QUANTITATIVE ANALYSIS

PART I EXAM WEIGHT | 20% (QA)

This area focuses on basic probability and statistics, regression and time series analysis, and various quantitative techniques useful in risk management. The broad knowledge points covered in Quantitative Analysis include the following:

- Discrete and continuous probability distributions
- Estimating the parameters of distributions
- Population and sample statistics
- Bayesian analysis
- Statistical inference and hypothesis testing
- Measures of correlation
- Linear regression with single and multiple regressors
- Time series analysis and forecasting
- Simulation methods
- Machine learning

The readings that you should focus on for this section and the specific learning objectives to achieve with each reading are:

Global Association of Risk Professionals. *Quantitative Analysis*. New York, NY: Pearson, 2022.

Chapter 1: Fundamentals of Probability [QA–1]

After completing this reading, you should be able to:

- Describe an event and an event space.
- Describe independent events and mutually exclusive events.
- Explain the difference between independent events and conditionally independent events.
- Calculate the probability of an event for a discrete probability function.
- Define and calculate a conditional probability.
- Distinguish between conditional and unconditional probabilities.
- Explain and apply Bayes' rule.

Chapter 2: Random Variables [QA–2]

After completing this reading, you should be able to:

- Describe and distinguish a probability mass function from a cumulative distribution function and explain the relationship between these two.
- Understand and apply the concept of a mathematical expectation of a random variable.
- Describe the four common population moments.
- Explain the differences between a probability mass function and a probability density function.

- Characterize the quantile function and quantile-based estimators.
- Explain the effect of a linear transformation of a random variable on the mean, variance, standard deviation, skewness, kurtosis, median, and interquartile range.

Chapter 3: Common Univariate Random Variables [QA–3]

After completing this reading, you should be able to:

- Distinguish the key properties and identify the common occurrences of the following distributions: uniform distribution, Bernoulli distribution, binomial distribution, Poisson distribution, normal distribution, lognormal distribution, Chi-squared distribution, Student's t- and F-distributions.
- Describe a mixture distribution and explain the creation and characteristics of mixture distributions.

Chapter 4: Multivariate Random Variables [QA–4]

After completing this reading, you should be able to:

- Explain how a probability matrix can be used to express a probability mass function.
- Compute the marginal and conditional distributions of a discrete bivariate random variable.
- Explain how the expectation of a function is computed for a bivariate discrete random variable.
- Define covariance and explain what it measures.
- Explain the relationship between the covariance and correlation of two random variables, and how these are related to the independence of the two variables.
- Explain the effects of applying linear transformations on the covariance and correlation between two random variables.
- Compute the variance of a weighted sum of two random variables.
- Compute the conditional expectation of a component of a bivariate random variable.
- Describe the features of an independent and identically distributed (iid) sequence of random variables.
- Explain how the iid property is helpful in computing the mean and variance of a sum of iid random variables.

Chapter 5: Sample Moments [QA–5]

After completing this reading, you should be able to:

- Estimate the mean, variance, and standard deviation using sample data.
- Explain the difference between a population moment and a sample moment.
- Distinguish between an estimator and an estimate.
- Describe the bias of an estimator and explain what the bias measures.
- Explain what is meant by the statement that the mean estimator is BLUE.
- Describe the consistency of an estimator and explain the usefulness of this concept.
- Explain how the Law of Large Numbers (LLN) and Central Limit Theorem (CLT) apply to the sample mean.
- Estimate and interpret the skewness and kurtosis of a random variable.
- Use sample data to estimate quantiles, including the median.
- Estimate the mean of two variables and apply the CLT.
- Estimate the covariance and correlation between two random variables.
- Explain how coskewness and cokurtosis are related to skewness and kurtosis.

Chapter 6: Hypothesis Testing [QA–6]

After completing this reading, you should be able to:

- Construct an appropriate null hypothesis and alternative hypothesis and distinguish between the two.
- Differentiate between a one-sided and a two-sided test and identify when to use each test.
- Explain the difference between Type I and Type II errors and how these relate to the size and power of a test.

- Understand how a hypothesis test and a confidence interval are related.
- Explain what the p-value of a hypothesis test measures.
- Construct and apply confidence intervals for one-sided and two-sided hypothesis tests and interpret the results of hypothesis tests with a specific confidence level.
- Identify the steps to test a hypothesis about the difference between two population means.
- Explain the problem of multiple testing and how it can lead to biased results.

Chapter 7: Linear Regression [QA–7]

After completing this reading, you should be able to:

- Describe the models which can be estimated using linear regression and differentiate them from those which cannot.
- Interpret the results of an ordinary least squares (OLS) regression with a single explanatory variable.
- Describe the key assumptions of OLS parameter estimation.
- Characterize the properties of OLS estimators and their sampling distributions.
- Construct, apply, and interpret hypothesis tests and confidence intervals for a single regression coefficient in a regression.
- Explain the steps needed to perform a hypothesis test in a linear regression.
- Describe the relationship among a t-statistic, its p-value, and a confidence interval.
- Estimate the correlation coefficient from the R2 measure obtained in linear regressions with a single explanatory variable.

Chapter 8: Regression with Multiple Explanatory Variables [QA–8]

After completing this reading, you should be able to:

- Distinguish between the relative assumptions of single and multiple regression.
- Interpret regression coefficients in a multiple regression.
- Interpret goodness-of-fit measures for single and multiple regressions, including R2 and adjusted-R2.
- Construct, apply, and interpret joint hypothesis tests and confidence intervals for multiple coefficients in a regression.
- Calculate the regression R2 using the three components of the decomposed variation of the dependent variable data: the explained sum of squares, the total sum of squares, and the residual sum of squares.

Chapter 9: Regression Diagnostics [QA–9]

After completing this reading, you should be able to:

- Explain how to test whether a regression is affected by heteroskedasticity.
- Describe approaches to using heteroskedastic data.
- Characterize multicollinearity and its consequences, as well as distinguish between multicollinearity and perfect collinearity.
- Describe the consequences of excluding a relevant explanatory variable from a model and contrast those with the consequences of including an irrelevant regressor.
- Explain two model selection procedures and how these relate to the bias-variance trade-off.
- Describe the various methods of visualizing residuals and their relative strengths.
- Describe methods for identifying outliers and their impact.
- Determine the conditions under which OLS is the best linear unbiased estimator.

Chapter 10: Stationary Time Series [QA–10]

After completing this reading, you should be able to:

- Describe the requirements for a series to be covariance stationary.
- Define the autocovariance function and the autocorrelation function.
- Define white noise and describe independent white noise and normal (Gaussian) white noise.

- Define and describe the properties of autoregressive (AR) processes.
- Define and describe the properties of moving average (MA) processes.
- Explain how a lag operator works.
- Explain mean reversion and calculate a mean-reverting level.
- Define and describe the properties of autoregressive moving average (ARMA) processes.
- Describe the application of AR, MA, and ARMA processes.
- Describe sample autocorrelation and partial autocorrelation.
- Describe the Box-Pierce Q statistic and the Ljung-Box Q statistic.
- Explain how forecasts are generated from ARMA models.
- Describe the role of mean reversion in long-horizon forecasts.
- Explain how seasonality is modeled in a covariance-stationary ARMA.

Chapter 11: Non-Stationary Time Series [QA–11]

After completing this reading, you should be able to:

- Describe linear and nonlinear time trends.
- Explain how to use regression analysis to model seasonality.
- Describe a random walk and a unit root.
- Explain the challenges of modeling time series containing unit roots.
- Describe how to test if a time series contains a unit root.
- Explain how to construct an h-step-ahead point forecast for a time series with seasonality.
- Calculate the estimated trend value and form an interval forecast for a time series.

Chapter 12: Measuring Returns, Volatility, and Correlation [QA–12]

After completing this reading, you should be able to:

- Calculate, distinguish, and convert between simple and continuously compounded returns.
- Define and distinguish between volatility, variance rate, and implied volatility.
- Describe how the first two moments may be insufficient to describe non-normal distributions.
- Explain how the Jarque-Bera test is used to determine whether returns are normally distributed.
- Describe the power law and its use for non-normal distributions.
- Define correlation and covariance and differentiate between correlation and dependence.
- Describe properties of correlations between normally distributed variables when using a one-factor model.
- Compare and contrast the different measures of correlation used to assess dependence.

Chapter 13: Simulation and Bootstrapping [QA–13]

After completing this reading, you should be able to:

- Describe the basic steps to conduct a Monte Carlo simulation.
- Describe ways to reduce Monte Carlo sampling error.
- Explain the use of antithetic and control variates in reducing Monte Carlo sampling error.
- Describe the bootstrapping method and its advantage over Monte Carlo simulation.
- Describe pseudo-random number generation.
- Describe situations where the bootstrapping method is ineffective.
- Describe the disadvantages of the simulation approach to financial problem solving.

Chapter 14: Machine Learning Methods [QA–14]

After completing this reading, you should be able to:

- Discuss the philosophical and practical differences between machine learning techniques and classical econometrics.
- Compare and apply the two methods utilized for rescaling variables in data preparation.
- Explain the differences among the training, validation, and test data sub-samples, and how each is used.
- Understand the differences between and consequences of underfitting and overfitting, andpropose potential remedies for each.
- Use principal components analysis to reduce the dimensionality of a set of features.
- Describe how the K-means algorithm separates a sample into clusters.
- Describe natural language processing and how it is used.
- Differentiate among unsupervised, supervised, and reinforcement learning models.
- Explain how reinforcement learning operates and how it is used in decision-making.

Chapter 15: Machine Learning and Prediction [QA–15]

After completing this reading, you should be able to:

- Explain the role of linear regression and logistic regression in prediction.
- Evaluate the predictive performance of logistic regression models.
- Understand how to encode categorical variables.
- Discuss why regularization is useful, and distinguish between the ridge regression and LASSO approaches.
- Show how a decision tree is constructed and interpreted.
- Describe how ensembles of learners are built.
- Explain the intuition and processes behind the K nearest neighbors and support vector machine methods for classification.
- Understand how neural networks are constructed and how their weights are determined.
- Compare the logistic regression and neural network classification approaches using a confusion matrix.