

Exercise 9. Time series analysis

1. The Darwin SLP data discussed in the lectures is in `darwin.txt`. Plot the data against time. What can you tell about any cycles or long-term time trends?
2. Apply moving average filters of different lengths to obtain a clearer view of the long-term trends.
3. Compute the autocorrelation function up to lag 36 months. The cyclical behaviour indicates the seasonality of the data. If X_t is the SLP for month t , plot X_t against X_{t-12} to visualise this dependence.
4. Apply the backward difference filter with lag 12 to remove the seasonality, then plot the differenced series to obtain a clearer view of the inter-annual variations.
5. (The final two questions are optional.) In addition, apply the backward difference filter with lag 1 to remove the month-to-month dependence and plot the resulting series, which we shall call Y_t . This looks like white noise. If this were so, what would you expect the autocorrelation function to look like? Compute the autocorrelation function up to lag 36 and interpret the result.
6. The autocorrelation function for Y_t suggests that a particular seasonal ARIMA model, with moving average components at lags 1 and 12, will be a good de-

scription of the Darwin SLP series. Fit this model to the Darwin SLP series by selecting a seasonal model with period 12, one nonseasonal and one seasonal difference, one nonseasonal and one seasonal moving average term, and no constant term. Plot the autocorrelation function and histogram of the residuals. Do the residuals look like white noise?