduce the end effects of computing a Fourier transform on a series of finite length. If we calculate the periodogram after tapering the series, the effect is to use the modified Dirichlet kernel shown on the right-hand side of Exhibit 14.7 for $n = 100$. Now the side lobes have essentially disappeared.
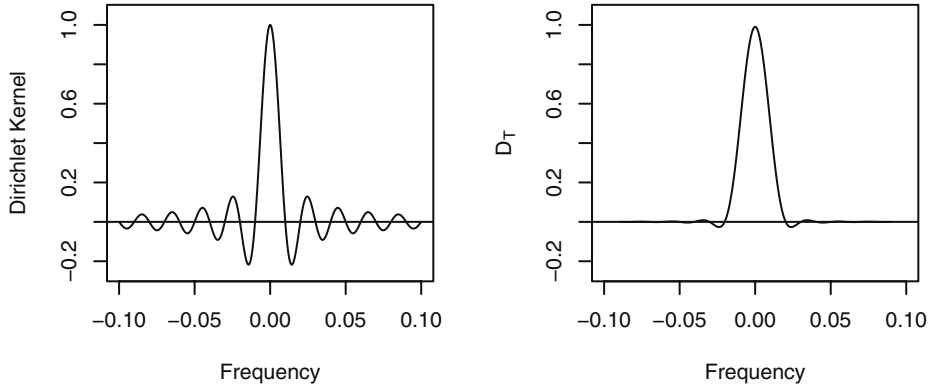

[†] Appendix K on page 381 gives some of the details.

Exhibit 14.7   Dirichlet Kernel and Dirichlet Kernel after Tapering

The most common form of tapering is based on a cosine bell. We replace the original series Y_t by $\tilde{Y}_t$, with

$$\tilde{Y}_t = h_t Y_t \tag{14.5.3}$$

where, for example, h_t is the cosine bell given by

$$h_t = \frac{1}{2}\left\{1 - \cos\left[\frac{2\pi(t - 0.5)}{n}\right]\right\} \tag{14.5.4}$$

A graph of the cosine bell with $n = 100$ is given on the left-hand side of Exhibit 14.8. A much more common taper is given by a split cosine bell that applies the cosine taper only to the extremes of the time series. The split cosine bell taper is given by

$$h_t = \begin{cases} \dfrac{1}{2}\left\{1 - \cos\left[\dfrac{\pi(t - 1/2)}{m}\right]\right\} & \text{for } 1 \le t \le m \\[2mm] 1 & \text{for } m + 1 \le t \le n - m \\[2mm] \dfrac{1}{2}\left\{1 - \cos\left[\dfrac{\pi(n - t + 1/2)}{m}\right]\right\} & \text{for } n - m + 1 \le t \le n \end{cases} \tag{14.5.5}$$

which is called a $100p\%$ cosine bell taper with $p = 2m/n$. A 10% split cosine bell taper is shown on the right-hand side of Exhibit 14.8 again with $n = 100$. Notice that there is a 10% taper on each end, resulting in a total taper of 20%. In practice, split cosine bell tapers of 10% or 20% are in common use.

Exhibit 14.8   Cosine Bell and 10% Taper Split Cosine Bell for _n_ = 100
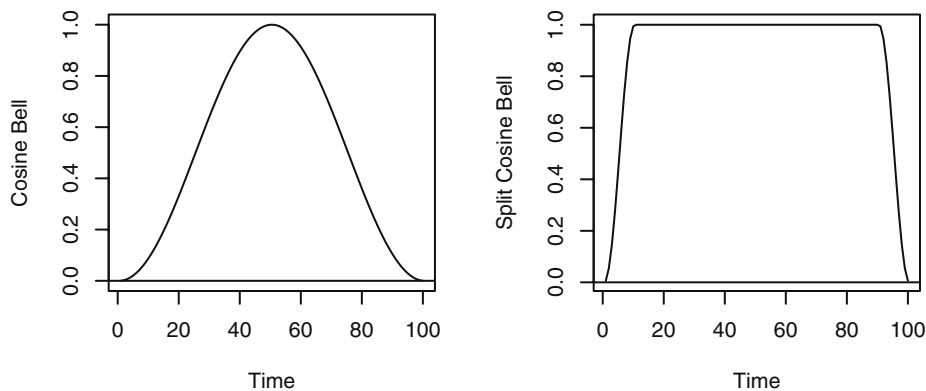
We return to the variable star brightness data first explored on page 325. Exhibit 14.9 displays four periodograms of this series, each with a different amount of tapering. Judging by the length of the 95% confidence intervals displayed in the respective "crosshairs", we see that the two peaks found earlier in the raw untapered periodogram at frequencies $f_1 = 21/600$ and $f_2 = 25/600$ are clearly real. A more detailed analysis of the minor peaks shown best in the bottom periodogram are all in fact harmonics of the frequencies f_1 and f_2. There is much more on the topic of leakage reduction and tapering in Bloomfield (2000).

Exhibit 14.9   Variable Star Spectra with Tapers of 0%, 10%, 20%, and 50%
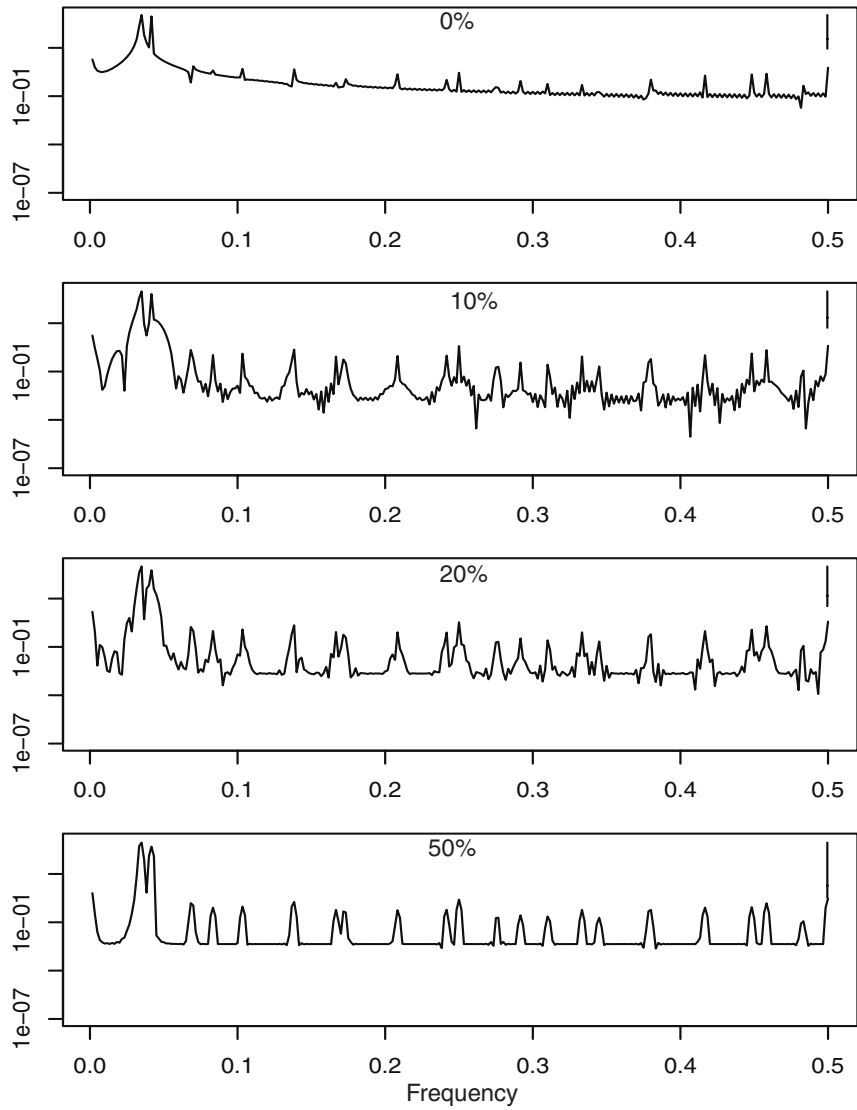
14.6 Autoregressive Spectrum Estimation

In the preceding sections on spectral density estimation, we did not make any assumptions about the parametric form of the true spectral density. However, an alternative method for estimating the spectral density would be to consider fitting an AR, MA, or ARMA model to a time series and then use the spectral density of that model with estimated parameters as our estimated spectral density. (Section 13.5, page 332, discussed the spectral densities of ARMA models.) Often AR models are used with possibly large order chosen to minimize the AIC criterion.

As an example, consider the simulated AR series with $\phi = -0.6$ and $n = 200$ that we used in Exhibits 13.20, 14.1, 14.2, and 14.5. If we fit an AR model, choosing the order to minimize the AIC, and then plot the estimated spectral density for that model, we obtain the results shown in Exhibit 14.10.

Exhibit 14.10 Autoregressive Estimation of the Spectral Density
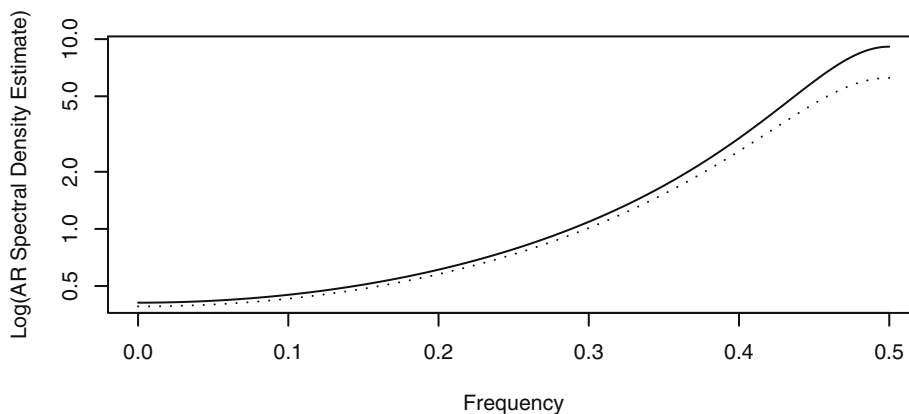

```
> sp=spec(y,method='ar',sub='',xlab='Frequency',
    ylab='Log(AR Spectral Density Estimate')
> lines(sp$freq,ARMAspec(model=list(ar=phi),freq=sp$freq,
    plot=F)$spec,lty='dotted')
```

Since these are simulated data, we also show the true spectral density as a dotted line. In this case, the order was chosen as $p = 1$ and the estimated spectral density follows the true density very well. We will show some examples with real time series in Section 14.8.

14.7 Examples with Simulated Data

A useful way to get a feel for spectral analysis is with simulated data. Here we know what the answers are and can see what the consequences are when we make choices of spectral window and bandwidth. We begin with an AR(2) model that contains a fairly strong peak in its spectrum.

AR(2) with $\phi_1 = 1.5$, $\phi_2 = -0.75$: A Peak Spectrum

The spectral density for this model contained a peak at about $f = 0.08$, as displayed in Exhibit 13.14 on page 336. We simulated a time series from this AR(2) model with normal white noise terms with unit variance and sample size $n = 100$. Exhibit 14.11 shows three estimated spectral densities and the true density as a solid line. We used the modified Daniell spectral window with three different values for $span = 2m + 1$ of 3, 9, and 15. A span of 3 gives the least amount of smoothing and is shown as a dotted line. A span of 9 is shown as a dashed line. With $span = 15$, we obtain the most smoothing, and this curve is displayed with a dot-dash pattern. The bandwidths of these three spectral windows are 0.018, 0.052, and 0.087, respectively. The confidence interval and bandwidth guide displayed apply only to the dotted curve estimate. The two others have wider bandwidths and shorter confidence intervals. The estimate based on $span = 9$ is probably the best one, but it does not represent the peak very well.

Exhibit 14.11 Estimated Spectral Densities
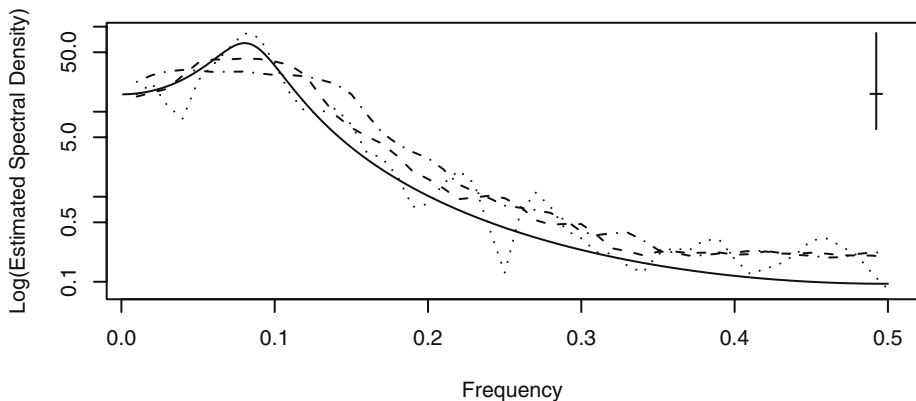

```
> win.graph(width=4.875,height=2.5,pointsize=8)
> set.seed(271435); n=100; phi1=1.5; phi2=-.75
> y=arima.sim(model=list(ar=c(phi1,phi2)),n=n)
> sp1=spec(y,spans=3,sub='',lty='dotted',xlab='Frequency',
    ylab='Log(Estimated Spectral Density)')
> sp2=spec(y,spans=9,plot=F); sp3=spec(y,spans=15,plot=F)
> lines(sp2$freq,sp2$spec,lty='dashed')
> lines(sp3$freq,sp3$spec,lty='dotdash')
```

```
> f=seq(0.001,.5,by=.001)
> lines(f,ARMAspec(model=list(ar=c(phi1,phi2)),freq=f,
    plot=F)$spec,lty='solid')
```

We also used the parametric spectral estimation idea and let the software choose the best AR model based on the smallest AIC. The result was an estimated AR(2) model with the spectrum shown in Exhibit 14.12. This is a very good representation of the underlying spectrum, but of course the model was indeed AR(2).

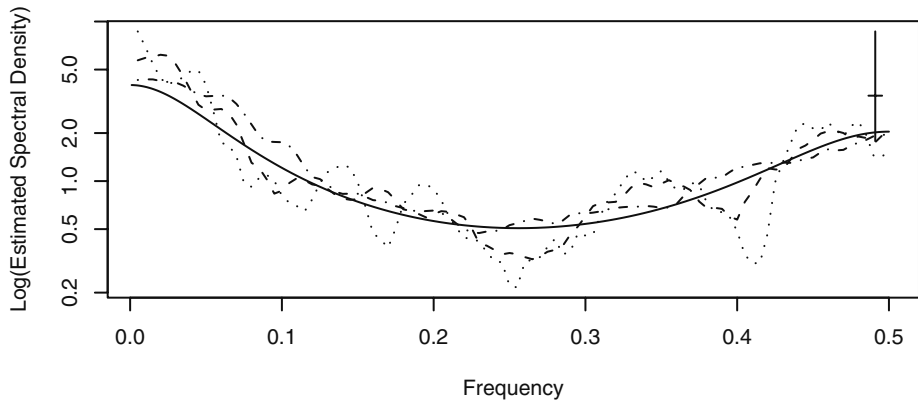**Exhibit 14.12 AR Spectral Estimation: Estimated (dotted), True (solid)**
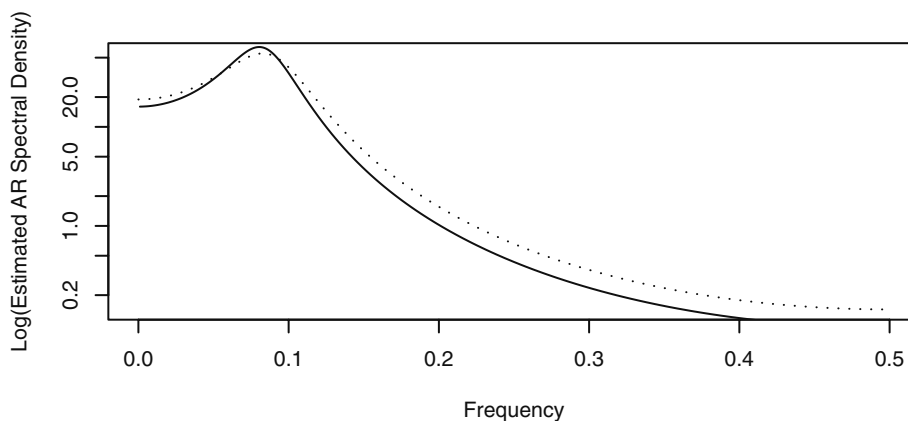

```
> sp4=spec(y,method='ar',lty='dotted',
    xlab='Frequency',ylab='Log(Estimated AR Spectral Density)')
> f=seq(0.001,0.5, by 0.001)
> lines(f,ARMAspec(model=list(ar=c(phi1,phi2)),freq=f,
    plot=F)$spec,lty='solid')
> sp4$method # This will tell you order of the AR model selected
```

AR(2) with $\phi_1 = 0.1$, $\phi_2 = 0.4$: A Trough Spectrum

Next we look at an AR(2) model with a trough spectrum and a larger sample size. The true spectrum is displayed in Exhibit 13.15 on page 337. We simulated this model with $n = 200$ and unit-variance normal white noise. The three smoothed spectral estimates shown are based on spans of 7, 15, and 31. As before, the confidence limits and bandwidth guide correspond to the smallest span of 7 and hence give the narrowest bandwidth and longest confidence intervals. In our opinion, the middle value of *span* = 15, which is about $\sqrt{n}$, gives a reasonable estimate of the spectrum.

**Exhibit 14.13 Estimated Spectrum for AR(2) Trough Spectrum Model**


```
> Use the R code for Exhibit 14.11 with new values for the
> parameters.
```

Exhibit 14.14 shows the AR spectral density estimate. The minimum AIC was achieved at the true order of the underlying model, AR(2), and the estimated spectral density is quite good.

Exhibit 14.14 AR Spectral Estimation: Estimated (dotted), True (solid)
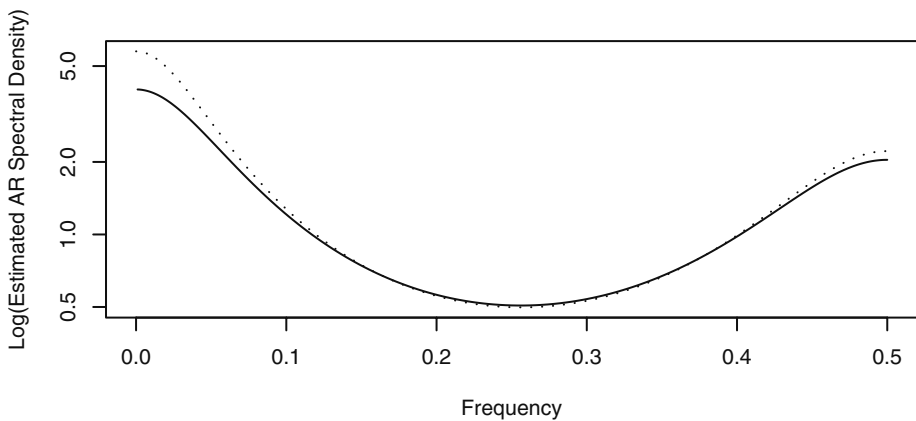


```
> Use the R code for Exhibits 14.11 and 14.12 with new values
> for the parameters.
```

ARMA(1,1) with $\phi = 0.5$, $\theta = 0.8$

The true spectral density of the mixed model ARMA(1,1) with $\phi = 0.5$ and $\theta = 0.8$ was shown in Exhibit 13.17 on page 338. This model has substantial medium- and high-frequency content but very little power at low frequencies. We simulated this model with a sample size of $n = 500$ and unit-variance normal white noise. Using $\sqrt{n} \approx 22$ as a guide for choosing m, we show three estimates with m of 11, 23, and 45 in Exhibit 14.15. The confidence interval guide indicates that the many peaks produced when $m = 11$ are likely spurious (which, in fact, they are). With such a smooth underlying spectrum, the maximum smoothing shown with $m = 45$ produces a rather good estimate.

Exhibit 14.15 Spectral Estimates for an ARMA(1,1) Process
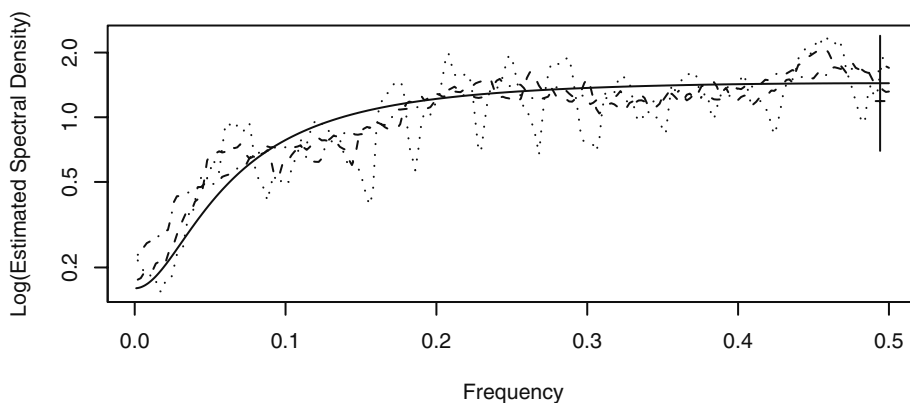

```
> win.graph(width=4.875,height=2.5,pointsize=8)
> set.seed(324135); n=500; phi=.5; theta=.8
> y=arima.sim(model=list(ar=phi,ma=-theta),n=n)
> sp1=spec(y,spans=11,sub='',lty='dotted',
    xlab='Frequency',ylab='Log(Estimated Spectral Density)')
> sp2=spec(y,spans=23,plot=F); sp3=spec(y,spans=45,plot=F)
> lines(sp2$freq,sp2$spec,lty='dashed')
> lines(sp3$freq,sp3$spec,lty='dotdash')
> f=seq(0.001,.5,by=.001)
> lines(f,ARMAspec(model=list(ar=phi,ma=-theta),f,
    plot=F)$spec,lty='solid')
```

In this case, a parametric spectral estimate based on AR models does not work well, as shown in Exhibit 14.16. The software selected an AR(3) model, but the resulting spectral density (dotted) does not reproduce the true density (solid) well at all.

Exhibit 14.16 AR Spectral Estimate for an ARMA(1,1) Process
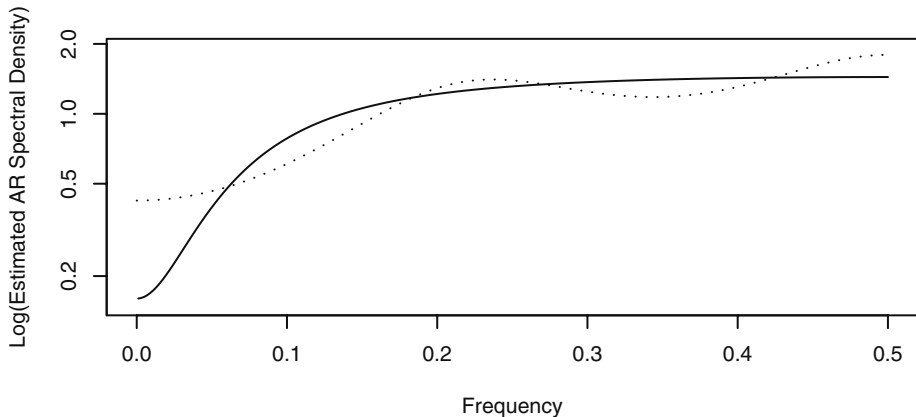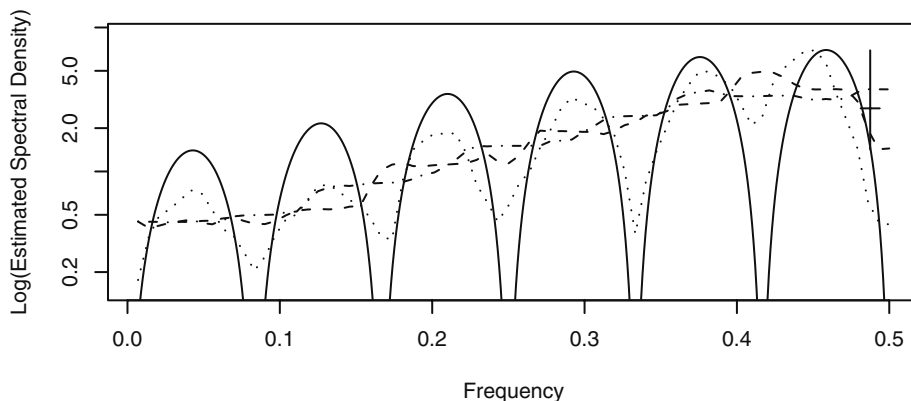


```
> sp4=spec(y,method='ar',lty='dotted',ylim=c(.15,1.9),
    xlab='Frequency',ylab='Log(Estimated AR Spectral Density)')
> f=seq(0.001,.5,by=.001)
> lines(f,ARMAspec(model=list(ar=phi,ma=-theta),f,
    plot=F)$spec,lty='solid')
```

Seasonal MA with θ = 0.4, Θ = 0.9, and *s* = 12

For our final example with simulated data, we choose a seasonal process. The theoretical spectral density is displayed in Exhibit 13.19 on page 340. We simulated $n = 144$ data points with unit-variance normal white noise. We may think of this as 12 years of monthly data. We used modified Daniell spectral windows with *span* = 6, 12, and 24 based on $\sqrt{n} \approx 12$.

This spectrum contains a lot of detail and is difficult to estimate with only 144 observations. The narrowest spectral window hints at the seasonality, but the two other estimates essentially smooth out the seasonality. The confidence interval widths (corresponding to $m = 6$) do seem to confirm the presence of real seasonal peaks.

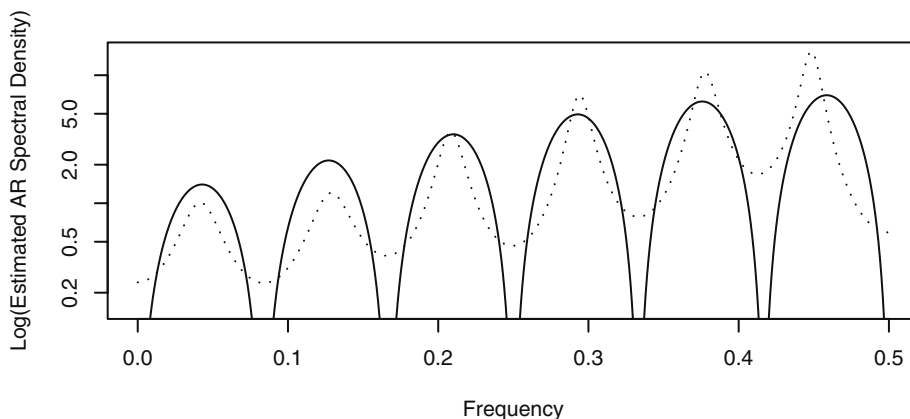**Exhibit 14.17 Spectral Estimates for a Seasonal Process**


```
> win.graph(width=4.875,height=2.5,pointsize=8)
> set.seed(247135); n=144; theta=.4;THETA=.9
> y=arima.sim(model=list(ma=c(-theta,rep(0,10),-THETA,theta*THETA
    )),n=n)
> sp1=spec(y,spans=7,sub='',lty='dotted',ylim=c(.15,9),
    xlab='Frequency',ylab='Log(Estimated Spectral Density)')
> sp2=spec(y,spans=13,plot=F); sp3=spec(y,spans=25,plot=F)
> lines(sp2$freq,sp2$spec,lty='dashed')
> lines(sp3$freq,sp3$spec,lty='dotdash')
> f=seq(0.001,.5,by=.001)
> lines(f,ARMAspec(model=list(ma=-theta,seasonal=list(sma=-THETA,
    period=12)),freq=f,plot=F)$spec,lty='solid')
```

Exhibit 14.18 AR Spectral Estimates for a Seasonal Process


```
> sp4=spec(y,method='ar',ylim=c(.15,15),lty='dotted',
    xlab='Frequency',ylab='Log(Estimated AR Spectral Density)')
```
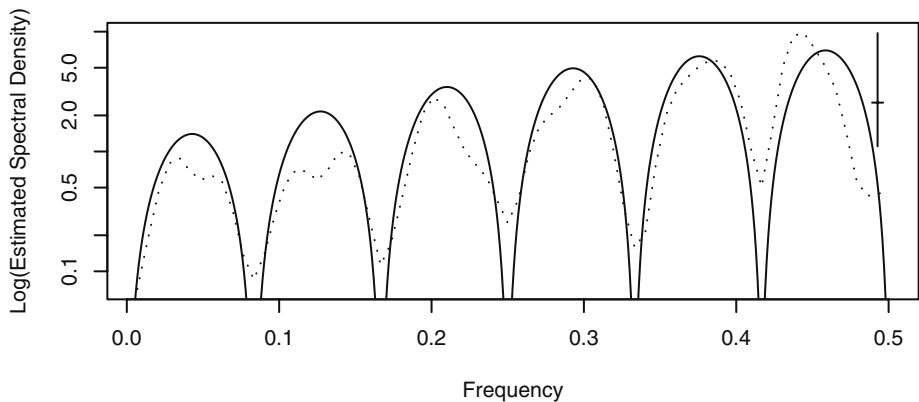
```
> f=seq(0.001,.5,by=.001)
> lines(f,ARMAspec(model=list(ma=-theta,seasonal=list(sma=-THETA,
    period=12)),freq=f,plot=F)$spec,lty='solid')
```

Exhibit 14.18 shows the estimated spectrum based on the best AR model. An order of 13 was chosen based on the minimum AIC, and the seasonality does show up quite well. However, the peaks are misplaced at the higher frequencies. Perhaps looking at both Exhibit 14.17 and Exhibit 14.18 we could conclude that the seasonality is real and that a narrow spectral window provides the best estimate of the underlying spectral density given the sample size available.

As a final estimate of the spectrum, we use a convolution of two modified Daniell spectral windows each with *span* = 3, as displayed in the middle of Exhibit 14.3 on page 354. The estimated spectrum is shown in Exhibit 14.19. This is perhaps the best of the estimates that we have shown.

Exhibit 14.19 Estimated Seasonal Spectrum with Convolution Window


```
> sp5=spec(y,spans=c(3,3),sub='',lty='dotted',
    xlab='Frequency',ylab='Log(Estimated Spectral Density)')
> f=seq(0.001,.5,by=.001)
> lines(f,ARMAspec(model=list(ma=-theta,seasonal=list(sma=-THETA,
    period=12)),freq=f,plot=F)$spec,lty='solid')
```
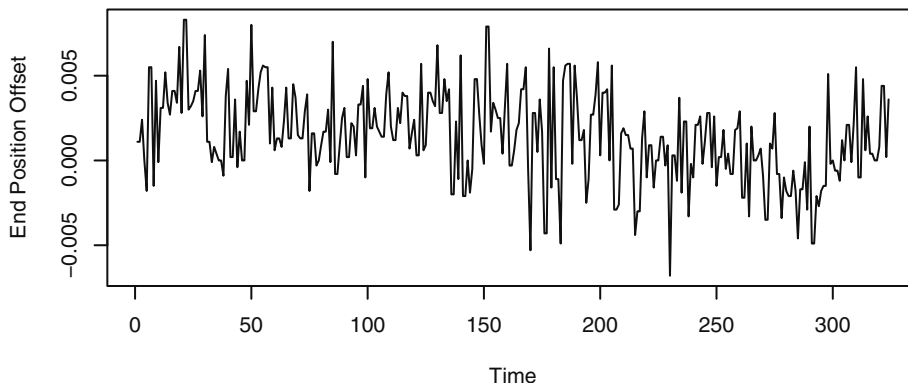
14.8 Examples with Actual Data

An Industrial Robot

An industrial robot was put through a sequence of maneuvers, and the distance from a desired target end position was recorded in inches. This was repeated 324 times to form the time series shown in Exhibit 14.20.

**Exhibit 14.20 Industrial Robot End Position Time Series**


```
> data(robot)
> plot(robot,ylab='End Position Offset',xlab='Time')
```

Estimates of the spectrum are displayed in Exhibit 14.21 using the convolution of two modified Daniell spectral windows with $m = 7$ (solid) and with a 10% taper on each end of the series. A plot of this spectral window is shown in the middle of Exhibit 14.3 on page 354. The spectrum was also estimated using a fitted AR(7) model (dotted), the order of which was chosen to minimize the AIC. Given the length of the 95% confidence interval shown, we can conclude that the peak at around a frequency of 0.15 in both estimates is probably real, but those shown at higher frequencies may well be spurious. There is a lot of power shown at very low frequencies, and this agrees with the slowly drifting nature of the series that may be seen in the time series plot in Exhibit 14.20.

Exhibit 14.21 Estimated Spectrum for the Industrial Robot
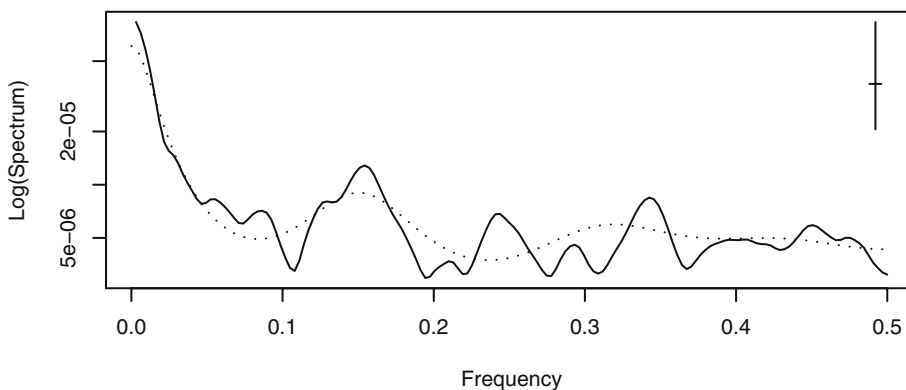
```
> spec(robot,spans=c(7,7),taper=.1,sub='',xlab='Frequency',
    ylab='Log(Spectrum)')
> s=spec(robot,method='ar',plot=F)
> lines(s$freq,s$spec,lty='dotted')
```
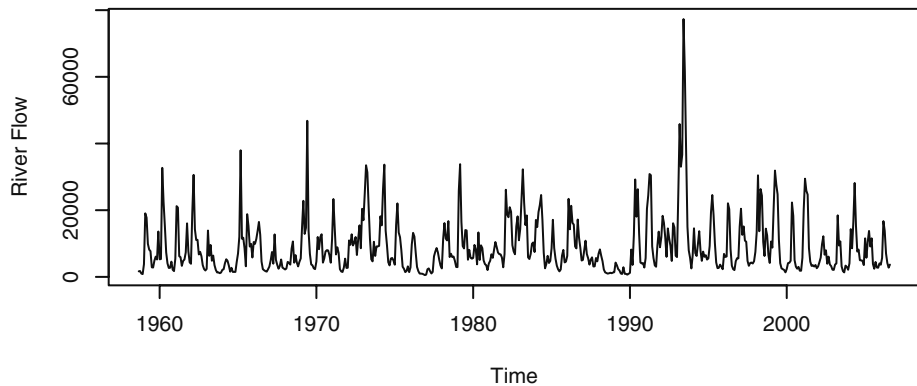
River Flow

Exhibit 14.22 shows monthly river flow for the Iowa River measured at Wapello, Iowa, for the period September 1958 through August 2006. The data are quite skewed toward the high values, but this was greatly improved by taking logarithms for the analysis.

Exhibit 14.22 River Flow Time Series


```
> data(flow); plot(flow,ylab='River Flow')
```

The sample size for these data is 576 with a square root of 24. The bandwidth of a modified Daniell spectral window is about 0.01. After some experimentation with several spectral window bandwidths, we decided that such a window smoothed too much and we instead used a convolution of two such windows, each with *span* = 7. The bandwidth of this convolved window is about 0.0044. The smoothed spectral density estimate is shown as a solid curve in Exhibit 14.23 together with an estimate based on an AR(7) model (dotted) chosen to minimize the AIC. The prominent peak at frequency 1/12 represents the strong annual seasonality. There are smaller secondary peaks at about $f \approx 0.17$ and $f \approx 0.25$ that correspond to multiples of the fundamental frequency of 1/12. They are higher harmonics of the annual frequency.

Exhibit 14.23 Log(Spectrum) of Log(Flow)
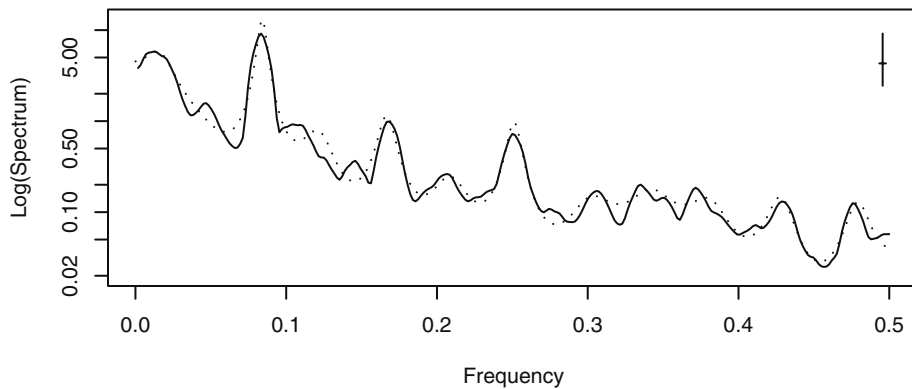

```
> spec(log(flow),spans=c(7,7),ylim=c(.02,13),sub='',
    ylab='Log(Spectrum)',xlab='Frequency')
> s=spec(log(flow),method='ar',plot=F)
> lines(s$freq,s$spec,lty='dotted')
```

Monthly Milk Production

The top portion of Exhibit 11.14 on page 264, showed U.S. monthly milk production from January 1994 through December of 2005. There is a substantial upward trend together with seasonality. We first remove the upward trend with a simple linear time trend model and consider the residuals from that regression—the seasonals. After trying several spectral bandwidths, we decided to use a convolution of two modified Daniell windows, each with *span* = 3. We believe that otherwise there was too much smoothing. This was confirmed by estimating an AR spectrum that ended up fitting an AR of order 15 with peaks at the same frequencies. Notice that the peaks shown in Exhibit 14.24 are located at frequencies 1/12, 2/12,…, 6/12, with the peak at 1/12 showing the most power.

Exhibit 14.24 Estimated Spectrum for Milk Production Seasonals
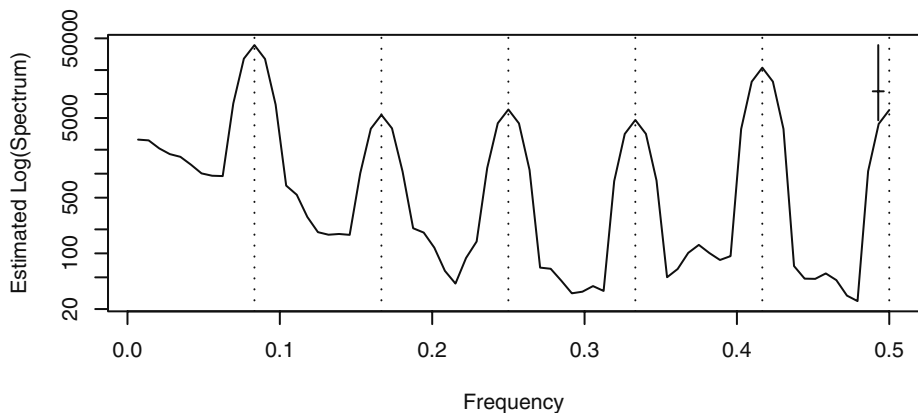

```
> data(milk)
> spec(milk,spans=c(3,3),detrend=T,sub='',
    ylab='Estimated Log(Spectrum)',xlab='Frequency')
> abline(v=seq(1:6)/12,lty='dotted')
```

For a final example in this section, consider the time series shown in Exhibit 14.25. These plots display the first 400 points of two time series of lengths 4423 and 4417, respectively. The complete series were created by recording a trombonist and a euphoniumist each sustaining a B flat (just below middle C) for about 0.4 seconds. The original recording produced data sampled at 44.1 MHz, but this was reduced by subsampling every fourth data point for the analysis shown. Trombones and euphonia are both brass wind instruments that play in the same range, but they have different sized and shaped tubing. The euphonium has larger tubing (a larger bore) that is mostly conical in shape, while the tenor trombone is mostly cylindrical in shape and has a smaller bore. The euphonium sound is considered more mellow than the bright, brassy sound of the trombone. When one listens to these notes being played, they sound rather similar. Our ques-

tion is: Does the tubing shape and size affect the harmonics (overtones) enough that the differences may be seen in the spectra of these sounds?

___

Exhibit 14.25 Trombone and Euphonium Playing Bb

Trombone Bb
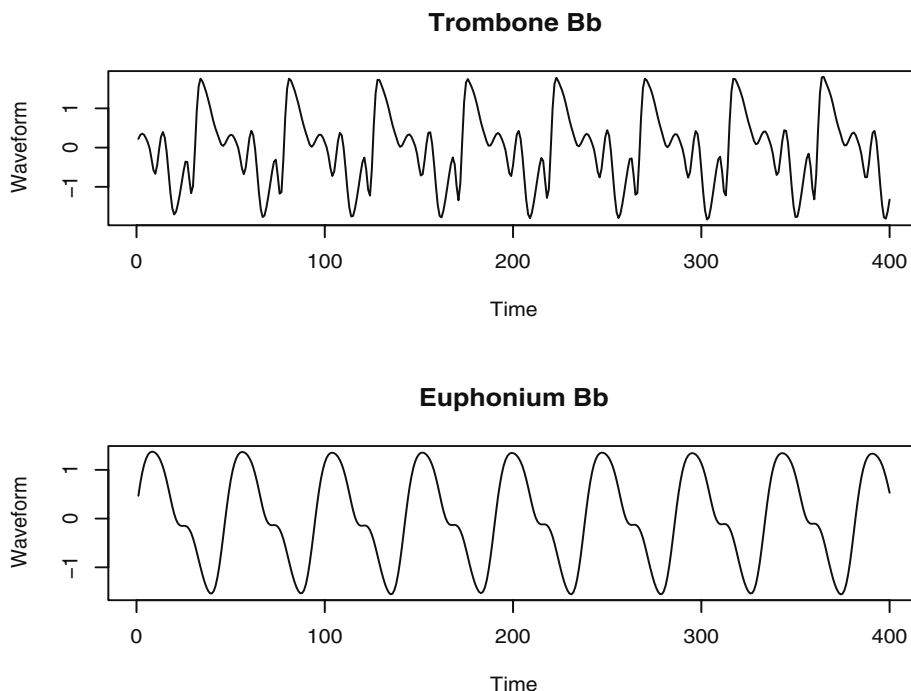
Euphonium Bb


```
> win.graph(width=4.875,height=4,pointsize=8)
> data(tbone); data(euph); oldpar=par; par(mfrow=(c(2,1)))
> trombone=(tbone-mean(tbone))/sd(tbone)
> euphonium=(euph-mean(euph))/sd(euph)
> plot(window(trombone,end=400),main='Trombone Bb',
    ylab='Waveform',yaxp=c(-1,+1,2))
> plot(window(euphonium,end=400),main='Euphonium Bb',
    ylab='Waveform',yaxp=c(-1,+1,2)); par=oldpar
```
___

Exhibit 14.26 displays the estimated spectra for the two waveforms. The solid curve is for the euphonium, and the dotted curve is for the trombone. We used the convolution of two modified Daniell spectral windows, each with *span* = 11, on both series. Since both series are essentially the same length, the bandwidths will both be about 0.0009 and barely perceptible on the bandwidth/confidence interval crosshair shown on the graph.
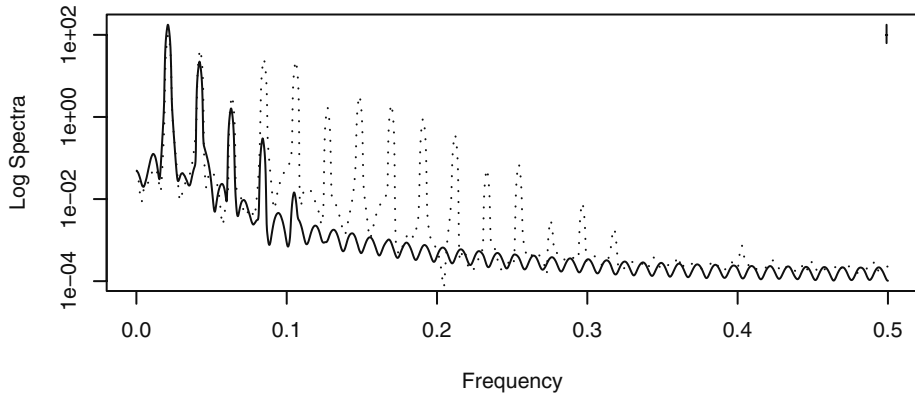
The first four major peaks occur at the same frequencies, but clearly the trombone has much more spectral power at distinct higher harmonic frequencies. It is suggested

that this may account for the more brassy nature of the trombone sound as opposed to
the more mellow sound of the euphonium.

Exhibit 14.26 Spectra for Trombone (dotted) and Euphonium (solid)


```
> win.graph(width=4.875,height=2.5,pointsize=8)
> spec(euph,spans=c(11,11),ylab='Log Spectra',
    xlab='Frequency',sub='')
> s=spec(tbone,spans=c(11,11),plot=F)
> lines(s$freq,s$spec,lty='dotted')
```

14.9 Other Methods of Spectral Estimation

Prior to widespread use of the fast Fourier transform, computing and smoothing the
sample spectrum was extremely intensive computationally—especially for long time
series. Lag window estimators were used to partially mitigate the computational diffi-
culties.

Lag Window Estimators

Consider the sample spectrum and smoothed sample spectrum. We have

$$\bar{S}(f) = \sum_{k=-m}^{m} W(k)\hat{S}\left(f+\frac{k}{n}\right)$$

$$= \sum_{k=-m}^{m} W(k)\left[\sum_{j=-n+1}^{n-1} \hat{\gamma}_j e^{-2\pi i\left(f+\frac{k}{n}\right)j}\right] \qquad (14.9.1)$$

$$= \sum_{j=-n+1}^{n-1} \hat{\gamma}_j\left[\sum_{k=-m}^{m} W(k)e^{-2\pi i\frac{k}{n}j}\right]e^{-2\pi ifj}$$

or

$$\bar{S}(f) = \sum_{j=-n+1}^{n-1} \hat{\gamma}_j w\left(\frac{j}{n}\right) e^{-2\pi ifj} \qquad (14.9.2)$$

where

$$w\left(\frac{j}{n}\right) = \sum_{k=-m}^{m} W(k) e^{-2\pi ik\left(\frac{j}{n}\right)} \qquad (14.9.3)$$

Equation (14.9.2) suggests defining and investigating a class of spectral estimators defined as

$$\tilde{S}(f) = \sum_{j=-n+1}^{n-1} w\left(\frac{j}{n}\right) \hat{\gamma}_j \cos(2\pi fj) \qquad (14.9.4)$$

where the function $w(x)$ has the properties

$$\begin{aligned} w(x) &= w(-x) \\ w(0) &= 1 \\ w(x) &\leq 1 \qquad \text{for } |x| \leq 1 \end{aligned} \qquad (14.9.5)$$

The function $w(x)$ is called a *lag window* and determines how much weight is given to the sample autocovariance at each lag.

The rectangular lag window is defined by

$$w(x) = 1 \qquad \text{for } |x| \leq 1 \qquad (14.9.6)$$

and the corresponding lag window spectral estimator is simply the sample spectrum. This estimator clearly gives too much weight to large lags where the sample autocovariances are based on too few data points and are unreliable.

The next simplest lag window is the truncated rectangular lag window, which simply omits large lags from the computation. It is defined as

$$w\left(\frac{j}{n}\right) = 1 \qquad \text{for } |j| \leq m \qquad (14.9.7)$$

where the computational advantage is achieved by choosing m much smaller than n.

The triangular, or Bartlett, lag window downweights higher lags linearly and is defined as

$$w\left(\frac{j}{n}\right) = 1 - \left|\frac{j}{m}\right| \qquad \text{for } |j| \leq m \qquad (14.9.8)$$

Other common lag windows are associated with the names of Parzen, Tukey-Hamming, and Tukey-Hanning. We will not pursue these further here, but much more information on the lag window approach to spectral estimation may be found in the books of Bloomfield (2000), Brillinger (2001), Brockwell and Davis (1991), and Priestley (1981).

Other Smoothing Methods

Other methods for smoothing the sample spectrum have been proposed. Kooperberg et al. (1995) proposed using splines to estimate the spectral distribution. Fan and Kreutzberger (1998) investigated local smoothing polynomials and Whittle's likelihood for spectral estimation. This approach uses automatic bandwidth selection to smooth the sample spectrum. See also Yoshihide (2006), Jiang and Hui (2004), and Fay et al. (2002).

14.10 Summary

Given the undesirable characteristics of the sample spectral density, we introduced the smoothed sample spectral density and showed that it could be constructed to improve the properties. The important topics of bias, variance, leakage, bandwidth, and tapering were investigated. A procedure for forming confidence intervals was discussed, and all of the ideas were illustrated with both real and simulated time series data.

EXERCISES

14.1 Consider the variance of $\bar{S}(f)$ with the Daniell spectral window. Instead of using Equation (14.2.4) on page 355, use the fact that $2\hat{S}(f)/S(f)$ has approximately a chi-square distribution with two degrees of freedom to show that the smoothed sample spectral density has an approximate variance of $S^2(f)/(2m+1)$.

14.2 Consider various convolutions of the simple Daniell rectangular spectral window.

    (a) Construct a panel of three plots similar to those shown in Exhibit 14.3 on page 354 but with the Daniell spectral window and with $m = 5$. The middle graph should be the convolution of two Daniell windows and the leftmost graph the convolution of three Daniell windows.

    (b) Evaluate the bandwidths and degrees of freedom for each of the spectral windows constructed in part (a). Use $n = 100$.

    (c) Construct another panel of three plots similar to those shown in Exhibit 14.3 but with the modified Daniell spectral window. This time use $m = 5$ for the first graph and convolve two with $m = 5$ and $m = 7$ for the second. Convolve three windows with m's of 5, 7, and 11 for the third graph.

    (d) Evaluate the bandwidths and degrees of freedom for each of the spectral windows constructed in part (c). Use $n = 100$.

14.3 For the Daniell rectangular spectral window show that

(a) $\dfrac{1}{n^2}\displaystyle\sum_{k=-m}^{m} k^2 W_m(k) = \dfrac{2}{n^2(2m+1)}\left(\dfrac{m^3}{3}+\dfrac{m^2}{2}+\dfrac{m}{6}\right)$

(b) Show that if m is chosen as $m = c\sqrt{n}$ for any constant c, then the right-hand side of the expression in part (a) tends to zero as n goes to infinity.

(c) Show that if $m = c\sqrt{n}$ for any constant c, then the approximate variance of the smoothed spectral density given by the right-hand side of Equation (14.2.4) on page 355 tends to zero as n tends to infinity.

14.4 Suppose that the distribution of $\bar{S}(f)$ is to be approximated by a multiple of a chi-square variable with degrees of freedom ν, so that $\bar{S}(f) \approx c\chi^2_\nu$. Using the approximate variance of $\bar{S}(f)$ given in Equation (14.2.4) on page 355 and the fact that $\bar{S}(f)$ is approximately unbiased, equate means and variances and find the values for c and ν (thus establishing Equation (14.4.2) on page 356).

14.5 Construct a time series of length $n = 48$ according to the expression

$$Y_t = \sin[2\pi(0.28)t]$$

Display the periodogram of the series and explain its appearance.

14.6 Estimate the spectrum of the Los Angeles annual rainfall time series. The data are in the file named larain. Because of the skewness in the series, use the logarithms of the raw rainfall values. The square root of the series length suggests a value for the *span* of about 11. Use the modified Daniell spectral window, and be sure to set the vertical limits of the plot so that you can see the whole confidence interval guide. Comment on the estimated spectrum.

14.7 The file named spots1 contains annual sunspot numbers for 306 years from 1700 through 2005.

(a) Display the time series plot of these data. Does stationarity seem reasonable for this series?

(b) Estimate the spectrum using a modified Daniell spectral window convoluted with itself and a *span* of 3 for both. Interpret the plot.

(c) Estimate the spectrum using an AR model with the order chosen to minimize the AIC. Interpret the plot. What order was selected?

(d) Overlay the estimates obtained in parts (b) and (c) above onto one plot. Do they agree to a reasonable degree?

14.8 Consider the time series of average monthly temperatures in Dubuque, Iowa. The data are in the file named tempdub and cover from January 1964 to December 1975 for an n of 144.

(a) Estimate the spectrum using a variety of *span* values for the modified Daniell spectral window.

(b) In your opinion, which of the estimates in part (a) best represents the spectrum of the process? Be sure to use bandwidth considerations and confidence limits to back up your argument.

380 Estimating the Spectrum

14.9 An EEG (electroencephalogram) time series is given in the data file named eeg. An electroencephalogram is a noninvasive test used to detect and record the electrical activity generated in the brain. These data were measured at a sampling rate of 256 per second and came from a patient suffering a seizure. The total record length is $n = 13,000$—or slightly less than one minute.

    (a) Display the time series plot and decide if stationarity seems reasonable.

    (b) Estimate the spectrum using a modified Daniell spectral window convolved with itself and a *span* of 51 for both components of the convolution. Interpret the plot.

    (c) Estimate the spectrum using an AR model with the order chosen to minimize the AIC. Interpret the plot. What order was selected?

    (d) Overlay the estimates obtained in parts (b) and (c) above onto one plot. Do they agree to a reasonable degree?

14.10 The file named electricity contains monthly U. S. electricity production values from January 1994 to December 2005. A time series plot of the logarithms of these values is shown in Exhibit 11.14 on page 264. Since there is an upward trend and increasing variability at higher levels in these data, use the first difference of the logarithms for the remaining analysis.

    (a) Construct a time series plot of the first difference of the logarithms of the electricity values. Does a stationary model seem warranted at this point?

    (b) Display the smoothed spectrum of the first difference of the logarithms using a modified Daniell spectral window and *span* values of 25, 13, and 7. Interpret the results.

    (c) Now use a spectral window that is a convolution of two modified Daniell windows each with *span* = 3. Also use a 10% taper. Interpret the results.

    (d) Estimate the spectrum using an AR model with the order chosen to minimize the AIC. Interpret the plot. What order was selected?

    (e) Overlay the estimates obtained in parts (c) and (d) above onto one plot. Do they agree to a reasonable degree?

14.11 Consider the monthly milk production time series used in Exhibit 14.24 on page 374. The data are in the file named milk.

    (a) Estimate the spectrum using a spectral window that is a convolution of two modified Daniell windows each with *span* = 7. Compare these results with those shown in Exhibit 14.24.

    (b) Estimate the spectrum using a single modified Daniell spectral window with *span* = 7. Compare these results with those shown in Exhibit 14.24 and those in part (a).

    (c) Finally, estimate the spectrum using a single modified Daniell spectral window with *span* = 11. Compare these results with those shown in Exhibit 14.24 and those in parts (a) and (b).

    (d) Among the four different estimates considered here, which do you prefer and why?

14.12 Consider the river flow series displayed in Exhibit 14.22 on page 372. An esti-
mate of the spectrum is shown in Exhibit 14.23 on page 373. The data are in the
file named flow.

    (a) Here n = 576 and $\sqrt{n}$ = 24. Estimate the spectrum using *span* = 25 with the
modified Daniell spectral window. Compare your results with those shown in
Exhibit 14.23.

    (b) Estimate the spectrum using *span* = 13 with the modified Daniell spectral
window and compare your results to those obtained in part (a) and in Exhibit
14.23.

14.13 The time series in the file named tuba contains about 0.4 seconds of digitized
sound from a tuba playing a B flat one octave and one note below middle C.

    (a) Display a time series plot of the first 400 of these data and compare your
results with those shown in Exhibit 14.25 on page 375, for the trombone and
euphonium.

    (b) Estimate the spectrum of the tuba time series using a convolution of two mod-
ified Daniell spectral windows, each with *span* = 11.

    (c) Compare the estimated spectrum obtained in part (b) with those of the trom-
bone and euphonium shown in Exhibit 14.26 on page 376. (You may want to
overlay several of these spectra.) Remember that the tuba is playing one
octave lower than the two other instruments.

    (d) Do the higher-frequency components of the spectrum for the tuba look more
like those of the trombone or those of the euphonium? (Hint: The euphonium
is sometimes called a tenor tuba!)

Appendix K: Tapering and the Dirichlet Kernel

Suppose $Y_t = \cos(2\pi f_0 t + \Phi)$ for t = 1, 2,..., n, where f_0 is not necessarily a Fourier
frequency. Since it will not affect the periodogram, we will actually suppose that

$$Y_t = e^{2\pi i f_0 t} \qquad\qquad (14.K.1)$$

in order to simplify the mathematics. Then the discrete-time Fourier transform of this
sequence is given by

$$\frac{1}{n}\sum_{t=1}^{n} Y_t e^{-2\pi i f t} = \frac{1}{n}\sum_{t=1}^{n} e^{2\pi i (f_0 - f) t} \qquad\qquad (14.K.2)$$

By Equations (13.J.7) and (13.J.8) on page 350, for any z,

$$\frac{1}{n}\sum_{t=1}^{n} e^{2\pi i z t} = \frac{1}{n} e^{2\pi i z}\frac{(e^{2\pi i n z} - 1)}{(e^{2\pi i z} - 1)}$$

$$= \frac{1}{n} e^{\pi i (n+1) z}\frac{(e^{\pi i n z} - e^{-\pi i n z})}{(e^{\pi i z} - e^{-\pi i z})}$$

so that

$$\frac{1}{n}\sum_{t=1}^{n}e^{2\pi izt}= e^{\pi i(n+1)z}\left[\frac{1}{n}\frac{\sin(\pi nz)}{\sin(\pi z)}\right] \tag{14.K.3}$$

The function

$$D(z) = \frac{1}{n}\frac{\sin(\pi nz)}{\sin(\pi z)} \tag{14.K.4}$$

is the Dirichlet kernel shown on the left-hand side of Exhibit 14.7 on page 360 for $n = 100$. These results lead to the following relationship for the periodogram of Y_t:

$$I(f) \propto |D(f-f_0)|^2 \tag{14.K.5}$$

Remember that for all Fourier frequencies $D(f) = 0$, so that this window has no effect at those frequencies. Leakage occurs when there is substantial power at non-Fourier frequencies. Now consider tapering Y_t with a cosine bell. We have

$$\tilde{Y}_t = \frac{1}{2}\left\{1 - \cos\left[\frac{2\pi(t-0.5)}{n}\right]\right\}Y_t$$

$$= \frac{1}{2}e^{2\pi if_0t} - \frac{1}{4}e^{2\pi if_0t + 2\pi i(t-\frac{1}{2})/n} - \frac{1}{4}e^{2\pi if_0t - 2\pi i(t-\frac{1}{2})/n} \tag{14.K.6}$$

and after some more algebra we obtain

$$\frac{1}{n}\sum_{t=1}^{n}\tilde{Y}_t e^{-2\pi ift}$$

$$= e^{\pi i(n+1)(f_0-f)t}\left[\frac{1}{4}D\left(f-f_0-\frac{1}{n}\right) + \frac{1}{2}D(f-f_0) + \frac{1}{4}D\left(f-f_0+\frac{1}{n}\right)\right] \tag{14.K.7}$$

The function

$$\tilde{D}(f) = \frac{1}{4}D\left(f-f_0-\frac{1}{n}\right) + \frac{1}{2}D(f-f_0) + \frac{1}{4}D\left(f-f_0+\frac{1}{n}\right) \tag{14.K.8}$$

is the tapered or modified Dirichlet kernel that is plotted on the right-hand side of Exhibit 14.7 on page 360 for $n = 100$. The periodogram of the tapered series is proportional to $|(\tilde{D}(f))|^2$, and the side lobe problem is substantially mitigated.