

Quantitative Finance

Lecturer: Daniil Esaulov, Sofya Budanova (part of Time Series Analysis)

Abstract

The course is elective and provides coverage of important topics in modern Quantitative Finance and Risk Management at the advanced undergraduate level. It is intended for the 4th-year undergraduate students of the International College of Economics and Finance, High School of Economics, Moscow. Particular attention is given to the topics such as the Efficient Market Hypothesis, financial markets microstructure and types of arbitrage, general principles of modelling the price dynamics of financial assets, market risk and other types of financial risks, Value-at-Risk (VaR) approach and applications, modelling of extreme market events, VaR analysis for financial derivatives using the Kolmogorov equations framework, modelling of periodic and quasiperiodic trends in time series in connection with technical analysis, and the foundations of high frequency arbitrage trading. The topics covered in this course will enable the students to develop the theoretical knowledge and practical skills required for successful working with multiple types of risks in modern financial markets, both Russian and international. The course is taught in English.

Prerequisites for the course are Elements of Econometrics and Microeconomics. Good command of methods of calculus, general probability theory and mathematical statistics are also required for the course.

Learning Objectives and Outcomes:

The goal of this course is to give students insights in the functioning of financial markets, understanding of measuring and forecasting financial risks. This course is aimed at giving students instruments required in order to analyze issues in asset pricing and market finance. After the course students should be familiar with recent empirical findings based on financial econometric models, have a good command of basic econometric techniques and understand practical issues in the forecasting of key financial market variables.

The student should be able to apply professional knowledge and skills acquired while studying the course in practical areas, including academic research, work in financial institutions, industry, state governance.

Methods of Instruction

The following methods and forms of study are used in the course:

- lectures (2 hours per week);
- practical classes, also in computer class (2 hours per week);
- self-study in computer class (doing home assignments using Excel, EViews, R, work with economic data, appliances in Internet);
- self-study with literature.

Special Equipment and Software Support

Laptop, projector, Internet connection

MS Excel, EViews, R

Grading System and Knowledge Assessment

- Midterm (in the 1st semester),
- December exam (at the end of the 1st semester),
- Home assignments,
- Quizzes,
- Final exam (at the end of the spring term).

Grade determination:

- The first term and final grades are derived from weighted averages of the results of assessment which are given the assigned weights. The formula to calculate these grades is the following:

$G_{\text{resultant}} = \omega_1 \cdot G_1 + \omega_2 \cdot G_2 + \omega_3 \cdot G_3$, where G_1, G_2, G_3 – grades for the types of assessment and $\omega_1, \omega_2, \omega_3$ – corresponding weights. • The first term grade will be determined by (weights are indicated in brackets): – December exam (70%), – Midterm test (20%), – Average grade for the home assignments (10%). • The final course grade is determined by (weights are indicated in brackets): – First semester grade (30%), – Average grade for the home assignments (10%), – Final exam (in April) (60%).

- Final grade is rounded up to the next larger integer number.
- Final grade determination in case of retakes or student's non-participation in different types of assessment and procedure for conducting retakes are regulated by the "guidelines for final grade determination for a course of the HSE and University of London Double Degree Bachelor Programme in Economics at the International College of Economics and Finance, the National Research University "Higher School of Economics".

Sample materials for knowledge assessment are available in ICEF Information system at <https://icef-info.hse.ru>.

Required reading:

- Patton A., Quantitative Finance, UoL Study Guide, 2015
- Enders W., Applied Econometric Time Series, 2nd ed., John Wiley & Sons, Inc., 2004
- Christoffersen P.F., Elements of Financial Risk Management, Academic Press, London, 2011.
- Diebold F.X. Elements of Forecasting, Thomson South-Western, Canada, 2006, fourth edition.

Optional reading:

- Ruey S. Tsay, Analysis of Financial Time Series, 2002. (RT) • Wilmott P., Quantitative Finance (selected chapters). 2nd ed. Wiley, 2006.
- McNeil A. J. et al, Quantitative Risk Management, Princeton University Press, 2005.
- Franke J., Haerdle W. K., Hafner C. M., Statistics of Financial Markets, Springer, 2004. • Hamilton, J., Time Series Analysis, Princeton University Press, Princeton, 1994.
- Brooks Ch., Introductory econometrics for finance, Cambridge University Press, 2002.

Course plan:

Note: topics 1-4 are covered by Sofya Budanova in the course Time Series and Panel Data Analysis. Check for further sources the corresponding syllabus. The midterm of the course counts as midterm for QF.

1. Basic time series concepts Many problems in quantitative finance involve the study of financial data. Such data most often comes in the form of 'time series', which is a sequence of random variables that are ordered through time. Before moving on to financial applications, we must first cover some fundamental topics in time series analysis, such as autocorrelation, white noise processes and ARMA processes. This topic is the most theoretical one in the course, and it may not appear too related to finance, but it lays the foundations for the (more interesting) topics we will cover later.

2. Testing for stationarity Testing for stationarity: graphical techniques and the formal unit root tests. (Augmented) DickeyFuller tests. Other tests of nonstationarity.

3. Modeling asset return volatility: introduction Risk plays a central role in financial decision making, and it is thus no surprise that a great deal of effort has been devoted to the study of the volatility of asset returns. This effort has paid large dividends: volatility modeling and forecasting methods have been shown to be very useful in many economic applications. In this topic we will cover some of the most widely-used models for modeling volatility, discuss the estimation of these models, and methods of testing for volatility predictability.

4. Vector Autoregression Dynamic interdependencies of financial variables can be uncovered using VAR analysis. Impulse response function. Granger causality.

— Midterm —

5. Empirical features of financial data Introducing some concepts of defining and describing financial data, we also discuss main stylized facts of returns.

6. Modeling asset return volatility: extensions In this chapter we discuss extensions of the basic ARCH/GARCH class of models, both univariate and multivariate. Univariate extensions have been proposed to capture more detailed features of asset return volatility, such as the so-called "leverage effect". Multivariate extensions of the GARCH model are used to assist with financial decisions that involve more than one risky asset, such as portfolio decisions or risk management.

7. Evaluating forecasts of risks and returns There are often many competing statistical models available for use in the forecasting of a particular financial variable. There are also many commercially available forecasts, issued by brokers or mutual funds. How do we determine whether the forecast is good or not? How do we determine which model or forecaster is best? These two questions relate to forecast evaluation and comparison. A third question that arises is whether we can take a collection of forecasts and combine them somehow to get an even better forecast. We will cover methods for answering these questions in this topic.

— Semester break—

8. The efficient market hypothesis and market predictability Much of modern quantitative finance relates to methods and models for predicting aspects of asset returns, and yet the classical theory of efficient markets may appear to suggest that asset returns should be completely unpredictable. In this topic we relate the concept of efficient markets, defined in various ways, to the evidence of predictability of financial variables and reconcile the empirical evidence for asset return predictability with the concept of an efficient market.

9. Risk management and Value-at-Risk: models Measuring and managing the exposure to risk generated by a trading desk, a structured product, or a traditional portfolio is one of the most important and interesting parts of quantitative finance. Modern risk management focuses heavily on a measure of risk known as "Value-at-Risk", or VaR. This is partly due to some advantages of this measure over variance,

and partly due to regulation (the Basel Accords are based on VaR as a measure of risk). In this topic we will introduce VaR formally and discuss some of the most common models for measuring VaR.

10. Risk management and Value-at-Risk: backtesting An important part of managing risk is testing how well your risk models are performing, a task known in the risk management literature as backtesting. Such tests can also be useful for indicating ways to improve risk models. This topic will cover some methods for backtesting VaR models.

11. Modeling high frequency financial data Traditionally, empirical studies in finance employed data at the daily and monthly frequencies. Many models and methods have been developed for the study of such data. Recently, high frequency (intraday) has become available to researchers, and empirical market microstructure is now an established sub-field within finance. Many of the methods developed for lower frequency data are applicable to high frequency data, but there are a few places where differences exist, and we will study two of these in this topic. How one treats the massive amounts of high frequency data available should depend on the problem the researcher wishes to address. In many cases, the question can be addressed by aggregating the 'tick' data up to a certain frequency and then analyze the sequence of aggregated returns. Doing so makes the data evenly spaced, and thus more similar to well-studied low frequency data. Many questions, however, are best addressed using tick data, meaning that we must find ways of dealing with the irregularly-spaced observations. Another problem that arises in certain analyses of high frequency data is seasonality. Seasonality is a well-studied problem in macro- and micro-econometrics, but is not generally a concern for financial econometricians. Intra-daily patterns (called 'diurnality' rather than 'seasonality') in certain measures are significant and must be dealt with. Three places where diurnality in high frequency returns has been found to be prominent are in the conditional variance, in bid-ask spreads and in trade durations.