UNIVERSITY OF LONDON | INTERNATIONAL PROGRAMMES

Course information 2015–16
FN3142 Quantitative finance

This course is aimed at students interested in obtaining a thorough grounding in market finance and related empirical methods.

Prerequisite
If taken as part of a BSc degree, courses which must be passed before this course may be attempted:

EC2020 Elements of econometrics ***and***
EC2066 Microeconomics.

Other rules
This course must be taken at the same time as or after FN3092 Corporate finance.

Aims and objectives
This course provides the econometric techniques, such as time-series analysis, required to analyse theoretical and empirical issues in finance. It provides applications in asset pricing, investments, risk analysis and management, market microstructure, and return forecasting.

Essential reading
For full details, please refer to the reading list

Christoffersen, P.F., 2011, *Elements of Financial Risk Management, Second Edition.* (Academic Press, London)

Diebold, F.X., 2007, *Elements of Forecasting, Fourth Edition.* (Thomson South-Western, Canada)

Learning outcomes
At the end of this course and having completed the essential reading and activities students should:

- ☑ To be able to demonstrate mastery of econometric techniques required in order to analyse issues in asset pricing and market finance
- ☑ To be able to demonstrate familiarity with recent empirical findings based on financial econometric models
- ☑ To understand and have gained valuable insights into the functioning of financial markets
- ☑ To understand some of the practical issues in the forecasting of key financial market variables, such as asset prices, risk and dependence.

Assessment
This course is assessed by a three hour unseen written examination.

Students should consult the *Programme Regulations for degrees and diplomas in Economics, Management, Finance and the Social Sciences* that are reviewed annually. Notice is also given in the *Regulations* of any courses which are being phased out and students are advised to check course availability.

Syllabus

This is a description of the material to be examined, as published in the *Programme handbook.* On registration, students will receive a detailed subject guide which provides a framework for covering the topics in the syllabus and directions to the essential reading.

Building on concepts introduced in course FN3092 Corporate finance and course EC2020 Elements of econometrics, this course will introduce students to some widely-used models used to study and forecast financial markets and familiarize them with the properties of financial data. Such data often comes in the form of time series, and thus much of the course will use methods from time series analysis. The models to be covered include autoregressive and ARMA models, GARCH models for volatility forecasting, and models using high frequency (intra-daily) asset prices. Students completing this course will have seen and applied many of the latest models used in financial econometrics and will understand some of the key features (both positive and negative) of these models.

Topics addressed by this course are:

- Concepts and measures of risk
- Time-series analysis
- Empirical features of financial asset returns
- Market risk models
- Models of financial market correlations
- Forecast evaluation methods
- Risk management
- Asset allocation decisions
- Market microstructure and high frequency data

This course is quantitative by nature. It aims however to investigate practical issues in the forecasting of key financial market variables and makes use of a number of real-world data sets and examples.

Examiners' commentaries 2015

FN3142 Quantitative finance

Important note

This commentary reflects the examination and assessment arrangements for this course in the academic year 2014–15. The format and structure of the examination may change in future years, and any such changes will be publicised on the virtual learning environment (VLE).

Information about the subject guide and the Essential reading references

Unless otherwise stated, all cross-references will be to the latest version of the subject guide (2015). You should always attempt to use the most recent edition of any Essential reading textbook, even if the commentary and/or online reading list and/or subject guide refer to an earlier edition. If different editions of Essential reading are listed, please check the VLE for reading supplements – if none are available, please use the contents list and index of the new edition to find the relevant section.

General remarks for Zone A

This year the format of the examination has been identical to last year's, so that candidates had to answer three out of four questions. In general, candidates answered the questions well, as both the number of firsts and of failures was around 19%, consistently with past years. All questions could be comfortably addressed after a careful reading of the subject guide and complementary references. Question 2 was the most difficult for candidates who opted to answer it.

Comments on specific questions – Zone A

Candidates should answer **THREE** of the following **FOUR** questions. All questions carry equal marks.

Question 1

(a) (*25 points*) **Assume that daily returns evolve as**

$$\begin{aligned} r_{t+1} &= \mu + \epsilon_{t+1} \\ \epsilon_{t+1} &\sim N(0, \sigma^2_{t+1}) \end{aligned}$$

Derive the GARCH(2,2) conditional variance model from the following ARMA(2,2) model for the squared residual:

$$\epsilon^2_{t+1} = \omega + \gamma_1 \epsilon^2_t + \gamma_2 \epsilon^2_{t-1} + \eta_{t+1} + \lambda_1 \eta_t + \lambda_2 \eta_{t-1}$$

where η_{t+1} is a Gaussian white noise process.

(b) **(*25 points*) Under the assumption of covariance stationarity, derive the unconditional variance of the GARCH(2,2) model in (a).**

(c) **(*25 points*) Show that the k steps ahead forecast $E_t[\sigma^2_{t+K}]$ can be written recursively as follows:**

$$E_t[\sigma^2_{t+K} - V] = (\alpha_1 + \beta_1)(E_t[\sigma^2_{t+K-1}] - V) + (\alpha_2 + \beta_2)(E_t[\sigma^2_{t+K-2}] - V)$$

where V is the unconditional variance.

(d) **(*25 points*) Explain how you can estimate the constant mean-GARCH(2,2) model by maximum likelihood, and write also the expression for the conditional likelihood corresponding to a sequence of observations $(r_1, r_2, \ldots, r_T)$. Recall that the probability density function of a normally distributed random variable with mean μ and variance σ^2 is**

$$f(x; \mu, \sigma^2) = (2\pi\sigma^2)^{-1/2} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$$

Reading for this question

- Subject guide, Chapter 8.
- Christoffersen, P.F. *Elements of Financial Risk Management.* Chapter 2.
- Brooks, C. *Introductory Econometrics for Finance.* Chapter 8.
- Taylor, S.J. *Asset Price Dynamics, Volatility and Prediction.* Chapters 8 and 9.
- Tsay, R.S. *Analysis of Financial Time Series.* Chapter 7.

Approaching the question

(a) Consider the following ARMA(2,2) model for the squared residual:

$$\epsilon^2_{t+1} = \omega + \gamma_1\epsilon^2_t + \gamma_2\epsilon^2_{t-1} + \eta_{t+1} + \lambda_1\eta_t + \lambda_2\eta_{t-1}. \quad (.1)$$

Taking expectations on both sides of (.1):

$$\begin{aligned}
E_t[\epsilon^2_{t+1}] &= \omega + \gamma_1\epsilon^2_t + \gamma_2\epsilon^2_{t-1} + \lambda_1\eta_t + \lambda_2\eta_{t-2}, \quad \text{because } \eta_{t+1} \sim N(0,1) \\
\sigma^2_{t+1} &= \omega + \gamma_1\epsilon^2_t + \gamma_2\epsilon^2_{t-1} + \lambda_1\left(\epsilon^2_t - E_{t-1}[\epsilon^2_t]\right) + \lambda_2\left(\epsilon^2_{t-1} - E_{t-2}[\epsilon^2_{t-1}]\right), \\
&\quad \text{substituting for } \eta_t \text{ and } \eta_{t-1} \\
\sigma^2_{t+1} &= \omega + (\gamma_1 + \lambda_1)\epsilon^2_t + (\gamma_2 + \lambda_2)\epsilon^2_{t-1} - \lambda_1\sigma^2_t - \lambda_2\sigma^2_{t-1}.
\end{aligned}$$

After setting $\beta_i = -\lambda_i$ and $\alpha_i = \gamma_i + \lambda_i$, $i = 1, 2$, the GARCH(2,2) follows.

(b) Take unconditional expectations on both sides of:

$$\sigma^2_{t+1} = \omega + \alpha_1\epsilon^2_t + \alpha_2\epsilon^2_{t-1} + \beta_1\sigma^2_t + \beta_2\sigma^2_{t-1} \quad (.2)$$

and applying the law of iterated expectations, we have:

$$\begin{aligned}
E[\sigma^2_{t+1}] &= \omega + \beta_1 E[\sigma^2_t] + \beta_2 E[\sigma^2_{t-1}] + \alpha_1 E[E_{t-1}[\epsilon^2_t]] + \alpha_2 E[E_{t-2}[\epsilon^2_{t-1}]] \\
E[\sigma^2_{t+1}] &= \omega + [(\alpha_1 + \beta_1) + (\alpha_2 + \beta_2)]E[\sigma^2_{t+1}], \quad \text{because } E[\sigma^2_{t+1}] = E[\sigma^2_t] = E[\sigma^2_{t-1}]
\end{aligned}$$

so that, provided $(\alpha_1 + \beta_1) + (\alpha_2 + \beta_2) < 1$, we obtain:

$$E[\sigma^2_{t+1}] = \frac{\omega}{1 - (\alpha_1 + \beta_1) - (\alpha_2 + \beta_2)}.$$

(c) Let $V = \omega/[1 - (\alpha_1 + \beta_1) - (\alpha_2 + \beta_2)]$ denote the unconditional variance. Starting from expression (.2) for the conditional variance and substituting $\omega = V[1 - (\alpha_1 + \beta_1) - (\alpha_2 + \beta_2)]$, we obtain, at time $t + K$:

$$\sigma^2_{t+K} - V = \alpha_1(\epsilon^2_{t+K-1} - V) + \alpha_2(\epsilon^2_{t+K-2} - V) + \beta_1(\sigma^2_{t+K-1} - V) + \beta_2(\sigma^2_{t+K-2} - V).$$

Taking expectations conditional on time t information and applying the law of iterated expectations ($E_t[\epsilon^2_{t+K-1}] = E_t[E_{t+K-2}[\epsilon^2_{t+K-1}]] = E_t[\sigma^2_{t+K-1}]$), we obtain the desired representation.

(d) The answer is similar to what is illustrated in Section 8.6 of the subject guide for the GARCH(1,1). The constant mean, GARCH(2,2) model for returns implies that the one-step ahead conditional distribution of returns is:

$$r_{t+1}|\mathcal{F}_t \sim N\left(\mu, \omega + \alpha_1\epsilon_t^2 + \alpha_2\epsilon_{t-1}^2 + \beta_1\sigma_t^2 + \beta_2\sigma_{t-1}^2\right).$$

Therefore, the likelihood for the model is:

$$\begin{aligned} L(\theta|r_1, r_2, \ldots, t_T) &= \prod_{t=2}^{T} \frac{1}{\sqrt{2\pi\sigma_t^2}} \exp\left(-\frac{\epsilon_t^2}{2\sigma_t^2}\right) \\ \epsilon_t &= r_t - \mu \\ \sigma_t^2 &= \omega + \alpha_1\epsilon_{t-1}^2 + \alpha_2\epsilon_{t-2}^2 + \beta_1\sigma_{t-1}^2 + \beta_2\sigma_{t-2}^2 \end{aligned}$$

where $\theta = (\mu, \omega, \alpha_1, \alpha_1, \beta_1, \beta_2)$. At this point candidates should take logarithms and write the log likelihood explicitly. They should also note that assumptions are needed about the two initial values of residuals ϵ_0, ϵ_1 and conditional variances σ_0^2, σ_1^2. Typically, these values are set at the corresponding unconditional levels. Since explicit expressions for the maximum likelihood estimates of the parameters are not available, candidates should just note that the log likelihood is maximised numerically.

Question 2

(a) You hold two different corporate bonds (bond A and bond B), each with a face value of 1m \$. The issuing firms have a 2% probability of defaulting on the bonds, and both the default events and the recovery values are independent of each other. Without default, the notional value is repaid, while in case of default, the recovery value is uniformly distributed between 0 and the notional value.

(a1) (*20 points*) Find the 1% VaR for bond A or bond B and report your calculations. (HINT: remember that the CDF of a uniformly distributed random variable on $[a, b]$ is $F(x) = (x - a)/(b - a)$.)

(a2) (*60 points*) Taking into account that the PDF of the sum of two independent uniformly distributed random variables on $[a, b]$ is:

$$f(x) = \begin{cases} x & 2a < x < 2a + b \\ 2b - x & 2a + b < x < 2b \end{cases},$$

explain how you would find the 1% VaR for a portfolio combining the two bonds (A + B). Report your calculations without finding the actual value.

(b) (*20 points*) The α% expected shortfall is defined as the expected loss given that the loss exceeds the α% VaR. Find the 1% expected shortfall for bond A and report your calculations.

Reading for this question

- Subject guide, Chapter 13.
- Christoffersen, P.F. *Elements of Financial Risk Management.* Chapter 2 and 6.
- Tsay, R.S. *Analysis of Financial Time Series.* Chapter 7.

Approaching the question

(a1) According to the definition of VaR in the subject guide (Section 13.2), we need to identify the 1%− ile of the PDF of the bond A (or B) value. Since the event of default has 2% probability, and in this event the loss is uniformly distributed between 0 and $1m$, a bond value of $0.5m$ has a probability of 1%, so that $VaR = 0.5m$. Although this simple reasoning suffices to answer the question, one might want to alternatively apply the following formal reasoning:

$$\text{prob}(x \leq \overline{x}|\text{default}) \times (\text{prob. of default}) = 0.01$$

where x is the bond value. Substituting:

$$\frac{1m - \overline{x}}{1m} 0.02 = 0.01 \quad \rightarrow \overline{x} = 0.5m.$$

(a2) This question is slightly more involved, but candidates are asked to report only their reasoning, and the calculations leading to the answer, rather than the numerical answer.

It should as a preliminary be pointed out that the formula suggested for the PDF of the sum of two independent uniform random variables (on the interval $[a, b]$) is not correct. The correct one is:

$$f(x) = \begin{cases} \frac{x-2a}{2(b-a)} & 2a < x < a+b \\ \frac{-x+2b}{2(b-a)} & a+b < x < 2b \end{cases}.$$

While an apology is due, it must be pointed out that this typo is inconsequential, and cannot affect your ability to answer for two reasons:

- If you substitute to a and b its values, $a = 0$ and $b = 1$ (units here are millions \$), the formula coincides with the suggested one.
- This question explicitly asks you to avoid performing calculations, but to report those leading to the answer, where you can use $f(x)$ for the density function, or report its form suggested in the text, which of course will be considered correct for the purpose of this exercise.

Now the solution. There are four mutually exclusive cases concerning bonds' default: 1) both A and B default, which occurs with probability 2% × 2%, because the events are independent. In this case the portfolio $A + B$ has a value which is distributed on $[0, 2m]$ with density $f(x)$. 2) A defaults but B doesn't, which occurs with probability 2% × 98%. In this case the portfolio $A + B$ has a value which is uniformly distributed on $[1m, 2m]$. 3) B defaults but A doesn't. The probability and the portfolio distribution are as in case 2. 4) No bond defaults, which occurs with probability 98% × 98%, and the portfolio value is $2m$. Intuitively, we need to consider the first three cases, because 98% × 98% < 99%, so the 1%-ile of the portfolio value is smaller than $2m$. In order to find the 1%-ile, we need to solve:

$$\text{prob}(x \leq \overline{x}|2 \text{ defaults}) \times (\text{prob. of 2 defaults}) + 2 \times \text{prob}(x \leq \overline{x}|1 \text{ default}) \times (\text{prob. of 1 default}) = 0.01.$$

At this stage the answer is considered complete, although candidates might go one step further and substitute, to obtain:

$$0.02 \times 0.02 \times \int_0^{\overline{x}} f(x)dx + 2 \times 0.98 \times 0.02 \times \frac{2m - \overline{x}}{2m - 1m} \times \mathbf{1}(\overline{x} > 1m) = 0.01.$$

where $\mathbf{1}()$ denotes the indicator function, which is 1 if $\overline{x} > 1m$ and 0 otherwise. If $\overline{x}$ solves the last equation, then $2m - \overline{x}$ is the 1% VaR of the portfolio including both bonds.

(b) Since the recovery value is uniformly distributed in case of default, the expected shortfall is just the average between the VaR ($0.5m$) and the maximum loss ($1m$), thus $0.75m$. This argument is sufficient to answer. A bit more formally, one might argue that (since the density of the uniform over $[a, b]$ is $f(x) = x/(b - a)$):

$$E[x|x < \overline{x}] = \frac{0.02 \times \int_0^{\overline{x}} \frac{x}{1m} dx}{0.02 \times \frac{\overline{x}}{1m}} = \frac{\overline{x}}{2} = 0.25m$$

where $\overline{x} = 0.5m$ is the 1%−ile of the portfolio value distribution. So the expected shortfall is $1m - 0.25m = 0.75m$.

Question 3

(a) **(*10 points*) What are the two main problems that a researcher encounters in multivariate volatility modeling?**

(b) **(*45 points*) Describe in detail the Constant Conditional Correlations - GARCH(1,1) model of multivariate volatility, making use of matrix notation. Discuss benefits and drawbacks of this model.**

(c) **(*45 points*) Suggest a two-stage estimation method for this model based on Maximum Likelihood. In particular, write explicitly one of the likelihood functions that you maximize at the first stage, for an individual returns series $(r_1^i, r_2^i, \ldots, r_T^i)$, i being one of the N assets. For simplicity, assume that the one period mean of returns is the constant μ^i. Recall that the probability density function of a normally distributed random variable with mean μ and variance σ^2 is**

$$f(x; \mu, \sigma^2) = (2\pi\sigma^2)^{-1/2} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$$

Reading for this question

- Subject guide, Chapter 10.
- Christoffersen, P.F. *Elements of Financial Risk Management.* Chapter 7.
- Tsay, R.S. *Elements of Financial Risk Management.* Chapter 10.

Approaching the question

(a) This question refers to Chapter 10 of the subject guide, where extensions to asset return volatility modelling are discussed. The two main challenges that a researcher encounters with multivariate volatilities are: 1) parameter proliferation, due to the need of modelling the dynamics of the $N \times (N+1)/2$ free components of a variance-covariance matrix of returns. This problem can lead to a curse of dimensionality in the estimation procedure and indeed to overfitting issues. 2) Guaranteeing that the variance-covariance matrix is positive definite. Modelling the dynamics of variances and covariances separately would not likely meet this requirement, thus an integrated approach is needed.

(b) Here candidates are asked to write down Bollerslev's CCC model. First, report the dynamics of returns in a multivariate model with time-varying volatilities:

$$r_{t+1} = \mu_{t+1} + \epsilon_{t+1}$$

$$\epsilon_{t+1} | \mathcal{F}_t \sim N(0, H_{t+1})$$

where r, μ, ϵ are $N \times 1$ vectors and H is the $N \times N$ variance-covariance matrix of returns. The CCC models the return variances h_{t+1}^{ii} (diagonal entries of H) as GARCH(1,1), while the covariances are:

$$h_{t+1}^{ij} = \overline{\rho}_{ij} \sqrt{h_{t+1}^{ii} h_{t+1}^{jj}}$$

where $\overline{\rho}_{ij}$ is a constant correlation, obtained as unconditional correlation of standardized residuals $\epsilon_{t+1}^i / \sqrt{h_{t+1}^{ii}}$ and $\epsilon_{t+1}^j / \sqrt{h_{t+1}^{jj}}$. The advantages of the model are its parsimony and the guarantee of positive-definiteness of the covariance matrix, while its main drawback is the lack of empirical evidence supporting the assumption that asset return correlations are constant.

(c) Though not explicitly mentioned in the subject guide, the description of the CCC model at point (b) naturally suggests a two stage estimation procedure: 1) first estimate the GARCH(1,1) model for the conditional variances on the univariate return series. At this point candidates are asked to describe explicitly the maximum likelihood procedure and write the likelihood function. See Section 4.5 of the subject guide or the answer to Question 1.d of this paper. 2) Build the standardized residuals $\widehat{\epsilon}^{i}_{t+1}/\sqrt{\widehat{h}^{ii}_{t+1}}$ for all assets, from the GARCH estimates at stage 1, and compute their sample correlation to obtain an estimate of $\overline{\rho}_{ij}$.

Question 4

Your aim is to compare two forecasts Y_t^1 and Y_t^2 of some financial variable Y_t. To this aim, consider some loss functions $l(Y_t^i, Y_t)$, $i = 1, 2$.

(a) ***(50 points)* Describe the Diebold-Mariano test to compare the forecasts in case the data are serially independent. Provide a numerical example where the loss function is the squared forecast error: $l(Y_t^i, Y_t) = (Y_t^i - Y_t)^2$, $i = 1, 2$.**

(b) ***(50 points)* Explain how you modify the procedure in (a) when the data is serially dependent. Describe the Newey-West estimator that has been proposed for this case.**

Reading for this question

- Subject guide, Chapter 12.
- Diebold, F.X. *Elements of Forecasting.* (Thomson South-Western, Canada, 2006) fourth edition [ISBN 9780324323597]. Chapter 12.

Approaching the question

(a) This is a descriptive question, where candidates are asked to concisely illustrate the procedure of Diebold-Mariano to compare two forecasts of a financial variable. Since there is no specific problem to be solved, I refer to Chapter 12.2 of the subject guide for a detailed treatment. Candidates can provide a simple numerical example by considering the squared error as loss functions, $l(Y_t^i, Y_t) = (Y_t^i - Y_t)^2$, and the time series of the differences:

$$d_t = (Y_t^1 - Y_t)^2 - (Y_t^2 - Y_t)^2$$

and then assigning numerical values to the Mean Squared Errors of the two forecasts, the sample variance of the squared errors and the sample size, to obtain a T-statistic for the asymptotic T-test. Candidates should be reminded that in the serially uncorrelated case, the denominator of the T-statistic is simply the sample variance divided by the square root of sample size.

(b) When data are serially correlated, the covariance between the differences d_t at different points in time does not vanish. Thus the denominator of the T-statistic ($V\left(\sum \frac{1}{T} d_t\right)$) needs to include these additional terms. The Newey-West estimator to be applied in this case is fully described in Chapter 12.2.