Advanced methods: Dynamic regression models

Dynamic regression models

- ARIMA models allow for the inclusion of information from past observations of a series, but not for the inclusion of other information that may also be relevant
- for example, the effects of holidays, competitor activity, changes in the law, the wider economy, or other external variables, may explain some of the historical variation and may lead to more accurate forecasts
- on the other hand, linear regression models allow for the inclusion of a lot of relevant information from predictor variables, but do not allow for the subtle time series dynamics that can be handled with ARIMA models
- Solution: we will allow the errors from a regression to contain autocorrelation.

Dynamic regression models

Regression with ARIMA errors

A linear regression model with ARIMA errors may be defined as

$$y_t = \beta_0 + \beta_1 x_{1,t} + \cdots + \beta_k x_{k,t} + \eta_t$$
$$(1 - \phi_1 B)(1 - B)\eta_t = (1 + \theta_1 B)e_t$$

where e_t is a white noise series.
In this case, η_t follows an ARIMA $(1, 1, 1)$ model.

Example

US quarterly changes in personal consumption and income

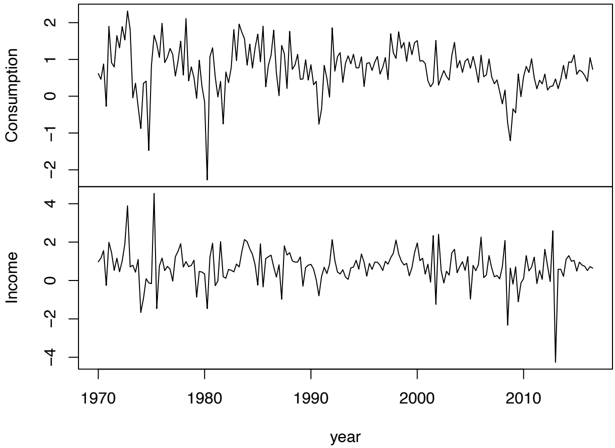

Example

US quarterly changes in personal consumption and income

```
Series: consumption
Regression with ARIMA(1,0,2) errors

Coefficients:
         ar1      ma1     ma2  intercept    xreg
      0.6922  -0.5758  0.1984     0.5990  0.2028
s.e.  0.1159   0.1301  0.0756     0.0884  0.0461

sigma^2 = 0.3219  log likelihood = -156.95
AIC=325.91   AICc=326.37   BIC=345.29
```

Example

US quarterly changes in personal consumption and income

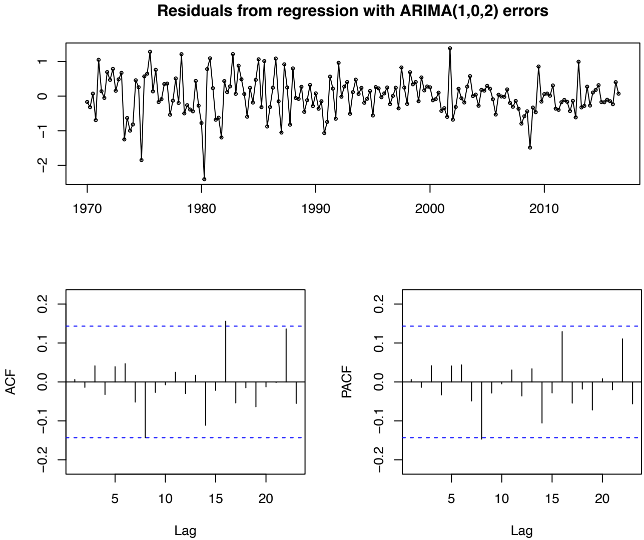

Example

US quarterly changes in personal consumption and income

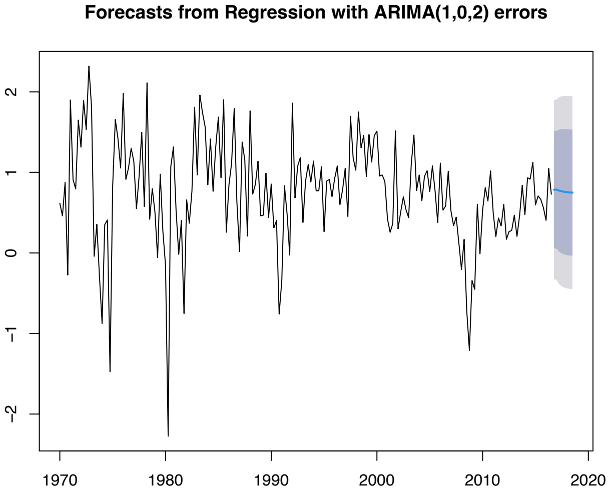

Deterministic trend

Total annual visitors in Australia

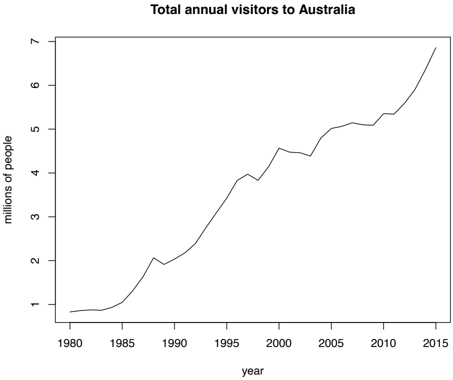

A deterministic trend can be obtained by using the model

$$y_t = \beta_0 + \beta_1 t + \eta_t$$

where η_t is an ARMA process.

Example

Total annual visitors in Australia

```
Series: austa
Regression with ARIMA(2,0,0) errors

Coefficients:
         ar1      ar2  intercept    xreg
      1.1127  -0.3805     0.4156  0.1710
s.e.  0.1600   0.1585     0.1897  0.0088

AIC=-17.2   AICc=-15.2   BIC=-9.28
```

The model can be written as

$$\begin{aligned} y_t &= 0.42 + 0.17t + \eta_t \\ \eta_t &= 1.11\eta_{t-1} - 0.38\eta_{t-2} + e_t \end{aligned}$$

Example

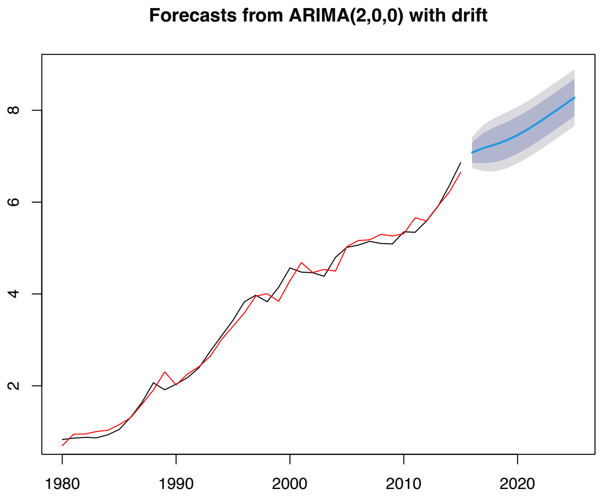
Forecasts from ARIMA(2,0,0) with drift
2
4
6
8
1980
1990
2000
2010
2020

Lagged predictors

- Sometimes the impact of a predictor included in a regression model will not be simple and immediate: for example an advertising campaign may impact sales for some time beyond the end of the campaign
- in these situations, we need to allow for lagged effects of the predictor

A model that allows for lagged effects can be written as

$$y_t = \beta_0 + \gamma_0 x_t + \gamma_1 x_{t-1} + \cdots + \gamma_k x_{t-k} + \eta_t$$

where η_t is an ARIMA process.
The value k can be selected using the AIC along with the values of p and q for the ARIMA.

Example

TV advertising and insurance quotations

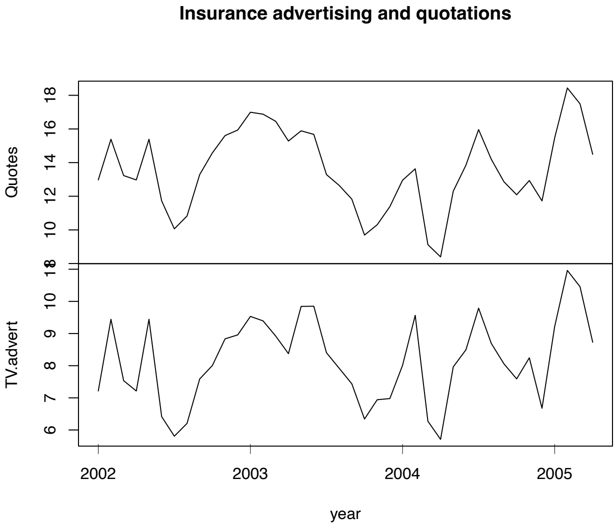

Example

TV advertising and insurance quotations

```
Series: insurance[, 1]
Regression with ARIMA(3,0,0) errors

Coefficients:
         ar1      ar2     ar3  intercept  AdLag0  AdLag1
      1.4117  -0.9317  0.3591     2.0393  1.2564  0.1625
s.e.  0.1698   0.2545  0.1592     0.9931  0.0667  0.0591

sigma^2 = 0.2165:  log likelihood = -23.89
AIC=61.78   AICc=65.4   BIC=73.43
```

The best model, according to the AIC, has two lagged predictors, and can be written as

$$y_t = 2.04 + 1.26x_t + 0.16x_{t-1} + \eta_t$$
$$\eta_t = 1.41\eta_{t-1} - 0.93\eta_{t-2} + 0.36\eta_{t-3} + e_t$$

Example

TV advertising and insurance quotations
we assume future advertising is 8 units in each future month

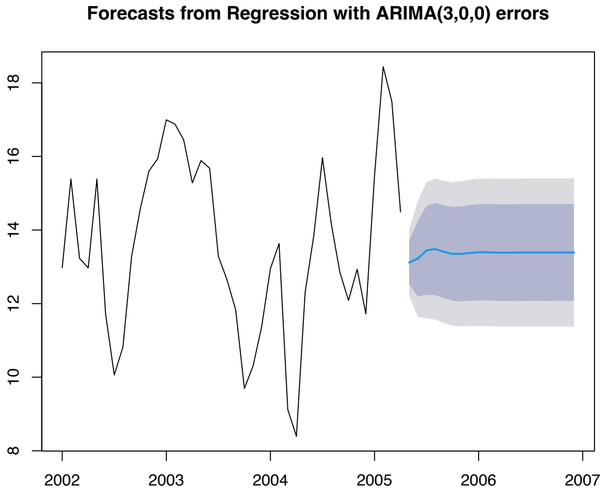

ARMAX refinement in diffusion models

An ARMAX model is defined as

$$y_t = \lambda x_t + \phi_1 y_{t-1} + \cdots + \phi_h y_{t-h} - \theta_1 z_{t-1} - \cdots - \theta_k z_{t-k} + z_t.$$

If we consider $x_t = f(\hat{\beta}, t)$, we can write

$$y_t = \lambda f(\hat{\beta}, t) + \phi_1 y_{t-1} + \cdots + \phi_h y_{t-h} - \theta_1 z_{t-1} - \cdots - \theta_k z_{t-k} + z_t$$

or

$$y_t - \lambda f(\hat{\beta}, t) = \phi_1 y_{t-1} + \cdots + \phi_h y_{t-h} - \theta_1 z_{t-1} - \cdots - \theta_k z_{t-k} + z_t.$$

If $\lambda = 1$ then

$$y_t - \lambda f(\hat{\beta}, t) = \hat{\varepsilon}(t) = \phi_1 y_{t-1} + \cdots + \phi_h y_{t-h} - \theta_1 z_{t-1} - \cdots - \theta_k z_{t-k} + z_t.$$

Residuals follow an ARMA(p, q).

Example

Apple iPod

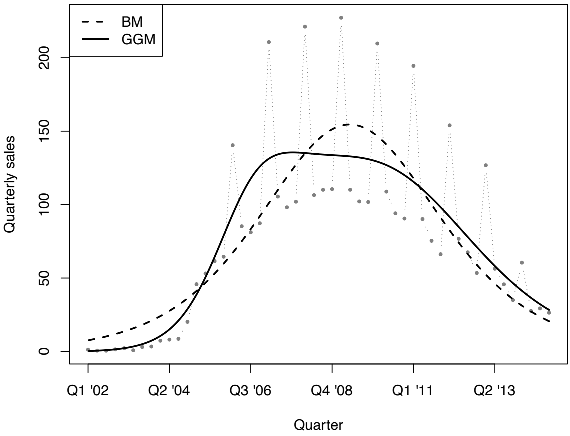

Example

Analysis of residuals

- We select the best nonlinear model (in this case GGM)
- we perform an analysis of residuals to detect possible autocorrelations: evidence of seasonality

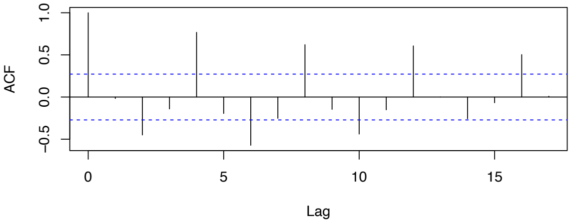

Example

SARMAX refinement after GGM

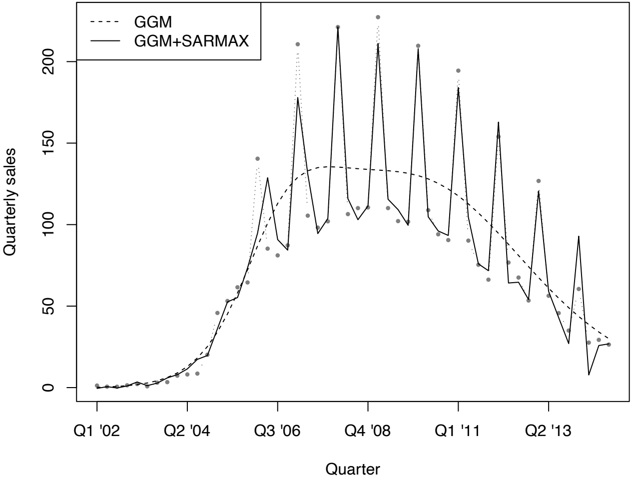

Example

parameter	ar1	ar2	ma1	ma2	sar1	sma1	λ
estimate	1.54	-0.81	-0.99	-0.006	0.85	0.38	1.00
s.e.	0.09	0.08	0.20	0.002	0.85	0.15	0.001

- Parameter $\lambda = 1$
- the ACF of residuals has no significant components

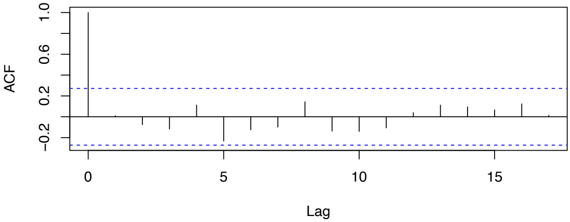

Example

Zoom share prices

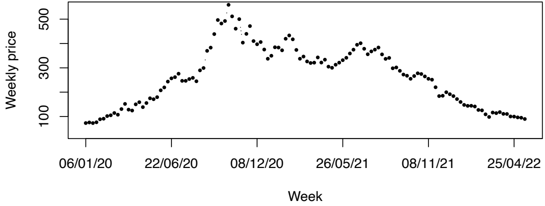

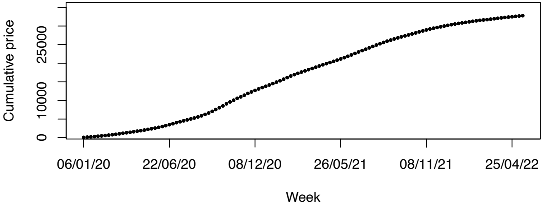

Example

Comparison between models: GBMe1 and GGM

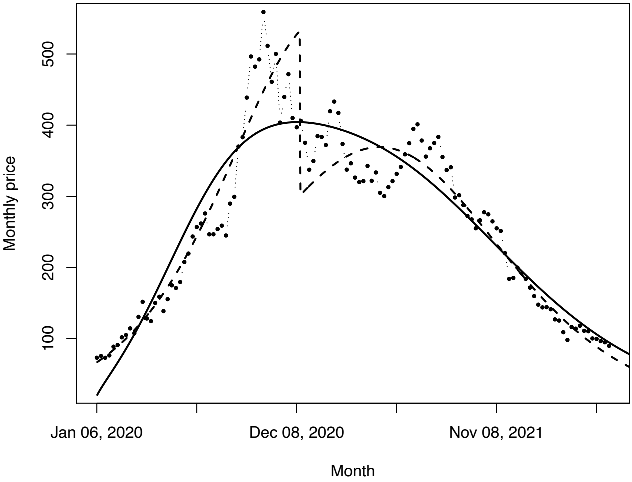

Example

SARMAX refinement after GBMe1