2024 FRM®
Exam Prep

SchweserNotes™
Market Risk Measurement and Management

PART II BOOK 1

KAPLAN SCHWESER

Kaplan Schweser's Path to Success

FRM® Exam Part II

FRM®

Welcome

As the head of Advanced Designations at Kaplan Schweser, I am pleased to have the opportunity to help you prepare for the FRM® exam. Kaplan Schweser has decades of experience in delivering the most effective FRM exam prep products in the market and I know you will find them to be invaluable in your studies.

Our products are designed to be an integrated study solution across print and digital media to provide you with the best learning experience, whether you are studying with a physical book, online, or on your mobile device.

Our core product, the SchweserNotes™, addresses all Topic Areas, Readings, and Learning Objectives in the FRM curriculum. Each reading in the SchweserNotes has been broken into smaller, bite-sized modules with Module Quizzes interspersed throughout to help you continually assess your comprehension. Topic Quizzes and Checkpoint Exams appear online to help you gauge your knowledge of the material before you move on to the next section.

All purchasers of the SchweserNotes receive online access to the Kaplan Schweser online platform (our learning management system or LMS) at www.Schweser.com. In the LMS, you will see a dashboard that tracks your overall progress and performance as well as an Activity Feed, which provides structure and organization to the tasks required to prepare for the FRM exam. You also have access to the online versions of the SchweserNotes and Module Quizzes. Look for the icons indicating where Module Quizzes are available online. I strongly encourage you to enter your Module Quiz answers online and use the dashboard to track your progress and stay motivated.

Again, thank you for trusting Kaplan Schweser with your FRM exam preparation. We're here to help you throughout your journey to become a certified Financial Risk Manager.

Regards,

Derek Burkett

Derek Burkett, CFA, FRM, CAIA
Vice President (Advanced Designations)

Book 1: Market Risk Measurement and Management

SchweserNotes™ 2024

FRM Part II

KAPLAN SCHWESER

SCHWESERNOTES™ 2024 FRM® PART II BOOK 1: MARKET RISK MEASUREMENT AND MANAGEMENT

Published in 2024 by Kaplan, Inc.

ISBN: 978-1-0788-4249-5

WELCOME TO THE 2024 SCHWESERNOTES™

Thank you for trusting Kaplan Schweser to help you reach your career and educational goals. We are very pleased to be able to help you prepare for the FRM Part II exam. In this introduction, I want to explain the resources included with the SchweserNotes, suggest how you can best use Kaplan Schweser materials to prepare for the exam, and direct you toward other educational resources you will find helpful as you study for the exam.

SchweserNotes™

The SchweserNotes consist of five volumes that include complete coverage of all FRM assigned readings and learning objectives (LOs), as well as module quizzes (multiple-choice questions for every reading) to help you master the material and check your retention of key concepts.

Practice Questions

To retain the material, it is important to quiz yourself often. We offer an online version of the SchweserPro™ QBank, which contains hundreds of Part II practice questions and explanations. We also offer Topic Quizzes and Checkpoint Exams online to further help you retain and apply what you have learned.

Mock Exams

Schweser offers four full 4-hour, 80-question practice exams. These online exams are important tools for gaining the speed and skills you will need to pass the exam. The Mock Exams contain answers with full explanations for self-grading and evaluation.

OnDemand Class

Our OnDemand Class provides comprehensive online instruction of every reading in the FRM curriculum. This video lecture series brings the personal attention of a classroom into your home or office with over 50 hours of instruction. The class offers in-depth coverage of difficult concepts as well as a discussion of sample exam questions. All videos are available for viewing at any time throughout the season. Candidates enrolled in the OnDemand Class also have the ability to email questions to the instructor at any time.

Late-Season Review

Late-season review and exam practice can make all the difference. Our OnDemand Review Package helps you evaluate your exam readiness with products specifically designed for late-season studying. This study package includes the OnDemand Review

(20-hour archived online workshop covering essential curriculum topics) and Schweser's Secret Sauce® (concise summary of the FRM curriculum).

Part II Exam Weightings

When preparing for the exam, be familiar with the weightings assigned to each topic area within the curriculum. The Part II exam weights and questions are as follows:

Book	Topic Area	Exam Weight	Exam Questions
1	Market Risk Measurement and Management	20%	16
2	Credit Risk Measurement and Management	20%	16
3	Operational Risk and Resiliency	20%	16
4	Liquidity and Treasury Risk Measurement and Management	15%	12
5	Risk Management and Investment Management	15%	12
5	Current Issues in Financial Markets	10%	8

How to Succeed

The FRM Part II exam is a formidable challenge (covering 103 assigned readings and almost 600 learning objectives), so you must devote considerable time and effort to be properly prepared. There are no shortcuts! You must learn the material, know the terminology and techniques, understand the concepts, and be able to answer 80 multiple-choice questions quickly and (at least 70%) correctly. A good estimate of the study time required is 300 hours on average, but some candidates will need more or less time, depending on their individual backgrounds and experience.

Expect the Global Association of Risk Professionals (GARP) to test your knowledge in a way that will reveal how well you know the Part II curriculum. You should begin studying early and stick to your study plan. You should first read the SchweserNotes and complete the practice questions for each reading. After completing each book, you should answer the provided online topic quiz questions to understand how concepts may be tested on the exam.

It is recommended that you finish your initial study of the entire curriculum at least two weeks (earlier if possible) prior to your exam date to allow sufficient time for practice and targeted review. During this period, you should take all of your Mock Exams. This final review period is when you will get a clear indication of how effective your study efforts have been and which readings require significant additional review. Answering exam-like questions across all readings and working on your exam time management skills will be important determinants of your success on exam day.

Best regards,

Eric Smith

Eric Smith, CFA, FRM, FDP
Director, Advanced Designations
Kaplan Schweser

CONTENTS

Readings and Learning Objectives

STUDY SESSION 1

1. **Estimating Market Risk Measures: An Introduction and Overview**
   Kevin Dowd, *Measuring Market Risk, 2nd Edition* (West Sussex, UK: John Wiley & Sons, 2005). Chapter 3.
   After completing this reading, you should be able to:
   a. estimate VaR using a historical simulation approach.
   b. estimate VaR using a parametric approach for both normal and lognormal return distributions.
   c. estimate the expected shortfall given profit and loss (P&L) or return data.
   d. estimate risk measures by estimating quantiles.
   e. evaluate estimators of risk measures by estimating their standard errors.
   f. interpret quantile-quantile (QQ) plots to identify the characteristics of a distribution.
2. **Non-Parametric Approaches**
   Kevin Dowd, *Measuring Market Risk, 2nd Edition* (West Sussex, UK: John Wiley & Sons, 2005). Chapter 4.
   After completing this reading, you should be able to:
   a. apply the bootstrap historical simulation approach to estimate coherent risk measures.
   b. describe historical simulation using non-parametric density estimation.
   c. compare and contrast the age-weighted, the volatility-weighted, the correlation-weighted, and the filtered historical simulation approaches.
   d. identify advantages and disadvantages of non-parametric estimation methods.
3. **Parametric Approaches (II): Extreme Value**
   Kevin Dowd, *Measuring Market Risk, 2nd Edition* (West Sussex, UK: John Wiley & Sons, 2005). Chapter 7.
   After completing this reading, you should be able to:
   a. explain the importance and challenges of extreme values in risk management.
   b. describe extreme value theory (EVT) and its use in risk management.
   c. describe the peaks-over-threshold (POT) approach.
   d. compare and contrast the generalized extreme value (GEV) and POT approaches to estimating extreme risks.
   e. discuss the application of the generalized Pareto (GP) distribution in the POT approach.
   f. explain the multivariate EVT for risk management.
4. **Backtesting VaR**
   Philippe Jorion, *Value at Risk: The New Benchmark for Managing Financial Risk, 3rd Edition* (New York, NY: McGraw Hill, 2007). Chapter 6.
   After completing this reading, you should be able to:
   a. describe backtesting and exceptions and explain the importance of backtesting VaR models.
   b. explain the significant difficulties in backtesting a VaR model.
   c. verify a model based on exceptions or failure rates.
   d. identify and describe Type I and Type II errors in the context of a backtesting process.
   e. explain the need to consider conditional coverage in the backtesting framework.
   f. describe the Basel rules for backtesting.
5. **VaR Mapping**
   Philippe Jorion, *Value at Risk: The New Benchmark for Managing Financial Risk, 3rd Edition* (New York, NY: McGraw Hill, 2007). Chapter 11.
   After completing this reading, you should be able to:
   a. explain the principles underlying VaR mapping and describe the mapping process.
   b. explain and demonstrate how the mapping process captures general and specific risks.
   c. differentiate among the three methods for mapping portfolios of fixed-income securities.
   d. summarize how to map a fixed-income portfolio into positions of standard instruments.

e. describe how mapping of risk factors can support stress testing.
f. explain how VaR can be computed and used relative to a performance benchmark.
g. describe the method of mapping forwards, forward rate agreements, interest rate swaps, and options.

6. Messages From the Academic Literature on Risk Measurement for the Trading Book

"Messages from the Academic Literature on Risk Measurement for the Trading Book," Basel Committee on Banking Supervision, Working Paper No. 19, Jan 2011.

After completing this reading, you should be able to:

a. explain the following lessons on VaR implementation: time horizon over which VaR is estimated, the recognition of time-varying volatility in VaR risk factors, and VaR backtesting.
b. describe exogenous and endogenous liquidity risk and explain how they might be integrated into VaR models.
c. compare VaR, expected shortfall, and other relevant risk measures.
d. compare unified and compartmentalized risk measurement.
e. compare the results of research on top-down and bottom-up risk aggregation methods.
f. describe the relationship between leverage, market value of asset, and VaR within an active balance sheet management framework.

STUDY SESSION 2

7. Correlation Basics: Definitions, Applications, and Terminology

Gunter Meissner, *Correlation Risk Modeling and Management, 2nd Edition* (Risk Books, 2019). Chapter 1.

After completing this reading, you should be able to:

a. describe financial correlation risk and the areas in which it appears in finance.
b. explain how correlation contributed to the global financial crisis of 2007–2009.
c. describe how correlation impacts the price of quanto options as well as other multi-asset exotic options.
d. describe the structure, uses, and payoffs of a correlation swap.
e. estimate the impact of different correlations between assets in the trading book on the VaR capital charge.
f. explain the role of correlation risk in market risk and credit risk.
g. relate correlation risk to systemic and concentration risk.

8. Empirical Properties of Correlation: How Do Correlations Behave in the Real World?

Gunter Meissner, *Correlation Risk Modeling and Management, 2nd Edition* (Risk Books, 2019). Chapter 2.

After completing this reading, you should be able to:

a. describe how equity correlations and correlation volatilities behave throughout various economic states.
b. calculate a mean reversion rate using standard regression and calculate the corresponding autocorrelation.
c. identify the best-fit distribution for equity, bond, and default correlations.

9. Financial Correlation Modeling—Bottom-Up Approaches

Gunter Meissner, *Correlation Risk Modeling and Management, 2nd Edition* (Risk Books, 2019). Chapter 5, pages 126–134.

After completing this reading, you should be able to:

a. explain the purpose of copula functions and how they are applied in finance.
b. describe the Gaussian copula and explain how to use it to derive the joint probability of default of two assets.
c. summarize the process of finding the default time of an asset correlated to all other assets in a portfolio using the Gaussian copula.

STUDY SESSION 3

10. Empirical Approaches to Risk Metrics and Hedging

Bruce Tuckman and Angel Serrat, *Fixed Income Securities: Tools for Today's Markets, 3rd Edition* (Hoboken, NJ: John Wiley & Sons, 2011). Chapter 6.

After completing this reading, you should be able to:

a. explain the drawbacks to using a DV01-neutral hedge for a bond position.
b. describe a regression hedge and explain how it can improve a standard DV01-neutral hedge.
c. calculate the regression hedge adjustment factor, beta.
d. calculate the face value of an offsetting position needed to carry out a regression hedge.
e. calculate the face value of multiple offsetting swap positions needed to carry out a two-variable regression hedge.
f. compare and contrast level and change regressions.
g. describe principal component analysis and explain how it is applied to constructing a hedging portfolio.

11. The Science of Term Structure Models

Bruce Tuckman and Angel Serrat, *Fixed Income Securities: Tools for Today's Markets, 3rd Edition* (Hoboken, NJ: John Wiley & Sons, 2011). Chapter 7.

After completing this reading, you should be able to:

a. calculate the expected discounted value of a zero-coupon security using a binomial tree.
b. construct and apply an arbitrage argument to price a call option on a zero-coupon security using replicating portfolios.
c. define risk-neutral pricing and apply it to option pricing.
d. distinguish between true and risk-neutral probabilities and apply this difference to interest rate drift.
e. explain how the principles of arbitrage pricing of derivatives on fixed-income securities can be extended over multiple periods.
f. define option-adjusted spread (OAS) and apply it to security pricing.
g. describe the rationale behind the use of recombining trees in option pricing.
h. calculate the value of a constant-maturity Treasury swap, given an interest rate tree and the risk-neutral probabilities.
i. evaluate the advantages and disadvantages of reducing the size of the time steps on the pricing of derivatives on fixed-income securities.
j. evaluate the appropriateness of the Black-Scholes-Merton model when valuing derivatives on fixed-income securities.

12. The Evolution of Short Rates and the Shape of the Term Structure

Bruce Tuckman and Angel Serrat, *Fixed Income Securities: Tools for Today's Markets, 3rd Edition* (Hoboken, NJ: John Wiley & Sons, 2011). Chapter 8.

After completing this reading, you should be able to:

a. explain the role of interest rate expectations in determining the shape of the term structure.
b. apply a risk-neutral interest rate tree to assess the effect of volatility on the shape of the term structure.
c. estimate the convexity effect using Jensen's inequality.
d. evaluate the impact of changes in maturity, yield, and volatility on the convexity of a security.
e. calculate the price and return of a zero-coupon bond incorporating a risk premium.

13. The Art of Term Structure Models: Drift

Bruce Tuckman and Angel Serrat, *Fixed Income Securities: Tools for Today's Markets, 3rd Edition* (Hoboken, NJ: John Wiley & Sons, 2011). Chapter 9.

After completing this reading, you should be able to:

a. construct and describe the effectiveness of a short-term interest rate tree assuming normally distributed rates, both with and without drift.
b. calculate the short-term rate change and standard deviation of the rate change using a model with normally distributed rates and no drift.
c. describe methods for addressing the possibility of negative short-term rates in term structure models.
d. construct a short-term rate tree under the Ho-Lee Model with time-dependent drift.
e. describe uses and benefits of the arbitrage-free models and assess the issue of fitting models to market prices.

f. describe the process of constructing a simple and recombining tree for a short-term rate under the Vasicek Model with mean reversion.
g. calculate the Vasicek Model rate change, standard deviation of the rate change, expected rate in T years and half-life.
h. describe the effectiveness of the Vasicek Model.

14. The Art of Term Structure Models: Volatility and Distribution

Bruce Tuckman and Angel Serrat, *Fixed Income Securities: Tools for Today's Markets, 3rd Edition* (Hoboken, NJ: John Wiley & Sons, 2011). Chapter 10.

After completing this reading, you should be able to:

a. describe the short-term rate process under a model with time-dependent volatility.
b. calculate the short-term rate change and determine the behavior of the standard deviation of the rate change using a model with time dependent volatility.
c. assess the efficacy of time-dependent volatility models.
d. describe the short-term rate process under the Cox-Ingersoll-Ross (CIR) and lognormal models.
e. calculate the short-term rate change and describe the basis point volatility using the CIR and lognormal models.
f. describe lognormal models with deterministic drift and mean reversion.

15. Volatility Smiles

John C. Hull, *Options, Futures, and Other Derivatives, 10th Edition* (New York, NY: Pearson, 2017). Chapter 20.

After completing this reading, you should be able to:

a. describe a volatility smile and volatility skew.
b. explain the implications of put-call parity on the implied volatility of call and put options.
c. compare the shape of the volatility smile (or skew) to the shape of the implied distribution of the underlying asset price and to the pricing of options on the underlying asset.
d. describe characteristics of foreign exchange rate distributions and their implications on option prices and implied volatility.
e. describe the volatility smile for equity options and foreign currency options and provide possible explanations for its shape.
f. describe alternative ways of characterizing the volatility smile.
g. describe volatility term structures and volatility surfaces and how they may be used to price options.
h. explain the impact of the volatility smile on the calculation of an option's Greek-letter risk measures.
i. explain the impact of a single asset price jump on a volatility smile.

16. Fundamental Review of the Trading Book

John C. Hull, *Risk Management and Financial Institutions, 5th Edition* (Hoboken, NJ: John Wiley & Sons, 2018). Chapter 18.

After completing this reading, you should be able to:

a. describe the changes to the Basel framework for calculating market risk capital under the Fundamental Review of the Trading Book (FRTB) and the motivations for these changes.
b. compare the various liquidity horizons proposed by the FRTB for different asset classes and explain how a bank can calculate its expected shortfall using the various horizons.
c. explain the FRTB revisions to Basel regulations in the following areas:
   - classification of positions in the trading book compared to the banking book.
   - backtesting, profit and loss attribution, credit risk, and securitizations.

The following is a review of the Market Risk Measurement and Management principles designed to address the learning objectives set forth by GARP®. Cross-reference to GARP assigned reading—Dowd, Chapter 3.

READING 1

ESTIMATING MARKET RISK MEASURES: AN INTRODUCTION AND OVERVIEW

Study Session 1

EXAM FOCUS

In this reading, the focus is on the estimation of market risk measures, such as value at risk (VaR). VaR identifies the probability that losses will be greater than a pre-specified threshold level. For the exam, be prepared to evaluate and calculate VaR using historical simulation and parametric models (both normal and lognormal return distributions). One drawback to VaR is that it does not estimate losses in the tail of the returns distribution. Expected shortfall (ES) does, however, estimate the loss in the tail (i.e., after the VaR threshold has been breached) by averaging loss levels at different confidence levels. Coherent risk measures incorporate personal risk aversion across the entire distribution and are more general than expected shortfall. Quantile-quantile (QQ) plots are used to visually inspect if an empirical distribution matches a theoretical distribution.

ESTIMATING RETURNS

To better understand the material in this reading, it is helpful to recall the computations of arithmetic and geometric returns. Note that the convention when computing these returns (as well as VaR) is to quote return losses as positive values. For example, if a portfolio is expected to decrease in value by $1 million, we use the terminology "expected loss is $1 million" rather than "expected profit is −$1 million."

Profit/loss data: Change in value of asset/portfolio, P_t, at the end of period t plus any interim payments, D_t.

$$P/L_t = P_t + D_t - P_{t-1}$$

Arithmetic return data: Assumption is that interim payments do not earn a return (i.e., no reinvestment). Hence, this approach is not appropriate for long investment horizons.

$$r_t = \frac{P_t + D_t - P_{t-1}}{P_{t-1}} = \frac{P_t + D_t}{P_{t-1}} - 1$$

Geometric return data: Assumption is that interim payments are continuously reinvested. Note that this approach ensures that asset price can never be negative.

$$R_t = \ln\left(\frac{P_t + D_t}{P_{t-1}}\right)$$

MODULE 1.1: HISTORICAL AND PARAMETRIC ESTIMATION APPROACHES

Historical Simulation Approach

LO 1.a: Estimate VaR using a historical simulation approach.

Estimating VaR with a historical simulation approach is by far the simplest and most straightforward VaR method. To make this calculation, you simply order return observations from largest to smallest. The observation that follows the threshold loss level denotes the VaR limit. We are essentially searching for the observation that separates the tail from the body of the distribution. More generally, the observation that determines VaR for *n* observations at the (1 − α) confidence level would be: (α × n) + 1.

PROFESSOR'S NOTE
Recall that the confidence level, (1 − α), is typically a large value (e.g., 95%) whereas the significance level, usually denoted as α, is much smaller (e.g., 5%).

To illustrate this VaR method, assume you have gathered 1,000 monthly returns for a security and produced the distribution shown in Figure 1.1. You decide that you want to compute the monthly VaR for this security at a confidence level of 95%. At a 95% confidence level, the lower tail displays the lowest 5% of the underlying distribution's returns. For this distribution, the value associated with a 95% confidence level is a return of −15.5%. If you have $1,000,000 invested in this security, the one-month VaR is $155,000 (= − 15.5% × $1,000,000).

Figure 1.1: Histogram of Monthly Returns

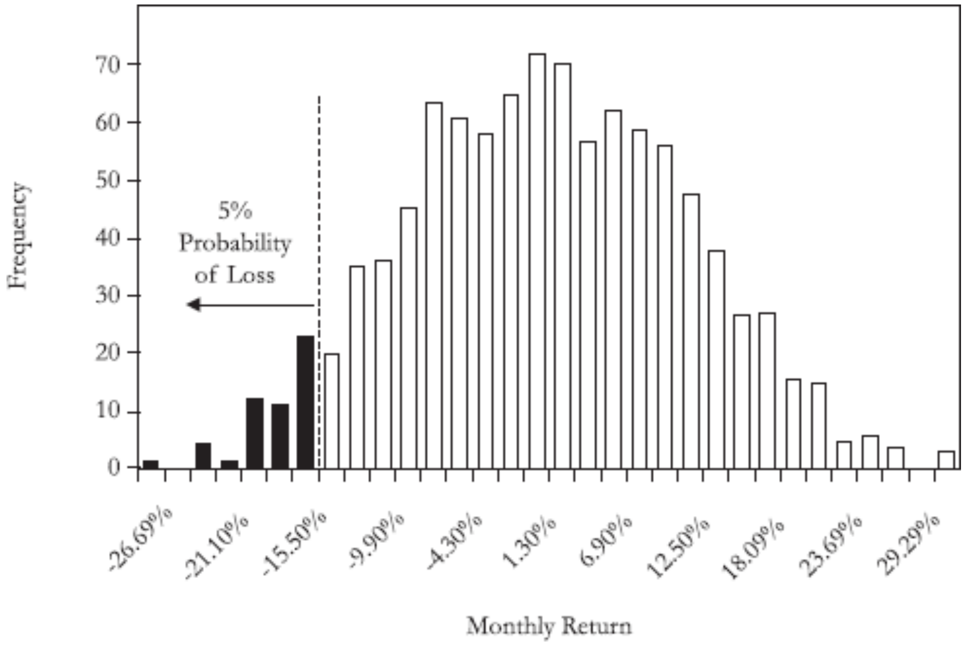

EXAMPLE: Identifying the VaR limit

Identify the ordered observation in a sample of 1,000 data points that corresponds to VaR at a 95% confidence level.

Answer:

Since VaR is to be estimated at 95% confidence, this means that 5% (i.e., 50) of the ordered observations would fall in the tail of the distribution. Therefore, the 51st ordered loss observation would separate the 5% of largest losses from the remaining 95% of returns.

PROFESSOR'S NOTE

VaR is the quantile that separates the tail from the body of the distribution. With 1,000 observations at a 95% confidence level, there is a certain level of arbitrariness in how the ordered observations relate to VaR. In other words, should VaR be the 50th observation (i.e., $\alpha \times n$), the 51st observation [i.e., $(\alpha \times n) + 1$], or some combination of these observations? In this example, using the 51st observation was the approximation for VaR, and the method used in the assigned reading. However, on past FRM exams, VaR using the historical simulation method has been calculated as just: $(\alpha \times n)$, in this case, as the 50th observation.

EXAMPLE: Computing VaR

A long history of profit/loss data closely approximates a standard normal distribution (mean equals zero; standard deviation equals one). **Estimate** the 5% VaR using the historical simulation approach.

Answer:

The VaR limit will be at the observation that separates the tail loss with area equal to 5% from the remainder of the distribution. Since the distribution is closely approximated by the standard normal distribution, the VaR is 1.65 (5% critical value from the *z*-table). Recall that since VaR is a one-tailed test, the entire significance level of 5% is in the left tail of the returns distribution.

From a practical perspective, the historical simulation approach is sensible only if you expect future performance to follow the same return generating process as in the past. Furthermore, this approach is unable to adjust for changing economic conditions or abrupt shifts in parameter values.

Parametric Estimation Approaches

LO 1.b: Estimate VaR using a parametric approach for both normal and lognormal return distributions.

In contrast to the historical simulation method, the parametric approach (e.g., the delta-normal approach) explicitly assumes a distribution for the underlying observations. In this section, we will analyze two cases: (1) VaR for returns that follow a normal distribution, and (2) VaR for returns that follow a lognormal distribution.

Normal VaR

Intuitively, the VaR for a given confidence level denotes the point that separates the tail losses from the remaining distribution. The VaR cutoff will be in the left tail of the returns distribution. Hence, the calculated value at risk is negative, but is typically reported as a positive value since the negative amount is implied (i.e., it is the value that is at risk). In equation form, the VaR at significance level α is:

$$\text{VaR}(\alpha\%) = -\mu_{P/L} + \sigma_{P/L}\, z_{\alpha}$$

where μ and σ denote the mean and standard deviation of the profit/loss distribution and *z* denotes the critical value (i.e., quantile) of the standard normal. In practice, the population parameters μ and σ are not likely known, in which case the researcher will use the sample mean and standard deviation.

EXAMPLE: **Computing VaR (normal distribution)**

Assume that the profit/loss distribution for XYZ is normally distributed with an annual mean of $15 million and a standard deviation of $10 million. **Calculate** the VaR at the 95% and 99% confidence levels using a parametric approach.

Answer:

VaR(5%) = −$15 million + $10 million × 1.65 = $1.5 million. Therefore, XYZ expects to lose at most $1.5 million over the next year with 95% confidence. Equivalently, XYZ expects to lose more than $1.5 million with a 5% probability.

VaR(1%) = −$15 million + $10 million × 2.33 = $8.3 million. Note that the VaR (at 99% confidence) is greater than the VaR (at 95% confidence) as follows from the definition of value at risk.

Now suppose that the data you are using is arithmetic return data rather than profit/loss data. The arithmetic returns follow a normal distribution as well. As you would expect, because of the relationship between prices, profits/losses, and returns, the corresponding VaR is very similar in format:

$$\text{VaR}(\alpha\%) = (-\mu_r + \sigma_r \times z_\alpha) \times P_{t-1}$$

EXAMPLE: Computing VaR (arithmetic returns)

A portfolio has a beginning period value of $100. The arithmetic returns follow a normal distribution with a mean of 10% and a standard deviation of 20%. **Calculate** VaR at both the 95% and 99% confidence levels.

Answer:

$$\text{VaR}(5\%) = (-10\% + 1.65 \times 20\%) \times 100 = \$23.0$$

$$\text{VaR}(1\%) = (-10\% + 2.33 \times 20\%) \times 100 = \$36.6$$

Lognormal VaR

The lognormal distribution is right-skewed with positive outliers and bounded below by zero. As a result, the lognormal distribution is commonly used to counter the possibility of negative asset prices (P_t). Technically, if we assume that geometric returns follow a normal distribution (μ_R, σ_R), then the natural logarithm of asset prices follows a normal distribution and P_t follows a lognormal distribution. After some algebraic manipulation, we can derive the following expression for **lognormal VaR**:

$$\text{VaR}(\alpha\%) = P_{t-1} \times (1 - e^{\mu_R - \sigma_R \times z_\alpha})$$

EXAMPLE: Computing VaR (lognormal distribution)

A diversified portfolio exhibits a normally distributed geometric return with mean and standard deviation of 10% and 20%, respectively. **Calculate** the 5% and 1% lognormal VaR assuming the beginning period portfolio value is $100.

Answer:

$$\begin{aligned}\text{Lognormal VaR}(5\%) &= 100 \times (1 - \exp[0.1 - 0.2 \times 1.65]) \\ &= 100 \times (1 - \exp[-0.23]) \\ &= \$20.55\end{aligned}$$

$$\begin{aligned}\text{Lognormal VaR}(1\%) &= 100 \times (1 - \exp[0.1 - 0.2 \times 2.33]) \\ &= 100 \times (1 - \exp[-0.366]) \\ &= \$30.65\end{aligned}$$

Note that the calculation of lognormal VaR (geometric returns) and normal VaR (arithmetic returns) will be similar when we are dealing with short time periods and practical return estimates.

MODULE QUIZ 1.1

1. The VaR at a 95% confidence level is estimated to be 1.56 from a historical simulation of 1,000 observations. Which of the following statements is most likely true?
   A. The parametric assumption of normal returns is correct.
   B. The parametric assumption of lognormal returns is correct.
   C. The historical distribution has fatter tails than a normal distribution.
   D. The historical distribution has thinner tails than a normal distribution.

2. Assume the profit/loss distribution for XYZ is normally distributed with an annual mean of $20 million and a standard deviation of $10 million. The 5% VaR is calculated and interpreted as which of the following statements?
   A. 5% probability of losses of at least $3.50 million.
   B. 5% probability of earnings of at least $3.50 million.
   C. 95% probability of losses of at least $3.50 million.
   D. 95% probability of earnings of at least $3.50 million.

MODULE 1.2: RISK MEASURES

Expected Shortfall

LO 1.c: Estimate the expected shortfall given profit and loss (P&L) or return data.

A major limitation of the VaR measure is that it does not tell the investor the amount or magnitude of the actual loss. VaR only provides the maximum value we can lose for a given confidence level. The **expected shortfall (ES)** provides an estimate of the tail loss by averaging the VaRs for increasing confidence levels in the tail. Specifically, the tail mass is divided into *n* equal slices and the corresponding n − 1 VaRs are computed. For example, if n = 5, we can construct the following table based on the normal distribution:

Figure 1.2: Estimating Expected Shortfall

Confidence Level	VaR	Difference
96%	1.7507	
97%	1.8808	0.1301
98%	2.0537	0.1729
99%	2.3263	0.2726
Average	2.003	
Theoretical true value	2.063	

Observe that the VaR increases (from *Difference* column) in order to maintain the same interval mass (of 1%) because the tails become thinner and thinner. The average of the four computed VaRs is 2.003 and represents the probability-weighted expected tail loss (a.k.a. expected shortfall). Note that as *n* increases, the expected shortfall will increase and approach the theoretical true loss [2.063 in this case; the average of a high number of VaRs (e.g., greater than 10,000)].

Estimating Coherent Risk Measures

LO 1.d: Estimate risk measures by estimating quantiles.

A more general risk measure than either VaR or ES is known as a coherent risk measure. A **coherent risk measure** is a weighted average of the quantiles of the loss distribution where the weights are user-specific based on individual risk aversion. ES (as well as VaR) is a special case of a coherent risk measure. When modeling the ES case, the weighting function is set to [1 / (1 − confidence level)] for all tail losses. All other quantiles will have a weight of zero.

Under expected shortfall estimation, the tail region is divided into equal probability slices and then multiplied by the corresponding quantiles. Under the more general coherent risk measure, the entire distribution is divided into equal probability slices weighted by the more general risk aversion (weighting) function.

This procedure is illustrated for n = 10. First, the entire return distribution is divided into nine (i.e., n − 1) equal probability mass slices at 10%, 20%, …, 90% (i.e., loss quantiles). Each breakpoint corresponds to a different quantile. For example, the 10% quantile (confidence level = 10%) relates to −1.2816, the 20% quantile (confidence level = 20%) relates to −0.8416, and the 90% quantile (confidence level = 90%) relates to 1.2816. Next, each quantile is weighted by the specific risk aversion function and then averaged to arrive at the value of the coherent risk measure.

This coherent risk measure is more sensitive to the choice of *n* than expected shortfall, but will converge to the risk measure's true value for a sufficiently large number of observations. The intuition is that as *n* increases, the quantiles will be further into the tails where more extreme values of the distribution are located.

Even though the risk measure estimate eventually converges to the true value as the number of observations is sufficiently large, knowing the exact value of *n* can be useful. One approach involves beginning with a small value of *n* and repeatedly doubling it until the risk measure estimates stabilize. Every time the number of observations is doubled, the width of the tail slides is cut in half. This process allows for the calculation of the "halving error," and the ideal number of tail slides is found when the halving error is near zero (i.e., the difference between the estimated risk measures as *n* increases is minimal).

LO 1.e: Evaluate estimators of risk measures by estimating their standard errors.

Sound risk management practice reminds us that estimators are only as useful as their precision. That is, estimators that are less precise (i.e., have large standard errors and wide confidence intervals) will have limited practical value. Therefore, it is best practice to also compute the standard error for all coherent risk measures.

PROFESSOR'S NOTE
The process of estimating standard errors for estimators of coherent risk measures is quite complex, so your focus should be on interpretation of this concept.

First, let's start with a sample size of *n* and arbitrary bin width of *h* around quantile, *q*. Bin width is just the width of the intervals, sometimes called "bins," in a histogram. Computing standard error is done by realizing that the square root of the variance of the quantile is equal to the standard error of the quantile. After finding the standard error, a confidence interval for a risk measure such as VaR can be constructed as follows:

$$[q + se(q) \times z_{\alpha}] > VaR > [q - se(q) \times z_{\alpha}]$$

EXAMPLE: Estimating standard errors

Construct a 90% confidence interval for 5% VaR (the 95% quantile) drawn from a standard normal distribution. Assume bin width = 0.1 and that the sample size is equal to 500.

Answer:

The quantile value, *q*, corresponds to the 5% VaR which occurs at 1.65 for the standard normal distribution. The confidence interval takes the following form:

$$[1.65 + 1.65 \times se(q)] > VaR > [1.65 - 1.65 \times se(q)]$$

PROFESSOR'S NOTE
Recall that a confidence interval is a two-tailed test (unlike VaR), so a 90% confidence level will have 5% in each tail. Given that this is equivalent to the 5% significance level of VaR, the critical values of 1.65 will be the same in both cases.

Since bin width is 0.1, *q* is in the range 1.65 ± 0.1/2 = [1.7, 1.6]. Note that the left tail probability, *p*, is the area to the left of −1.7 for a standard normal distribution.

Next, calculate the probability mass between [1.7, 1.6], represented as *f(q)*. From the standard normal table, the probability of a loss *greater* than 1.7 is 0.045 (left tail). Similarly, the probability of a loss *less* than 1.6 (right tail) is 0.945. Collectively, f(q) = 1 − 0.045 − 0.945 = 0.01.

The standard error of the quantile is derived from the variance approximation of *q* and is equal to:

$$se(q) = \frac{\sqrt{p(1-p)/n}}{f(q)}$$

Now we are ready to substitute in the variance approximation to calculate the confidence interval for VaR:

$$\left[1.65 + 1.65\frac{\sqrt{0.045(1-0.045)/500}}{0.01}\right]$$

$$> \text{VaR} >$$

$$\left[1.65 - 1.65\frac{\sqrt{0.045(1-0.045)/500}}{0.01}\right] = 3.18 > \text{VaR} > 0.1$$

Let's return to the variance approximation and perform some basic comparative statistics. What happens if we increase the sample size holding all other factors constant? Intuitively, the larger the sample size the smaller the standard error and the narrower the confidence interval.

Now suppose we increase the bin size, h, holding all else constant. This will increase the probability mass $f(q)$ and reduce p, the probability in the left tail. The standard error will decrease and the confidence interval will again narrow.

Lastly, suppose that p increases indicating that tail probabilities are more likely. Intuitively, the estimator becomes less precise and standard errors increase, which widens the confidence interval. Note that the expression $p(1 - p)$ will be maximized at $p = 0.5$.

The above analysis was based on one quantile of the loss distribution. Just as the previous section generalized the expected shortfall to the coherent risk measure, we can do the same for the standard error computation. Thankfully, this complex process is not the focus of the LO.

Quantile-Quantile Plots

LO 1.f: Interpret quantile-quantile (QQ) plots to identify the characteristics of a distribution.

A natural question to ask in the course of our analysis is, "From what distribution is the data drawn?" The truth is that you will never really know since you only observe the realizations from random draws of an unknown distribution. However, visual inspection can be a very simple but powerful technique.

In particular, the **quantile-quantile (QQ) plot** is a straightforward way to visually examine if empirical data fits the reference or hypothesized theoretical distribution (assume standard normal distribution for this discussion). The process graphs the quantiles at regular confidence intervals for the empirical distribution against the theoretical distribution. As an example, if both the empirical and theoretical data are drawn from the same distribution, then the median (confidence level = 50%) of the empirical distribution would plot very close to zero, while the median of the theoretical distribution would plot exactly at zero.

Continuing in this fashion for other quantiles (40%, 60%, and so on) will map out a function. If the two distributions are very similar, the resulting QQ plot will be linear.

Let us compare a theoretical standard normal distribution relative to an empirical *t*-distribution (assume that the degrees of freedom for the *t*-distribution are sufficiently small and that there are noticeable differences from the normal distribution). We know that both distributions are symmetric, but the *t*-distribution will have fatter tails. Hence, the quantiles near zero (confidence level = 50%) will match up quite closely. As we move further into the tails, the quantiles between the *t*-distribution and the normal will diverge (see Figure 1.3). For example, at a confidence level of 95%, the critical *z*-value is −1.65, but for the *t*-distribution, it is closer to −1.68 (degrees of freedom of approximately 40). At 97.5% confidence, the difference is even larger, as the *z*-value is equal to −1.96 and the *t*-stat is equal to −2.02. More generally, if the middles of the QQ plot match up, but the tails do not, then the empirical distribution can be interpreted as symmetric with tails that differ from a normal distribution (either fatter or thinner).

Figure 1.3: QQ Plot

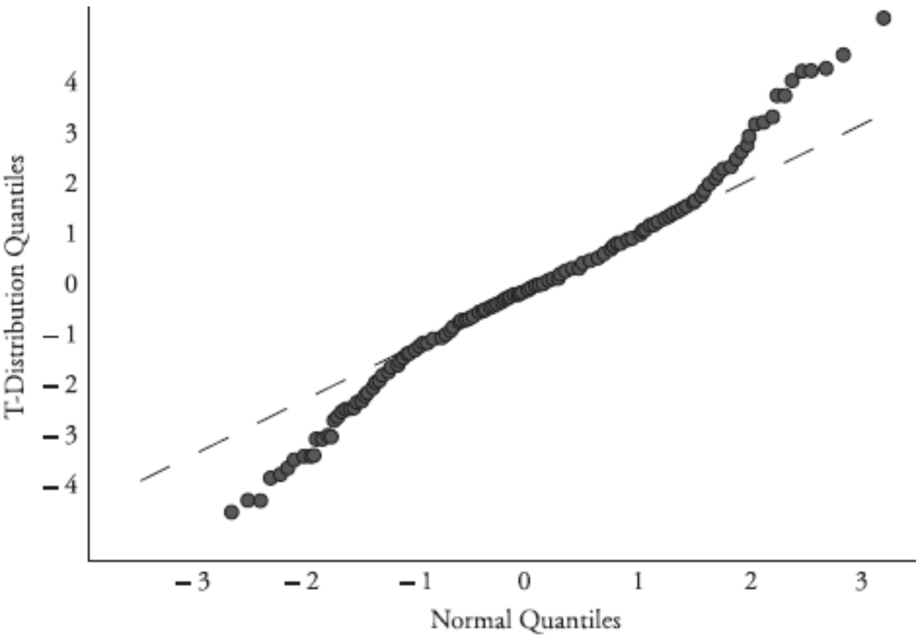

MODULE QUIZ 1.2

1. Which of the following statements about expected shortfall estimates and coherent risk measures are true?
   A. Expected shortfall and coherent risk measures estimate quantiles for the entire loss distribution.
   B. Expected shortfall and coherent risk measures estimate quantiles for the tail region.
   C. Expected shortfall estimates quantiles for the tail region and coherent risk measures estimate quantiles for the non-tail region only.
   D. Expected shortfall estimates quantiles for the entire distribution and coherent risk measures estimate quantiles for the tail region only.

2. Which of the following statements most likely increases standard errors from coherent risk measures?
   A. Increasing sample size and increasing the left tail probability.
   B. Increasing sample size and decreasing the left tail probability.
   C. Decreasing sample size and increasing the left tail probability.

D. Decreasing sample size and decreasing the left tail probability.

3. The quantile-quantile plot is best used for what purpose?
   A. Testing an empirical distribution from a theoretical distribution.
   B. Testing a theoretical distribution from an empirical distribution.
   C. Identifying an empirical distribution from a theoretical distribution.
   D. Identifying a theoretical distribution from an empirical distribution.

KEY CONCEPTS

LO 1.a

Historical simulation is the easiest method to estimate value at risk. All that is required is to reorder the profit/loss observations in increasing magnitude of losses and identify the breakpoint between the tail region and the remainder of distribution.

LO 1.b

Parametric estimation of VaR requires a specific distribution of prices or equivalently, returns. This method can be used to calculate VaR with either a normal distribution or a lognormal distribution.

Under the assumption of a normal distribution, VaR (i.e., delta-normal VaR) is calculated as follows:

$$\text{VaR} = -\mu_{P/L} + \sigma_{P/L} \times z_{\alpha}$$

Under the assumption of a lognormal distribution, lognormal VaR is calculated as follows:

$$\text{VaR} = P_{t-1} \times (1 - e^{\mu_R - \sigma_R \times z_{\alpha}})$$

LO 1.c

VaR identifies the lower bound of the profit/loss distribution, but it does not estimate the expected tail loss. Expected shortfall overcomes this deficiency by dividing the tail region into equal probability mass slices and averaging their corresponding VaRs.

LO 1.d

A more general risk measure than either VaR or ES is known as a coherent risk measure. A coherent risk measure is a weighted average of the quantiles of the loss distribution where the weights are user-specific based on individual risk aversion. A coherent risk measure will assign each quantile (not just tail quantiles) a weight. The average of the weighted VaRs is the estimated loss.

LO 1.e

Sound risk management requires the computation of the standard error of a coherent risk measure to estimate the precision of the risk measure itself. The simplest method creates a confidence interval around the quantile in question. To compute standard error, it is necessary to find the variance of the quantile, which will require estimates from the underlying distribution.

LO 1.f

The quantile-quantile (QQ) plot is a visual inspection of an empirical quantile relative to a hypothesized theoretical distribution. If the empirical distribution closely matches the theoretical distribution, the QQ plot would be linear.

ANSWER KEY FOR MODULE QUIZZES

Module Quiz 1.1

1. **D** The historical simulation indicates that the 5% tail loss begins at 1.56, which is less than the 1.65 predicted by a standard normal distribution. Therefore, the historical simulation has thinner tails than a standard normal distribution. (LO 1.a)

2. **D** The value at risk calculation at 95% confidence is: −20 million + 1.65 × 10 million = −$3.50 million. Since the expected loss is negative and VaR is an implied negative amount, the interpretation is that XYZ will earn less than +$3.50 million with 5% probability, which is equivalent to XYZ earning at least $3.50 million with 95% probability. (LO 1.b)

Module Quiz 1.2

1. **B** ES estimates quantiles for n − 1 equal probability masses in the tail region only. The coherent risk measure estimates quantiles for the entire distribution including the tail region. (LO 1.c)

2. **C** Decreasing sample size clearly increases the standard error of the coherent risk measure given that standard error is defined as:

   $$se(q) = \frac{\sqrt{p(1-p)/n}}{f(q)}$$

   As the left tail probability, *p,* increases, the probability of tail events increases, which also increases the standard error. Mathematically, p(1 − p) increases as *p* increases until *p* = 0.5. Small values of *p* imply smaller standard errors. (LO 1.e)

3. **C** Once a sample is obtained, it can be compared to a reference distribution for possible identification. The QQ plot maps the quantiles one to one. If the relationship is close to linear, then a match for the empirical distribution is found. The QQ plot is used for visual inspection only without any formal statistical test. (LO 1.f)

The following is a review of the Market Risk Measurement and Management principles designed to address the learning objectives set forth by GARP®. Cross-reference to GARP assigned reading—Dowd, Chapter 4.

READING 2

NON-PARAMETRIC APPROACHES

Study Session 1

EXAM FOCUS

This reading introduces non-parametric estimation and bootstrapping (i.e., resampling). The key difference between these approaches and parametric approaches discussed in the previous reading is that with non-parametric approaches the underlying distribution is not specified, and it is a data driven, not assumption driven, analysis. For example, historical simulation is limited by the discreteness of the data, but non-parametric analysis "smooths" the data points to allow for any VaR confidence level between observations. For the exam, pay close attention to the description of the bootstrap historical simulation approach as well as the various weighted historical simulations approaches.

MODULE 2.1: NON-PARAMETRIC APPROACHES

Non-parametric estimation does not make restrictive assumptions about the underlying distribution like parametric methods, which assume very specific forms such as normal or lognormal distributions. Non-parametric estimation lets the data drive the estimation. The flexibility of these methods makes them excellent candidates for VaR estimation, especially if tail events are sparse.

Bootstrap Historical Simulation Approach

LO 2.a: Apply the bootstrap historical simulation approach to estimate coherent risk measures.

The **bootstrap historical simulation** is a simple and intuitive estimation procedure. In essence, the bootstrap technique draws a sample from the original data set, records the VaR from that particular sample and "returns" the data. This procedure is repeated over and over and records multiple sample VaRs. Since the data is always "returned" to the data set, this procedure is akin to sampling with replacement. The best VaR estimate from the full data set is the average of all sample VaRs.

This same procedure can be performed to estimate the expected shortfall (ES). Each drawn sample will calculate its own ES by slicing the tail region into *n* slices and averaging the VaRs at each of the n − 1 quantiles. This is exactly the same procedure described in the previous reading. Similarly, the best estimate of the expected shortfall for the original data set is the average of all of the sample expected shortfalls.

Empirical analysis demonstrates that the bootstrapping technique consistently provides more precise estimates of coherent risk measures than historical simulation on raw data alone.

Applying Non-Parametric Estimation

LO 2.b: Describe historical simulation using non-parametric density estimation.

The clear advantage of the traditional historical simulation approach is its simplicity. One obvious drawback, however, is that the discreteness of the data does not allow for estimation of VaRs between data points. If there were 100 historical observations, then it is straightforward to estimate VaR at the 95% or the 96% confidence levels, and so on. However, this method is unable to incorporate a confidence level of 95.5%, for example. More generally, with *n* observations, the historical simulation method only allows for *n* different confidence levels.

One of the advantages of non-parametric density estimation is that the underlying distribution is free from restrictive assumptions. Therefore, the existing data points can be used to "smooth" the data points to allow for VaR calculation at all confidence levels. The simplest adjustment is to connect the midpoints between successive histogram bars in the original data set's distribution. See Figure 2.1 for an illustration of this **surrogate density function**. Notice that by connecting the midpoints, the lower bar "receives" area from the upper bar, which "loses" an equal amount of area. In total, no area is lost, only displaced, so we still have a probability distribution function, just with a modified shape. The shaded area in Figure 2.1 represents a possible confidence interval, which can be utilized regardless of the size of the data set. The major improvement of this non-parametric approach over the traditional historical simulation approach is that VaR can now be calculated for a continuum of points in the data set.