MONASH BUSINESS SCHOOL

Rob J Hyndman

Making forecasting easier

forecast v7 for R

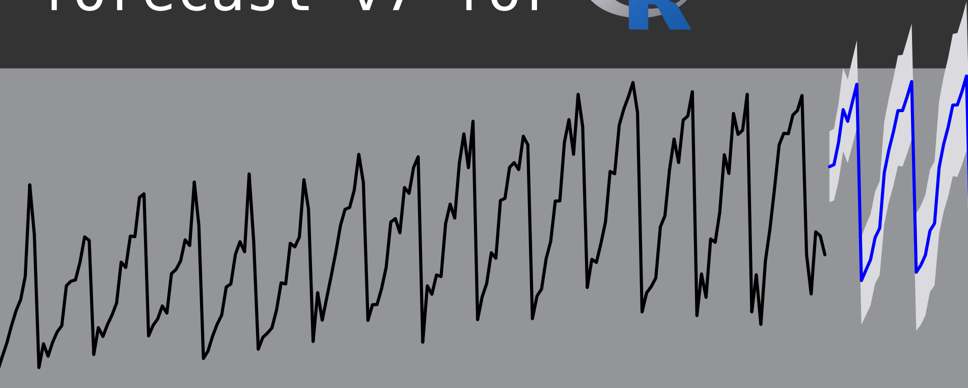

Outline

1. **Motivation and history**
2. Automatic forecasting in R
3. ggplot2 graphics
4. Bias adjustment

Motivation

Australian Government

Department of Health and Ageing

Motivation

FOXTEL digital

Australian Government

Department of Health and Ageing

Motivation

FOXTEL digital

Incitec Pivot

Australian Government

Department of Health and Ageing

Motivation

FOXTEL digital

CAT

Incitec Pivot

Australian Government

Department of Health and Ageing

Motivation

FOXTEL digital

CAT

hp

Incitec Pivot

Australian Government

Department of Health and Ageing

Motivation

1. Common in business to have thousands of products that need forecasting at least monthly.
2. Forecasts are often required by people who are untrained in time series analysis.

Specifications

Automatic forecasting algorithms must:

- ➥ determine an appropriate time series model;
- ➥ estimate the parameters;
- ➥ compute the forecasts with prediction intervals.

Motivation

1. Common in business to have thousands of products that need forecasting at least monthly.
2. Forecasts are often required by people who are untrained in time series analysis.

Specifications

Automatic forecasting algorithms must:

- ➥ determine an appropriate time series model;
- ➥ estimate the parameters;
- ➥ compute the forecasts with prediction intervals.

Forecast package history

Pre 2003	Collection of functions used for consulting projects
July/August 2003	**ets** and **thetaf** added
August 2006	v1.0 available on CRAN
May 2007	**auto.arima** added
July 2008	JSS paper (Hyndman & Khandakar)
September 2009	v2.0. Unbundled.
May 2010	**arfima** added
Feb/March 2011	**tslm**, **stlf**, **naive**, **snaive** added
August 2011	v3.0. Box Cox transformations added
December 2011	**tbats** added
April 2012	Package moved to github
November 2012	v4.0. **nnetar** added
June 2013	Major speed-up of **ets**
January 2014	v5.0. **tsoutliers** and **tsclean** added
May 2015	v6.0. Added several new plots
December 2015	264,000 package downloads in one month!
February 2016	v7.0. Added ggplot2 graphics & bias adjustment

Outline

Automatic methods in forecast package

Automatic model selection

- auto.arima + forecast
- ets + forecast
- tbats + forecast
- bats + forecast
- arfima + forecast
- ar + forecast
- nnetar + forecast
- stlm + forecast

Automatic forecasting

- forecast.ts
- stlf
- thetaf
- dshw, hw, holt, ses
- splinef
- rwf, naive
- croston

All produce an object of class “forecast”

ets algorithm in R

Based on Hyndman, Koehler, Snyder & Grose (IJF 2002):

- For each model, optimize parameters and initial values of underlying state space model using MLE.
- Select best method using AICc.
- Produce forecasts and prediction intervals using best method.

ets algorithm in R

Based on Hyndman, Koehler, Snyder & Grose (IJF 2002):

- For each model, optimize parameters and initial values of underlying state space model using MLE.
- Select best method using AICc.
- Produce forecasts and prediction intervals using best method.

ets algorithm in R

Based on Hyndman, Koehler, Snyder & Grose (IJF 2002):

- For each model, optimize parameters and initial values of underlying state space model using MLE.
- Select best method using AICc.
- Produce forecasts and prediction intervals using best method.

Exponential smoothing

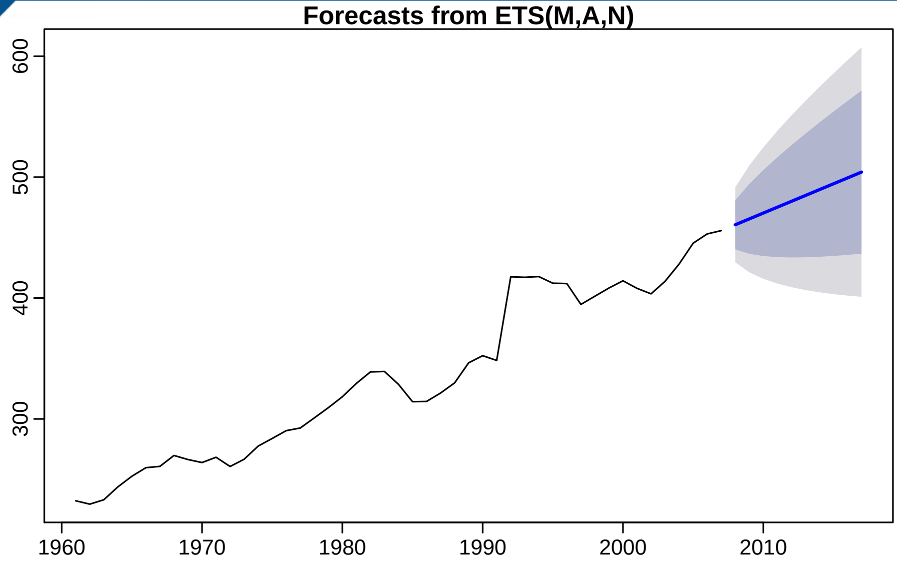

Exponential smoothing

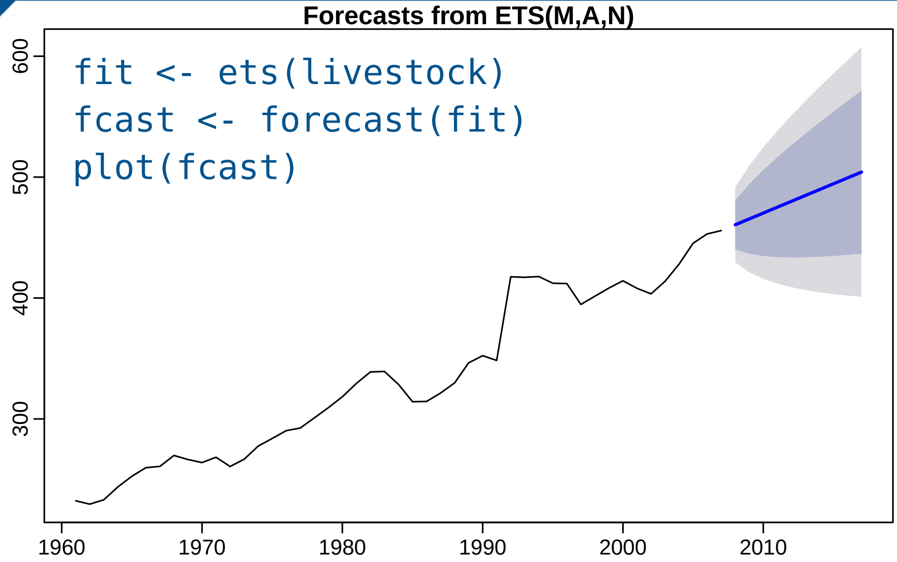


```r
fit <- ets(livestock)
fcast <- forecast(fit)
plot(fcast)
```

Exponential smoothing

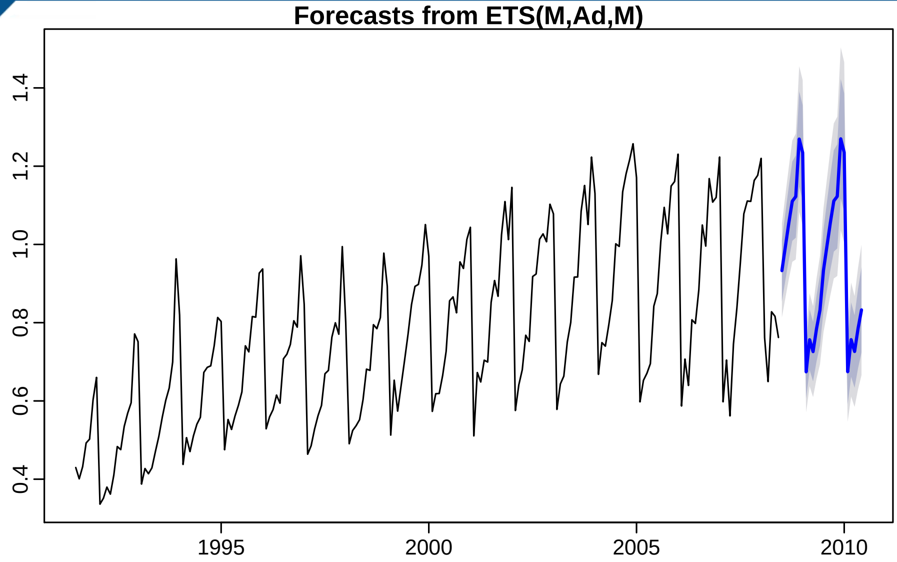

Exponential smoothing

```
fit <- ets(h02)
fcast <- forecast(fit)
plot(fcast)
```

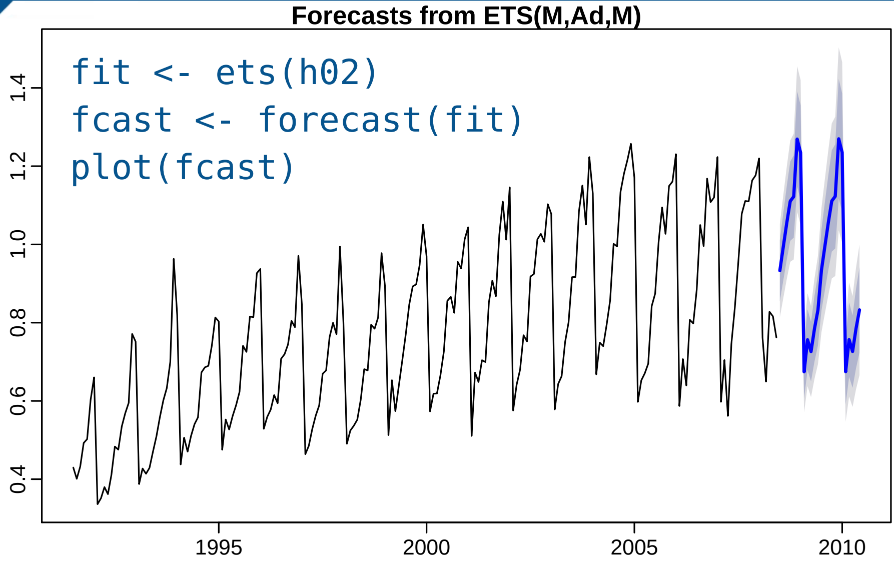

auto.arima algorithm in R

Based on Hyndman and Khandakar (JSS 2008):

- Select no. differences via unit root tests.
- Use stepwise search to traverse model space, starting with a simple model and considering nearby variants.
- For each model, optimize parameters using MLE.
- Select best method using AICc.
- Produce forecasts and prediction intervals using best method.

auto.arima algorithm in R

Based on Hyndman and Khandakar (JSS 2008):

- Select no. differences via unit root tests.
- Use stepwise search to traverse model space, starting with a simple model and considering nearby variants.
- For each model, optimize parameters using MLE.
- Select best method using AICc.
- Produce forecasts and prediction intervals using best method.

auto.arima algorithm in R

Based on Hyndman and Khandakar (JSS 2008):

- Select no. differences via unit root tests.
- Use stepwise search to traverse model space, starting with a simple model and considering nearby variants.
- For each model, optimize parameters using MLE.
- Select best method using AICc.
- Produce forecasts and prediction intervals using best method.

auto.arima algorithm in R

Based on Hyndman and Khandakar (JSS 2008):

- Select no. differences via unit root tests.
- Use stepwise search to traverse model space, starting with a simple model and considering nearby variants.
- For each model, optimize parameters using MLE.
- Select best method using AICc.
- Produce forecasts and prediction intervals using best method.

auto.arima algorithm in R

Based on Hyndman and Khandakar (JSS 2008):

- Select no. differences via unit root tests.
- Use stepwise search to traverse model space, starting with a simple model and considering nearby variants.
- For each model, optimize parameters using MLE.
- Select best method using AICc.
- Produce forecasts and prediction intervals using best method.

Auto ARIMA

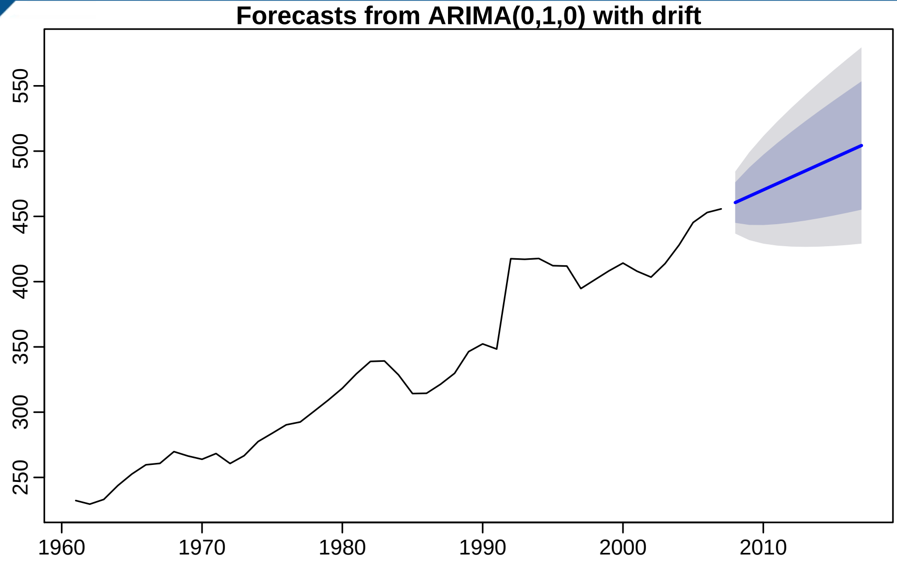

Forecasts from ARIMA(0,1,0) with drift
550
500
450
400
350
300
250
1960
1970
1980
1990
2000
2010

Auto ARIMA

```
fit <- auto.arima(livestock)
fcast <- forecast(fit)
plot(fcast)
```

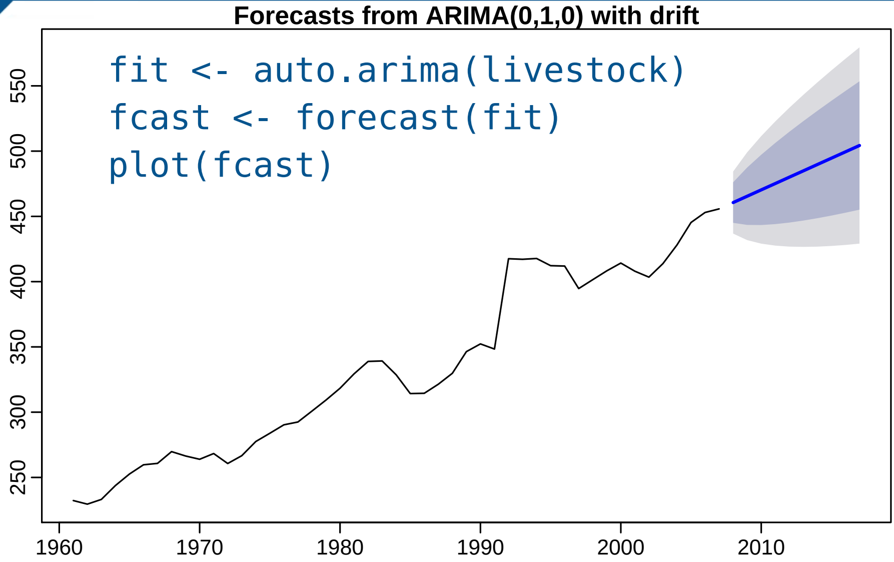

Auto ARIMA

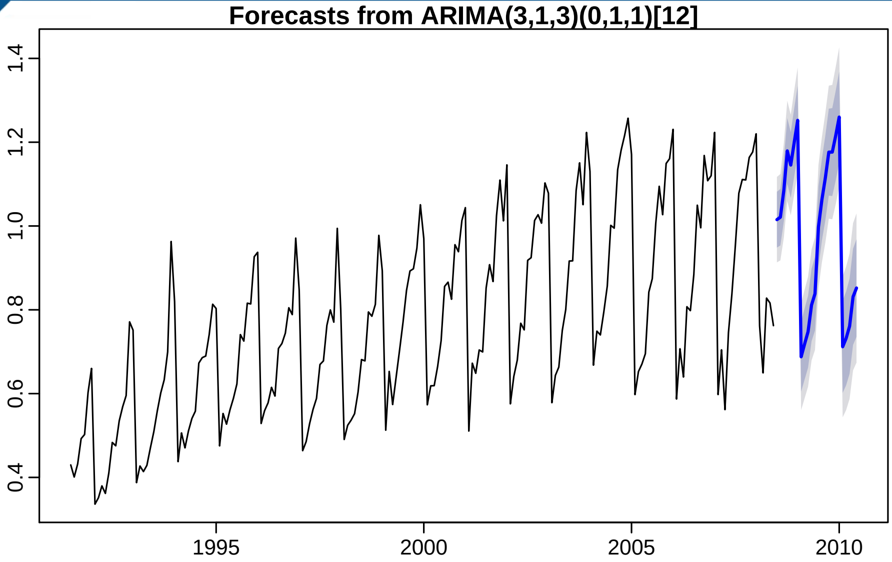

Forecasts from ARIMA(3,1,3)(0,1,1)[12]
0.4
0.6
0.8
1.0
1.2
1.4
1995
2000
2005
2010

Auto ARIMA

```
fit <- auto.arima(h02)
fcast <- forecast(fit)
plot(fcast)
```

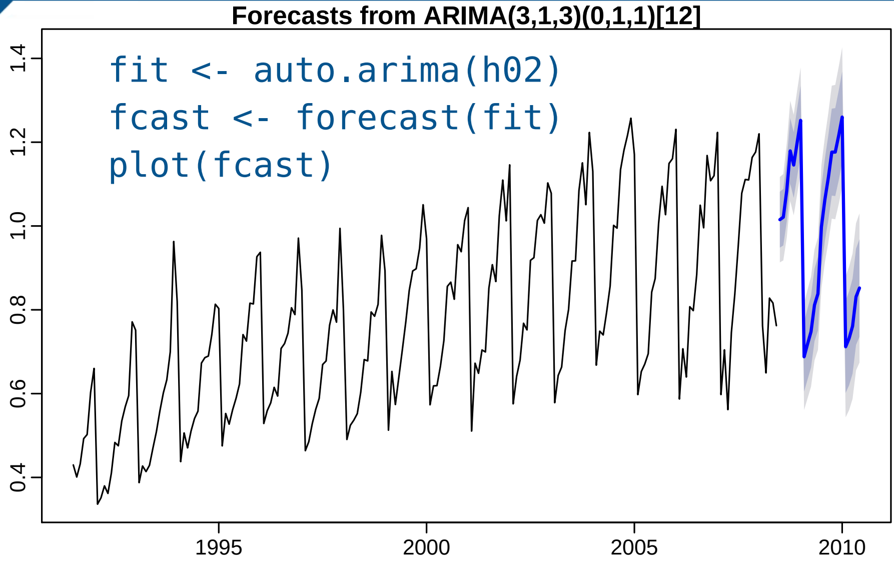

TBATS model

TBATS

Trigonometric terms for seasonality
Box-Cox transformations for heterogeneity
ARMA errors for short-term dynamics
Trend (possibly damped)
Seasonal (including multiple and non-integer periods)

Automatic algorithm described in
De Livera, Hyndman and Snyder (JASA 2011).

Examples

```r
fit <- tbats(gas)
fcast <- forecast(fit)
plot(fcast)
```

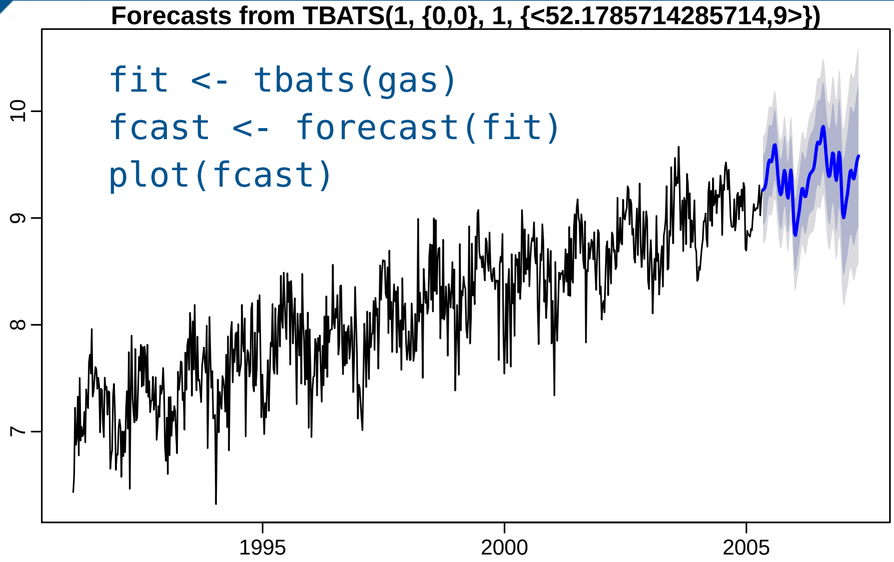

Examples

```
fit <- tbats(callcentre)
fcast <- forecast(fit)
plot(fcast)
```

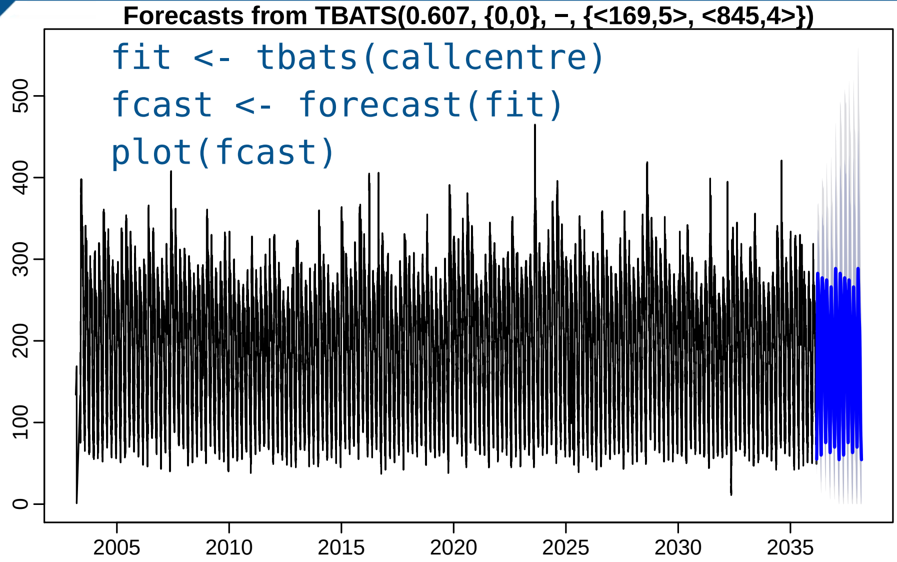

Outline

ggplot2 graphics

```r
plot(WWWusage, xlab="Minutes",
     ylab="Internet users",
     main="Internet Usage per Minute")
```

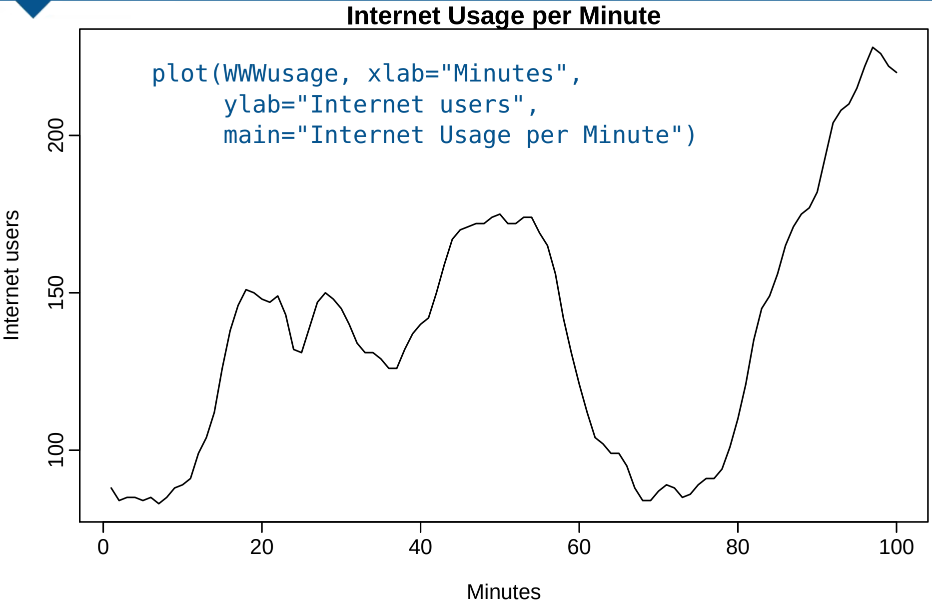

ggplot2 graphics

```
autoplot(WWWusage, xlab="Minutes",
         ylab="Internet users",
         main="Internet Usage per Minute")
```

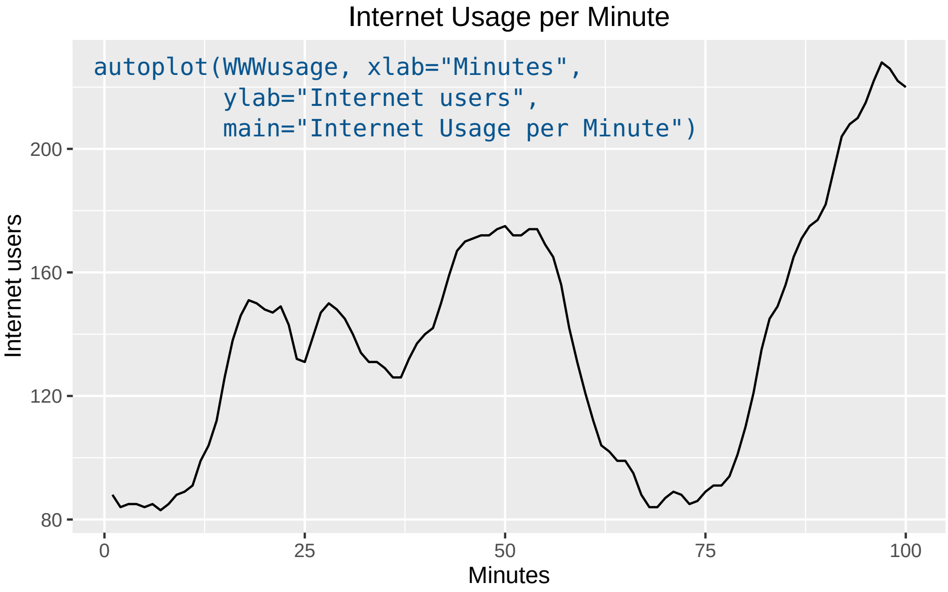

ggplot2 graphics

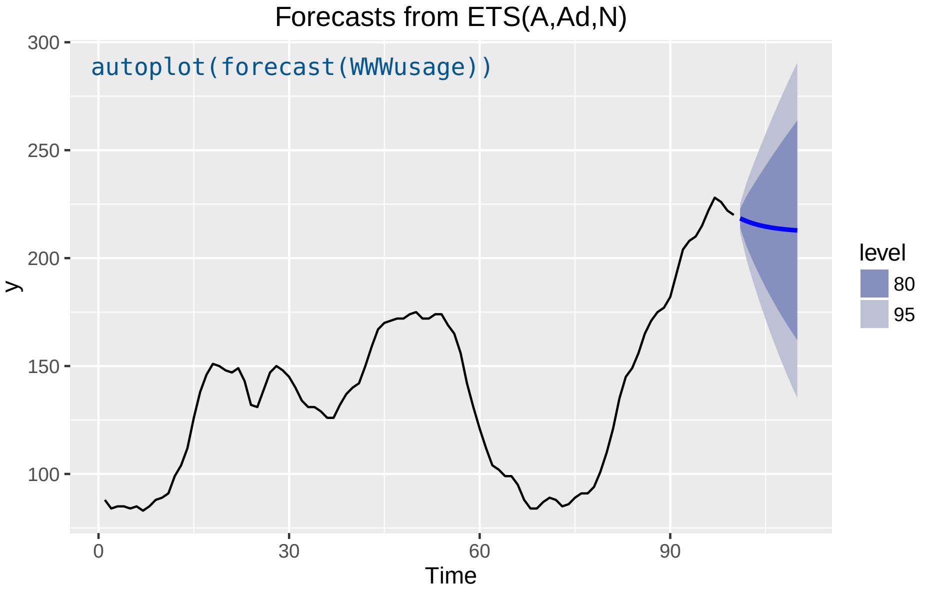

ggplot2 graphics

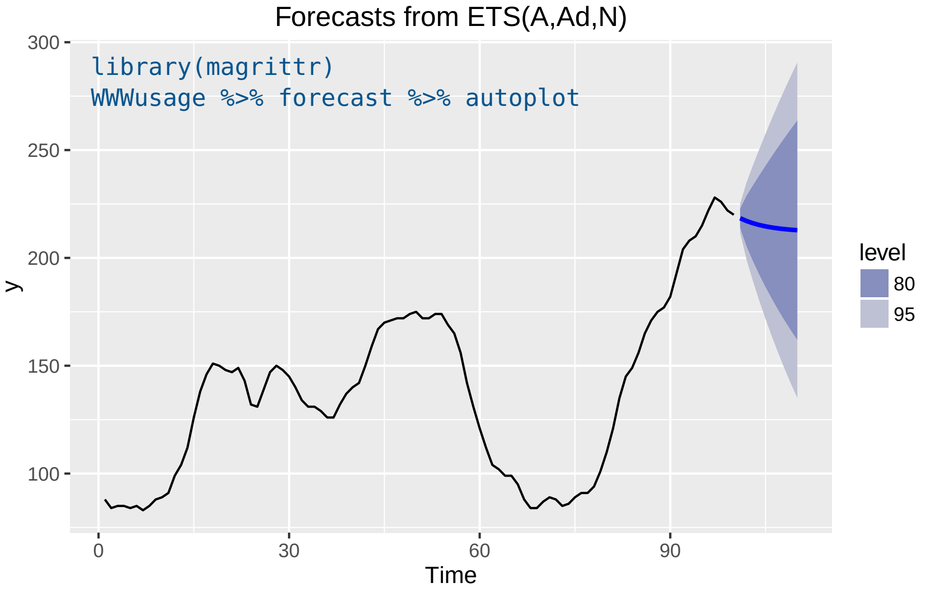

ggplot2 graphics

```
WWWusage %>% auto.arima %>% forecast(level=c(50,80,95)) -> fc
autoplot(WWWusage) +
  geom_forecast(fc, color='#ffcccc', show.legend=FALSE) +
  ggtitle("Forecasts from ARIMA model") +
  labs(x="Minute", y="Number of users")
```

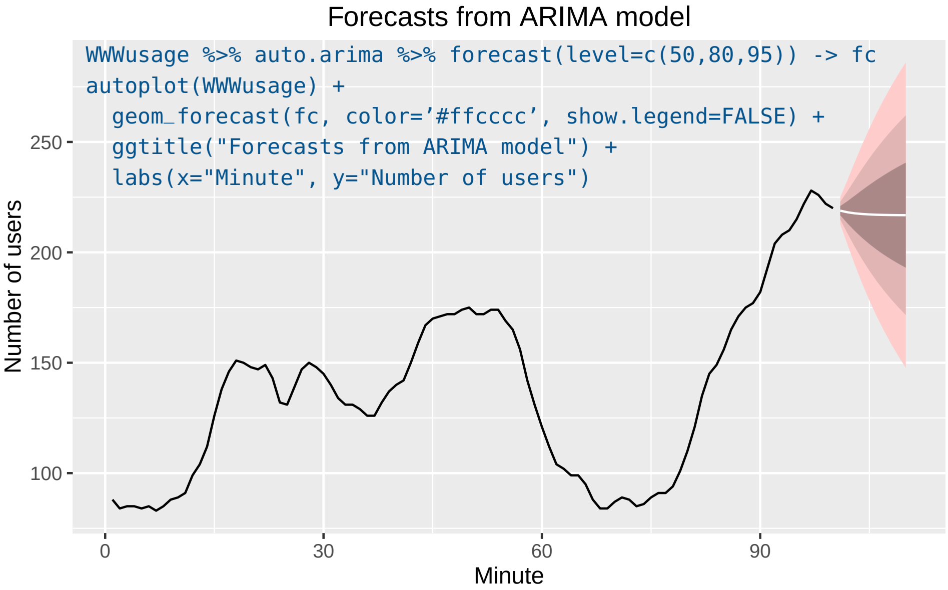

ggplot2 graphics

autoplot methods

- ts
- forecast
- acf
- stl
- Arima
- ets
- ...

Other ggplot2 graphics

- ggseasonplot
- ggmonthplot
- ggtsdisplay

Outline

Box-Cox transformations

$$w_t = f_\lambda(y_t) = \begin{cases} \log(y_t), & \lambda = 0; \\ (y_t^\lambda - 1)/\lambda, & \lambda \neq 0. \end{cases}$$

- $\lambda = 1$: (No substantive transformation)
- $\lambda = \frac{1}{2}$: (Square root plus linear transformation)
- $\lambda = 0$: (Natural logarithm)
- $\lambda = -1$: (Inverse plus 1)

Most modelling and forecasting functions in the forecast package have a lambda argument allowing Box-Cox transformations.

Box-Cox transformations

$$w_t = f_\lambda(y_t) = \begin{cases} \log(y_t), & \lambda = 0; \\ (y_t^\lambda - 1)/\lambda, & \lambda \neq 0. \end{cases}$$

- $\lambda = 1$: (No substantive transformation)
- $\lambda = \frac{1}{2}$: (Square root plus linear transformation)
- $\lambda = 0$: (Natural logarithm)
- $\lambda = -1$: (Inverse plus 1)

Most modelling and forecasting functions in the forecast package have a lambda argument allowing Box-Cox transformations.

Box-Cox transformations

$$w_t = f_\lambda(y_t) = \begin{cases} \log(y_t), & \lambda = 0; \\ (y_t^\lambda - 1)/\lambda, & \lambda \neq 0. \end{cases}$$

- $\lambda = 1$: (No substantive transformation)
- $\lambda = \frac{1}{2}$: (Square root plus linear transformation)
- $\lambda = 0$: (Natural logarithm)
- $\lambda = -1$: (Inverse plus 1)

Most modelling and forecasting functions in the forecast package have a lambda argument allowing Box-Cox transformations.

Box-Cox transformations

$$w_t = f_\lambda(y_t) = \begin{cases} \log(y_t), & \lambda = 0; \\ (y_t^\lambda - 1)/\lambda, & \lambda \neq 0. \end{cases}$$

- $\lambda = 1$: (No substantive transformation)
- $\lambda = \frac{1}{2}$: (Square root plus linear transformation)
- $\lambda = 0$: (Natural logarithm)
- $\lambda = -1$: (Inverse plus 1)

Most modelling and forecasting functions in the forecast package have a `lambda` argument allowing Box-Cox transformations.

Box-Cox transformations

$$w_t = f_\lambda(y_t) = \begin{cases} \log(y_t), & \lambda = 0; \\ (y_t^\lambda - 1)/\lambda, & \lambda \neq 0. \end{cases}$$

- $\lambda = 1$: (No substantive transformation)
- $\lambda = \frac{1}{2}$: (Square root plus linear transformation)
- $\lambda = 0$: (Natural logarithm)
- $\lambda = -1$: (Inverse plus 1)

Most modelling and forecasting functions in the forecast package have a `lambda` argument allowing Box-Cox transformations.

Box-Cox transformations

$$w_t = f_\lambda(y_t) = \begin{cases} \log(y_t), & \lambda = 0; \\ (y_t^\lambda - 1)/\lambda, & \lambda \neq 0. \end{cases}$$

- $\lambda = 1$: (No substantive transformation)
- $\lambda = \frac{1}{2}$: (Square root plus linear transformation)
- $\lambda = 0$: (Natural logarithm)
- $\lambda = -1$: (Inverse plus 1)

Most modelling and forecasting functions in the forecast package have a `lambda` argument allowing Box-Cox transformations.

Box-Cox transformations

$$w_t = f_\lambda(y_t) = \begin{cases} \log(y_t), & \lambda = 0; \\ (y_t^\lambda - 1)/\lambda, & \lambda \neq 0. \end{cases}$$

- $\lambda = 1$: (No substantive transformation)
- $\lambda = \frac{1}{2}$: (Square root plus linear transformation)
- $\lambda = 0$: (Natural logarithm)
- $\lambda = -1$: (Inverse plus 1)

Most modelling and forecasting functions in the forecast package have a `lambda` argument allowing Box-Cox transformations.

Box-Cox transformations

Back-transformation

We must reverse the transformation (or *back-transform*) to obtain forecasts on the original scale. The reverse Box-Cox transformations are given by

$$y_t = f_\lambda^{-1}(w_t) = \begin{cases} \exp(w_t), & \lambda = 0; \\ (\lambda w_t + 1)^{1/\lambda}, & \lambda \neq 0. \end{cases}$$

```
fit <- auto.arima(elec, lambda=1/3)
fc <- forecast(fit)
plot(fc, include=120)
```

Back-transformation

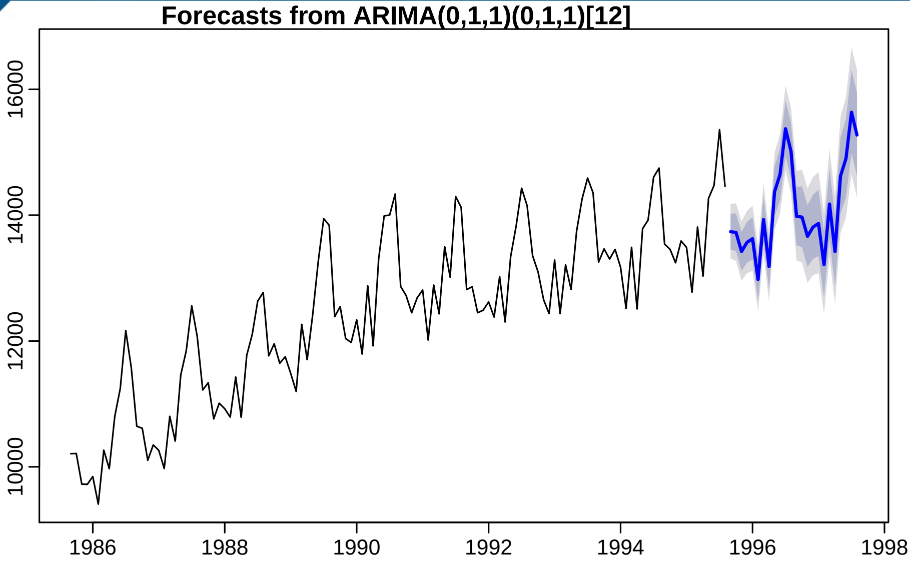

Bias and Box-Cox

- If the forecast is the **mean** on the transformed scale, it is *not* the **mean** on the original scale.
- If the forecast is the **median** on the transformed scale, it *is* the **median** on the original scale.
- Quantiles are preserved because the transformation is monotonically increasing.

If $\text{E}(W) = \mu$ and $\text{Var}(W) = \sigma^2$, then

Bias

$$\text{E}(Y) - e^{\mu} \approx \begin{cases} (\lambda\mu + 1)^{1/\lambda}\left[1 + \frac{\sigma^2(1-\lambda)}{2(\lambda\mu+1)^2}\right] - e^{\mu} & \text{if } \lambda \neq 0; \\ \frac{1}{2}e^{\mu}\sigma^2 & \text{if } \lambda = 0. \end{cases}$$

Bias and Box-Cox

- If the forecast is the **mean** on the transformed scale, it is *not* the **mean** on the original scale.
- If the forecast is the **median** on the transformed scale, it *is* the **median** on the original scale.
- Quantiles are preserved because the transformation is monotonically increasing.

If $\text{E}(W) = \mu$ and $\text{Var}(W) = \sigma^2$, then

Bias

$$\text{E}(Y) - e^{\mu} \approx \begin{cases} (\lambda\mu + 1)^{1/\lambda}\left[1 + \frac{\sigma^2(1-\lambda)}{2(\lambda\mu+1)^2}\right] - e^{\mu} & \text{if } \lambda \neq 0; \\ \frac{1}{2}e^{\mu}\sigma^2 & \text{if } \lambda = 0. \end{cases}$$

Bias and Box-Cox

- If the forecast is the **mean** on the transformed scale, it is *not* the **mean** on the original scale.
- If the forecast is the **median** on the transformed scale, it *is* the **median** on the original scale.
- Quantiles are preserved because the transformation is monotonically increasing.

If $\text{E}(W) = \mu$ and $\text{Var}(W) = \sigma^2$, then

Bias

$$\text{E}(Y) - e^{\mu} \approx \begin{cases} (\lambda\mu + 1)^{1/\lambda}\left[1 + \frac{\sigma^2(1-\lambda)}{2(\lambda\mu+1)^2}\right] - e^{\mu} & \text{if } \lambda \neq 0; \\ \frac{1}{2}e^{\mu}\sigma^2 & \text{if } \lambda = 0. \end{cases}$$

Bias and Box-Cox

- If the forecast is the **mean** on the transformed scale, it is *not* the **mean** on the original scale.
- If the forecast is the **median** on the transformed scale, it *is* the **median** on the original scale.
- Quantiles are preserved because the transformation is monotonically increasing.

If $\text{E}(W) = \mu$ and $\text{Var}(W) = \sigma^2$, then

Bias

$$\text{E}(Y) - e^{\mu} \approx \begin{cases} (\lambda\mu + 1)^{1/\lambda}\left[1 + \frac{\sigma^2(1-\lambda)}{2(\lambda\mu+1)^2}\right] - e^{\mu} & \text{if } \lambda \neq 0; \\ \frac{1}{2}e^{\mu}\sigma^2 & \text{if } \lambda = 0. \end{cases}$$

Bias and Box-Cox

- If the forecast is the **mean** on the transformed scale, it is *not* the **mean** on the original scale.
- If the forecast is the **median** on the transformed scale, it *is* the **median** on the original scale.
- Quantiles are preserved because the transformation is monotonically increasing.

If $\text{E}(W) = \mu$ and $\text{Var}(W) = \sigma^2$, then

Bias

$$\text{E}(Y) - e^{\mu} \approx \begin{cases} (\lambda\mu + 1)^{1/\lambda}\left[1 + \frac{\sigma^2(1-\lambda)}{2(\lambda\mu+1)^2}\right] - e^{\mu} & \text{if } \lambda \neq 0; \\ \frac{1}{2}e^{\mu}\sigma^2 & \text{if } \lambda = 0. \end{cases}$$

Bias adjustment

```
fit <- ets(eggs, lambda=0)
fc <- forecast(fit, h=50, level=95)
fc2 <- forecast(fit, h=50, level=95, biasadj=TRUE)
autoplot(fc, main="Price of a dozen eggs", ylab="cents") +
  geom_forecast(fc2, plot.conf=FALSE, color="red") +
  guides(fill=FALSE)
```

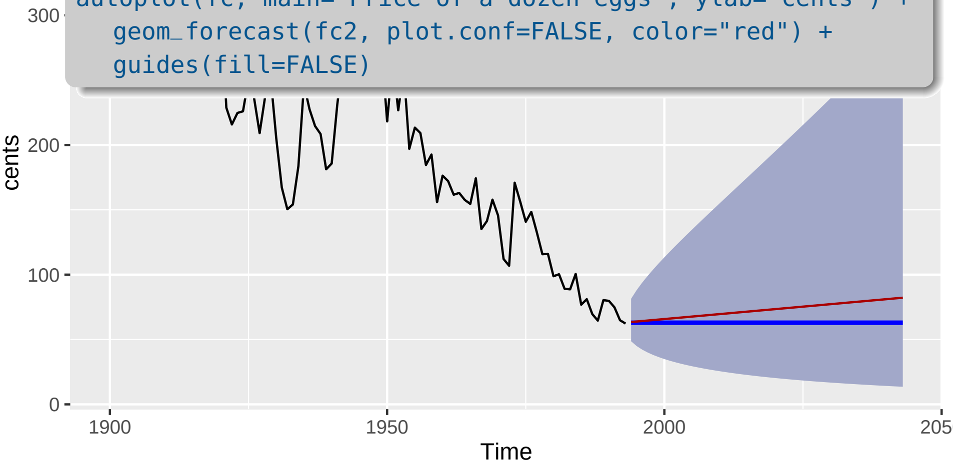

Bias adjustment

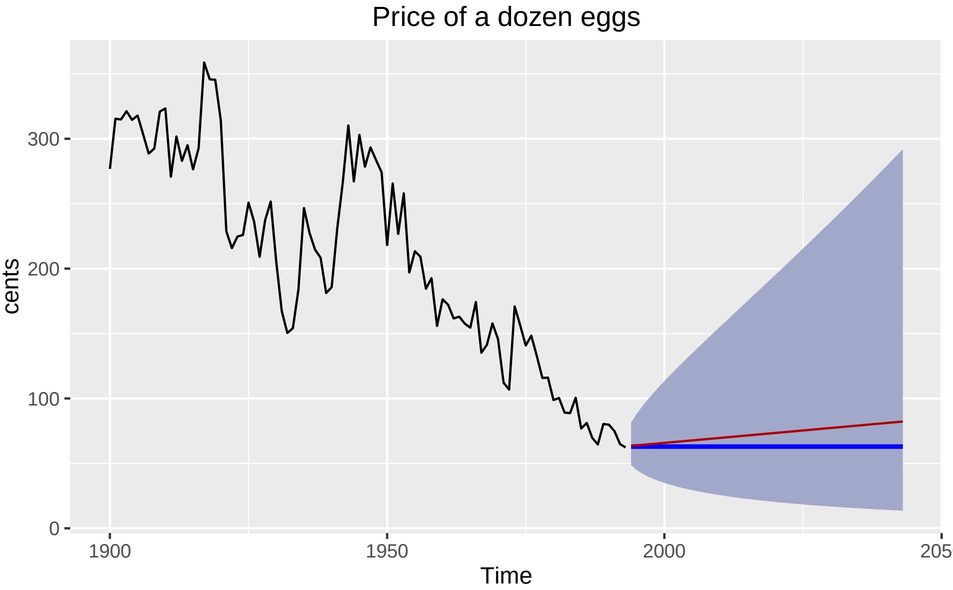

For further information

robjhyndman.com

- Slides for this talk.
- Links to all papers and books.
- Links to R packages.
- A blog about forecasting research.

OTexts.org/fpp

- Free online book based on forecast package for R.