LAM Sample Detailed Study Manual

You have downloaded a sample of our ILA-LAM detailed study manual. The full version covers the entire syllabus and is included with the online seminar.

Each portion of the detailed study manual is available in PDF with a clickable table of contents. Each reading (and sub-chapters if applicable) are also bookmarked in the PDF for ease of navigation in your favorite desktop, tablet, or smartphone PDF viewer.

If you have additional questions about the detailed study manual or any aspect of the exam, please email me.

Zak Fischer, FSA

zak@theinfiniteactuary.com

LAM Detailed Study Guide - Sample

LAM-117: Key Rate Durations

Thomas S.Y. Ho (1992)

Overview of This Reading

This was a seminal paper in the world of ALM and was the first to formalize the concept of key rate duration

Key rate durations involve relatively simple math, but they can be difficult to understand conceptually, especially within the relatively limited universe of information on the LAM syllabus

Even though the concepts may seem fairly academic and difficult to parse in the reading itself—especially if you don't have a background in ALM work—the framework is fairly straightforward

The video lesson for this reading takes a much more example-based approach to teaching the key concepts of this reading. If you can understand the basic idea of key rate durations through visuals and examples, there is actually very little to memorize, but the concepts are very far-reaching in terms of testability, so it's very important to spend time with this material and have a sound understanding of it.

Key topics for the exam include:

- Why effective duration is a poor measure of interest rate risk
- Definition of key rate duration
- 3 propositions for key rate durations
- Interpolating between key rates
- Using KRDs to build specific shift attributes
- 3 classic yield curve shifts—level, steepness, and curvature
- Analyzing the distribution of interest rate rate exposure
- KRD ratios in a portfolio context
- KRD profiles of various securities
    - Call provision
    - Callable bond coupon level
    - Sinking fund
    - Call vs. put
    - European call vs. American call
    - Mortgage pass-through
    - Principal- vs. interest-only strips

Introduction

Different types of fixed income securities are affected differently by changes in the yield curve

- Bonds with embedded call or put options
  - Sensitive to any shift of the yield curve, but tend to be more sensitive to non-parallel than parallel shifts
- Zero-coupon bond prices
  - Sensitive only to the shift in the yield curve rate at the maturity date
- Principal-only securities
  - More sensitive to rotational shifts of the yield curve

Understanding the interest rate risk exposure of different instruments is central to active and structured portfolio management

- Enables portfolio managers to better evaluate interest rate bets and control risk

Effective Duration

Historically, effective duration has been the most commonly used measure of interest rate risk exposure for interest rate-contingent claims

- Effective duration = ratio of the proportional change in the security value to the infinitesimal parallel shift of the spot curve
- Example:
  - If a bond has an effective duration of ten years, a 1 bps drop in the spot curve leads to a 0.10% rise in the bond value
- If a security's effective duration is high, the interpretation is that the security is exposed to significant interest rate risk

Note: Effective duration is not well covered on the current syllabus, and that makes learning KRD even harder. The video lesson begins with an example of calculating effective duration.

Effective duration is often not precise enough in many bond portfolio strategies

- Effective duration is based on a parallel shift of the yield curve ⇒ this rarely happens
- When hedging or immunization strategies based on effective duration fail, it is usually because the yield curve did not make a parallel shift

Solutions to address the effective duration disadvantage

- Estimate price sensitivity to various duration measures for specific different types of yield curve shifts
  - Some have developed durations for 3 types of common yield curve shifts in historical data:
    1. Level – parallel shift up or down

2. Steepness – long-term rates move more than short-term rates (or vice versa)
3. Curvature – long- and short-term rates move more than medium-term rates

Note: as the paper progresses, it's clear that there are no absolute definitions of the steepness and curvature movements, and the author acknowledges this. In general, "steepness" seems to refer to more of a change in the slope of the curve from short to long rates. "Curvature" can refer to up/down bends in middle portion of the curve or twists about the medium-term rate. Obviously, there are an infinite number of shapes a yield curve could take on, but these 3 are the primary shifts historically. This is discussed more in the video lesson.

- Use a duration vector for option-free bonds
  - This indicates that the interest rate risk exposure of a bond is better measured by a vector of sensitivities, not just a single one

Overview of Key Rate Durations (KRDs)

- These are new measures of interest rate risk exposure
- It is a vector representing the price sensitivity of a security to each key rate change

Advantages of Key Rate Durations

1. Identify the price sensitivity of an option-embedded bond to each segment of the spot yield curve
2. Recognizes that the yield curve movement is driven by multiple market factors
3. Easy to use to create a replicating portfolio of a bond with embedded options using zero-coupon bonds

Modeling Non-Parallel Yield Curve Shifts

Main assumption of effective duration: spot yield curve shift is parallel

- Can be misleading because the returns of 2 securities with the same effective duration can be significantly different if the yield curve undergoes non-parallel shifts, such as steepness or curvature

Key rate durations define the risk of the changing shape of the spot curve

- Not a single measure as is effective duration
- Define the price sensitivity of a security over the entire domain of possible movements of the yield curve
- $\sum$ KRDs = effective duration

Defining Key Rate Durations

1. Typically 11 key rates on the spot curve are chosen
   - 3-mo, 1-yr, 2-yr, 3-yr, 5-yr, 7-yr, 10-yr, 15-yr, 20-yr, 25-yr, and 30-yr

2. For each key rate, KRD = effective duration for that specific key rate
   - *For example, the KRD for year 10 predicts the % change in PV of a cash flow that will occur at year 10. If a security has additional cash flows at other years, their sensitivity is ignored in the 10th year KRD, which measures ONLY the sensitivity to a change in the 10-year rate. Therefore, a security with cash flows at different key rates, will have a KRD for each key rate, and the sum of all KRDs = total effective duration.*
3. Shifts for rates between the key rates (e.g. 4-year rate) can be approximated by linear interpolation of the shifts of each key rate
   - Linear interpolation is used for simplicity (more complex methods could be used but the author says it is unnecessary)
   - Enables the shifted yield curve to capture a curvature similar to the underlying yield curve
   - *Process is described next and also better illustrated in the video lesson*

Interpolating Between Key Rates

Basic key rate shift – for the ith key rate shift, the shift is zero for both maturities shorter than the $(i-1)$th key rate and longer than the $(i+1)$th key rate

> *This simply says, for example, when the 2-year rate shifts, the shift in all other key rates (before and after year 2) are zero*

For rates between the $(i-1)$th key rate and the ith key rate itself, the shift will grade linearly from zero and peak at the ith key rate

- Similarly, the shift will grade down linearly from the ith key rate shift to zero at key rate $i+1$

In formal notation, let $b[t;i,d(i)]$ be the ***i*th basic key rate shift** of term t (i.e. a specific yield curve shape at a specific point in time) and let $d(i)$ = the level of the shift to the new curve

$$b[t;1,d(1)] = \begin{cases} d(1) & \text{for } 0 < t < 0.25 \\ d(1)\frac{(1-t)}{(1.0-0.25)} & \text{for } 0.25 < t < 1.0 \\ 0 & \text{for } 1.0 < t < 30 \end{cases}$$

$$b[t;2,d(2)] = \begin{cases} 0 & \text{for } 0 < t < 0.25 \\ d(2)\frac{(t-0.25)}{(1.0-0.25)} & \text{for } 0.25 < t < 1.0 \\ d(2)\frac{(2-t)}{(2.0-1.0)} & \text{for } 1 < t < 2 \\ 0 & \text{for } 2.0 < t < 30 \end{cases}$$

A total yield curve shift is the vector sum of the 11 basic key rate shifts (including the interpolation affect on the neighboring "between" rates)

- In other words, the 11 independent basic key rate shifts together can approximate all the small yield curve shifts

- If $S[t; d(1), \ldots, d(11)]$ represents the vector of key rate shifts (the $d(i)$'s), then

$$S[t; d(1), \ldots, d(11)] = b[t; 1, d(1)] + \cdots + b[t; 11, d(11)]$$

Interpolation Example

Suppose there is a +100 bps shift in the 5-year key rate:

Maturity	Shift
0 – 3 yr	0
4 yr	+50
5 yr	**+100**
6 yr	+50
7 yr – 30 yr	0

The shift grades linearly from zero at key rate 3 to a peak of +100 at key rate 5, then grades down linearly to zero at key rate 7

You can also fill in at the monthly level. For example, the shift at maturity 3.5 (3 years and 6 months) would be +25 bps.

Note that a shift in key rate 3 would also impact rates for maturities between year 3 and 4. For a shift in multiple rates, the combined interpolation effects on rates between year 3 and 4 would be added to determine the final shift in the 4-year rate. This is the adding of the $b[\cdot]$ terms described above. This simply "connects the dots" between the key rates on the total shifted curve. For example, if the shock to key rate 3 is negative, this will pull the 3.5 rate down more, essentially putting it on a straight line between the shocked 3 and 5 year key rates. This is described more in the video lesson, which emphasizes learning the process more than memorizing the formulas.

Key Rate Durations

Defining the key rate duration $D(i)$:

- Let P be the initial security price
- Suppose the price is changed to P^* with the shift of the key rate
- Then the ith key rate duration $D(i)$ is defined as:

  $P^* - P = -P \times D(i) \times d(i)$

  where $d(i)$ is the shift defining the ith basic key rate shift

Note that the math here is identical to the way you calculated the change in price for effective duration

Approximating the shift of the yield curve

- You can approximate it by the sum of all the basic key rate shifts

$$P^* - P = -P \times D(1) \times d(1) - \cdots - P \times D(11) \times d(11)$$

- In other words, the total proportional change in price is the sum of the effect of each key rate change on the security

Proposition 1 on Key Rate Durations

- Key rate durations are a linear decomposition of effective duration
- Let D be the effective duration of a bond
- Let $D(i)$ be the ith key rate duration of the bond
- Then: $D = D(1) + \cdots + D(11)$

Numerical Example

There is a numerical example on p. 33 of the study note. This example is helpful if you can accept that the KRDs presented are not really KRDs but rather Macauley durations. I believe the author presented them this way for simplicity, but it causes a lot of confusion. Just accept that the KRDs presented are correct. The video lesson has a TIA-unique example involving KRDs that are more technically correct and internally consistent.

Key Rate Durations and Alternative Schemes

Suppose we define level, steepness, and curvature shifts in terms of a proportional change to the 30-year rate[1]

1. Level: $\{1, 1, 1\}$
2. Steepness: $\{-1, 0, 1\}$
3. Curvature: $\{1, 0, 1\}$

For example, if the 30-year rate increases 50 bps, key short-term rate must decrease 50 bps under the steepness definition above

We can express these 3 shifts as linear combinations of the key rate durations for the 3 key rates selected

$$\text{Duration (level shift)} = D(1) + D(2) + D(3)$$
$$\text{Duration (steepening shift)} = -D(1) + D(3)$$
$$\text{Duration (curvature shift)} = D(1) + D(3)$$

Since the steepness and curvature patterns defined above do not impact the medium-term rate, there is zero interest rate sensitivity to the medium-term rate

The best way to grasp this is to think of each $D(i)$ above as the value-weighted KRD of a portfolio with total effective duration D. The total effective duration corresponding to each shape shift is the weighted average

[1] As noted earlier, these are arbitrary definitions for steepness and curvature. There is obviously a lot of subjectivity in how these movements can be defined.

of the KRDs. If you wanted to develop total effective durations for specific types of non-parallel shifts, this would be a way to do it. Each effective duration would predict the change in the total portfolio for that shift.

The Link Between Key Rate Durations and Pricing Models

Pricing interest rate-contingent claims requires a model of term structure movements

Arbitrage-free (AR) models – Models that directly specify subsequent arbitrage-free rate movements relative to any initial spot yield curve

Key Rate Duration Profiles

On p. 35 of the study note, the author begins presenting a series of key rate duration profiles for various fixed income securities and options on fixed income securities. The key takeaways for each are noted below, but I highly recommend watching the video lesson, which takes a much more visual approach and makes key comparisons among the different profiles of the different securities. Once you understand the basic nature of securities with and without embedded options, the KRD profile becomes much more intuitive—as well as the implications for the different profiles. This means less to memorize for the exam.

Key rate duration profile: – a plot of the key rate durations against the terms of all the key rates

Zero-Coupon Bond (20 years)

- The simplest type of bond
- All the key rate durations are zero except for the key rate on the maturity date
- The KRD in this case equals the effective duration of the bond
- The key rate duration of a zero-coupon bond must lie on the 45° line

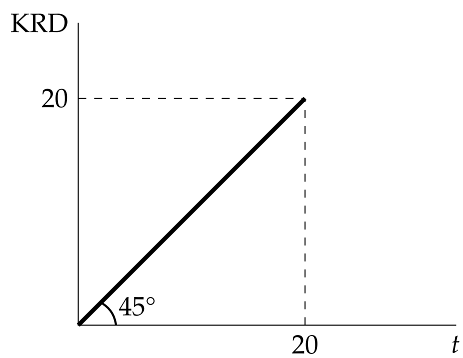

Another thing that is not terribly important for the exam, but creates a lot of confusion among first-time readers of this paper: The author is trying to simplify the presentation and assumes that the KRD of a zero coupon bond = maturity date. This is technically the Macauley duration, not an effective duration measure, however. The KRD would technically = modified duration = 20v, which would result in something a little smaller than a 45° angle.

Coupon Bond (30 year with 9% coupon Treasury bond)

- KRDs increase and decrease with the term, with a significantly large key rate duration at maturity
- Reason: Each coupon payment that is paid farther in the future has a higher duration but less present value

Callable Corporate Bond (30 yr with 9% coupon and 30 year with 8% coupon)

- Callability greatly changes the interest rate risk exposure of the bond
- Callability shortens the effective duration
- The callable bonds are more sensitive to the shorter-term key rate changes
- The callable bond with a lower coupon rate has less probability of being called → This means that the effective duration is higher and the bond also tends to be more sensitive to the longer-term key rate risk
- The increase in interest rate risk exposure is not uniform across the yield curve but is more substantial for the long-term rate

Callable Bond with a Sinking Fund (30 yr, 9% coupon with sinking fund provision)

- The sinking fund reduces the effective duration of the bond
- The sinking fund makes the bond more sensitive to the shorter-term key rates
- The sinking fund significantly reduces the long-term rate exposure since much of the portion of the bonds has been sunk over the years

European Call Option (30 yr, 9% coupon, expires in 10 years)

- The option is insensitive to the changes of any key rates with a term before the expiration date
- The key rate duration for the expiration date is <u>negative</u>
- A call option has a <u>negative</u> key rate duration
  - A call option is exposed much more to a curvature yield curve movement than to a parallel movement

European Put Option (30 year, 9% coupon, expires in 5 years)

- Has a positive key rate duration at expiration but negative key rate durations beyond the expiration date
- The key rate duration profile of a European put option is almost the mirror image of that of a European call option, but the magnitude of the key rate durations is not the same
- The relationship between the key rate durations of the call and the put can be derived from the put/call parity

American call option embedded in a callable bond (30 yr, 9% coupon)

- Can be viewed as a combination of a 9% coupon, 30 year option-free bond with shorting a call option
- The bond can be called any time
- The option is sensitive to the short-term key rate movements
- For American options there is no longer one expiration date, but instead a range of expiration dates
  - As a result, some of the key rate durations are negative for the key rates with terms less than ten years
  - The option is likely to be exercised before ten years to purchase the remaining life of a thirty-year coupon bond

GNMA mortgage pass-through (30 yr fixed rate 10% pool)

- The prepayment option significantly affects the key rate durations
- An unseasoned current coupon GNMA pass-through is most sensitive to the yield curve movement in the medium-term range
- An embedded option does not necessarily result in any negative key rate durations

Note that profile and optionality of this security is extremely similar to a callable bond with a sinking fund. The repayment of mortgage principal effectively "sinks" the mortgage, and prepayments are essentially "calls" to refinance.

Principal-Only and Interest-Only GNMAs

- The GNMA Pass-Through can be made up of the principal-only and interest-only strips
  - I*n other words, the principal portion of the mortgage payments is sold to investors separately from the interest portion of the mortgage payments. When added together, IOs and POs equal the performance of the total mortgage pool, but in isolation, IOs and POs behave very differently.*
- Principal-only
  - Pays the investor the principal part of the mortgage payments
  - Has negative KRDs up to the seventh year
  - Then the security has very high positive ten and fifteen year key rate durations
  - The key rate duration profile show the high interest rate risk exposure of a PO
  - Early payment of principal combined with the lower discount rate of the cash flow results in a much higher value of the PO
  - When interest rates rise, the prepayments slow down with the rise in the discount rate → the PO value falls rapidly which explains the high duration of a PO
  - A PO has high convexity so the PO can yield a high upside return with a fall in interest rates, while protecting against downside risk

- Interest-only strip
  - KRDs are positive up to the tenth year and then become negative
  - The KRD profile of an IO is the mirror image of the PO except for the magnitude
  - The IO value enhances with a rise in interest rates if the yield curve moves in a parallel fashion
  - The IO will behave like a regular bond if a curvature movement results in little or no yield curve movement beyond the 10 year maturity range

Summary of KRDs

- The KRD profiles of different types of securities show how each security type can be exposed to different interest rate risks
- Principal-only can have negative KRDs while it has high effective duration
- The KRD profiles of option-embedded bonds may change dramatically with changing interest rates

Key Rate Durations in a Portfolio Context

Extending Key Rate Durations from a single security to a portfolio:

- Understand the link between the KRDs of a single security vs. KRD of a portfolio
- The key rates are the zero-coupon rates determined from Treasury coupon issues while portfolio managers are concerned with yields of coupon issues

Key Rate Durations of a Portfolio

- Consider a portfolio of option-embedded bonds and option-free bonds
- The ith key rate duration of a portfolio is the sensitivity of the portfolio value to the shift in the ith key rate

Notation for a bond portfolio

- Let the bond portfolio consist of m bond positions
- Let the value of the jth bond position be $V(j)$
- The sum of all the bond position values must equal the portfolio value V
- The weight $w(j)$ is the proportion of the bond position to the portfolio value where:

$$w(j) = \frac{V(j)}{V} \text{ for } j = 1 \ldots m$$

Calculating the key rate durations of a portfolio
The KRD of a portfolio is the weighted sum of the KRDs of each bond position where you use the weights $w(j)$

$$D(i) = w(1)D(1, i) + w(2)D(2, i) + \cdots + w(m)D(m, i)$$

for each $i = 1 \ldots 11$

$D(i)$ is the portfolio's ith key rate duration

$D(j, i)$ is the jth bond position's ith key rate duration

Identifying Interest Rate Bets using Key Rate Durations

- Effective duration of the portfolio can be viewed as the total interest rate risk exposure
- Each key rate duration contributes to the total risk exposure
- You can think of the proportion of a KRD of the effective duration as the weight allocated to the risk exposure to that particular key rate

Holding a barbell portfolio[2]

- This is a portfolio position with high key rate durations for the short and long rates and low key rate durations for the medium term
- This is a bet that the short rate and the long rate fall relative to the medium-term rate

On p. 41, the study note provides a numerical example showing how you can construct a portfolio with different proportions of exposure so that the portfolio will out-perform a portfolio that simply invests equal proportions in all maturities. This is also covered through example in the video lesson.

Constructing Hedge Ratios for Option-Embedded Bonds

Constructing a portfolio of zero-coupon bonds

Do this to replicate option-embedded bonds using dynamic strategies

1. Let $D(i)$ be the portfolio's ith key rate duration
2. Calculate the weights as $W(i) = \frac{D(i)}{T(i)}$ for $i = 1 \ldots 11$ where $T(i)$ is the term of the ith key rate
3. Now define $W(0)$ such that $1 = W(0) + W(1) + \cdots + W(11)$
   - $W(0)$ can be seen as the balancing item
4. Now invest $\$VW(i)$ in the zero-coupon bond maturing on the term of the ith key rate and hold $\$VW(0)$ in cash

This results in a portfolio that has the same key rate durations and value to a benchmark

You now have $VW(0) + \cdots + VW(11) = V$

The Portfolio Duration $= \frac{VW(i)}{V} \times T(i) = W(i) \times T(i)$

See the video lesson for an example.

[2] A barbell portfolio is a portfolio made up of investments in very short and very long maturities, with few or no investments in medium-term maturities.

Main Uses of Key Rate Durations

- Can be used to measure the interest rate risk exposure of an interest rate-contingent claim
- Can be a useful tool in portfolio management where portfolios of option-embedded bonds can be replicated using option-free bonds by continually revising the hedge ratios
- Can be used with structured portfolio management techniques such as enhanced indexation and immunization
- Can quantify the interest rate bets and provide a procedure to control interest rate risk exposure
- Can be used to identify all sources of risk and return of a security or portfolio over a holding period
  - *Keep in mind that even though the author makes this statement, it is very much contradicted by other readings that cover concepts like convexity. A first-order duration measure can never tell the whole story of interest rate risk for insurance liabilities and assets, which can be highly convex.*

Appendix: Adjusted Key Rate Durations

Key rate durations can also be calculated relative to changes in specific coupon rates (e.g. on-the-run Treasury issues)—i.e. you don't have to just use the spot curve

The author calls these "adjusted key rate durations," but the math is identical to the framework in the main body

Economic Scenario Generators: A Practical Guide, Ch 9

Conning (July 2016)

Overview of This Reading

This chapter discusses the concept of arbitrage, and discusses the implications on pricing and arbitrage-free ESGs. In particular, it is seen that the price of a cash flow stream can be calculated as the risk-neutral discounted expected value of these cash flows.

To be mathematically rigorous, the concept of a trading strategy (seen in detail already in the Chin textbook) is used. It is seen that the price of a cash flow stream is equal to the price of any self-financing trading strategy that generates the same cash flow stream.

Key topics for the exam include:

- Describe an arbitrage-free model
- State how to price a security with well-defined cash flows using an arbitrage-free ESG
- Describe trading strategies, and explain their importance for risk-neutral pricing models
- Explain what is done when performing the *martingale test*

Chapter 9: Arbitrage-Free Modeling Considerations

Arbitrage Definitions

- The reading looks at a number of ways to define arbitrage:
  1. **Arbitrage** - the opportunity to make a riskless profit by exploiting price differences of identical or similar financial instruments, on different markets, or in different forms
  2. Arbitrage is an explicit trading strategy involving the simultaneous buying and selling of securities, leading to either of the following outcomes:
     - The trader would receive a positive cash flow today and no further liabilities during the rest of the trade horizon
     - The trader could enter into the trade at zero net cost and receive only nonnegative cash flows during the rest of the trade horizon, with a positive cash flow occurring with probability greater than zero
  3. A market is arbitrage-free if there do not exist any self-financing trading strategies that produce a risk-free gain

The existence of arbitrage is sometimes referred[3] to as a "free lunch"

[3] You may be familiar with the phrase "there's no such thing as a free lunch", the title of Milton Friedman's book.

9.1: Arbitrage-Free Financial Models

Arbitrage-Free Models & Connections to Pricing

- The theory of an arbitrage-free model requires that all assets in the model have **well-defined cash flow structures**
- Two types of trading strategies:
  - Modeled trading strategies - establish a mechanism of buying and selling securities
  - Self-financing trading strategies - do not require the injection of additional capital during their execution
- **Key Idea: Two identical cash flow streams ought to have the same price**
  - The price of a security can be determined based on a replicating portfolio. The replicating portfolio is comprised of securities with known prices and well-defined cash flows
    - In other words, pricing can be determined by first identifying a trading strategy that generates the same cash flows as the cash flows that need to be priced. The price of this cash flow stream is equal to the known prices of the securities that replicate that cash flow
  - If a model is arbitrage-free, then one can assign a price to a cash flow stream according to the principle that the price of a cash flow stream is equal to the price of any self-financing trading strategy that generates the same cash flow stream
  - In other words, pricing of a derivative security can be done as follows:
    1. Determine a trading strategy that generates the cash flows of the derivative security
    2. Use the known prices of traded securities to infer the arbitrage-free price of the derivative security
- A model is arbitrage-free if and only if there exists a change of measure[4] under which the discounted gains process for all traded assets, for any trading strategy, is a martingale
  - The reading does not prove this point, but it should be intuitive. In an arbitrage-free model, pricing can be obtained by discounting the future risk-neutral expected cash flows at the risk-free rate

Next, we will build the mathematical framework for pricing

First, we will introduce the key notation:

- $P(0, T)$: Time 0 price of a zero-coupon bond maturing for unit value at time T
- $B(t)$: Accumulated value of a bank account at time t
  - $B(t) = e^{\int_0^t r_u du}$

[4] Change of measure refers to a change in probability measure (e.g. switching to risk-neutral probabilities). P-measure uses real-world probabilities while Q-measure uses risk neutral probabilities.

- S_t: Price of non-dividend paying stock

The First Fundamental Theorem of Asset Pricing: Price must be given as the expected discounted future cash flows, where the expectation is taken with respect to a risk-neutral measure

$$\text{Bond Price: } P(0,T) = E^*\left[\frac{1}{B_T}\right] = E^*\left[e^{-\int_0^T r_u du}\right]$$

$$\text{Stock Price: } S_0 = E^*\left[\frac{1}{B_T}S_T\right] = E^*\left[e^{-\int_0^T r_u du}S_T\right]$$

Where in this reading, the (*) denotes that the expectation is taken using the risk-neutral measure[5]

9.2: A Practical Application

Consider a position of v_u shares of stock. Then the gain on the stock position is given by:

$$\text{Gains } = \int_0^t v_u dS_u + \int_0^t v_u dD_u$$

where the notation is defined as follows:

- v_u: number of units of stock in the trading strategy
- dS : change in value per unit stock position
- dD : cash flow (i.e. dividend) paid per unit stock position

Next, we will additionally allow a bond position in the portfolio. Let ϕ_u denote the number of units of bond in the trading strategy

A cash flow stream C_u is said to be **hedged** if there exists a trading strategy (ϕ_u, v_u) for all $u \in [0,T]$ such that both:

1. For all $t \in [0,T]$, cash flows can be exactly replicated by dynamically trading in the securities market:

$$\underbrace{\phi_t B_t + v_t S_t - \phi_0 B_0 - v_0 S_0}_{\text{Self-Financing Portfolio}} = \underbrace{\int_0^t \phi_u dB_u}_{\text{Bond}} + \underbrace{\int_0^t v_u dS_u + \int_0^t v_u dD_u}_{\text{Stock}} = \underbrace{\int_0^t dC_u}_{\text{Cash Flow Stream}}$$

2. All capital has been used up when the cash flow stream expires, so that $\phi_T B_T + v_T S_T = 0$

[5] Note that * and Q are both common ways to denote the use of a risk-neutral measure.

Now, remember that *the price of a cash flow stream is equal to the price of any self-financing trading strategy that generates the same cash flow stream*. Thus, the value of the cash flow stream can be equated to the initial cost of the trading strategy:

$$\text{Price}(\{C_u\}) = \phi_0 B_0 + v_0 S_0$$

However, it is often mathematically easier to write the price as the risk-neutral present value of cash flows:

$$\text{Price}(\{C_u\}) = E^{\mathbb{Q}}\left[\int_0^T \frac{1}{B_u} dC_u\right]$$

State Price Density

- The *state price density* values cash flow streams according to the following formula:

$$\text{Price } = E^{\mathbb{P}}\left[\int_0^T \rho_u dC_u\right]$$

- The state price density is a density function over states, in order to price assets under an arbitrage-free system and using real-world probabilities
- Defines security dynamics in a manner consistent with actual, real-world experience
- However, pricing is typically more complex, and so this method is used much less often than the risk-neutral approach

Completeness

- Not all cash flows can be created through dynamic trading
- Example: Catastrophe Bonds depend on variables that are not present in traditional bond/equity markets

9.3: Is the Arbitrage-Free Framework Necessary?

- Arbitrage-free models are a necessary requirement when one wants to apply an ESG to risk-neutral pricing problems. However, as the example below illustrates, there are situations where the real-world model is preferred
- Interest Rate Behavior Example:
  - Many interesting dynamics could be useful in better capturing ultra-low interest rates, but they are not easily brought into the arbitrage-free framework
  - Arbitrage-free models are subject to greater mathematically rigorous requirements[6] than real-world models.
  - This can give real-world models more flexibility to adequately model features such as ultra-low interest rates

[6] In particular, the requirement of no-arbitrage constrains relationships on how security prices should behave based on risk-neutral cash flow forecasts.

- If ESG models are applied in real-world modeling contexts, then these models should be able to capture return behavior and market dynamics, but do not necessarily have to be technically arbitrage-free
- Broadly speaking, one could construct a defensible ESG that is not technically arbitrage-free, so long as the model produces returns that relate to actual market behavior and respects general relationships between financial variables

- Arbitrage-free pricing constrains dynamics that can be used to drive risk, and may not always lead to tractable pricing formulas

The bottom line is that real-world models have their place too, and can be more useful for asking "what-if" questions. Real-world ESGs can provide further insight into risk management based on an analysis of the realistic distribution of interest rates

9.4: Martingale Tests

Martingale Test

- A common test that is applied to risk-neutral models is the *martingale test*
- The martingale test is based on the fact that in a risk-neutral framework, there is an explicit relationship between the price of a cash flow stream and its expected discounted future cash flows
- To perform the martingale test, the model needs to include sufficient output detail
- The martingale test only applies for risk-neutral / arbitrage-free models
- **Key Idea: The average discounted cash flow should (approximately, up to simulation error) equal the time zero value**

Examples

1. First, consider a security with no intermittent cash flows, such as a zero-coupon bond or non-dividend paying stock. The martingale test should confirm that the simulated price, discounted by the risk-neutral short rate, is approximately equal to the time zero asset price
2. Second, let us generalize to the case that there are intermittent cash flows. In this case, the martingale test will confirm that the expected value of the discounted dividend stream over the horizon plus the discounted value of the asset at the horizon equals the time zero asset price
3. If a total return index is based on a portfolio of traded assets, then the expected discounted value of that index at a given horizon will also be equal to the time zero value of the index

Martingale Test

You are given that the yield curve is flat and that the continuous interest rate is 2%.
Suppose that ABC Total Return Index currently equals $251.

Your model outputs the following statistics for ABC Total Return Index simulated two years ahead from today, based on 1000 simulations:

Statistic	Value
1st percentile	$97
50th percentile	$265
99th percentile	$387
Mean	$261

Assess whether the martingale test is passed based on the model output.

Solution:

We will test that the expected discounted value of that index at a given horizon will also be equal to the time zero value of the index:

$$261 \cdot e^{-.02 \cdot 2} = 250.77 \approx 251$$

Thus, we can see that the martingale test is passed. ✓

9.5: Chapter Summary

1. The concept of an arbitrage-free financial model is both reasonable and desirable, so long as we believe that markets are fairly efficient. Models that are not arbitrage-free can still be very useful but need to be interpreted carefully
2. The practitioner view on whether ESG models should be arbitrage-free has evolved:
   - Prior to the financial crisis, most practitioners would generally agree that ESG models should be arbitrage-free
   - Now, some practitioners want "flexibility" in their interest rate models, and one can build all sorts of interesting real-world interest rate dynamics, as long as one is unfettered by the absence of arbitrage
3. The First Fundamental Theorem of Asset Pricing implies that price must be given as the risk-neutral expected discounted future cash flows
4. The technical meaning of *arbitrage-free* is expressed in terms of trading strategies. This led to the conclusion that the price of a cash flow stream is equal to the price of any self-financing trading strategy that generates the same cash flow stream
5. An alternative approach is based on the concept of the *state price density*, which operates under the real-world measure, which means that the probability distributions that define all security dynamics are consistent with the way we actually experience them
   - However, in practice, the risk-neutral measure is more commonly applied, because the formula is simpler to use
6. Formulas for asset prices in terms of discounted expectations can be very useful for model building

- A standard example is the expression of zero-coupon bond prices in terms of the expectation of the integral of the "short rate"

7. The formula for asset prices in terms of discounted expectations also imposes strong restrictions on ESG models, so that an actuary cannot arbitrarily invent and impose a realistic asset price dynamic and still have an arbitrage-free model

8. The formula for asset prices in terms of discounted expectations leads to a martingale condition for asset prices. This martingale condition can be used as a test of ESG scenarios as a necessary condition for the scenarios to be risk-neutral. This test is widely applied by regulators and is generally referred to as "the martingale test"

    - Setting up a martingale test requires care to properly account for the nature of the cash flows for the asset being tested. The construction of the test differs for treasury bonds, equity returns and corporate bonds

Economic Scenario Generators: A Practical Guide, Ch 10

Conning (July 2016)

Overview of This Reading

This chapter largely repeats points made in prior chapters, but also goes into more detail on some points, such as risk-neutral scenarios and martingale tests

Key topics for the exam include:

- Explain what is meant by the "1=1" martingale test
- State when to use risk-neutral vs real-world scenarios
- State an example of when both risk-neutral and real-world scenarios may be used together

Chapter 10: The Role of Risk-Neutral Scenarios

10.1: What Does "Risk-Neutral" Mean?

The price of a traded security at any point in the simulation must be equal to the expected value of its future discounted gains process under the risk-neutral measure

Risk-Neutral Scenarios

- Difficult to interpret, not intended to conform to historical distributions
- Primary criterion is the *ability to fit today's market prices*
- Risk-neutral means the parameters are calibrated so that the expected discounted gains process is the same for all time horizons
- A risk-neutral scenario set can be interpreted as pricing scenarios, because it permits a series of cash flows, under varying economic scenarios, to be discounted at a risk-free rate to produce prices of the cash flows

The price of a zero-coupon bond is given by:

$$\text{Zero-Coupon Bond Price} = E^{\mathbb{Q}}\left[\exp\left(-\int_0^T r_u\,du\right)\right] = E^{\mathbb{Q}}\left[\frac{1}{B_T}\right]$$

Where, similar to Ch 9, B_T is the accumulated value of a bank account where $B_T = \exp\left(\int_0^T r_u\,du\right)$.

Next, denote V_t as the value of a portfolio generated by a self-financing trading strategy. Then, under the appropriate technical conditions, we have that:

$$\frac{V_0}{B_0} = \underbrace{\int_0^t v_u dM_u}_{\text{Stock Gains}} + \frac{V_t}{B_t}$$

Suppose we have that:

$$E^{\mathbb{Q}}\left(\int_0^t v_u dM_u\right) = 0$$

Then,

$$\boxed{\frac{V_0}{B_0} = E^{\mathbb{Q}}\left[\frac{V_t}{B_t}\right]}$$

This mathematically justifies the point that the price of a traded security at any point in the simulation must be equal to the expected value of its future discounted gains process under the risk-neutral measure. Note from the equation above, we can see the close connections with martingales[7]

Martingale Tests

- Key martingale property - the conditional expectation of the next value in the sequence is equal to the present observed value, given knowledge of all prior observed values
- Regulators often require martingale tests in order to validate model output

$$X_0 = E^{\mathbb{Q}}\left[\frac{1}{B_T} Q_T X_T\right] \approx \frac{1}{M}\sum_{i=1}^{M} \frac{1}{B_T(i)} Q_T(i) X_T(i)$$

where the notation is defined as follows:

- X : ex-dividend price process
- Q : units of the asset that are held at time T

Dividing the above equation by X_0 gives that:

$$1 = \frac{1}{M}\sum_{i=1}^{M} \frac{1}{B_T(i)} Q_T(i) \frac{X_T(i)}{X_0}$$

For the martingale test, the right hand side of the above equation is computed based on simulation results. The results should be close to 1 for the martingale test to pass

[7] See the end of the DSM for this reading for an overview of martingales.

Total Return Index Example

Consider performing the martingale test on a total return index. To perform the martingale test, one must have available a risk-neutral scenario set of the total return series for the index

The total return index I can be computed from the total returns on assets j as follows:

$$1 + R(t, t+s) = 1 + \sum_j \omega_j R^j(t, t+s)$$

$$I(t+s) = I(t) \cdot (1 + R(t, t+s))$$

where the notation is defined as follows:

- ω_j: portfolio weight in asset j
- $R^j(t, t+s)$: total return on asset j from time t to $t+s$
- $I(0) = 1$ is the initial value of the total return index

The martingale test equation is then:

$$1 = \frac{1}{M} \sum_{i=1}^{M} \frac{1}{B_T(i)} \frac{I_T(i)}{I_0} = \frac{1}{M} \sum_{i=1}^{M} \frac{I_T(i)}{B_T(i)}$$

Again, for the martingale test, the right hand side of the above equation is computed based on simulation results. The results should be close to 1 for the martingale test to pass

10.2: Applications of Martingale Tests in Practice

Suppose that the current price of an asset $X_0 = 1$. The martingale test will analyze the average projected value, discounted at the risk-free rate:

$$\bar{X}_0 = e^{-rT} \frac{1}{M} \sum_{i=1}^{M} X_T(i) \approx X_0$$

where $X_T(i)$ is the asset price at time T under simulation i. M is the total number of simulations

If the martingale test passes, we expect:

- The average discounted value should approximately equal 1, but does not need to exactly equal 1 because of sampling error

It can be helpful to run the martingale test a few different times, particularly for high-volatility assets. This will minimize the impact of unusual/biased draws

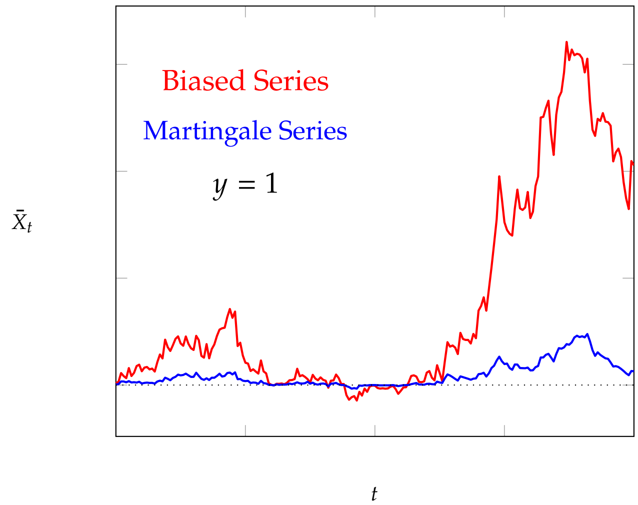

For example, in the above graph, we see that the martingale series (in blue) is close to 1 for all time horizons. These results seem reasonable and the martingale test passes.

Note that often, as in this example, the martingale test is set up as a "1=1" test, meaning that we are looking for the blue series to be fairly close to 1 to validate the martingale test

10.3: What is the Purpose of Risk-Neutral Scenarios?

The main purpose of risk-neutral scenarios is to *value cash flows*. Remember that the current price of an asset is the risk-neutral expected discounted value of the future cash flows. This holds even for more complex assets, as seen in the example below

Insurance Guarantee Example

Consider an insurance company that offers an investment guarantee with performance depending on a portfolio of several equity indices. Suppose the investment guarantee is path-dependent, and so it might depend on some history of the index returns

Then the price can be determined as:

$$E^{\mathbb{Q}}\left[\exp\left(-\int_0^T r_u\,du\right)\underbrace{\varphi(S_{t_1}, S_{t_2}, \ldots, S_{t_n})}_{\text{Value of Guarantee at Expiration}}\right]$$

Monte Carlo simulation can be used to compute the expression under the expectation for each scenario, and then average the results to approximate the current value of the insurance guarantee

Calibration

- When an ESG is utilized for risk-neutral simulation, the focus of the calibration will be on matching equity and swaption volatility surfaces, the initial yield curve, and reproduction of market prices in general

- The ESG should be calibrated to market price data as tightly as possible, so that the pricing of insurance liabilities and other contingent claims that are not priced in the market will have prices that reflect liquid market prices as accurately as possible

The reading makes the point that risk-neutral scenarios aren't necessary for deterministic modeling. This should be obvious, because the measure choice is irrelevant in the deterministic case.

10.4: Deciding Between Real-World and Risk-Neutral Scenarios

The table below summarizes key differences between risk-neutral and real-world scenarios:

Scenario Type	Typical Use	Calibration Method
Risk-Neutral	Pricing*	Pricing Data
Real-World	Risk Management	Historical Data or Specific View

*In particular, when no closed-form solution exists. Risk-neutral scenarios are useful for pricing complicated cash flows that depend upon stochastic financial variables

Risk-Neutral vs Real-World Scenarios

- The difference between risk-neutral scenarios and real-world scenarios is not the individual scenarios themselves; it is *the probability of those scenarios occurring*
- A risk-neutral scenario set weights the bad scenarios more than a real-world scenario set
- Risk-neutral scenarios often look unreasonable if viewed relative to historical interest rates or returns
- Risk-neutral scenarios embed the risk premiums[8] necessary to get appropriate market prices
- Real-world scenarios are applicable for measuring returns, volatilities and other risk exposures

[8] We understand this might be flipped from what you're traditionally used to, and we agree. But this is how the reading describes this point. They are viewing it from the perspective that risk neutral scenarios don't look like historical behavior, and there is an embedded difference that they are labeling as the risk premium. The full language from the reading is as follows: "Risk-neutral scenarios often look unreasonable if viewed relative to historical interest rates or returns. This is largely because risk-neutral scenarios embed the risk premiums necessary to get appropriate market prices."

10.5: When are Real-World and Risk-Neutral Scenarios Used Together?

Real-world and risk-neutral scenarios will often be applied together in applications that **combine risk management metrics with pricing**.

Closed-Form Pricing

- However, risk-neutral scenarios will **not** be necessary for risk management applications with closed-form pricing. In this case, a closed-form solution can be used in place of risk-neutral scenarios
- Risk-neutral scenarios will enter the picture when the real-world risk measurement exercises requires prices for liabilities or assets that do not have closed-form formulas

Examples

- Risk management of mortgage-backed securities (MBS)
- Nested stochastic projections for a variable annuity block
  - The risk management and hedging of variable annuities is a standard example in which real-world and risk-neutral scenarios are applied together
  - Computationally intensive
  - Real-world simulation is used to assess the overall risk of the variable annuity book and to measure the effectiveness of the hedging strategy
  - Since the liabilities associated with the variable annuity book are complicated, closed-form formulas are not available, and risk-neutral scenarios must be used to price the variable annuity book at each node of the real-world simulation

VA Example

A variable annuity block typically has exposure to equities, and equity risk is quantified by deltas (e.g. S&P delta, EAFE delta, etc.)

An insurer may want to simulate future paths of the S&P delta of its VA block (e.g. 2018 S&P delta, 2019 S&P delta, 2020 S&P delta) for risk management purposes.

The calculation *between* time steps would be real-world. For example, simulating the variable annuity block from 2018 to 2019 to 2020 would involve a real-world projection.

The calculation *at* each time step would be risk-neutral. For example, consider the calculation of the 2020 S&P delta. First, the VA block is projected from 2018 to 2020 using a real-world projection. Then, because delta calculations are closely related to pricing and Black-Scholes, the delta calculation at each particular time step is a risk-neutral calculation.

Thus, VA's are a common example of a product that relies on projections that use both risk-neutral and real-world scenarios together.

10.6: Chapter Summary

1. Risk neutral has a very specific meaning, which is inherently technical in nature but has an important practical manifestation that is the basis for generating pricing scenarios
2. Risk-neutral scenarios are typically used when one must price cash flows for which closed formulas are not available.
   - This situation is usually thought of in the context of Monte Carlo simulation
   - Although Monte Carlo simulation may be used to generate the risk-neutral scenarios, there is no other link between Monte Carlo and risk neutral
   - Market-consistent embedded value (MCEV), variable annuity risk management (i.e. exposure assessment and hedging) and the pricing of complicated cash flows are all examples in which risk-neutral scenarios are applied
3. The martingale test is a check to confirm that a set of scenarios are risk neutral
   - Often required by regulators
   - The implementation of the martingale test requires a careful assembly of the relevant risk-neutral variables
   - Volatile asset returns such as equity returns can produce martingale test results that are not perfect passes, even though the scenarios may be legitimate risk-neutral scenarios
4. Real-World vs Risk-Neutral Scenarios
   - Real-world scenarios are applicable for measuring returns, volatilities and other risk exposures. Risk-neutral scenarios are applicable for pricing cash flows
   - Real-world scenarios are parameterized to reflect historical benchmarks or clients' own views. Risk-neutral scenarios are calibrated to pricing data such as swaptions and index options
   - Risk-neutral scenarios often look unreasonable if viewed relative to historical interest rates or returns. This is largely because risk-neutral scenarios embed the risk premiums necessary to get appropriate market prices
   - Risk-neutral scenarios have a reputation of being difficult. This is largely the result of failing to understand the fundamental differences between real-world and risk-neutral scenarios
5. Using Both Real-World and Risk-Neutral Scenarios
   - Real-world and risk-neutral scenarios will often be applied together in applications that combine risk management metrics with pricing
   - Such applications typically involve nested stochastic calculations where the state of the world or node is simulated under the real-world measure and cash flows are priced in that state of the world using risk-neutral scenarios

Martingale Background

The reading does not explicitly include this section, but I thought it might be helpful to take a step back and give a basic overview of martingales.

What is a Martingale?

- A martingale is a process which is neither expected to grow nor diminish over time
  - The expected future value equals the current value[9]
  - Another way to define a martingale is "a stochastic process in which the conditional expectation of the next value, given the current and preceding values, is the current value"

Examples

- **Age:** Your age would be an example of something that is clearly not a martingale - your age grows over time! The expected future value is greater than the current value in this case, and so age is not a martingale
- **Constant:** Anything that is constant, by default, is going to be a martingale. For example, if you stick $100 in a lockbox undisturbed, the amount in your lockbox will always be a constant $100 and the amount of money in the lockbox is a martingale
- **World Population:** Assuming the world population grows over time, then the world population is not a martingale
- **Coin Flipping:** If you have W dollars currently and gain x dollars if a fair coin is flipped heads, or lose y dollars if it is flipped tails, your wealth is a martingale if x equals y. But if x is greater than y, this is not a martingale, since your wealth is expected to grow
- **Weather:** Temperature is an example of something that is not a martingale – if it is currently winter, we expect the temperature to be higher in the upcoming summer
- **Stocks:** Generally speaking, the value of a portfolio stocks is typically not a martingale because it is expected to grow over time (e.g. 5-7% expected growth per year)
  - For example, if you start with a portfolio of $100,000 and have an expected growth rate of 6% - then the expected portfolio value in a year (let's just use simple interest here) is approximately $106,000 – which is greater than the current value of $100,000

[9] Warning: Sometimes students get confused and think the "value" in the martingale definition is the stock value. Keep in mind it is the value of whichever process you are analyzing, which could have nothing to do with stocks. For example, you could analyze the population of the world over time – in this case, the expected future value is likely greater than the current value due to population growth. So, the value here is not necessarily the stock value. It could also even be a function of the stock value - so you could analyze a process defined as the discounted stock price.

- **Discounted Stocks:** The process of a discounted stock portfolio[10] under a risk-neutral measure[11] is a martingale
  - To understand this, first ignore the discounting piece. As we just saw, the stock portfolio is expected to grow at the risk-free rate under a risk-neutral setting
  - Then, applying discounting on top of this will result in a process that is neither expected to grow nor diminish over time
    - Since stocks are expected to grow at the risk-free rate and then we are also discounting at the risk-free rate, these two impacts "cancel out" (as seen mathematically below)
    - The future expected value equals the current value, where value is defined as the value of the discounted stock portfolio
  - Mathematically, the discounted stock portfolio value is $e^{-rt}S_t$, and if we are currently at time 0 and the future is time t, then:

$$\text{Expected Future Value} = E_0\left(e^{-rt}S_t\right) = e^{-rt}E_0\left(S_t\right) = e^{-rt}e^{rt}S_0 = S_0 = \text{Current Value}$$

[10] One of the most common examples of a martingale in finance is the risk-neutral discounted value of a portfolio.

[11] Measure refers to probability measure, which can, for example, be real-world or risk-neutral.