Chapter 3

Case Studies of Linear Time Series

3.2 Global Temperature Anomalies

Global warming is a topic of considerable importance and has attracted much attention in recent years, ranging from environmental engineers to scientists to economists. If the rise in global temperature continues, it will have a major impact on the global economy. In this section, we analyze the monthly global temperature anomalies from January 1880 to August 2010. Our goal is not to debate the evidence of global warming, but to demonstrate empirical time series analysis. Specifically, our goals are (a) to illustrate the methods discussed in Chapter 2 about time-series modeling and forecasting, (b) to compare different models, (c) to see the limitation of time series models in long-term prediction, and (d) to show the difficulty in distinguishing trend-stationarity from unit-root stationarity based purely on the data. Global climate changes involve many other factors.

There are several data sets available for global temperature anomalies. See the websites http://data.giss.nasa.gov/gistemp/ of the Goddard Institute for Space Studies (GISS), National Aeronautics and Space Administration (NASA) and http://www.ncdc.noaa.gov/cmb-faq/anomalies.html of the National Climatic Data Center, National Oceanic and Atmospheric Administration (NOAA). We employ the series of monthly means based on land-surface air temperature anomalies of GISS, NASA. However, we obtained similar results from the data of NOAA. In fact, the same models apply to both series.

Figure 3.1 shows the time plot of the global temperature anomalies from January 1880 to August 2010 for 1568 observations over 131 years. The GISS data are in 0.01 degrees of Celsius. An upward trend is clearly seen from the plot. In particular, the slope of the trend seems to increase in the early 1980s. On the other hand, the variability of the temperature is relatively stable over the 131 years.

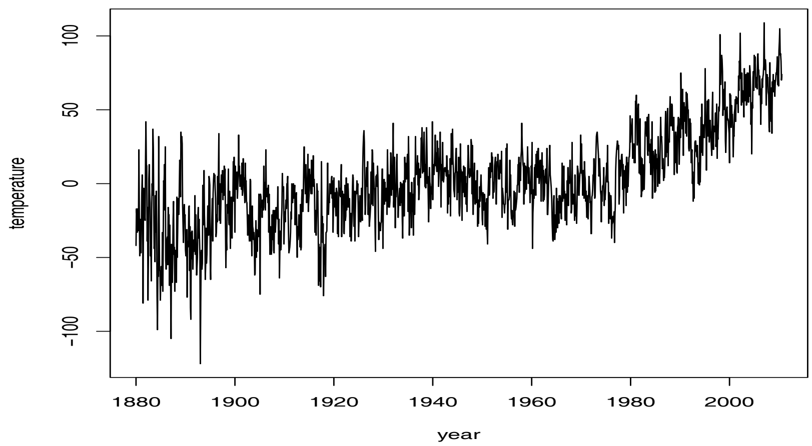

Figure 3.1: Monthly global temperature anomalies from January 1880 to August 2010.

3.2.1 Unit-root Stationarity

Let G_t denote the monthly global temperature anomalies. To specify a model for G_t, we start by examining the dynamic dependence of the series. Figure 3.2 gives the sample ACFs of G_t. As expected, the ACFs are high and decay slowly. A careful inspection also shows that the ACFs exhibit a cyclic pattern with peaks occurring around lags 24 and 36. This latter feature is not surprising because temperature often has a seasonal pattern.

Because of the strong serial dependence, we consider the differenced data $x_t = (1-B)G_t$. Figure 3.3 shows the sample ACFs and PACFs of x_t. Both ACFs and PACFs become small, indicating that x_t can be approximated by a stationary time series model. Since the sample size 1568 is large, we can entertain a relatively more complicated model. A careful examination of the sample ACFs of x_t shows that (a) ACF at lags 1, 2, 4, 5, and 8 are either significant or marginally significant, and (b) ACF at lag 24 is significant. The sample PACF, on the other hand, has several significant values. In particular, the lag-1 PACF is much larger than others and the PACF does not decay exponentially. Putting information together and following the general guidelines of the previous section, we start with the simple ARIMA(1,1,2) model

$$(1-\phi B)(1-B)G_t = (1-\theta_1 B - \theta_2 B^2)a_t, \tag{3.1}$$

for G_t. Here we use $p = 1$ because the differenced series x_t has a large lag-1 PACF and $q = 2$ because the first two ACFs of x_t are significant. Since the specified MA(2) model can have significant PACFs at lower order lags, we decide

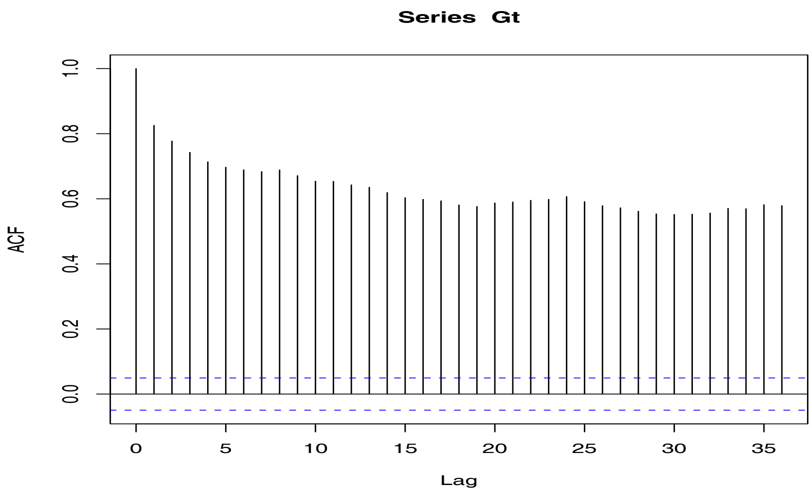

Figure 3.2: Sample autocorrelation function of the monthly global temperature anomalies.

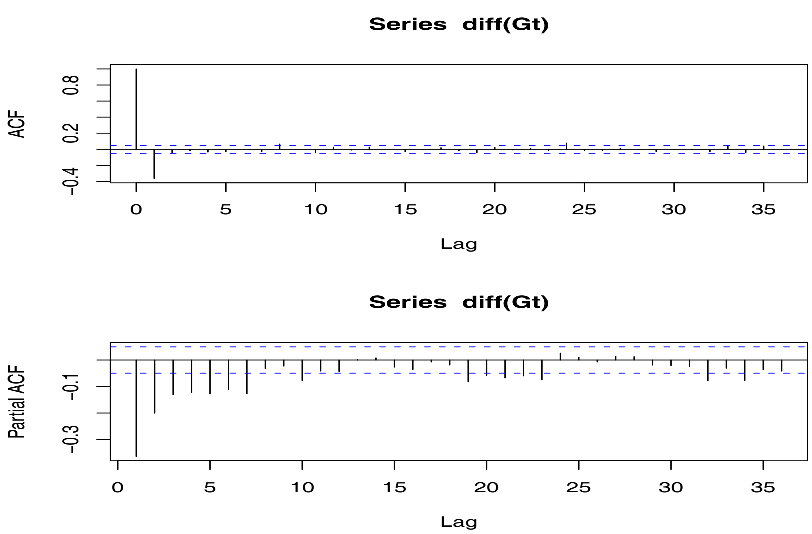

Figure 3.3: Sample autocorrelation and partial autocorrelation functions of the differenced global temperature anomalies.

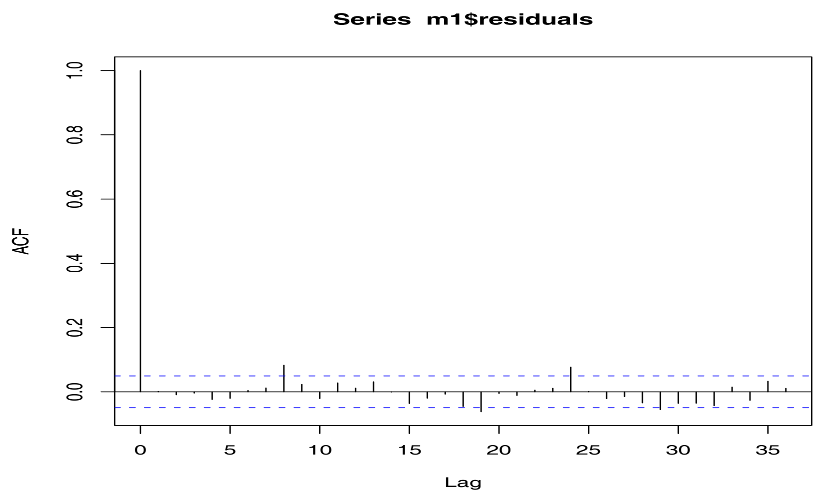

Figure 3.4: Sample autocorrelation function of the residuals of ARIMA(1,1,2) model for global temperature anomalies.

to keep $p = 1$. The high order ACFs are temporarily ignored because we like to keep the model simple. The fitted model is

$$(1 - 0.739B)(1 - B)G_t = (1 - 1.297B + 0.318B^2)a_t, \quad \sigma_a^2 = 272.1. \tag{3.2}$$

All estimates are highly significant. Figure 3.4 shows the sample ACF of the residuals of the ARIMA(1,1,2) model in Eq. (3.2). Based on the residual ACFs, the model is inadequate because the ACFs are significant at lags 8 and 24.

Turn to model refinement. The significance of ACF at lag 24 is understandable because of the seasonal nature of temperature. On the other hand, it is not easy to explain the serial correlation at lag 8. Consequently, we refine the model as

$$(1 - \phi B)(1 - B)G_t = (1 - \theta_1 B - \theta_2 B^2)(1 - \theta_{24} B^{24})a_t. \tag{3.3}$$

The fitted model is

$$(1-0.761B)(1-B)G_t = (1-1.324B+0.342B^2)(1-0.072B^{24})a_t, \quad \sigma_a^2 = 270.6, \tag{3.4}$$

where the standard errors of the estimates are, in order, 0.038, 0.052, 0.049, and 0.024, respectively. These estimates are statistically significant. Figure 3.5 gives the diagnostic plots for the seasonal model in Eq. (3.4). The residual plot looks reasonable and the p values of Ljung-Box statistics are above 0.05 except for $Q(8)$ and $Q(19)$. As expected, the residual ACFs show marginally significant values at lags 8 and 19. As mentioned earlier, it is hard to explain

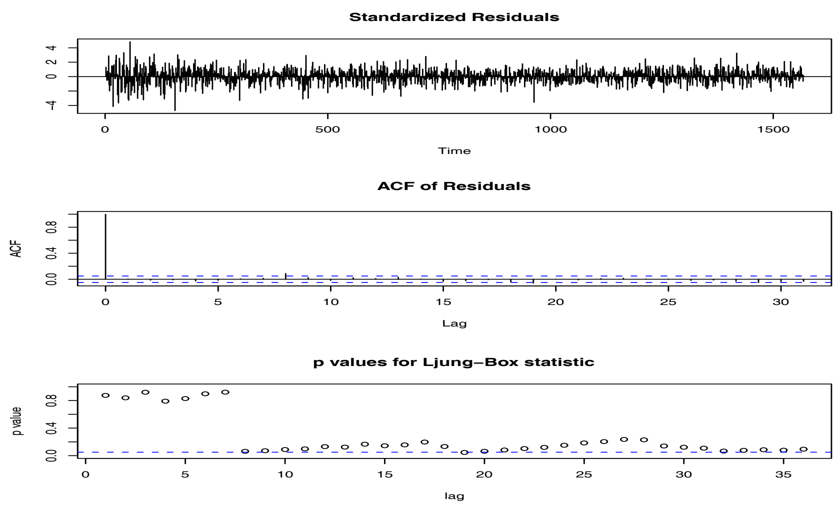

Figure 3.5: Diagnostic checking plots of the model in Eq. (3.4) for the monthly global temperature anomalies.

the lag-8 serial correlation and the magnitude of the ACF is small, we terminate the modeling process and treat the model in Eq. (3.4) as an adequate model. The AIC of model (3.4) is 13234.4 which is smaller than 13241.1 of model (3.2).

The model in Eq. (3.4) is called a unit-root nonstationary model because it uses the first difference to transform the global temperature into a stationary series. The time series G_t is said to be difference-stationary.

R Demonstration

```
> Gt=scan(file='m-GLBTs.txt')
> Gtemp=ts(Gt,frequency=12,start=c(1880,1))
> plot(Gtemp,xlab='year',ylab='temperature',type='l') % Plot the data
> acf(diff(Gt),lag=36)
> pacf(diff(Gt),lag=36)

> m1=arima(Gt,order=c(1,1,2))
> m1
arima(x = Gt, order = c(1, 1, 2))

Coefficients:
         ar1      ma1     ma2
      0.7387  -1.2973  0.3183
s.e.  0.0406   0.0533  0.0492

sigma^2 estimated as 272.1: log likelihood=-6616.6, aic=13241.1
> acf(m1$residuals,lag=36)
```

```
> m1=arima(Gt,order=c(1,1,2),seasonal=list(order=c(0,0,1),period=24))
> m1
arima(x=Gt,order=c(1,1,2),seasonal=list(order=c(0,0,1),period=24))

Coefficients:
         ar1      ma1     ma2    sma1
      0.7612  -1.3241  0.3416  0.0717
s.e.  0.0379   0.0519  0.0485  0.0243

sigma^2 estimated as 270.6: log likelihood=-6612.2, aic=13234.4
> tsdiag(m1,gof=36)
```

3.2.2 Trend-nonstationarity

In the literature, some analysts and scientists use *time trend* to model the global temperature anomalies. By time trend, we mean using time index as an explanatory variable. Consider the model

$$G_t = \beta_0 + \beta_1 t + z_t, \tag{3.5}$$

where z_t is an innovation series denoting the deviation of the global temperature anomalies from a time trend. If z_t is a stationary time series, then G_t is called a trend-stationary time series, meaning that it can be transformed into a stationary series by removing the effect of a time trend. In Model (3.5), β_1 is the slope of the time trend. A positive β_1 indicates that G_t will increase with time and eventually goes to positive infinity as t approaches infinity. For monthly data, β_1 is the monthly growth rate of G_t. Conceptually, a trend-stationary time series is very different from a difference-stationary one because the latter does not contain a fixed trend. We shall discuss further the difference between the two models when we consider long-term prediction.

For the global temperature anomalies, the fitted linear regression model is

$$G_t = -38.04 + 0.05156t + z_t, \tag{3.6}$$

where the standard errors of the coefficient estimates are 1.135 and 0.0013, respectively, and the standard deviation of z_t is 22.46. The time slope is positive and highly significant. Figure 3.6 shows the sample ACFs and PACFs of the innovation series z_t of Eq. (3.6). The PACFs decay quickly and the ACFs do not show any high value. Therefore, it is reasonable to assume that z_t does not have a unit root. That is, z_t is stationary and, hence, G_t is trend-stationary.

Next, we specify a model for the innovation series z_t. Because its ACFs in Figure 3.6 do not cut-off at any finite lag, z_t does not follow a simple MA model. In other words, some AR component is needed. That is, $p > 0$. The PACFs of Figure 3.6 have two discernible features. First, the first eight lags of PACFs are significant, indicating that z_t does not follow a low order AR model. This implies that $q > 0$. Second, the PACFs do not follow a simple exponentially decaying pattern. This means that $p > 1$. Putting information together and

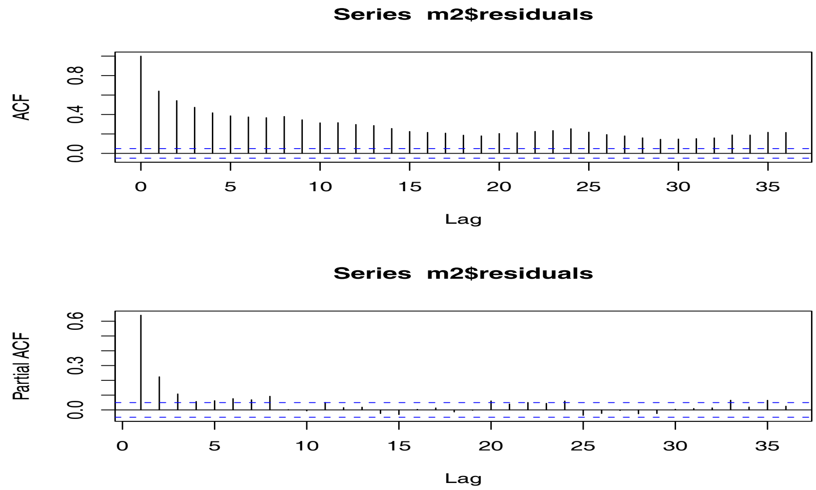

Figure 3.6: Sample autocorrelations and partial autocorrelations of the innovation series z_t in Eq. (3.6) for monthly global temperature anomalies.

keeping the order simple, we start with an ARMA(2,1) model for z_t, i.e.

$$(1 - \phi_1 B - \phi_2 B^2) z_t = (1 - \theta_1 B) a_t.$$

The model for G_t then becomes

$$(1 - \phi_1 B - \phi_2 B^2)(G_t - \beta_0 - \beta_1 t) = (1 - \theta_1 B) a_t. \tag{3.7}$$

The fitted model is

$$(1-1.239B+0.272B^2)(G_t+38.72-0.053t) = (1-0.78B)a_t, \quad \sigma_a^2 = 272.9, \tag{3.8}$$

where all estimates are highly significant. However, the residual ACFs of the model shows a significant value at lag 24. This is not surprising based on the ARIMA model used in the previous section. Consequently, we further refine the model and obtain

$$(1 - 1.196B + 0.239B^2)(G_t + 38.72 - 0.0529t) = (1 - 0.745B)(1 - 0.0856B^{24})a_t, \tag{3.9}$$

where $\sigma_a^2 = 270.8$ and all estimates are statistically significant at the 5% level. The standard errors of the coefficient estimates are, in order, 0.059, 0.048, 5.18, 0.006, 0.049, and 0.024, respectively. Figure 3.7 gives the diagnostic checking plots of the model in Eq. (3.9). Except for a minor residual ACF at lag 8, the checking statistics fail to indicate any inadequacy of the fitted model. Consequently, we select the model in Eq. (3.9) as the final model for G_t under trend stationarity.

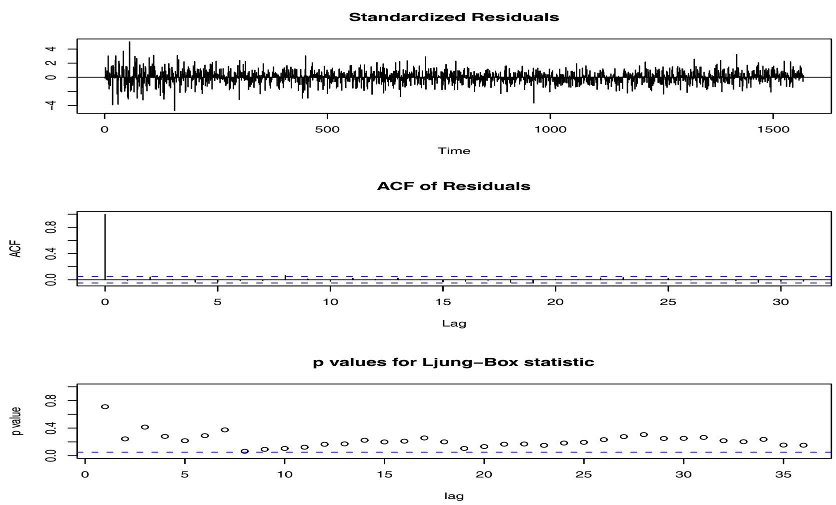

Figure 3.7: Diagnostic checking of the fitted model in Eq. (3.9) for the monthly global temperature anomalies.

Based on the model in Eq. (3.9), the global temperature increases on average 0.0529/100 degrees in Celsius per month. That is, the global temperature increases 0.00635 degrees in Celsius per year. This is very significant because it implies that the global temperature, on average, will increase 1 degree in Celsius every 157 years. The AIC of the model in Eq. (3.9) is 13247.5, which is larger than that of the difference-stationary model in Eq. (3.4).

R Demonstration

```
> time=c(1:1568) % time index
> m2=lm(Gt~time)
> summary(m2)
lm(formula = Gt ~ time)

Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept) -38.039763   1.134960  -33.52   <2e-16 ***
time          0.051560   0.001253   41.15   <2e-16 ***
---
Residual standard error: 22.46 on 1566 degrees of freedom
Multiple R-squared: 0.5195,    Adjusted R-squared: 0.5192
> par(mfcol=c(2,1))
> acf(m2$residuals,lag=36)
> pacf(m2$residuals,lag=36)
> m2=arima(Gt,order=c(2,0,1),xreg=time)
> m2
```

```
arima(x = Gt, order = c(2, 0, 1), xreg = time)

Coefficients:
         ar1      ar2      ma1  intercept    time
      1.2385  -0.2719  -0.7802   -38.8493  0.0530
s.e.  0.0567   0.0477   0.0460     5.3548  0.0059

sigma^2 estimated as 272.9: log likelihood=-6623.0, aic=13257.97
> tsdiag(m2,gof=36)  % Significant ACF at lag 24.
> m2=arima(Gt,order=c(2,0,1),seasonal=list(order=c(0,0,1),
  period=24),xreg=time)
> m2
arima(x=Gt,order=c(2,0,1),seasonal=list(order=c(0,0,1),period=24),
    xreg = time)

Coefficients:
         ar1      ar2      ma1    sma1  intercept    time
      1.1960  -0.2394  -0.7451  0.0856   -38.7150  0.0529
s.e.  0.0587   0.0482   0.0486  0.0241     5.1843  0.0057

sigma^2 estimated as 270.8: log likelihood=-6616.7, aic=13247.5
> tsdiag(m2,gof=36) % model checking
```

3.2.3 Model Comparison

We have obtained two models for the monthly global temperature anomalies from January 1880 to August 2010. The first model in Eq. (3.4) is difference-stationary whereas the second one in Eq. (3.9) is trend-stationary. Both models are adequate because they passed the rigorous model checking. A question arises naturally is which model should one choose. We shall address this question in this section.

In-sample comparison
As mentioned in Chapter 2, one approach to compare models for a given time series is to consider their in-sample goodness of fit. A commonly used criterion here is the Akaike information criterion or the Bayesian information criterion. For the global temperature data, the difference-stationary model in Eq. (3.4) is selected based on the AIC because it has a smaller value at 13,234. The two models, however, are close in many ways. For instance, Figure 3.8 shows time plots of the residuals of the two competing models with the top panel for the difference-stationary model. The two residuals series are essentially the same.

Another approach to in-sample model comparison is to study the implications of the competing models. Here the two fitted models for the global temperature anomalies differ dramatically. As mentioned before, the trend-stationary model in Eq. (3.9) imposes *a priori* a time trend for the temperature. The estimated time slope is 0.0529, which is positive and highly significant. Thus, the model implies that the gobal temperature will continue to increase at a pace

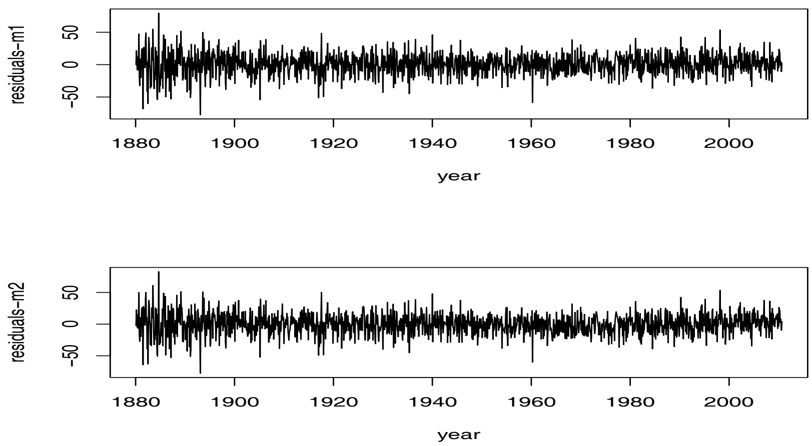

Figure 3.8: Residual plots for the models in Eqs. (3.4) and (3.9) for the monthly global temperature anomalies. The upper panel is for model (3.4).

of 0.00635 degrees in Celsius per year under the current environment. On the other hand, the difference-stationary model in Eq. (3.4) does not provide a definite support for the global warming. Similar to the randon walk model for stock prices, the future temperature may increase or decrease because a random walk has no fixed direction. What the model implies is that there exists substantial uncertainty in the future global temperature. It can go anywhere from positive infinity to negative infinity. The observed data, which were over 131 years, provide little guidance about the future global temperature. This appears to be unconceivable at the first glance, but 131 years are not sufficiently long when we are making inference about hundreds or thousands of years into the future. In short, the data we have do not contain sufficient evidence to distinguish between trend-stationarity and difference-stationarity.

Out-of-sample comparison

If the goal of time series analysis is forecasting, one can use out-of-sample prediction to compare competing models. Again, we use the backtesting of Chapter 2 to evaluate out-of-sample prediction. For the global temperature data, we divide the sample into modeling and forecasting subsamples with the latter consisting of the last 200 observations. We then apply the backtesting method to compute the 1-step ahead prediction of the two competing models in Eqs. (3.4) and (3.9). For the global temperature data with 200 1-step ahead out-of-sample predictions, we obtain the following results:

Model	RMSFE	MAFE
Difference-stationary model in Eq. (3.4)	14.526	11.167
Trend-stationary model in Eq. (3.9)	15.341	11.966

Clearly, the difference-stationary model is preferred based on the 1-step ahead prediction. The drop in RMSFE is about $(15.341 - 14.526)/14.526 = 5.6\%$, a moderate amount. This exercise also shows that time series models are useful in short-term prediction because the RMSFEs of the two models are smaller than the unconditional standard error 16.43 of the innovations of the fitted models. For the difference-stationary model, the reduction in RMSFE is approximately $(16.43 - 14.53)/14.53 = 13.1\%$.

R Demonstration

```
> source("backtest.R")
> pm1=backtest(m1,Gt,1368,1)
[1] "RMSE of out-of-sample forecasts"
[1] 14.52598
[1] "Mean absolute error of out-of-sample forecasts"
[1] 11.16746
> time=as.matrix(time)
> pm2=backtest(m2,Gt,1368,1,xre=time)
[1] "RMSE of out-of-sample forecasts"
[1] 15.34131
[1] "Mean absolute error of out-of-sample forecasts"
[1] 11.96595
```

3.2.4 Long-term Prediction

Global warming is concerned with long-term prediction. In this section we consider and compare the performance of the two competing models in Eqs. (3.4) and (3.9) using long-term prediction. Specifically, using August 2010, which gives the last data point, as the forecast origin, we compute 1-step to 1200-step ahead predictions of the monthly global temperature anomalies. In other words, we use the models built based on data of the past 131 years to predict the global temperatures for the next 100 years. To compute the predictions of the trend-stationary model in Eq. (3.9), the time index is given. Figure 3.9 shows the point predictions and the corresponding 95% interval forecasts of the global temperature anomalies based on the difference-stationary model in Eq. (3.4). The plot highlights several features of the model. First, similar to other unit-root models, the long-term forecasts converge to a constant represented by a horizontal line in the plot. The level of this horizontal line depends on the forecast origin. Second, the length of the 95% interval forecasts continues to grow with the forecast horizon. In fact, the length of the interval diverges to infinity eventually. These two features have important implications in forecasting. First, they indicate that the long-term forecasts are rather uncertain. This makes intuitive sense because long-term predictions of the model are dominated

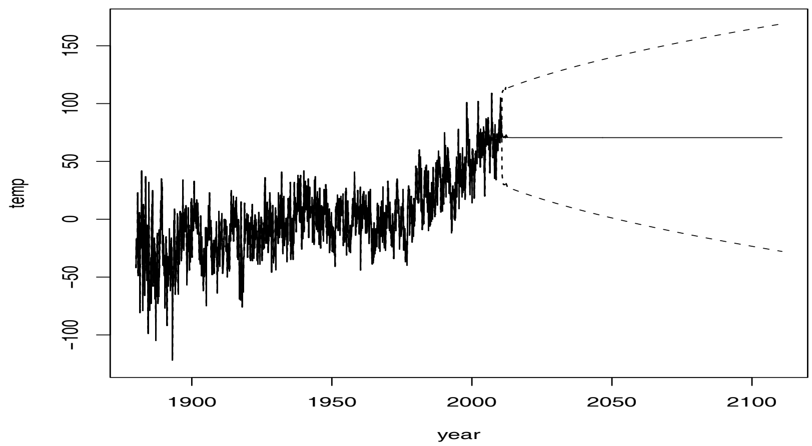

Figure 3.9: Long-term point and interval forecasts of the monthly global temperature anomalies based on the difference-stationary model in Eq. (3.4). The forecast origin is August 2010 and the forecast horizon is 100 years.

by its random walk component and for a random walk the current value contains little information about the future. Second, they demonstrate clearly that the model is only informative in short-term prediction. To a great degree, this is true for most time series models discussed in the book.

Turn to the trend-stationary model in Eq. (3.9). Figure 3.10 shows the point predictions and the corresponding 95% interval forecasts of global temperature anomalies of the model. Again, the forecast origin is August 2010 and the forecast horizon is 100 years. The plot also highlights the key features of the model. First, because of the positive time slope, the predictions grow with the forecast horizon. Second, the lengths of the interval forecasts are stable over time. In fact, the lengths converge to a constant quickly with the constant being approximately $4 \times \sigma_z$, where σ_z is the sample standard error of the innovation series z_t. As stated in Eq. (3.6), the innovation series z_t is stationary. As such the variances of the forecasts of z_t converge to its variance, σ_z^2, when the forecast horizon increases. For the trend-stationary model in (3.9), the prediction of the time trend is certain conditioned on the coefficients β_0 and β_1 being fixed. The uncertainty in forecasts is determined by that of z_t. Consequently, the variances of forecast errors of the model in Eq. (3.9) converge to that of z_t. As a matter of fact, the lengths of the forecast intervals match well with the range of the data. These two features also have important implications about the model. First, by imposing a time trend, the forecasts of the model diverge to infinity as the forecast horizon increases. Is this reasonable? How certain are we that the model is the true model for the global temperature anomalies? This type of

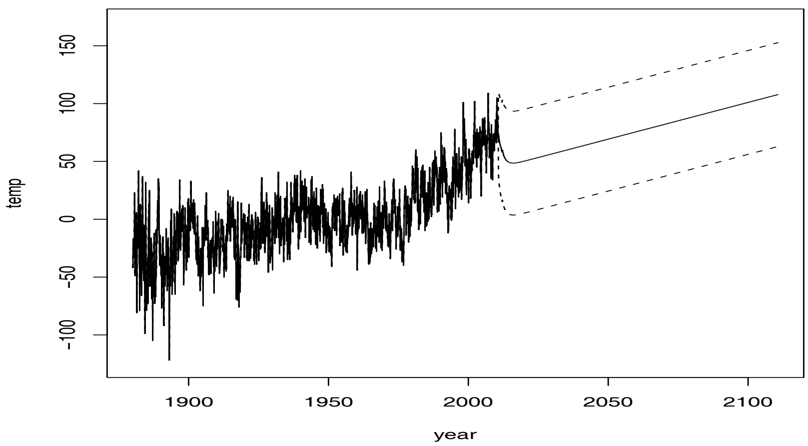

Figure 3.10: Long-term point and interval forecasts of the monthly global temperature anomalies based on the trend-stationary model in Eq. (3.9). The forecast origin is August 2010 and the forecast horizon is 100 years.

uncertainty is not shown by the model, nor by its forecasts. Second, the finite interval forecasts further confirm that the global temperature will continue to rise under the model. This is very different from that of the difference-stationary model in Eq. (3.4).

Finally, we provide a direct comparison by plotting the long-term predictions of the two competing models together. See Figure 3.11. The contrast between the two models is clearly seen. On the other hand, the plot contains no information about which model is more appropriate for the data, leading to the assertion that the two models are not useful in long-term prediction.

3.2.5 Discussion

In this case study, we employed two models to show several important issues of linear time series analysis. We demonstrated the process of model building and model selection. The models built are useful in short-term prediction because their out-of-sample root mean squares of forecast errors are smaller than their residual standard errors. Both the in-sample AIC criterion and the 1-step ahead backtesting select the difference-stationary model for the global temperature anomalies. On the other hand, the selected models are not informative in long-term prediction. Other important issues raised by the case study are summarized below.

Model Uncertainty

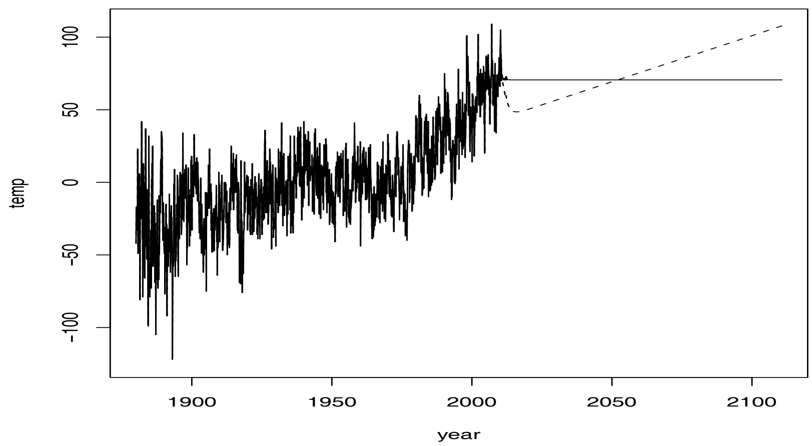

Figure 3.11: Point forecasts of the monthly global temperature anomalies based on two competing models in Eqs. (3.4) and (3.9). The forecast origin is August 2010 and the forecast horizon is 100 years.

All statistical models are wrong; they simply provide approximation to the underlying process. As such, statistical models contain uncertainty not only in parameter estimates but also in the choice of the models themselves. One should consider model uncertainty in making inference whenever possible. For time series analysis, model uncertainty can be addressed by model averaging, e.g., combining forecasts. A simple approach to model averaging in forecasting is to use the simple average among predictions of the competing models. For the monthly global temperature anomalies, we can use the average of two predictions to obtain a combined forecast. Again, this combination is useful mainly in short-term prediction.

Short versus Long-term Forecasts
The difference between short-term and long-term predictions shows that in time series analysis the choice of a model may depend on the forecast horizon. The best model for 1-step ahead prediction might be different from that for 2-step ahead prediction. This type of ideas has been explored in the literature and leads to the development of *adaptive forecast*. See Tiao and Tsay (1994) and the references therein.

Trend-shift Model
To demonstrate the arbitrariness (or subjectivity) involved in using linear time trend model, we consider an alternative specification. Recall that the time plot

in Figure 3.1 shows an increase in time slope around 1980. This feature has led some analysts to employ a *trend-shift* model. Specifically, consider the model

$$G_t = \beta_0 + \beta_1 t + \beta_2 x_t + n_t, \tag{3.10}$$

where x_t is defined as

$$x_t = \begin{cases} 0 & \text{if } t \leq 1212, \\ t & \text{if } 1212 < t \leq 1568, \end{cases} \tag{3.11}$$

where $t = 1212$ corresponds to December 1980. This simple model allows the time slope to change at the beginning of 1981. Specifically, the time slope is β_1 before January 1981 and is $\beta_1 + \beta_2$ starting in 1981. The parameter β_2 thus denotes the change in time slope. If β_2 is positive, the time slope jumps from β_1 to $\beta_1 + \beta_2$ in January 1981. The time plot of $\beta_1 t + \beta_2 x_t$ would show a jump and a change of direction in January 1981 so that the model in Eq. (3.10) is referred to as a trend-shift model.

For the global temperature anomalies, we obtain the model

$$G_t = -28.92 + 0.0313t + 0.0214x_t + n_t,$$

where all estimates are highly significant. This regression model says that the time slope for the monthly global temperature was 0.0313 before January 1981 and it increased to 0.0527 thereafter. The increase was statistical significant. The sample ACFs and PACFs of the innovation series z_t of the prior regression model are similar to those of the model in Eq. (3.6). Therefore, we employ the same ARMA model for z_t. The resulting model for the monthly global temperature anomalies is

$$(1 - 1.122B + 0.197B^2)(1 - B)z_t = (1 - 0.684B)(1 + 0.0823B^{24})a_t, \tag{3.12}$$

where $z_t = G_t + 29.263 - 0.317t - 0.0219x_t$, $\sigma_a^2 = 267.5$, and all coefficient estimates are significantly different from zero. The AIC of the model is 13,230, which is smaller than those of the two competing models built before. The diagnostic checking plots of this trend-shift model are similar to those of Eq. (3.9), indicating that model (3.12) is also adequate for the monthly global temperature anomalies. Note that this refined model shows some improvements over that of Eq. (3.9) because its residual variance reduces to 267.5.

The significance of the slope change 0.0219 in Janaury 1981 suggests that, under the postulated trend-shift model, the global warming has accelerated in recent years. The monthly temperature change increased from 0.0317 to 0.0536 starting in January 1981.

Finally, based on the AIC, the trend-shift model in Eq. (3.12) is preferred over the two competing models discussed before. This is not surprising because we impose a trend shift after examining the data. One can further improve the in-sample fit by employing a more flexible polynomial trend for the data. This type of improvement is not recommended in general because it can easily lead to over-fitting in a real application. Furthermore, long-term forecasts of such models do not take into account the possibility of slope changes in the future.

R Demonstration

```
> Gt=scan(file='m-GLBTs.txt')
> time=c(1:1568)
> time1=c(rep(0,1212),time[1213:1568])
> mm1=lm(Gt~time+time1)
> summary(mm1)
lm(formula = Gt ~ time + time1)

Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept) -28.924419   1.191061  -24.29   <2e-16 ***
time          0.031330   0.001702   18.40   <2e-16 ***
time1         0.021397   0.001318   16.23   <2e-16 ***
---
Residual standard error: 20.79 on 1565 degrees of freedom
Multiple R-squared: 0.5887,    Adjusted R-squared: 0.5882
F-statistic:  1120 on 2 and 1565 DF,  p-value: < 2.2e-16

> x1=cbind(time,time1)
> mm1=arima(Gt,order=c(2,0,1),seasonal=list(order=c(0,0,1),period=24),
  xreg=x1)
> mm1
arima(x=Gt,order=c(2,0,1),seasonal=list(order=c(0,0,1),period=24),xreg=x1)

Coefficients:
         ar1      ar2      ma1    sma1  intercept    time   time1
      1.1220  -0.1973  -0.6835  0.0823   -29.2630  0.0317  0.0219
s.e.  0.0727   0.0542   0.0643  0.0239     4.1411  0.0058  0.0044

sigma^2 estimated as 267.5:  log likelihood = -6607,  aic = 13230
> tsdiag(mm1,gof=36)
> Box.test(mm1$residuals,lag=8,type='Ljung')
        Box-Ljung test

data:  mm1$residuals
X-squared = 15.4598, df = 8, p-value = 0.0508
```

Other Data Set

There exist other time series measuring the monthly global temperature anomalies. See, for instance, the series given in NCDC, NOAA mentioned before. This series measured in degrees of Celsius is close, but not identical, to that of GISS used in our case study. The basic characteristics of the data, however, are the same. As a matter of fact, the same models apply to this new series. Thus, the choice of data sets does not affect the general conclusions of the analysis. For example, consider the trend-stationary model. We have

$$(1-1.298B+0.306B^2)(y_t+0.495-0.000661t) = (1-0.934B)(1-0.083B^{24})a_t,$$

where y_t denotes the monthly global temperature anomalies, $\sigma_a^2 = 0.0883$ and all coefficient estimates are statistically significant at the 5% level. The time slope is 0.000661, implying that the global temperature increases on average 0.00793 degrees in Celsius per year. This is slightly greater than 0.00635 degrees suggested by the GISS data set. The difference, however, appears to be not statistically significant.

R Demonstration

```
> da=read.table("m-ncdc-noaa-glbtemp.txt")
> head(da)
    V1 V2      V3
1 1880  1 -0.0405
2 1880  2 -0.6112
 ...
> tail(da)
       V1 V2        V3
1568 2010  8    0.8970
1569 2010  9 -999.0000
 ....
1572 2010 12 -999.0000
> da=da[1:1568,]
> temp=da[,3]
> m3=arima(temp,order=c(1,1,2),seasonal=list(order=c(0,0,1),period=24))
> m3
arima(x=temp,order=c(1,1,2),seasonal=list(order=c(0,0,1),period=24))

Coefficients:
         ar1      ma1     ma2    sma1
      0.5817  -1.2414  0.2639  0.0854
s.e.  0.0704   0.0827  0.0781  0.0243

sigma^2 estimated as 0.0881:  log likelihood = -321.11,  aic = 652.21
> tsdiag(m3,gof=36)
> m4=arima(temp,order=c(2,0,1),seasonal=list(order=c(0,0,1),period=24),
  xreg=time)
> m4
arima(x=temp,order=c(2,0,1),seasonal=list(order=c(0,0,1),period=24),
    xreg = time)

Coefficients:
         ar1      ar2      ma1    sma1  intercept   time
      1.2975  -0.3057  -0.9344  0.0827    -0.4952  7e-04
s.e.  0.0562   0.0480   0.0430  0.0255     0.1178  1e-04

sigma^2 estimated as 0.08825: log likelihood=-322.2, aic=658.4
```

```
> m4$coef
          ar1           ar2           ma1         sma1   intercept
 1.2974716910 -0.3056922789 -0.9343898897  0.0826971  -0.4952302
         time
 0.0006613375
> sqrt(diag(m4$var.coef))
      ar1       ar2        ma1       sma1  intercept       time
0.0561629  0.048026  0.0429865 0.02545154 0.11775435 0.00013652
> m4$coef/sqrt(diag(m4$var.coef))  % Compute t-ratios
       ar1        ar2        ma1       sma1  intercept     time
 23.101928  -6.365140 -21.736831   3.249198  -4.205621   4.8441
> tsdiag(m4,gof=36)
> %%% Backtesting
> pm3=backtest(m3,temp,1368,1)
[1] "RMSE of out-of-sample forecasts"
[1] 0.3160872
[1] "Mean absolute error of out-of-sample forecasts"
[1] 0.2361009
> pm4=backtest(m4,temp,1368,1,xre=time)
[1] "RMSE of out-of-sample forecasts"
[1] 0.3270271
[1] "Mean absolute error of out-of-sample forecasts"
[1] 0.241655
```