2024 FRM®

Exam Prep

Schweser's Secret Sauce®

PART II

KAPLAN SCHWESER

Schweser's Secret Sauce®

FRM Part II

2024

KAPLAN SCHWESER

SCHWESER'S SECRET SAUCE®: 2024 FRM® PART II

Published in 2024 by Kaplan, Inc.

ISBN: 978-1-0788-4261-7

CONTENTS

FOREWORD

This review book is a valuable addition to the study tools of any FRM exam candidate. It offers concise coverage of exam topics to enhance your retention of the FRM curriculum.

We suggest that you use this book as a companion to your other, more comprehensive, study materials. It is easier to carry with you and will allow you to study these key concepts, definitions, and techniques over and over, which is a crucial part of mastering the material. For a majority of you, there are no shortcuts to learning the broad array of subject matter covered by the FRM curriculum, but this book should be a very valuable tool for learning and reviewing the material as you progress in your studies over the weeks leading up to exam day.

Previous Part II exam pass rates have been close to 60%, and many FRM candidates have commented on the high difficulty level of the exam. This is an indication that you should not underestimate the task at hand. Our SchweserNotes, Mock Exams, SchweserPro QBank, OnDemand Class, and Schweser's Secret Sauce are all designed to help you study as efficiently as possible, grasp and retain the material, and apply your knowledge with confidence on the exam.

As a reminder, the 2024 FRM Part II topic area coverage and weightings assigned by GARP 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 Resilience	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

MARKET RISK MEASUREMENT AND MANAGEMENT

Study Sessions 1–3

Weight on Exam	20%
SchweserNotes™ Reference	Book 1

STUDY SESSION 1: MARKET RISK ANALYSIS

ESTIMATING MARKET RISK MEASURES: AN INTRODUCTION AND OVERVIEW

Cross-reference to GARP assigned reading – Dowd, Chapter 3.

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.

Parametric Estimation Approaches

In contrast to the historical simulation method, the parametric approach (e.g., the delta-normal approach) explicitly assumes a distribution for the underlying observations.

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:

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

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
>
> 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

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:**
>
> $$\begin{aligned}\text{Lognormal VaR}(5\%) &= 100 \times (1 - \exp[0.1 - 0.2 \times 1.65]) \\ &= 100 \times (1 - \exp[-0.23]) \\ &= \$20.55 \\ \text{Lognormal VaR}(1\%) &= 100 \times (1 - \exp[0.1 - 0.2 \times 2.33]) \\ &= 100 \times (1 - \exp[-0.366]) \\ &= \$30.65\end{aligned}$$

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.

Expected Shortfall

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.

Estimating Coherent Risk Measures

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.

Quantile-Quantile Plots

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.

NON-PARAMETRIC APPROACHES

Cross-reference to GARP assigned reading – Dowd, Chapter 4.

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

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.

Applying Non-Parametric 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.

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.

Weighted Historical Simulation Approaches

Age-Weighted Historical Simulation

The obvious adjustment to the equal-weighted assumption used in historical simulation is to weight recent observations more and distant observations less. In this approach, the weight for an observation that is *i* days old is equal to:

$$w(i) = \frac{\lambda^{i-1}(1-\lambda)}{1-\lambda^n}$$

The implication of the age-weighted simulation is to reduce the impact of ghost effects and older events that may not reoccur. Note that this more general weighting scheme suggests that historical simulation is a special case where $\lambda = 1$ (i.e., no decay) over the estimation window.

Volatility-Weighted Historical Simulation

Another approach is to weight the individual observations by volatility rather than proximity to the current date. The intuition is that if recent volatility has increased, then using historical data will underestimate the current risk level. Similarly, if current volatility is markedly reduced, the impact of older data with higher periods of volatility will overstate the current risk level.

There are several advantages of the volatility-weighted method. First, it explicitly incorporates volatility into the estimation procedure in contrast to other historical methods. Second, the near-term VaR estimates are likely to be more sensible in light of current market conditions. Third, the volatility-adjusted returns allow for VaR estimates that are higher than estimates with the historical data set.

Correlation-Weighted Historical Simulation

This procedure is more complicated than the volatility-weighting approach, but it follows the same basic principles. Intuitively, the historical correlation (or equivalently variance-covariance) matrix needs to be adjusted to the new information environment. This is accomplished, loosely speaking, by “multiplying” the historic returns by the revised correlation matrix to yield updated correlation-adjusted returns.