

# CHAPTER 8

## MODEL DIAGNOSTICS

We have now discussed methods for specifying models and for efficiently estimating the parameters in those models. Model diagnostics, or model criticism, is concerned with testing the goodness of fit of a model and, if the fit is poor, suggesting appropriate modifications. We shall present two complementary approaches: analysis of residuals from the fitted model and analysis of overparameterized models; that is, models that are more general than the proposed model but that contain the proposed model as a special case.

### 8.1 Residual Analysis

---

We already used the basic ideas of residual analysis in Section 3.6 on page 42 when we checked the adequacy of fitted deterministic trend models. With autoregressive models, residuals are defined in direct analogy to that earlier work. Consider in particular an AR(2) model with a constant term:

$$Y_t = \phi_1 Y_{t-1} + \phi_2 Y_{t-2} + \theta_0 + e_t \quad (8.1.1)$$

Having estimated  $\phi_1$ ,  $\phi_2$ , and  $\theta_0$ , the residuals are defined as

$$\hat{e}_t = Y_t - \hat{\phi}_1 Y_{t-1} - \hat{\phi}_2 Y_{t-2} - \hat{\theta}_0 \quad (8.1.2)$$

For general ARMA models containing moving average terms, we use the inverted, infinite autoregressive form of the model to define residuals. For simplicity, we assume that  $\theta_0$  is zero. From the inverted form of the model, Equation (4.5.5) on page 80, we have

$$Y_t = \pi_1 Y_{t-1} + \pi_2 Y_{t-2} + \pi_3 Y_{t-3} + \cdots + e_t$$

so that the residuals are defined as

$$\hat{e}_t = Y_t - \hat{\pi}_1 Y_{t-1} - \hat{\pi}_2 Y_{t-2} - \hat{\pi}_3 Y_{t-3} - \cdots \quad (8.1.3)$$

Here the  $\pi$ 's are not estimated directly but rather implicitly as functions of the  $\phi$ 's and  $\theta$ 's. In fact, the residuals are not calculated using this equation but as a by-product of the estimation of the  $\phi$ 's and  $\theta$ 's. In Chapter 9, we shall argue, that

$$\hat{Y}_t = \hat{\pi}_1 Y_{t-1} + \hat{\pi}_2 Y_{t-2} + \hat{\pi}_3 Y_{t-3} + \cdots$$

is the best forecast of  $Y_t$  based on  $Y_{t-1}, Y_{t-2}, Y_{t-3}, \dots$ . Thus Equation (8.1.3) can be rewritten as

$$\text{residual} = \text{actual} - \text{predicted}$$

in direct analogy with regression models. Compare this with Section 3.6 on page 42.


If the model is correctly specified and the parameter estimates are reasonably close to the true values, then the residuals should have nearly the properties of white noise. They should behave roughly like independent, identically distributed normal variables with zero means and common standard deviations. Deviations from these properties can help us discover a more appropriate model.

### Plots of the Residuals

Our first diagnostic check is to inspect a plot of the residuals over time. If the model is adequate, we expect the plot to suggest a rectangular scatter around a zero horizontal level with no trends whatsoever.

Exhibit 8.1 shows such a plot for the standardized residuals from the AR(1) model fitted to the industrial color property series. Standardization allows us to see residuals of unusual size much more easily. The parameters were estimated using maximum likelihood. This plot supports the model, as no trends are present.

**Exhibit 8.1 Standardized Residuals from AR(1) Model of Color**




```
> win.graph(width=4.875,height=3,pointsize=8)
> data(color)
> m1.color=arima(color,order=c(1,0,0)); m1.color
> plot(rstandard(m1.color),ylab='Standardized Residuals',
      type='o'); abline(h=0)
```

As a second example, we consider the Canadian hare abundance series. We estimate a subset AR(3) model with  $\phi_2$  set to zero, as suggested by the discussion following Exhibit 7.8 on page 166. The estimated model is

$$\sqrt{Y_t} = 3.483 + 0.919\sqrt{Y_{t-1}} - 0.5313\sqrt{Y_{t-3}} + e_t \quad (8.1.4)$$

and the time series plot of the standardized residuals from this model is shown in Exhibit 8.2. Here we see possible reduced variation in the middle of the series and increased variation near the end of the series—not exactly an ideal plot of residuals.<sup>†</sup>

**Exhibit 8.2 Standardized Residuals from AR(3) Model for Sqrt(Hare)**




---

```

> data(hare)
> m1.hare=arima(sqrt(hare),order=c(3,0,0)); m1.hare
> m2.hare=arima(sqrt(hare),order=c(3,0,0),fixed=c(NA,0,NA,NA))
> m2.hare
> # Note that the intercept term given in R is actually the mean
  in the centered form of the ARMA model; that is, if
  y(t)=sqrt(hare)-intercept, then the model is
  y(t)=0.919*y(t-1)-0.5313*y(t-3)+e(t)
> # So the 'true' intercept equals 5.6889*(1-0.919+0.5313)=3.483
> plot(rstandard(m2.hare),ylab='Standardized Residuals',type='o')
> abline(h=0)

```

---

Exhibit 8.3 displays the time series plot of the standardized residuals from the IMA(1,1) model estimated for the logarithms of the oil price time series. The model was fitted using maximum likelihood estimation. There are at least two or three residuals early in the series with magnitudes larger than 3—very unusual in a standard normal distribution.<sup>‡</sup> Ideally, we should go back to those months and try to learn what outside factors may have influenced unusually large drops or unusually large increases in the price of oil.


<sup>†</sup> The seemingly large negative standardized residuals are not outliers according to the Bonferroni outlier criterion with critical values  $\pm 3.15$ .

<sup>‡</sup> The Bonferroni critical values with  $n = 241$  and  $\alpha = 0.05$  are  $\pm 3.71$ , so the outliers do appear to be real. We will model them in Chapter 11.

---

**Exhibit 8.3 Standardized Residuals from Log Oil Price IMA(1,1) Model**

---



---


```
> data(oil.price)
> m1.oil=arima(log(oil.price),order=c(0,1,1))
> plot(rstandard(m1.oil),ylab='Standardized residuals',type='l')
> abline(h=0)
```

---

**Normality of the Residuals**

As we saw in Chapter 3, quantile-quantile plots are an effective tool for assessing normality. Here we apply them to residuals.

A quantile-quantile plot of the residuals from the AR(1) model estimated for the industrial color property series is shown in Exhibit 8.4. The points seem to follow the straight line fairly closely—especially the extreme values. This graph would not lead us to reject normality of the error terms in this model. In addition, the Shapiro-Wilk normality test applied to the residuals produces a test statistic of  $W = 0.9754$ , which corresponds to a  $p$ -value of 0.6057, and we would not reject normality based on this test.


**Exhibit 8.4 Quantile-Quantile Plot: Residuals from AR(1) Color Model**


---

```
> win.graph(width=2.5,height=2.5,pointsize=8)
> qqnorm(residuals(m1.color)); qqline(residuals(m1.color))
```

---

The quantile-quantile plot for the residuals from the AR(3) model for the square root of the hare abundance time series is displayed in Exhibit 8.5. Here the extreme values look suspect. However, the sample is small ( $n = 31$ ) and, as stated earlier, the Bonferroni criteria for outliers do not indicate cause for alarm.

**Exhibit 8.5 Quantile-Quantile Plot: Residuals from AR(3) for Hare**


---

```
> qqnorm(residuals(m1.hare)); qqline(residuals(m1.hare))
```


---

Exhibit 8.6 gives the quantile-quantile plot for the residuals from the IMA(1,1) model that was used to model the logarithms of the oil price series. Here the outliers are quite prominent, and we will deal with them in Chapter 11.

---

**Exhibit 8.6 Quantile-Quantile Plot: Residuals from IMA(1,1) Model for Oil**


---




---

```
> qqnorm(residuals(m1.oil)); qqline(residuals(m1.oil))
```

---

### Autocorrelation of the Residuals

To check on the independence of the noise terms in the model, we consider the sample autocorrelation function of the residuals, denoted  $\hat{r}_k$ . From Equation (6.1.3) on page 110, we know that for true white noise and large  $n$ , the sample autocorrelations are approximately uncorrelated and normally distributed with zero means and variance  $1/n$ . Unfortunately, even residuals from a correctly specified model with efficiently estimated parameters have somewhat different properties. This was first explored for multiple-regression models in a series of papers by Durbin and Watson (1950, 1951, 1971) and for autoregressive models in Durbin (1970). The key reference on the distribution of residual autocorrelations in ARIMA models is Box and Pierce (1970), the results of which were generalized in McLeod (1978).

Generally speaking, the residuals are approximately normally distributed with zero means; however, for small lags  $k$  and  $j$ , the variance of  $\hat{r}_k$  can be substantially less than  $1/n$  and the estimates  $\hat{r}_k$  and  $\hat{r}_j$  can be highly correlated. For larger lags, the approximate variance  $1/n$  does apply, and further  $\hat{r}_k$  and  $\hat{r}_j$  are approximately uncorrelated.

As an example of these results, consider a correctly specified and efficiently estimated AR(1) model. It can be shown that, for large  $n$ ,

$$\text{Var}(\hat{r}_1) \approx \frac{\phi^2}{n} \quad (8.1.5)$$

$$\text{Var}(\hat{r}_k) \approx \frac{1 - (1 - \phi^2)\phi^{2k-2}}{n} \quad \text{for } k > 1 \quad (8.1.6)$$

$$\text{Corr}(\hat{r}_1, \hat{r}_k) \approx -\text{sign}(\phi) \frac{(1 - \phi^2)\phi^{k-2}}{1 - (1 - \phi^2)\phi^{2k-2}} \quad \text{for } k > 1 \quad (8.1.7)$$

where

$$\text{sign}(\phi) = \begin{cases} 1 & \text{if } \phi > 0 \\ 0 & \text{if } \phi = 0 \\ -1 & \text{if } \phi < 0 \end{cases}$$

The table in Exhibit 8.7 illustrates these formulas for a variety of values of  $\phi$  and  $k$ . Notice that  $\text{Var}(\hat{r}_1) \approx 1/n$  is a reasonable approximation for  $k \geq 2$  over a wide range of  $\phi$ -values.

**Exhibit 8.7 Approximations for Residual Autocorrelations in AR(1) Models**


$\phi$	0.3	0.5	0.7	0.9	$\phi$	0.3	0.5	0.7	0.9
$k$	Standard deviation of $\hat{r}_k$ times $\sqrt{n}$				Correlation $\hat{r}_1$ with $\hat{r}_k$				
1	0.30	0.50	0.70	0.90	1.00	1.00	1.00	1.00	1.00
2	0.96	0.90	0.87	0.92	-0.95	-0.83	-0.59	-0.21	
3	1.00	0.98	0.94	0.94	-0.27	-0.38	-0.38	-0.18	
4	1.00	0.99	0.97	0.95	-0.08	-0.19	-0.26	-0.16	
5	1.00	1.00	0.99	0.96	-0.02	-0.09	-0.18	-0.14	
6	1.00	1.00	0.99	0.97	-0.01	-0.05	-0.12	-0.13	
7	1.00	1.00	1.00	0.97	-0.00	-0.02	-0.09	-0.12	
8	1.00	1.00	1.00	0.98	-0.00	-0.01	-0.06	-0.10	
9	1.00	1.00	1.00	0.99	-0.00	-0.00	-0.03	-0.08	

If we apply these results to the AR(1) model that was estimated for the industrial color property time series with  $\hat{\phi} = 0.57$  and  $n = 35$ , we obtain the results shown in Exhibit 8.8.

**Exhibit 8.8 Approximate Standard Deviations of Residual ACF values**

Lag $k$	1	2	3	4	5	> 5
$\sqrt{\text{Var}(\hat{r}_k)}$	0.096	0.149	0.163	0.167	0.168	0.169

A graph of the sample ACF of these residuals is shown in Exhibit 8.9. The dashed horizontal lines plotted are based on the large lag standard error of  $\pm 2/\sqrt{n}$ . There is no evidence of autocorrelation in the residuals of this model.

**Exhibit 8.9 Sample ACF of Residuals from AR(1) Model for Color**

```
> win.graph(width=4.875,height=3,pointsize=8)
> acf(residuals(m1.color))
```

For an AR(2) model, it can be shown that

$$\text{Var}(\hat{r}_1) \approx \frac{\phi_2^2}{n} \quad (8.1.8)$$

and

$$\text{Var}(\hat{r}_2) \approx \frac{\phi_2^2 + \phi_1^2(1 + \phi_2)^2}{n} \quad (8.1.9)$$

If the AR(2) parameters are not too close to the stationarity boundary shown in Exhibit 4.17 on page 72, then

$$\text{Var}(\hat{r}_k) \approx \frac{1}{n} \quad \text{for } k \geq 3 \quad (8.1.10)$$

If we fit an AR(2) model<sup>†</sup> by maximum likelihood to the square root of the hare abundance series, we find that  $\hat{\phi}_1 = 1.351$  and  $\hat{\phi}_2 = -0.776$ . Thus we have

$$\sqrt{\text{Var}(\hat{r}_1)} \approx \frac{|-0.776|}{\sqrt{35}} = 0.131$$

$$\sqrt{\text{Var}(\hat{r}_2)} \approx \sqrt{\frac{(-0.776)^2 + (1.351)^2(1 + (-0.776))^2}{35}} = 0.141$$

$$\sqrt{\text{Var}(\hat{r}_k)} \approx 1/\sqrt{35} = 0.169 \quad \text{for } k \geq 3$$


<sup>†</sup> The AR(2) model is not quite as good as the AR(3) model that we estimated earlier, but it still fits quite well and serves as a reasonable example here.



Exhibit 8.10 displays the sample ACF of the residuals from the AR(2) model of the square root of the hare abundance. The lag 1 autocorrelation here equals  $-0.261$ , which is close to 2 standard errors below zero but not quite. The lag 4 autocorrelation equals  $-0.318$ , but its standard error is  $0.169$ . We conclude that the graph does not show statistically significant evidence of nonzero autocorrelation in the residuals.<sup>†</sup>

---

**Exhibit 8.10 Sample ACF of Residuals from AR(2) Model for Hare**




---

```
> acf(residuals(arima(sqrt(hare), order=c(2, 0, 0))))
```

---

With monthly data, we would pay special attention to possible excessive autocorrelation in the residuals at lags 12, 24, and so forth. With quarterly series, lags 4, 8, and so forth would merit special attention. Chapter 10 contains examples of these ideas.

It can be shown that results analogous to those for AR models hold for MA models. In particular, replacing  $\phi$  by  $\theta$  in Equations (8.1.5), (8.1.6), and (8.1.7) gives the results for the MA(1) case. Similarly, results for the MA(2) case can be stated by replacing  $\phi_1$  and  $\phi_2$  by  $\theta_1$  and  $\theta_2$ , respectively, in Equations (8.1.8), (8.1.9), and (8.1.10). Results for general ARMA models may be found in Box and Pierce (1970) and McLeod (1978).

### The Ljung-Box Test

In addition to looking at residual correlations at individual lags, it is useful to have a test that takes into account their magnitudes as a group. For example, it may be that most of the residual autocorrelations are moderate, some even close to their critical values, but, taken together, they seem excessive. Box and Pierce (1970) proposed the statistic

$$Q = n(\hat{r}_1^2 + \hat{r}_2^2 + \dots + \hat{r}_K^2) \quad (8.1.11)$$

to address this possibility. They showed that if the correct ARMA( $p, q$ ) model is estimated, then, for large  $n$ ,  $Q$  has an approximate chi-square distribution with  $K - p - q$

---

<sup>†</sup> Recall that an AR(3) model fits these data even better and has even less autocorrelation in its residuals, see Exercise 8.7.

degrees of freedom. Fitting an erroneous model would tend to inflate  $Q$ . Thus, a general “portmanteau” test would reject the ARMA( $p, q$ ) model if the observed value of  $Q$  exceeded an appropriate critical value in a chi-square distribution with  $K - p - q$  degrees of freedom. (Here the maximum lag  $K$  is selected somewhat arbitrarily but large enough that the  $\psi$ -weights are negligible for  $j > K$ .)

The chi-square distribution for  $Q$  is based on a limit theorem as  $n \rightarrow \infty$ , but Ljung and Box (1978) subsequently discovered that even for  $n = 100$ , the approximation is not satisfactory. By modifying the  $Q$  statistic slightly, they defined a test statistic whose null distribution is much closer to chi-square for typical sample sizes. The modified Box-Pierce, or **Ljung-Box**, statistic is given by

$$Q_* = n(n+2) \left( \frac{\hat{r}_1^2}{n-1} + \frac{\hat{r}_2^2}{n-2} + \cdots + \frac{\hat{r}_K^2}{n-K} \right) \quad (8.1.12)$$

Notice that since  $(n+2)/(n-k) > 1$  for every  $k \geq 1$ , we have  $Q_* > Q$ , which partly explains why the original statistic  $Q$  tended to overlook inadequate models. More details on the exact distributions of  $Q_*$  and  $Q$  for finite samples can be found in Ljung and Box (1978), see also Davies, Triggs, and Newbold (1977).

Exhibit 8.11 lists the first six autocorrelations of the residuals from the AR(1) fitted model for the color property series. Here  $n = 35$ .

---

**Exhibit 8.11 Residual Autocorrelation Values from AR(1) Model for Color**

---

Lag $k$	1	2	3	4	5	6
Residual ACF	-0.051	0.032	0.047	0.021	-0.017	-0.019

---

```
> acf(residuals(m1.color), plot=F)$acf
> signif(acf(residuals(m1.color), plot=F)$acf[1:6], 2)
> # display the first 6 acf values to 2 significant digits
```


---

The Ljung-Box test statistic with  $K = 6$  is equal to

$$Q_* = 35(35+2) \left( \frac{(-0.051)^2}{35-1} + \frac{(0.032)^2}{35-2} + \frac{(0.047)^2}{35-3} + \frac{(0.021)^2}{35-4} + \frac{(-0.017)^2}{35-5} + \frac{(-0.019)^2}{35-6} \right) \approx 0.28$$

This is referred to a chi-square distribution with  $6 - 1 = 5$  degrees of freedom. This leads to a  $p$ -value of 0.998, so we have no evidence to reject the null hypothesis that the error terms are uncorrelated.

Exhibit 8.12 shows three of our diagnostic tools in one display—a sequence plot of the standardized residuals, the sample ACF of the residuals, and  $p$ -values for the Ljung-Box test statistic for a whole range of values of  $K$  from 5 to 15. The horizontal dashed line at 5% helps judge the size of the  $p$ -values. In this instance, everything looks very good. The estimated AR(1) model seems to be capturing the dependence structure of the color property time series quite well.

**Exhibit 8.12 Diagnostic Display for the AR(1) Model of Color Property**

```
> win.graph(width=4.875,height=4.5)
> tsdiag(m1.color,gof=15,omit.initial=F)
```

As in Chapter 3, the runs test may also be used to assess dependence in error terms via the residuals. Applying the test to the residuals from the AR(3) model for the Canadian hare abundance series, we obtain expected runs of 16.09677 versus observed runs of 18. The corresponding  $p$ -value is 0.602, so we do not have statistically significant evidence against independence of the error terms in this model.

## 8.2 Overfitting and Parameter Redundancy

Our second basic diagnostic tool is that of **overfitting**. After specifying and fitting what we believe to be an adequate model, we fit a slightly more general model; that is, a model “close by” that contains the original model as a special case. For example, if an AR(2) model seems appropriate, we might overfit with an AR(3) model. The original AR(2) model would be confirmed if:

1. the estimate of the additional parameter,  $\phi_3$ , is not significantly different from zero, and
2. the estimates for the parameters in common,  $\phi_1$  and  $\phi_2$ , do not change significantly from their original estimates.

As an example, we have specified, fitted, and examined the residuals of an AR(1) model for the industrial color property time series. Exhibit 8.13 displays the output from the R software from fitting the AR(1) model, and Exhibit 8.14 shows the results from fitting an AR(2) model to the same series. First note that, in Exhibit 8.14, the estimate of  $\phi_2$  is not statistically different from zero. This fact supports the choice of the AR(1) model. Secondly, we note that the two estimates of  $\phi_1$  are quite close—especially when we take into account the magnitude of their standard errors. Finally, note that while the AR(2) model has a slightly larger log-likelihood value, the AR(1) fit has a smaller AIC value. The penalty for fitting the more complex AR(2) model is sufficient to choose the simpler AR(1) model.

---

### Exhibit 8.13 AR(1) Model Results for the Color Property Series

Coefficients: <sup>†</sup>	ar1	Intercept <sup>‡</sup>
	0.5705	74.3293
<b>s.e.</b>	0.1435	1.9151

sigma^2 estimated as 24.83: log-likelihood = -106.07, AIC = 216.15

---

<sup>†</sup> `m1.color # R code to obtain table`

<sup>‡</sup> Recall that the intercept here is the estimate of the process mean  $\mu$ —not  $\theta_0$ .

---

### Exhibit 8.14 AR(2) Model Results for the Color Property Series

Coefficients:	ar1	ar2	Intercept
	0.5173	0.1005	74.1551
<b>s.e.</b>	0.1717	0.1815	2.1463

sigma^2 estimated as 24.6: log-likelihood = -105.92, AIC = 217.84

---

`> arima(color, order=c(2, 0, 0))`

---

A different overfit for this series would be to try an ARMA(1,1) model. Exhibit 8.15 displays the results of this fit. Notice that the standard errors of the estimated coefficients for this fit are rather larger than what we see in Exhibits 8.13 and 8.14. Regardless, the estimate of  $\phi_1$  from this fit is not significantly different from the estimate in Exhibit 8.13. Furthermore, as before, the estimate of the new parameter,  $\theta$ , is not significantly different from zero. This adds further support to the AR(1) model.

**Exhibit 8.15 Overfit of an ARMA(1,1) Model for the Color Series**

Coefficients:	ar1	ma1	Intercept
	0.6721	-0.1467	74.1730
<b>s.e.</b>	0.2147	0.2742	2.1357

sigma^2 estimated as 24.63: log-likelihood = -105.94, AIC = 219.88

```
> arima(color, order=c(1, 0, 1))
```

As we have noted, any ARMA( $p, q$ ) model can be considered as a special case of a more general ARMA model with the additional parameters equal to zero. However, when generalizing ARMA models, we must be aware of the problem of **parameter redundancy** or **lack of identifiability**.

To make these points clear, consider an ARMA(1,2) model:

$$Y_t = \phi Y_{t-1} + e_t - \theta_1 e_{t-1} - \theta_2 e_{t-2} \quad (8.2.1)$$

Now replace  $t$  by  $t-1$  to obtain

$$Y_{t-1} = \phi Y_{t-2} + e_{t-1} - \theta_1 e_{t-2} - \theta_2 e_{t-3} \quad (8.2.2)$$

If we multiply both sides of Equation (8.2.2) by *any constant*  $c$  and then subtract it from Equation (8.2.1), we obtain (after rearranging)

$$Y_t - (\phi + c)Y_{t-1} + \phi c Y_{t-2} = e_t - (\theta_1 + c)e_{t-1} - (\theta_2 - \theta_1 c)e_{t-2} + c\theta_2 e_{t-3}$$

This apparently defines an ARMA(2,3) process. But notice that we have the factorizations

$$1 - (\phi + c)x + \phi c x^2 = (1 - \phi x)(1 - cx)$$

and

$$1 - (\theta_1 + c)x - (\theta_2 - c\theta_1)x^2 + c\theta_2 x^3 = (1 - \theta_1 x - \theta_2 x^2)(1 - cx)$$

Thus the AR and MA characteristic polynomials in the ARMA(2,3) process have a common factor of  $(1 - cx)$ . Even though  $Y_t$  does satisfy the ARMA(2,3) model, clearly the parameters in that model are not unique—the constant  $c$  is completely arbitrary. We say that we have **parameter redundancy** in the ARMA(2,3) model.<sup>†</sup>

The implications for fitting and overfitting models are as follows:

1. Specify the original model carefully. If a simple model seems at all promising, check it out before trying a more complicated model.
2. When overfitting, do not increase the orders of both the AR and MA parts of the model simultaneously.

<sup>†</sup> In backshift notation, if  $\phi(B)Y_t = \theta(B)e_t$  is a correct model, then so is  $(1 - cB)\phi(B)Y_t = (1 - cB)\theta(B)e_t$  for *any* constant  $c$ . To have unique parameterization in an ARMA model, we must cancel any common factors in the AR and MA characteristic polynomials.

3. Extend the model in directions suggested by the analysis of the residuals. For example, if after fitting an MA(1) model, substantial correlation remains at lag 2 in the residuals, try an MA(2), not an ARMA(1,1).

As an example, consider the color property series once more. We have seen that an AR(1) model fits quite well. Suppose we try an ARMA(2,1) model. The results of this fit are shown in Exhibit 8.16. Notice that even though the estimate of  $\sigma_e^2$  and the log-likelihood and AIC values are not too far from their best values, the estimates of  $\phi_1$ ,  $\phi_2$ , and  $\theta$  are way off, and none would be considered different from zero statistically.

---

### Exhibit 8.16 Overfitted ARMA(2,1) Model for the Color Property Series

Coefficients:	ar1	ar2	ma1	Intercept
	0.2189	0.2735	0.3036	74.1653
s.e.	2.0056	1.1376	2.0650	2.1121

sigma^2 estimated as 24.58: log-likelihood = -105.91, AIC = 219.82

---

```
> arima(color, order=c(2, 0, 1))
```

---

## 8.3 Summary

---

The ideas of residual analysis begun in Chapter 3 were considerably expanded in this chapter. We looked at various plots of the residuals, checking the error terms for constant variance, normality, and independence. The properties of the sample autocorrelation of the residuals play a significant role in these diagnostics. The Ljung-Box statistic portmanteau test was discussed as a summary of the autocorrelation in the residuals. Lastly, the ideas of overfitting and parameter redundancy were presented.

## EXERCISES

---

- 8.1 For an AR(1) model with  $\phi \approx 0.5$  and  $n = 100$ , the lag 1 sample autocorrelation of the residuals is 0.5. Should we consider this unusual? Why or why not?
- 8.2 Repeat Exercise 8.1 for an MA(1) model with  $\theta \approx 0.5$  and  $n = 100$ .
- 8.3 Based on a series of length  $n = 200$ , we fit an AR(2) model and obtain residual autocorrelations of  $\hat{r}_1 = 0.13$ ,  $\hat{r}_2 = 0.13$ , and  $\hat{r}_3 = 0.12$ . If  $\hat{\phi}_1 = 1.1$  and  $\hat{\phi}_2 = -0.8$ , do these residual autocorrelations support the AR(2) specification? Individually? Jointly?

- 8.4** Simulate an AR(1) model with  $n = 30$  and  $\phi = 0.5$ .
- (a) Fit the correctly specified AR(1) model and look at a time series plot of the residuals. Does the plot support the AR(1) specification?
  - (b) Display a normal quantile-quantile plot of the standardized residuals. Does the plot support the AR(1) specification?
  - (c) Display the sample ACF of the residuals. Does the plot support the AR(1) specification?
  - (d) Calculate the Ljung-Box statistic summing to  $K = 8$ . Does this statistic support the AR(1) specification?
- 8.5** Simulate an MA(1) model with  $n = 36$  and  $\theta = -0.5$ .
- (a) Fit the correctly specified MA(1) model and look at a time series plot of the residuals. Does the plot support the MA(1) specification?
  - (b) Display a normal quantile-quantile plot of the standardized residuals. Does the plot support the MA(1) specification?
  - (c) Display the sample ACF of the residuals. Does the plot support the MA(1) specification?
  - (d) Calculate the Ljung-Box statistic summing to  $K = 6$ . Does this statistic support the MA(1) specification?
- 8.6** Simulate an AR(2) model with  $n = 48$ ,  $\phi_1 = 1.5$ , and  $\phi_2 = -0.75$ .
- (a) Fit the correctly specified AR(2) model and look at a time series plot of the residuals. Does the plot support the AR(2) specification?
  - (b) Display a normal quantile-quantile plot of the standardized residuals. Does the plot support the AR(2) specification?
  - (c) Display the sample ACF of the residuals. Does the plot support the AR(2) specification?
  - (d) Calculate the Ljung-Box statistic summing to  $K = 12$ . Does this statistic support the AR(2) specification?
- 8.7** Fit an AR(3) model by maximum likelihood to the square root of the hare abundance series (filename hare).
- (a) Plot the sample ACF of the residuals. Comment on the size of the correlations.
  - (b) Calculate the Ljung-Box statistic summing to  $K = 9$ . Does this statistic support the AR(3) specification?
  - (c) Perform a runs test on the residuals and comment on the results.
  - (d) Display the quantile-quantile normal plot of the residuals. Comment on the plot.
  - (e) Perform the Shapiro-Wilk test of normality on the residuals.
- 8.8** Consider the oil filter sales data shown in Exhibit 1.8 on page 7. The data are in the file named oilfilters.
- (a) Fit an AR(1) model to this series. Is the estimate of the  $\phi$  parameter significantly different from zero statistically?
  - (b) Display the sample ACF of the residuals from the AR(1) fitted model. Comment on the display.

- 8.9** The data file named `robot` contains a time series obtained from an industrial robot. The robot was put through a sequence of maneuvers, and the distance from a desired ending point was recorded in inches. This was repeated 324 times to form the time series. Compare the fits of an AR(1) model and an IMA(1,1) model for these data in terms of the diagnostic tests discussed in this chapter.
- 8.10** The data file named `deere3` contains 57 consecutive values from a complex machine tool at Deere & Co. The values given are deviations from a target value in units of ten millionths of an inch. The process employs a control mechanism that resets some of the parameters of the machine tool depending on the magnitude of deviation from target of the last item produced. Diagnose the fit of an AR(1) model for these data in terms of the tests discussed in this chapter.
- 8.11** Exhibit 6.31 on page 139, suggested specifying either an AR(1) or possibly an AR(4) model for the difference of the logarithms of the oil price series. (The file-name is `oil.price`).
- (a) Estimate both of these models using maximum likelihood and compare the results using the diagnostic tests considered in this chapter.
  - (b) Exhibit 6.32 on page 140, suggested specifying an MA(1) model for the difference of the logs. Estimate this model by maximum likelihood and perform the diagnostic tests considered in this chapter.
  - (c) Which of the three models AR(1), AR(4), or MA(1) would you prefer given the results of parts (a) and (b)?