Econometria Financeira (ARIMA + GARCH)

Análise de Dados Financeiros e Econômicos com o R

Marcelo S. Perlin

EA-UFRGS

18/11/2026

Sumário

- Introduction
- Arima Models
- Garch Models

Introduction

Most popular models:

- ~~Linear models (OLS)~~
- ~~Generalized linear models (GLS)~~
- ~~Panel data models~~
- Arima models (Integrated Autoregressive Moving Averages)
- Garch models (generalized autoregressive conditional heteroskedasticity)
- ~~Regime switching models~~

Please read the references!

- The econometrics of financial markets (Campbell et al. 1997)
- Introductory econometrics for finance (Brooks 2014)
- Time series analysis (Hamilton 1994)
- Econometric analysis (Greene 2003)

Arima Models

Introduction to Arima Models

> Arima is a special type of model that uses the past to explain its own future.

Example of a ARIMA(AR = 1, D = 0, MA = 1):

$$y_t = 0.5y_{t-1} - 0.2\epsilon_{t-1} + \epsilon_t$$

- *AR* indicates the maximum *lag* used in y_t in the right hand side of the equation.
- D indicates the level of differentiation of the time series (Hamilton 1994)
- *MA* defines the maximum *lag* used for the error of the model.

Interesting properties

- A AR(1) is simply a ARIMA(1, 0, 0) model

$$y_t = 0.15y_{t-1} + \epsilon_t$$

- if $\theta = 1$ in a AR(1) model, then we have the simplest case of a Random walk model: $y_t = y_{t-1} + \epsilon_t$
- if $abs(\theta) > 1$, the stochastic process is explosive:

$$y_t = 1.5y_{t-1} + \epsilon_t$$

- A ARMA model is a parsimonious representation of a AR model with infinite lags $AR(\infty)$

Simulating Arima Models

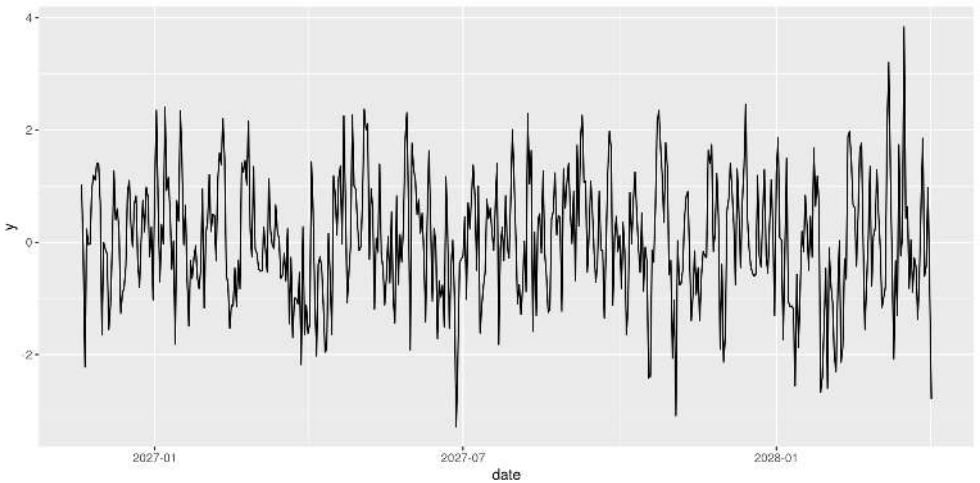
```r
library(ggplot2)

set.seed(1)

# set number of observati
n_T <- 500

# set model's parameters
arima_param <- list(ar =
                    ma =
my_sd <- 1

# simulate model
```



Estimating Arima Models

```r
my_arima <- stats::arima(
  my_ts,
  order = c(1,0,1))
```

Simulated Model:

$$y_t = 0.5 y_{t-1} + -0.1 \epsilon_{t-1} + \epsilon_t$$

Estimated values

```
Call:
stats::arima(x = my_ts,
order = c(1, 0, 1))

Coefficients:
          ar1      ma1
intercept
      0.4284  -0.0527
0.0248
s.e.  0.1006   0.1102
```

Estimating ARIMA for real data

1. make sure data is stationary (dick fuller test) and whether it needs differentiation
2. differentiate the time series (or not) – this sets parameter D
3. estimate ARMA model on resulting series
4. make predictions and applications

Example for SP500

```
df_SP500 <- yfR::yf_get(
   "^GSPC",
   "2015-01-01") |>
   na.omit()
```

The lazy man approach – Auto-ARIMA

> We can use computer power to find and estimate the best ARIMA model for the current data. It uses a fitting criteria to select models

```
# estimate arima model with automatic identification
my_autoarima <- forecast::auto.arima(
  x = df_SP500$ret_adjusted_prices)

my_autoarima
```

```
Series:
ARIMA(2,0,2) with non-zero mean

Coefficients:
          ar1      ar2     ma1     ma2    mean
      -1.7329  -0.8726  1.6394  0.7615  0.0005
s.e.   0.0255   0.0242  0.0340  0.0321  0.0002
```

```
sigma^2 = 0.0001174:  log likelihood = 9276.35
AIC=-18540.7   AICc=-18540.67   BIC=-18504.69
```

Forecasting Arima Models

```
1 # forecast model
2 forecast::forecast(my_autoarima, h = 5)
```

```
     Point Forecast        Lo 80       Hi 80        Lo 95
Hi 95
2987   0.0008412800 -0.01304730 0.01472986 -0.02039947
0.02208203
2988  -0.0009921824 -0.01494137 0.01295700 -0.02232562
0.02034125
2989   0.0027434305 -0.01122367 0.01671053 -0.01861740
0.02410427
2990  -0.0021302266 -0.01609763 0.01183718 -0.02349153
0.01923107
2991   0.0020557259 -0.01091016 0.01702062 -0.01831702
```

An empirical application of ARIMA Models

> “Can We Predict the Financial Markets Based on Google’s Search Queries?” – Journal of Forecasting

Garch Models

Introduction

> Garch models relate to the seminal work of Engle (1982) and Bollerslev (1986). The main innovation in this class of models is that the variance (uncertainty) of the residual can change over time.

Benefits:

- directly models the clustering of volatility, allowing for a better representation of the data generating process
- allows for the forecast of future volatility
- main R libraries: `fGarch` (**R-fGarch?**), `rugarch` (**R-rugarch?**)

A simple GARCH model

> A GARCH model is modular. We have a process that sets the conditional mean, and another that defines the variance of the error.

Example for AR(1)-GARCH(1, 1):

$$
\begin{aligned}
y_t &= \mu + \theta y_{t-1} + \epsilon_t \\
\epsilon_t &\sim N\left(0, h_t\right) \\
h_t &= \omega + \alpha \epsilon_{t-1}^2 + \beta h_{t-1}
\end{aligned}
$$

Simulating Garch Models

$$
\begin{aligned}
y_t &= 0.02 + 0.1 y_{t-1} + \epsilon_t \\
\epsilon_t &\sim N\left(0, h_t\right) \\
h_t &= 0.001 + 0.15 \epsilon_{t-1}^2 + 0.8 h_{t-1}
\end{aligned}
$$

Step 01: Set up model's specifications

```
# set list with model spe
my_sim_garch = list(
  omega=0.001,
  alpha=0.15,
  beta=0.8,
  mu=0.02,
  ar = 0.1)

# set garch spec
spec = fGarch::garchSpec(
  model = my_sim_garch
)

```

```
Formula:
 ~ ar(1) + garch(1, 1)
Model:
 ar:    0.1
 mu:    0.02
 omega: 0.001
 alpha: 0.15
 beta:  0.8
Distribution:
 norm
```

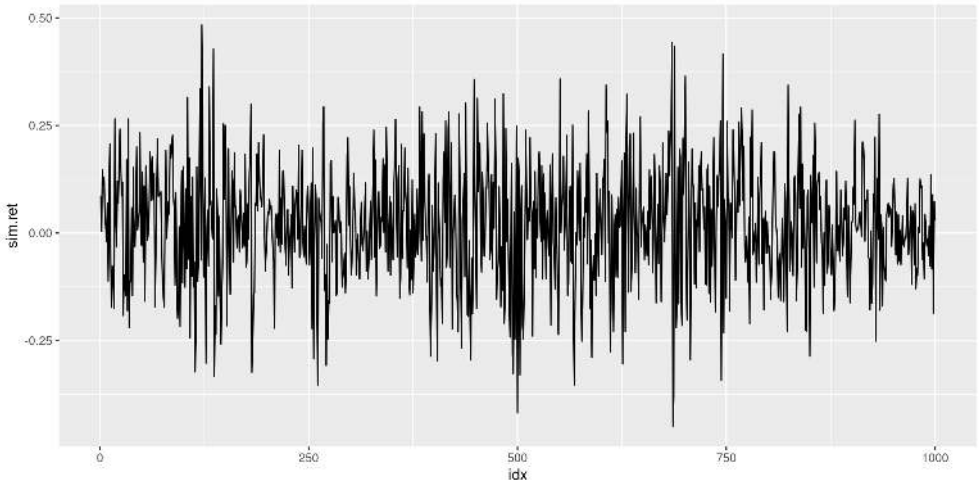
Step 02: Simulate it!

```r
set.seed(20)

# simulate garch model
n_T <- 1000
sim_garch = fGarch::garch

# set df for ggplot
temp_df <- dplyr::tibble(
  sim.ret = sim_garch$gar
  idx=seq_along(sim_garch

library(ggplot2)
p <- ggplot(temp_df, aes(
```



Estimating Garch Models

```r
# estimate garch model
my_garchfit <- fGarch::garchFit(
  data = sim_garch,
  formula = ~ arma(1,0) + garch(1,1),
  trace = FALSE
)

print(my_garchfit)
```

```
Title:
 GARCH Modelling

Call:
 fGarch::garchFit(formula = ~arma(1, 0) + garch(1, 1),
data = sim_garch,
    trace = FALSE)
```

Mean and Variance Equation:

Example for SP500

Steps:

1. Check for GARCH effects using LM Arch test

```
# test for Arch effects
MTS::archTest(
  rt = df_SP500$ret_adjus
)
```

```
Q(m) of squared series(LM
test):
Test statistic:  3287.009
p-value:  0
Rank-based Test:
Test statistic:  864.3169
p-value:  0
```

2. Estimate model

```
# set object for estimati
df_est <- fBasics::as.tim

# estimate garch model fo
my_garchfit_sp500 <- fGar
  data = df_est ,
  formula = ret_adjusted_
  trace = FALSE)

print(my_garchfit_sp500)
```

```
Title:
 GARCH Modelling

Call:
 fGarch::garchFit(formula =
ret_adjusted_prices ~
arma(1, 0) +
    garch(1, 1), data =
df_est, trace = FALSE)
```

Forecasting Garch Models

```
1 # static forecast for gar
2 my_forecast <- fGarch::pr
3   my_garchfit_sp500, n.ah
4
5 # print df
6 print(my_forecast)
```

```
  meanForecast  meanError
standardDeviation
1 0.0013943239 0.01415492
0.01415492
2 0.0007455040 0.01407883
0.01405674
3 0.0007816394 0.01398327
0.01396127
```

An empirical application of Garch Models

> “A GARCH Tutorial with R” – RAC

Exercises

1. Using package `yfR` (Perlin 2023), download data for the SP500 (ticker ^GSPC) between 2015 and today. Estimate a ARIMA (1, 0, 0) for the adjusted returns. What is the value of the autoregressive parameter?
2. Using same package, download data for the FTSE between 2010 and today. However, download data for daily, monthly and yeary observations. Estimate a ARIMA(1, 0, 1) for each of the time series of returns. Can you find any difference in the estimated coefficients?
3. Using same data from exercise 1, use function **forecast::auto.arima()** to estimate a ARIMA model with automatic selection of lags. What is the chosem model?
4. Using package `fGarch` (**R-fGarch?**), simulate 10000 observations of a GARCH(1,1) model with the following specification:

$$
\begin{aligned}
y_t &= 0.015 - 0.15 y_{t-1} + \epsilon_t \\
\epsilon_t &\sim N\left(0, h_t\right) \\
h_t &= 0.001 + 0.1 \epsilon_{t-1}^2 + 0.85 h_{t-1}
\end{aligned}
$$

5. Using the simulated data from previous exercise, estimate a GARCH(1, 1) model with **fGarch::garchFit()**. Do the estimated coefficients match the ones from the simulation?
6. Using package `yfR`, download data for the IBOVESPA index (ticker ^BVSP) between 2015 and today. Test the existence of arch effects in 1) prices series and 2) in the returns series. Report an discuss both results.
7. For the same data used in 4, estimate a GARCH(1, 1) for the adjusted returns. Report the model and its coefficients in the screen of R.

References

Bollerslev, Tim. 1986. “Generalized Autoregressive Conditional Heteroskedasticity.” *Journal of Econometrics* 31 (3): 307–27.

Brooks, Chris. 2014. *Introductory Econometrics for Finance*. Cambridge university press.

Campbell, John Y, Andrew Wen-Chuan Lo, Archie Craig MacKinlay, et al. 1997. *The Econometrics of Financial Markets*. Vol. 2. Princeton University press Princeton, NJ.

Engle, Robert F. 1982. “Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation.” *Econometrica: Journal of the Econometric Society*, 987–1007.

Greene, William H. 2003. *Econometric Analysis*. Pearson Education India.

Hamilton, James Douglas. 1994. *Time Series Analysis*. Vol. 2. Princeton university press Princeton.

Perlin, Marcelo. 2023. *yfR: Downloads and Organizes Financial Data from Yahoo Finance*. https://github.com/ropensci/yfR.