


Solution to Series 4


1. a) 1. AR(2): The ordinary autocorrelations describe a damped sine curve, and the partial autocorrelations are cut off at lag $k = 2$.
2. MA(3): The ordinary autocorrelations are cut off at lag $k = 3$, and the partial autocorrelations show a damped harmonic.
3. ARMA(1, 2): The ordinary autocorrelations coincide with that of an AR(1) process for lag $k > 2$ (that is, decay exponentially in absolute values); the partial autocorrelations coincide with that of an MA(2) process for lag $k > 1$ (that is, describe a damped harmonic).
- b) We show exemplary R code for the AR model; for the other models, we only show the resulting plots.

1. AR(2):


```
> par(mfrow = c(1, 2))
> ar.acf <- ARMAacf(ar = c(0.9, -0.5), lag.max = 30)
> plot(0:30, ar.acf, type = "h", ylab = "ACF")
> ar.pacf <- ARMAacf(ar = c(0.9, -0.5), lag.max = 30, pacf = TRUE)
> plot(1:30, ar.pacf, type = "h", ylab = "PACF")
```




2. MA(3):




3. ARMA(1, 2):




c) 1. AR(2):



2. MA(3):




3. ARMA(1,2):




In all models (perhaps with the exception of the PACF at lag 4 of the MA model), ACF and PACF coincide well with the expected values (plotted as red circles). Note, however, that you can find different values depending on the seed you used for the simulation.

2. 1. First time series, ts1:



The PACF shows—up to alternating signs—an exponential decay; the ACF describes a damped harmonic. It is therefore plausible to fit the time series as an $\text{ARMA}(p, q)$ process. The alternating signs of the PACF indicate the order $q = 1$. The order p of the AR part is hard to determine by eye; we therefore try out different orders and compare their AIC values:

```
> aic1 <- sapply(0:5, function(p)
  arima(ts1, order = c(p, 0, 1), method = "ML")$aic)
> plot(0:5, aic1, type = "l", xlab = "Order p", ylab = "AIC")
```



We see that the AIC is minimized at $p = 2$, hence we fit an $\text{ARMA}(2, 1)$ model:

```
> (arma1 <- arima(ts1, order = c(2, 0, 1), method = "ML"))
```

Call:

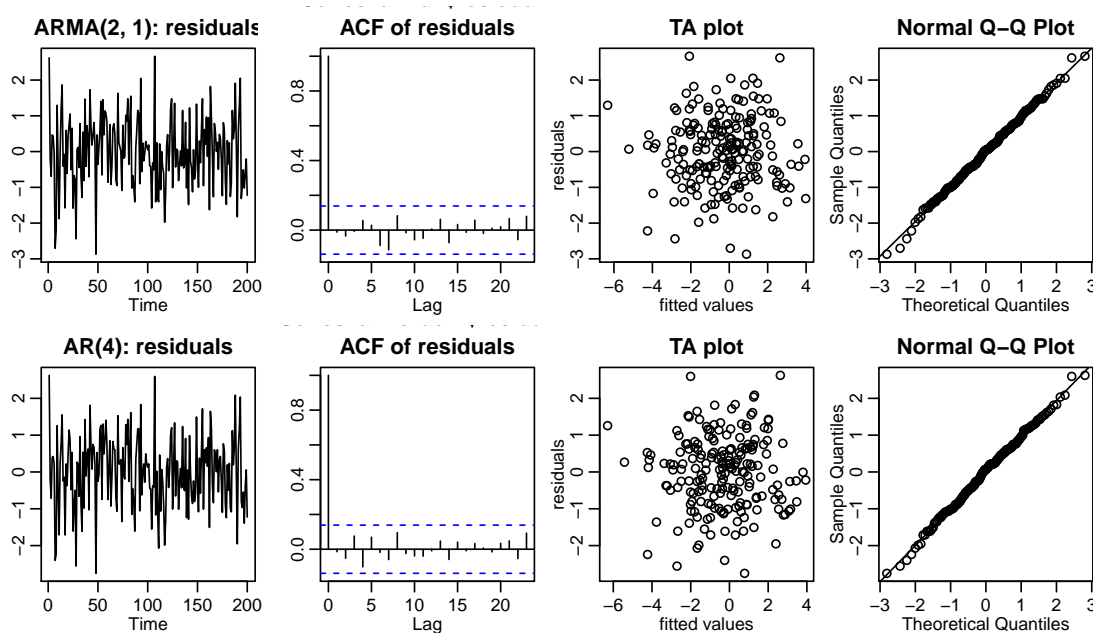
```
arima(x = ts1, order = c(2, 0, 1), method = "ML")
```

Coefficients:

	ar1	ar2	ma1	intercept
	0.8394	-0.2758	0.6776	-0.2967
s.e.	0.0845	0.0835	0.0636	0.2653


```
sigma^2 estimated as 0.9629: log likelihood = -281.21, aic = 572.41
```

Another possibility would be to interpret, based on the plots of ACF and PACF, the process as an $\text{AR}(4)$ (or possibly $\text{AR}(3)$) process. A residual analysis of such a model (second row in the following plot) shows that this model fits the time series as well as the $\text{ARMA}(2, 1)$ model from above (first row in the plots):



An AR(4) model, has, however, a slightly higher AIC value (575.2) as the ARMA(2, 1) model (572.4). Indeed, the time series was simulated from an ARMA(2, 1) model with parameters $\alpha_1 = 0.8$, $\alpha_2 = -0.2$ and $\beta_1 = 0.7$.

2. Second time series, ts2:



Candidate models are MA(2), or ARMA(1, q) ($p = 1$ because of the alternating sign of the ACF values, see above). The first one has an AIC value of 589.2. For the second ones, we determine the order q similarly as for the first time series:

```
> aic2 <- sapply(0:5, function(q)
+   arima(ts2, order = c(1, 0, q), method = "ML")$aic)
> round(aic2, 1)
```

```
[1] 686.4 658.5 590.6 591.4 593.0 594.5
```

MA(2) still has a lower AIC value than ARMA(1, 2), hence we fit the time series with an MA(2) model:

Call:

```
arima(x = ts2, order = c(0, 0, 2), method = "ML")
```


Coefficients:

	ma1	ma2	intercept
	-1.1828	0.9331	-0.0083
s.e.	0.0229	0.0321	0.0540

σ^2 estimated as 1.046: log likelihood = -290.59, aic = 589.18

This time series was simulated from an ARMA(1, 3) model with coefficients $\alpha_1 = -0.6$, $\beta_1 = -0.7$, $\beta_2 = 0.4$, $\beta_3 = 0.6$. The influence of the higher order coefficients (and of the AR component) is not visible any more in the time series.

3. Third time series, ts3:



The PACF has a clear cut-off at lag $k = 4$. We therefore fit an AR(4) model:

```
> arima(ts3, order = c(4, 0, 0), method = "ML")
```

Call:

```
arima(x = ts3, order = c(4, 0, 0), method = "ML")
```

Coefficients:

	ar1	ar2	ar3	ar4	intercept
	0.4620	-0.4908	0.5279	-0.6756	-0.0938
s.e.	0.0519	0.0488	0.0480	0.0515	0.0569

sigma² estimated as 0.8842: log likelihood = -273.02, aic = 558.04

The time series was indeed simulated from an AR(4) model with coefficients $\alpha_1 = 0.4$, $\alpha_2 = -0.5$, $\alpha_3 = 0.5$, $\alpha_4 = -0.7$.