2026

FRM®

Exam Prep

Schweser's
Secret Sauce®

Part II

KAPLAN SCHWESER

Schweser's Secret Sauce®

FRM Part II

2026

KAPLAN SCHWESER

SCHWESER'S SECRET SAUCE®: 2026 FRM® PART II

Published in 2026 by Kaplan, Inc.

ISBN: 978-1-0788-5469-6

CONTENTS

ABOUT THIS REVIEW GUIDE

Schweser's Secret Sauce® is intended to serve as a focused and practical companion after your initial studies. This concise book distills the most critical concepts, definitions, and exam strategies from the extensive FRM Part II curriculum. It is designed for efficient review, portability, and quick reference during the final phase of exam preparation.

It is important to emphasize, however, that not every learning objective in the curriculum is addressed in detail in this book. To build a comprehensive understanding, we strongly encourage you to utilize this guide as a review of the material from the SchweserNotes™, the official FRM curriculum, and additional practice resources. Repetition and consistent review are important components of successful mastery on exam day.

As a reminder, the 2026 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

Earning the FRM designation is a significant undertaking. The breadth and depth of the material require diligence, discipline, and perseverance. There are no shortcuts to genuine understanding of the material. Your preparation should be thorough and systematic, making full use of the suite of study tools available to you, including the SchweserNotes™, OnDemand Classes, SchweserPro™ QBank, and Mock Exams.

By pursuing the FRM credential, you are joining a distinguished global community of investment professionals dedicated to advancing the practice and principles of risk management. Whether you are seeking career advancement, a transition to a new role, or an enhancement of your investment acumen, the knowledge and skills you acquire in this program will serve you well throughout your professional journey.

On behalf of the entire Kaplan Schweser team, we thank you for selecting us as your educational partner. We extend our best wishes for your studies and look forward to your future achievements as a certified Financial Risk Manager.

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 with 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. 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 is a special case of a coherent risk measure. When modeling ES, 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.

Nonparametric estimation does not make restrictive assumptions about the underlying distribution like parametric methods, which assume very specific forms such as normal or lognormal distributions. Nonparametric 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 dataset, 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 dataset, this procedure is akin to sampling with replacement. The best VaR estimate from the full dataset is the average of all sample VaRs.

Applying Nonparametric 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 nonparametric 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 dataset'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 dataset.

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.

Correlation-weighted simulation is an even richer analytical tool than volatility-weighted simulation because it allows for updated variances (volatilities) as well as covariances (correlations).

Filtered Historical Simulation

The filtered historical simulation is the most comprehensive, and hence most complicated, of the nonparametric estimators. The process combines the historical simulation model with conditional volatility models (like GARCH or asymmetric GARCH). Thus, the method contains both the attractions of the traditional historical simulation approach and the sophistication of models that incorporate changing volatility. In simplified terms, the model is flexible enough to capture conditional volatility and volatility clustering as well as a surprise factor that could have an asymmetric effect on volatility.

Advantages and Disadvantages of Nonparametric Methods

Advantages of nonparametric methods include the following:

- Intuitive and often computationally simple.
- Not hindered by parametric violations of skewness, fat-tails, et cetera.
- Avoids complex variance-covariance matrices and dimension problems.
- Data is often readily available and does not require adjustments.
- Can accommodate more complex analysis.

Disadvantages of nonparametric methods include the following:

- Analysis depends critically on historical data.
- Volatile data periods lead to VaR and ES estimates that are too high.
- Quiet data periods lead to VaR and ES estimates that are too low.
- Difficult to detect structural shifts/regime changes in the data.
- Cannot accommodate plausible large impact events if they did not occur within the sample period.
- Difficult to estimate losses significantly larger than the maximum loss within the dataset.
- Need sufficient data, which may not be possible for new instruments or markets.

Parametric Approaches (II): Extreme Value

Cross-reference to GARP assigned reading – Dowd, Chapter 7.

Managing Extreme Values

The occurrence of extreme events is rare; however, it is crucial to identify these extreme events for risk management since they can prove to be very costly. Extreme values are the result of large market declines or crashes, the failure of major institutions, the outbreak of financial or political crises, or natural catastrophes. The challenge of analyzing and modeling extreme values is that there are only a few observations on which to build a model, and there are ranges of extreme values that have yet to occur.

Extreme Value Theory

Extreme value theory (EVT) is a branch of applied statistics that has been developed to address problems associated with extreme outcomes. EVT focuses on the unique aspects of extreme values and is different from “central tendency” statistics, in which the central-limit theorem plays an important role. Extreme value theorems provide a template for estimating the parameters used to describe extreme movements.

Peaks-Over-Threshold

The peaks-over-threshold (POT) approach is an application of extreme value theory to the distribution of excess losses over a high threshold. The POT approach generally requires fewer parameters than approaches based on extreme value theorems. The POT approach provides a natural way to model values that are greater than a high threshold, and in this way, it corresponds to the generalized extreme value (GEV) theory by modeling the maxima or minima of a large sample.

Generalized Pareto (GP) Distribution

The GP distribution exhibits a curve that dips below the normal distribution prior to the tail. It then moves above the normal distribution until it reaches the extreme tail. The GP distribution then provides a linear approximation of the tail, which more closely matches empirical data.

Generalized Extreme Value and Peaks-Over-Threshold

Extreme value theory is the source of both the GEV and POT approaches. These approaches are similar in that they both have a tail parameter denoted ξ. There is a subtle difference in that GEV theory focuses on the distributions of extremes, whereas POT focuses on the distribution of values that exceed a certain threshold.

Multivariate EVT

Multivariate EVT is important because we can easily see how extreme values can be dependent on each other. A terrorist attack on oil fields will produce losses for oil companies, but it is likely that the value of most financial assets will also be affected. We can imagine similar relationships between the occurrence of a natural disaster and a decline in financial markets as well as markets for real goods and services.

Backtesting VaR

Cross-reference to GARP assigned reading – Jorion, Chapter 6.

Backtesting VaR Models

Backtesting is the process of comparing losses predicted by a value at risk (VaR) model to those actually experienced over the testing period. It is an important tool for providing *model validation*, which is a process for determining whether a VaR model is adequate. The main goal of backtesting is to ensure that actual losses do not exceed expected losses at a given confidence level. The observations that fall outside a given confidence level are known as *exceptions*. The number of exceptions falling outside of the VaR confidence level should not exceed one minus the confidence level. For example, exceptions should occur less than 5% of the time if the confidence level is 95%.

VaR models are based on static portfolios, while actual portfolio compositions are constantly changing as relative prices change and positions are bought and sold. Multiple risk factors affect actual profit and loss, but they are not included in the VaR model. For example, the actual returns are complicated by intraday changes as well as profit-and-loss factors that result from commissions, fees, interest income, and bid-ask spreads. These effects can be minimized by backtesting over a short horizon, such as a daily holding period. The backtesting period constitutes a limited sample, and a challenge for risk managers is to find an acceptable level of exceptions.

Using Failure Rates in Model Verification

An unbiased measure of the number of exceptions as a proportion of the number of samples is called the **failure rate**. The probability of exception, *p*, equals one minus the confidence level ($p = 1 - c$). If we use *N* to represent the number of exceptions and *T* to represent the sample size, the failure rate is computed as *N/T*. This failure rate is unbiased if the computed *p* approaches 1 minus the confidence level as the sample size increases. Nonparametric tests can then be used to see if the number of times a VaR model fails is acceptable or not.

Testing that the model is correctly calibrated requires the calculation of a *z*-score, where *x* is the number of actual exceptions observed. This *z*-score is then compared to the critical value at the chosen level of confidence (e.g., 1.96 for the 95% confidence level) to determine whether the VaR model is unbiased.

$$z = \frac{x - pT}{\sqrt{p(1 - p)T}}$$

Type I and Type II Errors

A sample cannot be used to determine with absolute certainty whether the model is accurate. However, we can determine the accuracy of the model and the probability of having the number of exceptions that we experienced. When determining a range for the number of exceptions that we would accept, we must strike a balance between the chance of *rejecting an accurate model* (**Type I error**) and the chance of *failing to reject an inaccurate model* (**Type II error**). The model verification test involves a tradeoff between Type I and Type II errors. The goal in backtesting is to create a VaR model with a low Type I error and to include a test for a very low Type II error rate.

Unconditional and Conditional Coverage

The term **unconditional coverage** refers to the fact that we are not concerned about independence of exception observations or the timing of when the exceptions occur. We are simply concerned with the total number of exceptions.

Conditional coverage considers the time variation of the data. In addition to having a predictable number of exceptions, we also anticipate the exceptions to be fairly equally distributed across time. A bunching of exceptions may indicate that market correlations have changed or that our trading positions have been altered. In the event that exceptions are not independent, the risk manager should incorporate models that consider time variation in risk.

Basel Committee Rules for Backtesting

The Basel Committee is primarily concerned with identifying whether exceptions are the result of bad luck (Type I error) or a faulty model (Type II error). The Basel Committee requires that market VaR be calculated at the 99% confidence level and backtested over the past year. At the 99% confidence level, we would expect to have 2.5 exceptions (= 250 × 0.01) each year, given approximately 250 trading days.

To mitigate the risk that banks willingly commit a Type II error and use a faulty model, the Basel Committee designed the **Basel penalty zones**. The committee established a scale of the number of exceptions and corresponding increases in the capital multiplier, *k*. Thus, banks are penalized for exceeding four exceptions per year. The multiplier is normally three but can be increased to as much as four, based on the accuracy of the bank's VaR model. Increasing *k* significantly increases the amount of capital a bank must hold and lowers the bank's performance measures, like return on equity.

VaR Mapping

Cross-reference to GARP assigned reading – Jorion, Chapter 11.

Mapping Process

Value at risk (VaR) mapping involves replacing the current values of a portfolio with risk factor exposures. The first step in the process is to measure all current positions within a portfolio. These positions are then mapped to **risk factors** by means of **factor exposures**. Mapping involves finding common risk factors among positions in a given portfolio. If we have a portfolio consisting of a large number of positions, it may be difficult and time consuming to manage the risk of each individual position. Instead, we can evaluate the value of these positions by mapping them onto common risk factors (e.g., changes in interest rates or equity prices). By reducing the number of variables under consideration, we greatly simplify the risk management process.

The principles for VaR risk mapping are summarized as follows:

- VaR mapping aggregates risk exposure when it is impractical to consider each position separately. For example, there may be too many computations needed to measure the risk for each individual position.
- VaR mapping simplifies risk exposures into primitive risk factors. For example, a portfolio may have thousands of positions linked to a specific exchange rate that could be summarized with one aggregate risk factor.
- VaR risk measurements can differ from pricing methods where prices cannot be aggregated. The aggregation of a number of positions to one risk factor is acceptable for risk measurement purposes.
- VaR mapping is useful for measuring changes over time, as with bonds or options. For example, as bonds mature, risk exposure can be mapped to spot yields that reflect the current position.
- VaR mapping is useful when historical data is not available.

In our choice of general risk factors for use in VaR models, we should be aware that the types and number of risk factors we choose will have an effect on the size of residual or specific risks. **Specific risks** arise from unsystematic risk or asset-specific risks of various positions in the portfolio. The more precisely we define risk, the smaller the specific risk. For example, a portfolio of bonds may include bonds of different ratings, terms, and currencies. If we use duration as our only risk factor, there will be a significant amount of variance among the bonds that we referred to as specific risk. If we add a risk factor for credit risk, we could expect that the amount of specific risk would be smaller. If we add another risk factor for currencies, we would expect that the specific risk would be even smaller. Thus, the definition of specific risk is a function of general market risk.

Mapping Approaches for Fixed-Income Portfolios

Principal mapping. This method includes only the risk of repayment of principal amounts. For principal mapping, we consider the average maturity of the portfolio. VaR is calculated using the risk level from the zero-coupon bond that equals the average maturity of the portfolio. This method is the simplest of the three approaches.

Duration mapping. With this method, the risk of the bond is mapped to a zero-coupon bond of the same duration. For duration mapping, we calculate VaR using the risk level of the zero-coupon bond that equals the duration of the portfolio. Note that it may be difficult to calculate the risk level that exactly matches the duration of the portfolio.

Cash flow mapping. With this method, the risk of the bond is decomposed into the risk of each of the bond's cash flows. Cash flow mapping is the most precise method because we map the present value of the cash flows (i.e., face amount discounted at the spot rate for a given maturity) onto the risk factors for zeros of the same maturities and include the inter-maturity correlations.

The primary distinction between these mapping techniques is how they account for both the timing and the amount of cash flows.

Stress Testing

If we assume that there is perfect correlation among maturities of the zeros, the portfolio VaR would be equal to the **undiversified VaR** (i.e., the sum of the VaRs). Instead of calculating the undiversified VaR directly, we could reduce each zero-coupon value by its respective VaR and then revalue the portfolio. The difference between the revalued portfolio and the original portfolio value should be equal to the undiversified VaR. Stressing each zero by its VaR is a simpler approach than incorporating correlations; however, this method ceases to be viable if correlations are anything but perfect (i.e., 1).

Benchmarking a Portfolio

It is often convenient to measure VaR relative to a benchmark portfolio. This is what is referred to as benchmarking a portfolio. Portfolios can be constructed that match the risk factors of a benchmark portfolio but have either a higher or a lower VaR. The VaR of the deviation between the two portfolios is referred to as a **tracking error VaR**. In other words, tracking error VaR is a measure of the difference between the VaR of the target portfolio and the benchmark portfolio.

Mapping Approaches for Linear Derivatives

The delta-normal method provides accurate estimates of VaR for portfolios and assets that can be expressed as linear combinations of normally distributed risk factors. Once a portfolio or financial instrument is expressed as a linear combination of risk factors, a covariance (correlation) matrix can be generated, and VaR can be measured using matrix multiplication.

Forwards are appropriate for the application of the delta-normal method. Their values are a linear combination of a few general risk factors, which have commonly available volatility and correlation data.

Mapping Approaches for Nonlinear Derivatives

As mentioned, the delta-normal VaR method is based on linear relationships between variables. Options, however, exhibit nonlinear relationships between movements of the value of the underlying instruments and the value of the options. In many cases, the delta-normal method may still be applied because the value of an option may be expressed linearly as the product of the option delta and the underlying asset.

Unfortunately, the delta-normal VaR cannot be expected to provide an accurate estimate of the true VaR over ranges where deltas are unstable. In other words, over longer periods of time, the delta is not constant, which makes linear methods inappropriate. Conversely, over short periods of time, such as one day, a linear approximation of the delta is more accurate. However, the accuracy of this approximation is dependent on parameter inputs (i.e., delta increases with the underlying spot price).

Validating Bank Holding Companies' Value-at-Risk Models for Market Risk

Cross-reference to GARP assigned reading – Lynch et al., Chapter 2.

VaR Model Validation

The Basel Accords of 1996 established **value at risk (VaR)** as a risk metric for determining regulatory capital for banks. **Exceedances** (or exceptions) are defined as actual losses that exceed the VaR threshold. A bank's regulatory capital is adjusted higher for those banks with a higher number of exceedances.

One problem with traditional VaR estimation is that using historical data does not account for dynamic volatility in portfolio returns. To address this limitation, **GARCH VaR** models were developed to allow for more recent volatility changes.

Conceptual Soundness of a VaR Model

A VaR model should be validated by

- determining the soundness of the methodology used,
- the quality of the input data, and
- the validity of the model assumptions.

The model should be designed to meet the specific risk management objectives of the bank. Understanding the intended use of the VaR model (e.g., regulatory capital calculation, internal risk assessment) is essential for evaluating its appropriateness. Quantitative tests that can be used to evaluate the conceptual soundness of a VaR model include conducting sensitivity analysis and applying confidence intervals.

Sensitivity Analysis for VaR

Sensitivity analysis is a quantitative test used to examine the validity of VaR estimates and the conceptual soundness of a VaR model. Performing this analysis involves the following steps:

1. *Identify key inputs*. Determine the key inputs and assumptions of the VaR model, such as the portfolio positions.
2. *Adjust key inputs and recalculate VaR*. Adjusting one input at a time (e.g., increasing or decreasing a position weight by a certain percentage) recalculates the VaR threshold. **Marginal VaR** is the change in portfolio value due to a small change in the weight of a particular portfolio position.
3. *Analyze results*. Compare the new VaR estimates to the original VaR. Evaluate how sensitive the VaR is to changes in each input.

Data omissions may compromise the quality of sensitivity analysis. Scarce data can be replaced with proxies, but the volatility of the proxy or the correlation of the proxy to portfolio value may be different from the scarce data.

Confidence Intervals for VaR

A statistical approach for evaluating the accuracy of VaR estimates is to place a **confidence interval** around the estimate such that the true VaR would fall within that confidence interval, for example, 99% of the time. Calculating confidence intervals for VaR presents several challenges for financial institutions in the areas of (1) data quality and availability, (2) model assumptions, and (3) historical findings.

Regarding historical data limitations, accurate confidence interval calculations rely on high-quality historical data. If the data is incomplete, outdated, or contains errors, it can lead to unreliable VaR estimates and confidence intervals.

Regarding the choice of distribution, selecting the correct probability distribution for modeling changes in portfolio values is critical. Applying an incorrect distribution can lead to significant errors in confidence interval estimates.

Confidence intervals are not symmetric, and using more data will lead to tighter confidence intervals (i.e., more precise estimates). Using GARCH VaR of the historical profit and loss (P&L) of banks also tends to produce tighter confidence intervals compared to historical simulation VaR. During stress periods, historical simulation VaR produces wider confidence intervals (i.e., less precise estimates).

Benchmarking VaR Models

Benchmarking VaR involves comparing the performance of the bank's VaR model to another VaR model (i.e., a benchmark). This is usually done for a short time period when the bank is planning on transitioning to a new model. Benchmarking VaR models is crucial for validating their performance and ensuring that they provide accurate risk assessments. Formal benchmarking evaluation involves statistical tests.

Issues with statistical tests that compare two VaR models suggest the following:

- Because trading portfolios change frequently, errors are not independent and identically distributed, especially for regression-based results.
- Banks often do not have two VaR models to compare. Instead, regular backtesting of VaR models against actual P&L outcomes can serve as a robust benchmarking tool.

Beyond Exceedance-Based Backtesting of Value-at-Risk Models: Methods for Backtesting the Entire Forecasting

Distribution Using Probability Integral Transform

Cross-reference to GARP assigned reading – Iercosan et al., Chapter 4.

Exceedance-Based Approach

The most prevalent backtesting framework is the **exceedance-based backtest**, which counts the number of times actual losses exceed the VaR threshold. For example, we would expect that a 95% VaR will be exceeded 5% of the time. In other words, given 100 observations, we would expect five **exceedances**.

For **unconditional coverage**, if the *actual* number of exceedances is substantially different from the *expected* number of exceedances, we might have reason for concern. For example, if there are very few exceedances over many observations, we might suspect that the VaR model is *overestimating risk*. On the other hand, if there are a relatively large number of exceedances, it might indicate that the VaR model is *underestimating risk*.

We also need to understand whether the exceedances are **independent** from each other or if clustering is occurring. Does the probability of having an exceedance increase if we just had an exceedance the day before? If so, this is problematic, because an accurate model would have adjusted the forecasted distribution to arrive at a VaR that would be exceeded with the expected probability (e.g., 5% probability for a 95% VaR).

Finally, we can combine the unconditional coverage property with the independence property to check if there is **conditional coverage**. In other words, are we correct in saying that our model predicts the rate of exceedances—and are those exceedances independent?

Probability Integral Transforms

One criticism of the exceedance-based approach is that it focuses exclusively on a preselected distribution tail cutoff. In contrast, the **probability integral transform (PIT) approach** analyzes how well the VaR model does at predicting the entire distribution of outcomes. The PIT approach moves beyond a simple counting of exceedances by using more sophisticated mathematical techniques and frameworks.

The PIT approach steps are outlined as follows:

Step 1: Obtain a forecast of the distribution in the period T.

Step 2: Collect the observation for period T.

Step 3: Convert the observation into an "equivalent" one from a uniform distribution [0,1] by locating where the observation falls on the cumulative distribution function (CDF). This uniform distribution location is the PIT value for the observation.

Step 4: Analyze the distribution of PITs for uniformity, independence, and conditional coverage.

Distribution of PITs

Given a large number of PITs, we construct a uniform distribution between 0 and 1 that would also have independence among observations. In other words, the same percentage of PITs would be present in each distribution range. This is a generalization of exceedance-based backtests. In contrast, PIT-based backtests view the distribution more broadly. An imperfect VaR model might indicate that the PITs are forecasting distributions that underestimate losses or gains in some places and overestimate losses or gains in other places.

It is often more useful to construct a **quantile-quantile (QQ) plot**, which contrasts two probability distributions against each other. In this instance, we plot the portion of the PITs that are less than or equal to each value from 0 to 1 and contrast this with a perfect distribution, which would just be a straight line on the QQ plot.

Evaluating PIT Distributions

Goodness-of-fit tests for evaluating the distribution of PITs include the Kolmogorov-Smirnov (KS) test, the Anderson-Darling (A-D) test, and the Cramér-von Mises (CVM) test. These approaches consider how the theoretical distribution differs from the actual distribution of PITs. They use QQ plots to compare the actual cumulative distributions to the theoretical straight line true uniform distribution.

The KS test calculates the maximum vertical distance between the theoretical distribution and the actual distribution. The A-D test places higher emphasis on tail observations, making it a more appealing option for risk managers. The CVM test is very similar to the KS test, but instead of looking at the maximum difference, it considers the mean squared deviation.

STUDY SESSION 2: CORRELATION RISK

Correlation Basics: Definitions, Applications, and Terminology

Cross-reference to GARP assigned reading – Meissner, Chapter 1.

Correlation Risk

Correlation risk measures the risk of financial loss resulting from adverse changes in correlations between financial or nonfinancial assets. An example of financial correlation risk is the negative correlation between interest rates and commodity prices. If interest rates rise, losses occur in commodity investments.

Nonfinancial assets can also be impacted by correlation risk. For example, the correlation of sovereign debt levels and currency values can result in financial losses for exporters.

Correlation in Trading With Multi-Asset Options

Correlation trading strategies involve trading assets that have prices determined by the co-movement of one or more assets over time. **Correlation options** have prices that are very sensitive to the correlation between two assets and are often referred to as *multi-asset options*. For almost all correlation option strategies, a lower correlation results in a higher option price.

Correlation Swap

A correlation swap is used to trade a fixed correlation between two or more assets with the correlation that actually occurs. The correlation that will actually occur is unknown and is referred to as the *realized* or *stochastic correlation*.

The following equation calculates the realized correlation that actually occurs over the time period of the swap for a portfolio of *n* assets, where $\rho_{i,j}$ is the correlation coefficient:

$$\rho_{\text{realized}} = \frac{2}{n^2 - n} \sum_{i>j} \rho_{i,j}$$

The payoff for an investor buying a correlation swap is calculated as follows:

$$\text{notional amount} \times (\rho_{\text{realized}} - \rho_{\text{fixed}})$$

Role of Correlation Risk in Other Types of Risk

Given that correlation risk refers to the risk that the correlation between assets changes over time, the concern is how the covariance matrix used for calculating value at risk (VaR) or expected shortfall (ES) changes over time due to changes in *market risk*.

Risk managers are also concerned with measuring *credit risk* with respect to migration risk and default risk. **Migration risk** is the risk that the quality of a debtor decreases following the lowering of quality ratings. Lower debt quality ratings imply higher default probabilities.

Default correlation is of critical importance to financial institutions in quantifying the degree to which defaults occur at the same time. A lower default correlation is associated with greater diversification of credit risk.

Concentration risk is the financial loss that arises from the exposure to multiple counterparties within a specific group. Concentration risk is measured by the **concentration ratio**. A lower (higher) concentration ratio reflects that the creditor has more (less) diversified default risk.

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

Cross-reference to GARP assigned reading – Meissner, Chapter 2.

Correlations During Different Economic States

An empirical investigation of the correlations among the 30 common stocks in the Dow Jones Industrial Average (Dow) was conducted over the period from 1972 to 2017.

As expected, correlations were highest during recessions, when common stocks in equity markets tend to go down together. The low correlation levels during an expansionary period suggest common stock valuations are determined more on industry and company-specific information rather than macroeconomic factors.

The correlation volatilities during a recession, normal period, and expansionary period were 80.5%, 83.0%, and 71.2%, respectively. These results may seem a little surprising at first as one might expect volatilities to be highest during a recession. However, there is perhaps slightly more uncertainty in a normal economy regarding the overall direction of the stock market.

Mean Reversion and Autocorrelation

Mean reversion implies that over time, variables or returns regress back to the mean or average return. Empirical studies reveal evidence that bond values, interest rates, credit spreads, stock returns, volatility, and other variables are mean reverting.

The **mean reversion rate** is the degree of the attraction back to the mean and is also referred to as the speed or gravity of mean reversion. The mean reversion rate, *a*, is expressed as follows:

$$S_t - S_{t-1} = a(\mu - S_{t-1})$$

When a regression is run where $S_t - S_{t-1}$ (i.e., the *Y* variable) is regressed with respect to S_{t-1} (i.e., the *X* variable), we can think of the slope coefficient of the regression as being equal to the negative of the mean reversion rate, *a*.

Autocorrelation measures the degree to which a current variable value is correlated to past values. Autocorrelation is often calculated using an *autoregressive conditional heteroskedasticity (ARCH) model* or a *generalized autoregressive conditional heteroskedasticity (GARCH) model*. An alternative approach to measuring autocorrelation is running a regression equation. Autocorrelation has the exact opposite properties of mean reversion.

For example, equity correlations show high mean reversion rates (78%) and low autocorrelations (22%). These two rates must sum to 100%. Bond correlations and default probability correlations show much lower mean reversion rates and higher autocorrelation rates.

Best-Fit Distributions for Correlations

Based on the results of the Kolmogorov-Smirnov, Anderson-Darling, and chi-squared distribution fitting tests, the **Johnson SB distribution** (which has two shape parameters, one location parameter, and one scale parameter) provided the best fit for equity correlations.

The best-fit distribution for bond correlations was found to be the **generalized extreme value (GEV) distribution**. However, the normal distribution is also a good fit for bond correlations.

The default probability correlation distribution was similar to equity distributions in that the Johnson SB distribution is the best fit for both distributions. Figure 1.1 summarizes the findings of the empirical correlation analysis.

Figure 1.1: Empirical Findings for Equity, Bond, and Default Correlations

Correlation Type	Average Correlation	Correlation Volatility	Reversion Rate	Best Fit Distribution
Equity	35%	80%	78%	Johnson SB
Bond	42%	64%	26%	Generalized Extreme Value
Default Probability	30%	88%	30%	Johnson SB

Financial Correlation Modeling—Bottom-Up Approaches

Cross-reference to GARP assigned reading – Meissner, Chapter 5.

Copula Functions

A **correlation copula** is created by converting two or more unknown distributions that may have unique shapes and mapping them to a known distribution with well-defined properties, such as the normal distribution. A copula creates a joint probability distribution between two or more variables while maintaining their individual marginal distributions. This is accomplished by mapping multiple distributions to a single multivariate distribution.

Gaussian Copula

A **Gaussian copula** maps the marginal distribution of each variable to the standard normal distribution, which, by definition, has a mean of zero and a standard deviation of one. The mapping of each variable to the new distribution is done on a percentile-to-percentile basis.

For example, if the unique marginal distributions of X and Y are not well-behaved structures, it is difficult to define a relationship between the two variables. However, the standard normal distribution is a well-behaved distribution. Therefore, a copula is a way to indirectly define a correlation relationship between two variables when it is not possible to directly define a correlation.

STUDY SESSION 3: TERM STRUCTURES AND VOLATILITY MODELING

Regression Hedging and Principal Component Analysis

Cross-reference to GARP assigned reading – Tuckman and Serrat, Chapter 6.

DV01-Neutral Hedge

A standard DV01-neutral hedge assumes that the yield on a bond and the yield on a hedging instrument rise and fall by the same number of basis points. However, a one-to-one relationship does not always exist in practice.

To improve this DV01-neutral hedge approach, we can apply regression analysis techniques. Using a regression hedge examines the volatility of historical rate differences and adjusts the DV01 hedge accordingly, based on historical volatility.

Regression Hedge

The advantage of a regression framework is that it provides an estimate of a hedged portfolio's volatility. An investor can gauge the expected gain in advance and compare it to historical volatility to determine whether the hedged portfolio is an attractive investment.

Hedge Adjustment Factor

The alpha and beta coefficients of a least squares regression line will be determined by the line of best fit through historical yield data points.

$$\Delta y_t^N = \alpha + \beta \Delta y_t^R + \varepsilon_t$$

where:

Δy_t^N = changes in the nominal yield

Δy_t^R = changes in the real yield

Defining F^R and F^N as the face amounts of the real and nominal bonds, respectively, and their corresponding DV01s as $DV01^R$ and $DV01^N$, a DV01 hedge is adjusted by the hedge adjustment factor, or beta, as follows:

$$F^R = F^N \times \left(\frac{DV01^N}{DV01^R}\right) \times \beta$$

Two-Variable Regression Hedge

Regression hedging can also be conducted with two independent variables. For example, assume a trader in euro interest rate swaps buys/receives the fixed rate in a relatively illiquid 20-year swap and wishes to hedge this interest rate exposure. In this case, a regression hedge with swaps of different maturities would be appropriate. Since it may be impractical to hedge this position by immediately selling 20-year swaps, the trader may choose to sell a combination of 10- and 30-year swaps.

Level and Change Regressions

Using a single-variable approach, the formula for a change-on-change regression with dependent variable *y* and independent variable *x* is as follows:

$$\Delta y_t = \alpha + \beta \Delta x_t + \Delta \varepsilon_t$$

where:

$\Delta y_t = y_t - y_{t-1}$

$\Delta x_t = x_t - x_{t-1}$

Alternatively, the formula for a level-on-level regression is as follows:

$$y_t = \alpha + \beta x_t + \varepsilon_t$$

Reverse Regressions

In regression analysis, reversing the dependent and independent variables results in a technique called **reverse regression**. For a regression hedge, the yield change of a target bond (e.g., XYZ corporate bond) is regressed on the yield change of a hedging instrument (e.g., T-bond). In contrast, a reverse regression regresses changes in yield of the T-bond on changes in yield of the XYZ corporate bond.

The main difference in these hedging approaches is the varying amount of XYZ bonds needed to hedge the same amount of T-bonds under both approaches. The choice of which hedge to employ depends on the trader's objectives.

Principal Component Analysis

The advantage of principal component analysis (PCA) is that we only really need to describe the volatility and structure of the first three principal components (PCs) since the sum of the variances of the first three PCs is a good approximation of the sum of the variances of all rates.

For example, if we consider the set of swap rates from 1 to 30 years, at annual maturities, the PCA approach creates 30 interest rate factors or components, and each factor describes a change in each of the 30 rates. Changes in the 30 rates can now be expressed with changes in just three factors, which is a much simpler approach.

Arbitrage Pricing with Term Structure Models

Cross-reference to GARP assigned reading – Tuckman and Serrat, Chapter 7.

Interest Rate Tree (Binomial) Model

A **binomial model** is a model that assumes that interest rates can take only one of two possible values in the next period.

This interest rate model makes assumptions about interest rate volatility, along with a set of paths that interest rates may follow over time. This set of possible interest rate paths is referred to as an **interest rate tree**.

Constructing the Binomial Interest Rate Tree

The construction of an interest rate tree, binomial or otherwise, is a tedious process. In practice, the interest rate tree is usually generated using specialized computer software. There is one underlying rule governing the construction of an interest rate tree: *The values for on-the-run issues generated using an interest rate tree should prohibit arbitrage opportunities*. This means that the value of an on-the-run issue produced by the interest rate tree must equal its market price. It should be noted that in accomplishing this, the interest rate tree must maintain the interest rate volatility assumption of the underlying model.

Backward Induction

Backward induction refers to the process of valuing a bond using a binomial interest rate tree. The term "backward" is used because in order to determine the value of a bond at node 0, you need to know the values that the bond can take on at node 1. But to determine the values of the bond at node 1, you need to know the possible values of the bond at node 2, and so on. Thus, for a bond that has N compounding periods, the current value of the bond is determined by computing the bond's possible values at period N and working "backward" to node 0.

Know that *the value of a bond at a given node in a binomial tree is the average of the present values of the two possible values from the next period*. The appropriate discount rate is the forward rate associated with the node under analysis.

Using the 0.5 probabilities for up and down states as shown in the previous example may not produce an expected discounted value that exactly matches the market price of the bond. This is because the 0.5 probabilities are the assumed **true probabilities** of price movements. In order to equate the discounted value using a binomial tree and the market price, we need to use what is known as **risk-neutral probabilities**. Any difference between the risk-neutral and true probabilities is referred to as the **interest rate drift**.

Using the Risk-Neutral Interest Rate Tree

There are actually two ways to compute bond and bond derivative values using a binomial model. These techniques are referred to as **risk-neutral pricing**.

- The first method is to start with spot and forward rates derived from the current yield curve and then *adjust the interest rates* on the paths of the tree so that the value derived from the model is equal to the current market price of an on-the-run bond (i.e., The tree is created to be "arbitrage free.").
- The second method is to take the rates on the tree as given and then *adjust the probabilities* so that the value of the bond derived from the model is equal to its current market price.

There are three basic steps to valuing an option on a fixed-income instrument using a binomial tree:

Step 1: Price the bond value at each node using the projected interest rates.

Step 2: Calculate the intrinsic value of the derivative at each node at maturity.

Step 3: Calculate the expected discounted value of the derivative at each node using the risk-neutral probabilities and working backward through the tree.

Fixed-Income Securities and Black-Scholes-Merton

The Black-Scholes-Merton model is the best-known equity option-pricing model. Unfortunately, the model is based on three assumptions that do not apply to fixed-income securities:

1. The model's main shortcoming is that it assumes there is no upper limit to the price of the underlying asset. However, bond prices do have a maximum value. This upper limit occurs when interest rates equal zero so that zero-coupon bonds are priced at par and coupon bonds are priced at the sum of the coupon payments plus par.
2. It assumes the risk-free rate is constant. However, changes in short-term rates do occur, and these changes cause rates along the yield curve and bond prices to change.
3. It assumes bond price volatility is constant. With bonds, however, price volatility decreases as the bond approaches maturity.

Bonds With Embedded Options

Fixed-income securities are often issued with **embedded options**, such as a call feature. In this case, the price-yield relationship will change, and so will the price volatility characteristics of the issue.

Callable Bonds

For an option-free noncallable bond, prices will fall as yields rise, and prices will rise unabated as yields fall—in other words, they'll move in line with yields. That's not the case, however, with **callable bonds**. The decline in callable bond yield will reach the point where the rate of increase in the price of the callable bond will start slowing down and eventually level off.

This is known as **negative convexity**. Such behavior is due to the fact that the issuer has the right to retire the bond prior to maturity at some specified call price. The call price, in effect, acts to hold down the price of the bond (as rates fall) and causes the price-yield curve to flatten.

Putable Bonds

The put feature in **putable bonds** is another type of embedded option. The put feature gives the bondholder the right to sell the bond back to the issuer at a set price (i.e., the bondholder can "put" the bond to the issuer).

Expectations, Risk Premium, Convexity, and the Shape of the Term Structure

Cross-reference to GARP assigned reading – Tuckman and Serrat, Chapter 8.

Interest Rate Expectations and Volatility

Expectations of future interest rates are based on uncertainty. For example, an investor may expect that interest rates over the next period will be 8%. However, this investor may realize there is also a high probability that interest rates could be 7% or 9% over the next period.

Expectations play an important role in determining the shape of the yield curve and can be illustrated by examining yield curves that are flat, upward-sloping, and downward-sloping.

If the expected 1-year spot rates for the next three years are r_1, r_2, and r_3, then the 2-year and 3-year spot rates are computed as:

$$\hat{r}(2) = \sqrt[2]{(1 + r_1)(1 + r_2)} - 1$$
$$\hat{r}(3) = \sqrt[3]{(1 + r_1)(1 + r_2)(1 + r_3)} - 1$$

Interest rate expectations can describe the level of interest rates for long-term horizons.

When there is uncertainty regarding expected rates, the volatility of expected rates causes the future spot rates to be lower. With the implied rate, we can compute the value of convexity.

Convexity Effect

The convexity effect can be measured by applying a special case of Jensen's inequality as follows:

$$E\left[\frac{1}{(1 + r)}\right] > \frac{1}{E[1 + r]}$$

All else held equal, the value of convexity increases with maturity. In other words, as the maturity of a bond increases, the price-yield relationship becomes more convex.

This convexity occurs due to volatility. Thus, we can also say that the value of convexity increases with volatility.

Decomposition of Bond Returns

The return on a bond can be decomposed into three components:

1. *Forward rate*. The forward rate captures the return from the passage of time.
2. *Impact of changes in interest rates*. Bond returns move inversely with rate changes in proportion to duration.
3. *Impact of changes in interest rate volatility*. Bond returns increase in proportion to convexity.

Mathematically, the bond return decomposition can be expressed as follows:

$$\text{return} = f(t) - D\Delta r + \frac{1}{2} C \sigma_r^2$$

where:
f(t) = forward rate
D = duration
Δr = change in interest rates
C = convexity
σ = volatility of interest rates

For *risk-neutral investors,* a bond's expected return is equal to the short-term risk-free rate, r_0. In contrast, *risk-averse investors* demand additional compensation (i.e., a risk premium, λ) for bearing interest rate risk. In this case, the bond's expected return is computed as follows:

$$E(\text{return}) = r_0 + \lambda D$$

For example, consider a bond with a short-term rate of 2.5%, a duration of 7, and a risk premium of 15 basis points per unit of duration risk. For risk-neutral investors, the expected return would simply be 2.5%. However, for risk-averse investors who require compensation for interest rate risk, the expected return would increase to 2.5% + (7 × 0.15%) = 3.55%.

The Art of Term Structure Models: Drift

Cross-reference to GARP assigned reading – Tuckman and Serrat, Chapter 9.

Term Structure Model With No Drift (Model 1)

Model 1 is used in cases where there is no drift and interest rates are normally distributed. The continuously compounded instantaneous rate, denoted r_t, will change (over time) according to the following relationship:

$$dr = \sigma dw$$

where:
dr = change in interest rates over small time interval, dt
dt = small time interval (measured in years) (e.g., one month = 1/12)
σ = annual basis-point volatility of rate changes
dw = normally distributed random variable with mean 0 and standard deviation $\sqrt{dt}$

In Model 1, since the expected value of *dw* is zero [i.e., E(dw) = 0], the drift will be zero. Also, since the standard deviation of dw = $\sqrt{dt}$, the volatility of the rate change = $\sigma\sqrt{dt}$. This expression is also referred to as the standard deviation of the rate.

Figure 1.2: Interest Rate Tree With No Drift

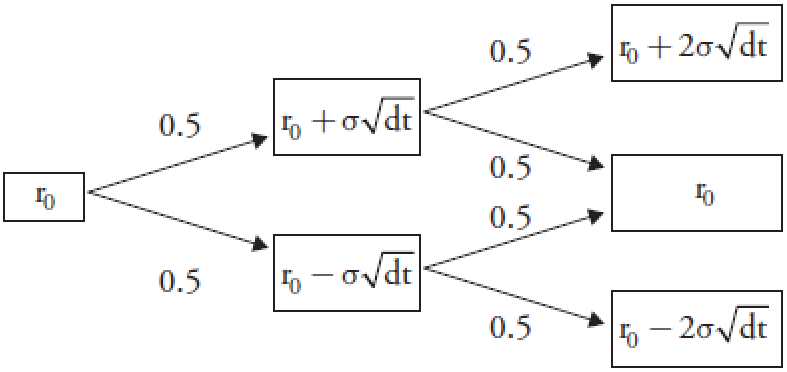

There are two reasonable solutions for negative interest rates. First, the model could use distributions that are always non-negative, such as lognormal or chi-squared distributions. Second, the interest rate tree can "force" negative interest rates to take a value of zero. This method may be preferred over the first method because it forces a change in the original distribution only in a very low interest rate environment whereas changing the entire distribution will impact a much wider range of rates.

Term Structure Model With Drift (Model 2)

A natural extension to Model 1 is to add a positive drift term that can be economically interpreted as a positive risk premium associated with longer time horizons. We can augment Model 1 with a constant drift term, which yields Model 2:

$$dr = \lambda dt + \sigma dw$$

Figure 1.3: Interest Rate Tree With Constant Drift

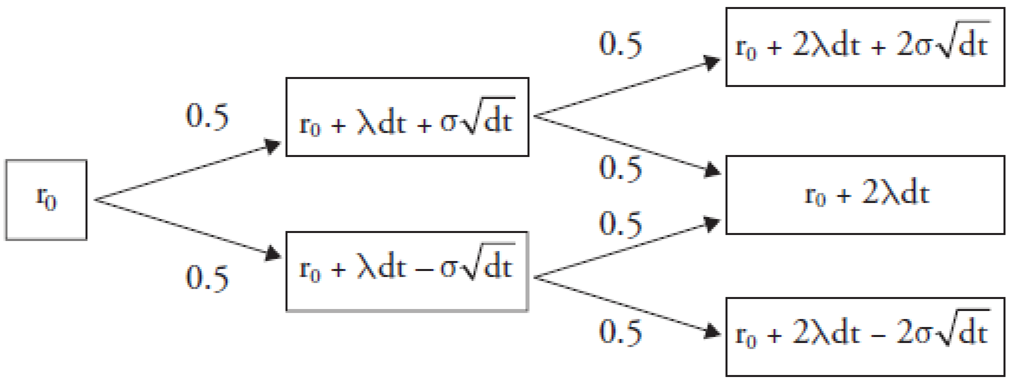

Ho-Lee Model

The Ho-Lee model further generalizes the drift to incorporate time-dependency. That is, the drift in Time 1 may be different from the drift in Time 2; additionally, each drift does not have to increase and can even be negative. Thus, the model is more flexible than the constant drift model.

Arbitrage-Free Models

Arbitrage models are often used to quote the prices of securities that are illiquid or customized. For example, an arbitrage-free tree is constructed to properly price on-the-run Treasury securities (i.e., the model price must match the market price). Then, the arbitrage-free tree is used to predict off-the-run Treasury securities and is compared to market prices to determine if the bonds are properly valued. These arbitrage models are also commonly used for pricing derivatives based on observable prices of the underlying security (e.g., options on bonds).

Vasicek Model

The **Vasicek model** assumes a mean-reverting process for short-term interest rates. The underlying assumption is that the economy has an equilibrium level based on economic fundamentals such as long-run monetary supply, technological innovations, and similar factors. Therefore, if the short-term rate is above the long-run equilibrium value, the drift adjustment will be negative to bring the current rate closer to its mean-reverting level.

The formal Vasicek model is expressed as follows:

$$dr = k(\theta - r)dt + \sigma dw$$

where:
k = parameter that measures the speed of reversion adjustment
θ = long-run value of the short-term rate assuming risk neutrality
r = current interest rate level

The risk neutrality assumption of the long-run value of the short-term rate allows θ to be approximated as:

$$\theta \approx r_l + \frac{\lambda}{k}$$

where:
r_l = long-run true rate of interest

The Art of Term Structure Models: Volatility and Distribution

Cross-reference to GARP assigned reading – Tuckman and Serrat, Chapter 10.

Term Structure Model With Time-Dependent Volatility

The relationships between volatility in each period could take on an almost limitless number of combinations. For example, the volatility of the short-term rate in one year, σ(1), could be 220 basis points and the volatility of the short-term rate in two years, σ(2), could be 260 basis points. It is also entirely possible that σ(1) could be 220 basis points and σ(2) could be 160 basis points. To make the analysis more tractable, it is useful to assign a specific parameterization of time-dependent volatility. Consider the following model, which is known as Model 3:

$$dr = \lambda(t)dt + \sigma e^{-\alpha t}dw$$

where:
σ = volatility at $t = 0$, which decreases exponentially to 0 for $\alpha > 0$

Cox-Ingersoll-Ross (CIR) and Lognormal Models

A popular model where the basis-point volatility (i.e., annualized volatility of *dr*) increases proportional to the square root of the rate (i.e., $\sigma\sqrt{r}$) is the **Cox-Ingersoll-Ross (CIR) model** where *dr* increases at a decreasing rate and σ is constant. The CIR model is shown as follows:

$$dr = k(\theta - r)dt + \sigma\sqrt{r}\, dw$$

A second common specification of a model where basis-point volatility increases with the short-term rate is the **lognormal model** (Model 4). An important property of the lognormal model is that the yield volatility, σ, is constant, but basis-point volatility increases with the level of the short-term rate. Specifically, basis-point volatility is equal to σr and the functional form of the model, where σ is constant and *dr* increases at σr, is:

$$dr = ardt + \sigma rdw$$

For both the CIR and the lognormal models, as long as the short-term rate is not negative then a positive drift implies that the short-term rate cannot become negative.

The Vasicek and Gauss+ Models

Cross-reference to GARP assigned reading - Tuckman and Serrat, Chapter 9.

Vasicek Model

The **Vasicek model** uses a single factor to model how short-term interest rates change over time. Its main characteristic is mean reversion where the short-term rate (r_t) tends to revert to its long-term average (θ).

The Vasicek model is useful for simple applications like pricing zero-coupon bonds and applying basic hedging techniques. However, it has the following drawbacks:

- It fails to replicate complex term structure shapes and real-world volatility patterns.
- It fails to account for observed changes in volatility across maturities.
- It fails to adequately incorporate macroeconomic trends and monetary policy shifts.

Gauss+ Model

Compared to the Vasicek model, the **Gauss+ model** is a more sophisticated multifactor framework that provides a better representation of real-world interest rate behaviors. The model consists of three interacting components: the short-term rate (r_t), a medium-term interest rate factor (m_t), and a long-term interest rate factor (l_t). These factors affect each other in a cascading manner in the following expressions:

$$dr_t = -\alpha_r(r_t - m_t)dt$$

$$dm_t = -\alpha_m(m_t - l_t)dt + \sigma_m(\rho dw_t^1 + \sqrt{1-\rho^2}dw_t^2)$$

$$dl_t = -\alpha_l(l_t - \mu)dt + \sigma_l dw_t^1$$

where:

$\alpha_r, \alpha_m, \alpha_l$ = positive mean reversion parameters

μ = constant term that includes rate expectations and risk premiums

Regarding the first of the three cascade equations, when $r_t < m_t$, the right-hand side of the equation suggests that the rate will be increasing as we approach the medium-term factor. When $r_t > m_t$, the rate will be decreasing as we approach the medium-term factor. When $r_t = m_t$, this is the equilibrium case, and there will be no change to the short-term rate. This process results in a cascade of convergence where r_t is seeking to converge to m_t, which is trying to converge to l_t, which is trying to converge to μ.

There are economic interpretations to these three rate factors. The short-term rate represents the policy interest rate controlled by central banks (e.g., the rate set by the Federal Reserve). The medium-term factor represents influences from the broader economy, such as business cycles and monetary policy effects. The long-term factor represents long-term expectations for economic trends, such as inflation and productivity growth. Accordingly, the short-term rate reacts quickly to the medium-term factor, the medium-term factor adjusts more gradually to the long-term factor, and the long-term factor shifts the slowest over time.

Vasicek Model vs. Gauss+ Model

Figure 1.4 summarizes some differences in dynamics, features, and applications between the Vasicek and Gauss+ models.

Figure 1.4: Differences Between Vasicek and Gauss+ Models

Aspect	Vasicek Model	Gauss+ Model
Dynamics	Single-factor model	Three-factor model
Mean reversion	Long-run average, θ	Cascading mean reversion to μ
Volatility	Constant	Hump-shaped volatility term structure
Historical alignment	Limited	High; reflects observed rate term structure shapes and behaviors
Applications	Pricing simple bonds; basic hedging	Interest rate derivatives, risk management, macroeconomic analysis

Volatility Smiles and Volatility Surfaces