Solution to Series 9

1. **a)** We use the log-transformation, we have seasonal data and a trend.

```
> d.airshort1 <- log(d.airshort)
> d.airshort2 <- diff(d.airshort1, lag=12)
> d.airshort3 <- diff(d.airshort2, 1)
> f.acf(d.airshort3)
```

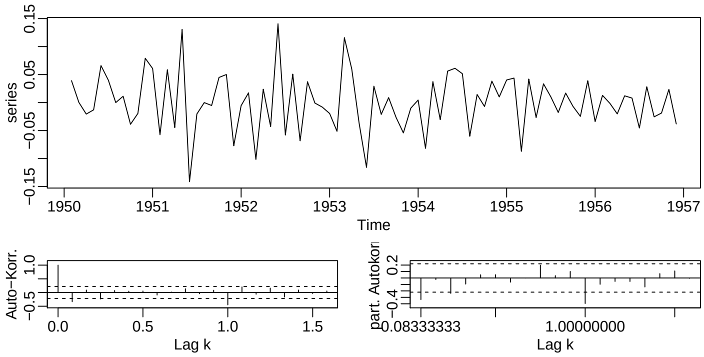

A $SARIMA(0,1,1)(0,1,1)^{12}$ could be an appropriate model.

```
> s1.air <- arima(d.airshort1,order=c(0,1,1),seasonal=c(0,1,1))
> s1.air

Call:
arima(x = d.airshort1, order = c(0, 1, 1), seasonal = c(0, 1, 1))

Coefficients:
         ma1     sma1
      -0.394   -0.613
s.e.   0.117    0.108

sigma^2 estimated as 0.00151:  log likelihood = 149,  aic = -292
> f.acf(s1.air$residuals)
```

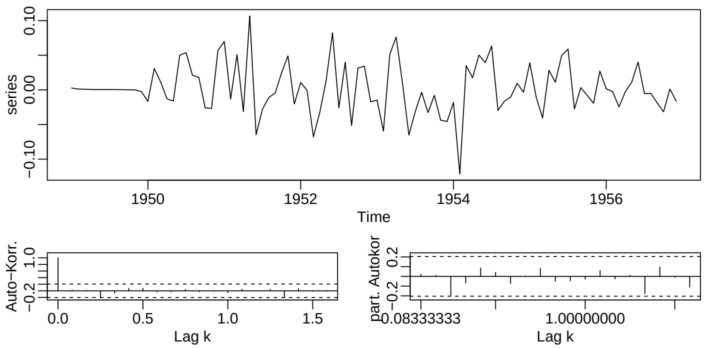

The residuals look ok.

```
> t.pr <- predict(s1.air, n.ahead=48)
> plot(log(d.air), xlim=c(1950,1961), ylim=c(4.5,7.5))
> t.u <- t.pr$pred+1.96*t.pr$se
> t.l <- t.pr$pred-1.96*t.pr$se
> lines(t.pr$pred, col="red")
> lines(t.u, col="green", lty=2)
> lines(t.l, col="green", lty=2)
```

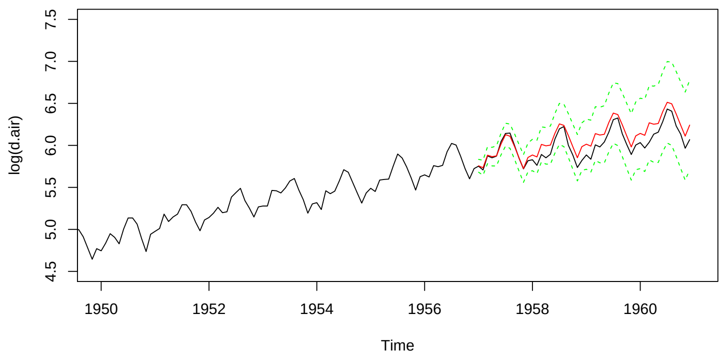

b)
```
> fit <- HoltWinters(d.airshort1)
> t.pr <- predict(fit, 48, prediction.interval=T)
> plot(log(d.air), xlim=c(1950,1961), ylim=c(4.5,7.5))
> lines(t.pr[,"fit"], col="red")
> lines(t.pr[,"upr"], col="green", lty=2)
> lines(t.pr[,"lwr"], col="green", lty=2)
```

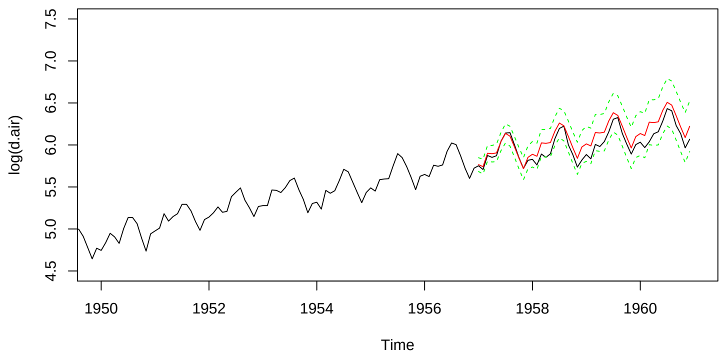


```
> f.acf(residuals(fit))
```

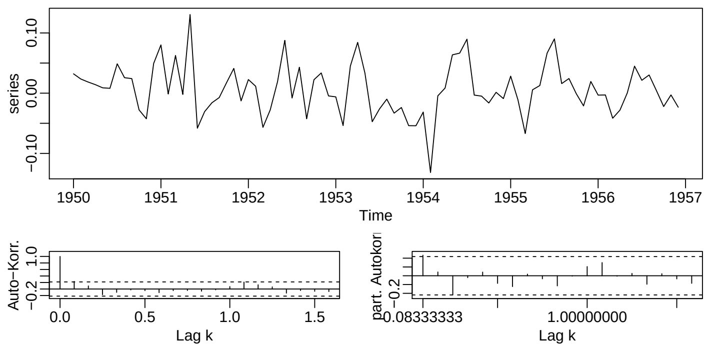

The residuals look ok.

Trend extrapolation:

c)
```
> fit <- stl(d.airshort1, s.window="periodic")
> trend <- fit$time.series[,2]
> ## Least Squares for trend over the last 4 years
> yy <- window(trend, start=c(1953,1))
> xx <- time(yy)
> reg <- lm(yy~xx)
> ## Trend Extrapolation
> kk <- 48
> trend.ex <- rev(trend)[1]+((1:kk)/12)*coef(reg)[2]
> trend.ex
 [1] 5.85 5.86 5.87 5.88 5.89 5.90 5.91 5.92 5.93 5.95 5.96
[12] 5.97 5.98 5.99 6.00 6.01 6.02 6.03 6.04 6.05 6.06 6.08
[23] 6.09 6.10 6.11 6.12 6.13 6.14 6.15 6.16 6.17 6.18 6.19
[34] 6.21 6.22 6.23 6.24 6.25 6.26 6.27 6.28 6.29 6.30 6.31
[45] 6.33 6.34 6.35 6.36
```

Saisonal effect:

```
> ## saisonality
> saison <- fit$time.series[,1]
> saisy <- window(saison, start=c(1953,1))
> sais.ex <- ts(saisy, start=c(1957,1), end=c(1960,12), freq=12)
```

Remainder:

```
> ## remainder
>
> remainder <- fit$time.series[,3]
> f.acf(remainder)
```

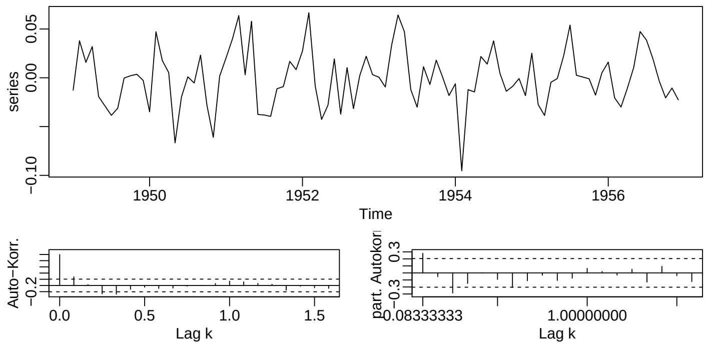

Remainder looks like an AR(3) model.

```
> fit.rem <- arima(remainder, order=c(3,0,0))
> f.acf(resid(fit.rem))
```

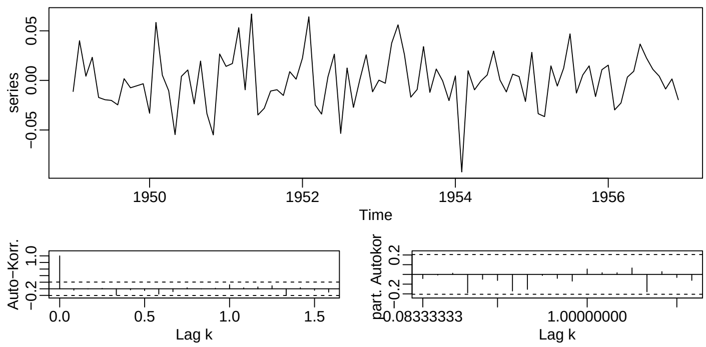

Residuals look good. Forecast of remainder:

```
> pred <- predict(fit.rem, n.ahead=48)
> rem.ex <- pred$pred
```

Putting everything together (trend, seasonality, remainder):

```
> ## everything together
> series.ex <- trend.ex+sais.ex+rem.ex
> ## Plotting
> plot(log(d.air), xlim=c(1950,1961), ylim=c(4.5,7.5))
> lines(series.ex, col="red", lwd=2)
```

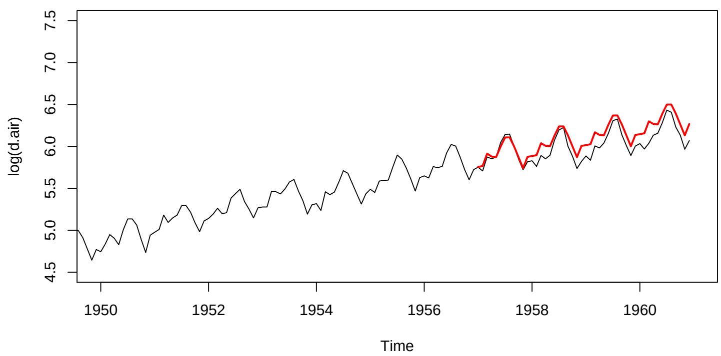

d) All of the three methods seem to overestimate this particular time series a little. Probably exponential smoothing works best.

e) No solution.