Analysis of Financial Time Series with EViews

Enrico Foscolo

Contents

1 Asset Returns

Most financial studies involve returns (more precisely, log return r_t of a stock) instead of prices.

Why?

Economic perspective Return of an asset is a scale–free summary of the investment opportunity.

Statistical perspective Return series are easier to handle than price series because of the more attractive statistical properties of the former (persistent effects in the second moment).

About our approch. One approach to forecast future values of an economic variable is to build a more or less structural econometric model, describing the relationship between the variable of interest with other economic quantities. Although this approach has the advantage of giving economic content to one's predictions, it is not always very useful. Here, we follow a pure time series approach. In this approach the current values of an economic variable are related to past values. The emphasis is purely on making use of the information in past values of a variable for forecasting its future.

1.1 Empirical Properties of Returns

- High–frequency of observation (cf. Figure 1).
- Non–normal empirical distribution (cf. Figure 2). The empirical density function has a higher peak around its mean, but fatter tails than that of the corresponding normal distribution.
- Daily returns of the market indexes and individual stocks tend to have high excess kurtosis (cf. Figure 2).
- The mean of log–return series is close to zero (cf. Figure 1).
- $p_t \sim I(1)$, so $r_t := \Delta \log p_t \sim I(0)$ (cf. Figure 1).
- We do not recognize autocorrelation in levels (r_t), but we do with squares of log–returns (cf. Figure 3).

***In–depth* 1 Stationary versus Nonstationary stochastic process**

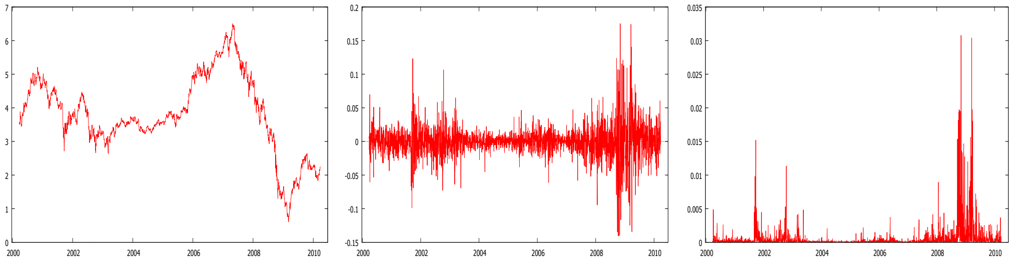

(a) (b) (c)

Figure 1: Plot of prices, log–returns, and squares of log–returns for Unicredit S.p.A. (ISIN code IT0004781412), respectively.

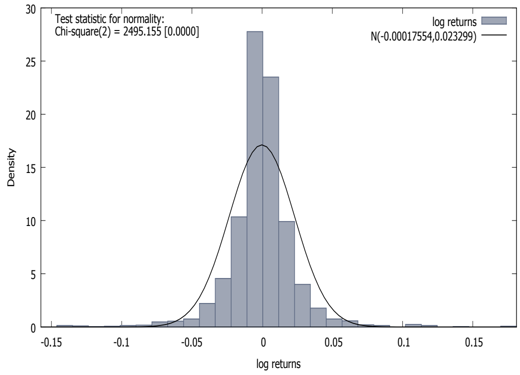

Figure 2: Frequency distribution of log–returns for Unicredit S.p.A. (ISIN code IT0004781412).

If the process for p_t has one unit root, we can eliminate the nonstationarity by transforming the series into first differences (changes). A series which becomes stationary after first differencing is said to be **integrated of order one**, denoted $I(1)$.

The main differences between $I(0)$ and $I(1)$ processes:

- $I(0)$ series fluctuates around its mean (it is said **mean reverting**) with a finite variance that does not depend on time, while $I(1)$ series wanders widely,
- $I(0)$ series has a limited memory of its past behavior (the effects of random innovations are only transitory), while a $I(1)$ series has an infinitely long memory (innovations permanently affect the process),
- autocorrelations of a $I(0)$ series rapidly decline as the lag increases,

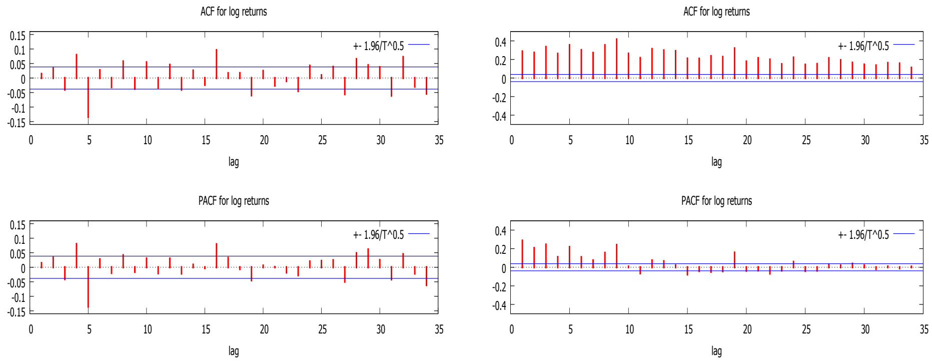

(a) (b)

Figure 3: Autocorrelation functions of log–returns and squares of log–returns for Unicredit S.p.A. (ISIN code IT0004781412), respectively.

while for a $I(1)$ process the estimated autocorrelation coefficients decay to zero very slowly.

Double click on the series name to open the series window, and choose `View / Unit Root Test...`.

The difference operator is given by function `d(.)`.

2 Heteroskedasticity and Autocorrelation

We discuss single equation regression techniques

$$r_t = \boldsymbol{x}_t^\top \boldsymbol{\beta} + u_t$$

The essential Gauss–Markov assumptions

$$E\{\boldsymbol{u}|\boldsymbol{X}\} = \boldsymbol{0}$$
$$Var\{\boldsymbol{u}|\boldsymbol{X}\} = \sigma^2 \boldsymbol{I}$$

A common finding in time series regressions:

- residuals are **correlated with their own lagged values**,
- error terms do **not have identical variances**.

Problems associated with serial correlation and heteroskedasticity:

- OLS is no longer efficient among linear estimators (**OLS is still unbiased**),
- standard errors computed using the OLS formula are not correct,
- t– and F–tests will no longer be valid and inference will be mislead.

We start with autocorrelation issue.

There are many forms of autocorrelation. Each one leads to a different structure for the error covariance matrix $Var\{\boldsymbol{u}\,|\boldsymbol{X}\}$.

Linear time series analysis provides a natural framework to study the dynamic structure of such a series. In general, we will be concerned with specifications of the form of higher–order autoregressive models; i.e., AR(p):

$$r_t = \boldsymbol{x}_t^\top \beta + u_t$$
$$u_t = \rho_1 u_{t-1} + \ldots + \rho_p u_{t-p} + \epsilon_t\,, \quad \epsilon_t \sim WN\left(0, \sigma_\epsilon^2\right)$$

The autocorrelations of a stationary AR(p) process gradually die out to zero, while the partial autocorrelations for lags larger than p are zero.

2.1 Testing for serial correlation

- **Correlogram and Ljung–Box Q–Statistic**

  `View / Residual Tests / Correlogram-Q-statistics`

- **Breusch–Godfrey Lagrange multiplier test statistic**

  `View / Residual Tests / Serial Correlation LM Test...`

2.2 AR estimation

Open an equation by selecting `Quick / Estimate Equation...` and enter your specification,

```
r c x1 ...  xk ar(1) ...  ar(p)
```

The stationarity condition for general AR(p) processes is that the inverted roots of the lag polynomial lie inside the unit circle. The **Inverted AR Roots** take place at the bottom of the regression output.

3 General ARMA Processes

ARIMA (Autoregressive Integrated Moving Average) models are generalizations of the simple AR model.

ARMA(p, q):

$$
\begin{aligned}
r_t &= \boldsymbol{x}_t^\top \beta + u_t \\
u_t &= \rho_1 u_{t-1} + \ldots + \rho_p u_{t-p} + \epsilon_t + \theta_1 \epsilon_{t-1} + \ldots + \theta_q \epsilon_{t-q}\,, \quad \epsilon_t \sim WN\left(0, \sigma_\epsilon^2\right)
\end{aligned}
$$

The goal of ARIMA analysis is a parsimonious representation of the process governing the residual.
The specification can also be applied directly to a series.

3.1 ARMA estimation

Open an equation by selecting `Quick / Estimate Equation...` and enter your specification,

```
y c x1 ...  xk ar(1) ...  ar(p) ma(1) ...  ma(q)
```

3.2 ARMA Equation Diagnostics

From the menu of an estimated equation `View / ARMA Structure...`.

Roots If the estimated ARMA process is (covariance) stationary, then all inverse AR roots should lie inside the unit circle. If the estimated ARMA process is invertible, then inverse all MA roots should lie inside the unit circle.

Correlogram If the ARMA model is correctly specified, the residuals from the model should be nearly white noise. More general tests for serial correlation in the residuals may be carried out with `View / Residual Tests / Correlogram-Q-statistic` and `View / Residual Tests / Serial Correlation LM Test...`.

4 GARCH Models

We turn into the heteroskedasticity issue.

Besides the return series, we also consider the volatility process. The volatility process is concerned with the evolution of **conditional variance** of the return over time, denoting by σ_t^2, namely the one–period ahead forecast variance based on past information.

Why?

The variabilities of returns vary over time and appear in clusters (i.e., volatility may be high for certain time periods and low for other periods). Plots of log–returns show the so–called **volatility clustering** phenomenon, where large changes in returns are likely to be followed by further large changes. Moreover, volatility seems to react differently to a big price increase or a big price drop, referred to as the **leverage effect**. In application, volatility plays an important role in pricing options, risk management (e.g., Value–at–Risk), and in asset allocation under the mean–variance framework.

The advantage of knowing about risks is that we can change our behavior to avoid them. Thus, we optimize our reactions and in particular our portfolio, to maximize rewards and minimize risks.

- A special feature of stock volatility is that it is not directly observable.
- Forecast confidence intervals may be time–varying, so that more accurate intervals can be obtained by modeling the variance of the errors.
- More efficient estimators can be obtained if heteroskedasticity in the errors is handled properly.
- The volatility index of a market has recently become a financial instrument. The VIX volatility index compiled by the Chicago Board of Option Exchange (CBOE) started to trade in futures on March 26, 2004.

***In–depth* 2 Option Pricing**

Consider the price of an *European call option*, which is a contract giving its holder the right, but not the obligation, to buy a fixed number of shares of a specified common stock at a fixed price on a given date. The fixed price is called the *strike price* and is commonly denoted by K. The given date is called the *expiration date*. The important time span here is the time to expiration, and we denote it by l. The well–known *Black–Scholes* option pricing formula states that the price of such a call option is

$$c_t = p_t \, \Phi\left(x\right) - K \, r^{-l} \, \Phi\left(x - \sigma_t \sqrt{l}\right) \quad \text{and} \quad x = \frac{\ln\left(p_t / K \, r^{-l}\right)}{\sigma_t \sqrt{l}} + 0.5 \, \sigma_t \sqrt{l} \, , \tag{1}$$

where p_t is the current price of the underlying stock, r is the risk–free interest rate, σ_t is the conditional standard deviation of the log return of the specified stock, and $\Phi\left(\cdot\right)$ is the cumulative distribution function of the standard normal random variable. The conditional standard deviation σ_t of the log return of the underlying stock plays an important role and it evolves over time.

4.1 Basic GARCH Specifications

In options markets, if one accepts the idea that the prices are governed by an econometric model such as the Black–Scholes formula, then one can use the price to obtain the **implied volatility**. This approach is often criticized because of its own assumptions that might not hold in practice (e.g., the assumed geometric Brownian motion for the price of the underlying asset). For instance, from the observed prices of a European call option, one can use the Black–Scholes formula in (1) to deduce the conditional standard deviation σ_t. It might be different from the actual volatility. Experience shows that implied volatility of an asset return tends to be larger than that obtained by using a GARCH type of volatility model. For instance, the VIX of CBOE is an implied volatility.

A simple approach, called **historical volatility**, is widely used. In this method, the volatility is estimated by the sample standard deviation of returns over a short period. If the period is too long, then it will not be so relevant for today and if the period is too short, it will be very noisy. It is logically inconsistent to assume, for example, that the variance is constant for a period.

Conditional heteroscedastic models can be classified into two general categories. Those in the first category use an exact function to govern the evolution of σ_t^2, whereas those in the second category use a stochastic equation to describe σ_t^2. The GARCH model belongs to the first category.

Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models are specifically designed to model and forecast conditional variances. The variance of the dependent variable is modeled as a function of past values of the dependent variable and independent, or exogenous variables.

GARCH(1, 1):

$$
\begin{aligned}
r_t &= \boldsymbol{x}_t^\top \boldsymbol{\gamma} + u_t \\
u_t &= \sigma_t \, \epsilon_t \, , \quad \epsilon_t \sim IID\,(0,1) \\
\sigma_t^2 &= \omega + \alpha \, u_{t-1}^2 + \beta \, \sigma_{t-1}^2
\end{aligned}
$$

An ordinary ARCH model is a special case of a GARCH specification in which there are no lagged forecast variances in the conditional variance equation; i.e., a GARCH(1, 0).

This specification is often interpreted in a financial context, where an agent or trader predicts this period's variance by forming a weighted average of a long term average (the constant), the forecasted variance from last period (the GARCH term), and information about volatility observed in the previous period (the ARCH term). If the asset return was unexpectedly large in either the upward or the downward direction, then the trader will increase the estimate of the variance for the next period. This model is also consistent with the volatility clustering.

Equivalent representations. Rearranging terms and defining $\upsilon_t = u_t^2 - \sigma_t^2$

$$u_t^2 = \omega + (\alpha + \beta)\, u_{t-1}^2 + \upsilon_t - \beta\, \upsilon_{t-1}$$

the squared errors follow a heteroskedastic ARMA(1, 1) process; the autoregressive root which governs the persistence of volatility shocks is $(\alpha + \beta)$ and, in many applied settings, this root is very **close to unity so that shocks die out rather slowly**.

Higher order GARCH models, denoted GARCH(p, q),

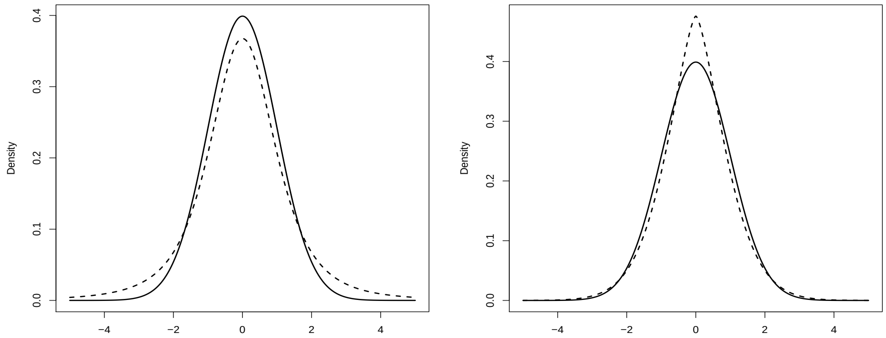

(a) (b)

Figure 4: Normal density (solid line) versus t–Student and GED density with 3 and 1.5 degree of freedom, respectively.

$$
\begin{aligned}
r_t &= \boldsymbol{x}_t^\top \boldsymbol{\gamma} + u_t \\
u_t &= \sigma_t \, \epsilon_t \,, \quad \epsilon_t \sim IID\,(0,1) \\
\sigma_t^2 &= \omega + \sum_{i=1}^{p} \alpha_i \, u_{t-i}^2 + \sum_{j=1}^{q} \beta_j \, \sigma_{t-j}^2
\end{aligned}
$$

The basic GARCH specification requires an assumption about the conditional distribution of the error term u_t.

- Gaussian distribution (cf. Figure 4).
- Student's t–distribution (cf. Figure 4).
- Generalized Error distribution (cf. Figure 4).

Given a distributional assumption, GARCH models are typically estimated by the method of **(quasi) maximum likelihood**.

`Quick / Estimate Equation ...` or by selecting `Object / New Object ... / Equation ...`. Select `ARCH` from the method combo box at the bottom of the dialog.

The presence of GARCH errors in a regression or autoregressive model does not invalidate OLS estimation. It does imply, however, that more efficient (nonlinear) estimators exist than OLS.

Estimation procedures are often computationally–intensive.

4.2 Diagnostic Checking

Building a volatility model for an asset return series consists of four steps.

1. Specify a mean equation by testing for serial dependence in the data and, if necessary, building an econometric model (e.g., an ARMA model) for the return series to remove any linear dependence.

2. Use the residuals of the mean equation to test for ARCH effects. To carry out the for autoregressive conditional heteroskedasticity (ARCH) in the residuals, push `View / Residual Tests / ARCH LM Test ...` on the equation toolbar and specify the order of ARCH to be tested against. It carries out Lagrange multiplier tests to test whether the standardized residuals exhibit additional ARCH. If the variance equation is correctly specified, there should be no ARCH left in the standardized residuals.

3. Specify a volatility model if ARCH effects are statistically significant and perform a joint estimation of the mean and volatility equations.

4. Check the fitted model carefully and refine it if necessary.

Actual, Fitted, Residual

GARCH Graph

Covariance Matrix

Coefficient Tests Likelihood ratio tests on the estimated coefficient.

Residual Tests / Correlogram–Q–statistics This view can be used to test for remaining serial correlation in the mean equation and to check the specification of the mean equation. If the mean equation is correctly specified, all Q–statistics should not be significant.

Residual Tests / Correlogram Squared Residuals This view can be used to test for remaining ARCH in the variance equation and to check the specification of the variance equation. If the variance equation is correctly specified, all Q–statistics should not be significant. See also `Residual Tests / ARCH LM Test`.

Residual Tests / Histogram–Normality Test

4.3 Regressors in the Variance Equation

$$\sigma_t^2 = \omega + \sum_{i=1}^{p} \alpha_i \, u_{t-i}^2 + \sum_{j=1}^{q} \beta_j \, \sigma_{t-j}^2 + \boldsymbol{z}_t^\top \boldsymbol{\pi}$$

Note that the forecasted variances from this model are not guaranteed to be positive.

4.4 The GARCH–M Model

Finance theory suggests that an asset with an higher perceived risk would pay an higher return on average. Then, the theory suggests that the mean return would be related to the conditional variance or standard deviation,

$$r_t = \boldsymbol{x}_t^\top \boldsymbol{\gamma} + \lambda \, \sigma_t^2 + u_t$$

The parameter λ is called the **risk premium** parameter. Not so used in practice; i.e., sometimes $\hat{\lambda}$ are not interpretable and catch effects different from the influence of volatility on conditional mean.

4.5 The Threshold GARCH (TARCH) Model

$$\sigma_t^2 = \omega + \sum_{i=1}^{p} \alpha_i \, u_{t-i}^2 + \sum_{j=1}^{q} \beta_j \, \sigma_{t-j}^2 + \sum_{k=1}^{r} \eta_k \, u_{t-k}^2 \, I_{t-k}$$

where $I_t = 1$ if $u_t < 0$ and 0 otherwise.

Good news, $u_{t-i} > 0$, and bad news, $u_{t-i} < 0$, have differential effects on the conditional variance; i.e., good news has an impact of α_i, while bad news has an impact of $\alpha_i + \eta_i$. If $\eta_i > 0$, bad news increases volatility, and we say that there is a **leverage effect** for the i–th order. If $\eta_i \neq 0$, the news impact is asymmetric.

4.6 The Exponential GARCH (EGARCH) Model

$$\log\left(\sigma_t^2\right) = \omega + \sum_{i=1}^{p} \alpha_i \left|\frac{u_{t-i}}{\sigma_{t-i}}\right| + \sum_{j=1}^{q} \beta_j \log\left(\sigma_{t-j}^2\right) + \sum_{k=1}^{r} \eta_k \frac{u_{t-k}}{\sigma_{t-k}}$$

This equation highlights the asymmetric responses in volatility to the past positive and negative shocks. The log of the conditional variance implies that the leverage effect is exponential and that forecasts of the conditional variance are guaranteed to be nonnegative. The presence of leverage effects can be tested by the hypothesis that $\eta_i < 0$. The impact is asymmetric if $\eta_i \neq 0$.

4.7 The Power ARCH (PARCH) Model

$$\sigma_t^\delta = \omega + \sum_{i=1}^{p} \alpha_i \left(|u_{t-i}| - \eta_i\, u_{t-i}\right)^\delta + \sum_{j=1}^{q} \beta_j\, \sigma_{t-j}^\delta$$

where $\delta > 0$, $\eta_i \leq 1$ for $i = 1, \ldots, r$, $\eta_i = 0$ for $i > r$ and $r \leq p$.

4.8 The Component GARCH (CGARCH) Model

It allows mean reversion to a varying level m_t,

$$\begin{aligned} \sigma_t^2 - m_t &= \bar{\omega} + \alpha\left(u_{t-1}^2 - \bar{\omega}\right) + \beta\left(\sigma_{t-1}^2 - \bar{\omega}\right) \\ m_t &= \omega + \rho\left(m_{t-1} - \omega\right) + \phi\left(u_{t-1}^2 - \sigma_{t-1}^2\right) \end{aligned}$$

σ_{t-1}^2 is still the volatility, while m_t takes the place of ω and is the time varying long–run volatility. The first equation describes the transitory component, $\sigma_t^2 - m_t$, which converges to zero with powers of $(\alpha + \beta)$. The second equation describes the long run component m_t, which converges to ω with powers of ρ.