Kamakura Corporation
2222 Kalakaua Avenue
Suite 1400
Honolulu, Hawaii 96815, USA
phone 1 808 791 9888
facsimile 1 808 791 9898
www.kamakuraco.com

Advanced Financial Risk Management:

Tools and Techniques for Integrated Credit Risk, Market Risk, Liquidity Risk and Interest Rate Risk Management

April 2012

Donald R. van Deventer
Kenji Imai
Mark Mesler

Part 2: Risk Management Techniques for Interest Rate Analytics

Chapter 12:
Legacy Approaches to Interest Rate Risk Management

In this chapter, we introduce a number of legacy interest rate risk concepts for one purpose only: as an aid in the reader's communication with colleagues who are still fond of these techniques. As for the authors, we believe that the time has come and gone for techniques that we discuss in this chapter: gap analysis, simple simulation methods, duration, and convexity. As we go through this chapter, we will be liberal in adding the phrase "MODEL RISK ALERT" when an analysis involves assumptions that are false and potentially very harmful from a career standpoint and an accuracy point of view. For readers who have no interest in the horse and buggy days of interest rate risk management, we recommend that you skip this chapter except for the section that compares a modern hedge using a multi-factor Heath Jarrow and Morton modeling approach with the legacy duration approach.

Gap Analysis and Simulation Models

Uyemura and van Deventer [1993] discuss the "80/20 rule" in the context of risk management. Getting 80% of the benefits for 20% of the work was the right strategy in 1993, when most serious risk analysis on an enterprise wide basis was still done on mainframe computers because of the severe limitations of the 16 bit operating systems of the personal computers of that era. This rule of thumb is probably still a good one in some senses, but it is simply an excuse for laziness when it involves the computer science aspects of risk management. The state of computer science and risk management software has evolved so far in the last decade that 98 percent of the benefits can now be achieved with 2 percent of the effort required for a perfect answer. This is particularly true with the rapid market penetration of advanced personal computers/servers running a 64 bit operating system which can address 4.3 billion times (2^{32}) more memory than the 32 bit operating systems they replace. We'll keep the distinction between the 98/2 rule and the 80/20 rule in mind through this chapter, where we review legacy interest rate risk management tools for only two distinct purposes: communication with senior management (essential) and exposition (not decision making) of risk issues in day to day interest rate risk operations.

Measuring Interest Rate Risk: A Review

In earlier chapters, we discussed how financial institutions can measure risk in a number of ways:

- The **Jarrow-Merton Put Option**: the value of a put option that would eliminate all of the interest rate risk (the downside risk only) on the financial institution's balance sheet for a specific time horizon. We expand this concept in later chapters to include all risks, with a special focus on credit risk.
- The **Interest Rate Sensitivity of the Market Value of the Financial Institution's Equity**, often proxied by the mark to model valuation of the financial institution's assets less the value of its liabilities
- **The Interest Rate Sensitivity of the Financial Institution's Net Income** (or Cash Flow) over a specific time horizon. MODEL RISK ALERT: The short time horizons of this analysis truncate visibility of risk, and generally accepted accounting principles drastically distort true cash flow and obscure risk. The authors recommend cash flow simulation over net income simulation, with the use of the latter recommended only if required by regulators or for communication with an unsophisticated audience (from a risk management perspective) like non-financial executives in senior management.

In Chapters 13 through 41, we emphasize a modern approach to each of these risk measures with respect to all risk: interest rate risk, market risk, liquidity risk, and credit risk. Our emphasis in this chapter is solely on tools in use since the 1938 introduction of the Macaulay duration concept. A sophisticated reader should review this chapter as if it is an exercise in archaeology, with the exception of the Heath Jarrow and Morton hedging example.

Legacy Rate Risk Tools: Interest Rate Sensitivity Gap Analysis

In earlier chapters, we discussed the concept of the interest rate risk management "safety zone" in which it is impossible for the financial institution to go bankrupt from interest rate risk. Every senior officer of a financial institution needs to know the answers to a number of questions related to the safety zone concept:

- How do we know whether the financial institution is in the safety zone or not?
- If we are out of the safety zone, how far out are we?
- How big should my interest rate risk hedge be to get back in the safety zone?
- How much of a move in interest rates would it take to cause bankruptcy from an interest rate risk move?

Interest rate sensitivity "gap" analysis was the first attempt that most financial institutions made to answer these questions. Today, financial institutions use this expository technique to communicate with senior management team members who do not have a background in finance or risk management. Gap analysis should be not used for decision making because its implications are too crude to be accurate. Gap analysis is too subject to arbitrary assumptions that can literally change the measured impact of higher interest rates on the financial institution from positive to negative or vice-versa. Gap analysis does have a 30,000 foot-level perspective, but there is no reason in the 21st century to use a Sopwith Camel, the British bi-plane fighter introduced in 1917, to hit risk management targets when risk management cruise missiles are available.

How was an interest rate sensitivity gap supposed to work? Consider a simple example in which \$1000 is invested in a floating rate loan with a floating interest rate that resets semiannually. We assume the maturity of the loan is three years. The loan is financed with \$800 in a matching maturity six month certificate of deposit and \$200 in equity, since we assume the bank was newly formed. Assume that the interest rate on the loan is set at 2% over the six month certificate of deposit (all-in) cost and that this cost changes from 6% in the first period to 4% to the second period. How would this balance sheet look in a standard interest rate sensitivity gap table?

Standard Interest Rate Sensitivity Gap Analysis

Balance Sheet Category	1 Day	2-7 Days	8-30 Days	2 Months	3 Months	4 Months	5 Months	6 Months	12 Months	2 years	3-5 years	Over 5	Other	Total
Assets														
Commercial Loan								1000						1000
	0	0	0	0	0	0	0	1000	0	0	0	0	0	1000
Liabilities														
Certificates of Deposit								800						800
	0	0	0	0	0	0	0	800	0	0	0	0	0	800
Interest Rate Sensitivity Gap	0	0	0	0	0	0	0	**200**	0	0	0	0	0	**200**
Cumulative Rate Sensivity Gap	0	0	0	0	0	0	0	200	200	200	200	200	200	

We show the \$1000 in loans in the six month "time zone," not its actual maturity, because the interest rate resets in approximately six months. The \$800 in certificates of deposit goes in the same time bucket.

What about the "effective maturity" of equity? While some large financial institutions believe they know what the effective maturity of equity is, independent of the actual balance sheet, most sophisticated practitioners and essentially all (if such a thing is possible) academics believe that the effective maturity of the financial institution's capital is derived from the net of the impact of assets less the impact of "non-equity" liabilities. Therefore, the mismatch we show in month six of \$200 is the effective maturity of the common stock of this financial institution.

If the financial institution had just issued \$200 in equity for cash, the financial institution could have made \$800 in loans funded with \$800 in CDs and then invested the \$200 in equity at any maturity. If the proceeds were invested in overnight funds, the interest rate mismatch would look like this:

Standard Interest Rate Sensitivity Gap Analysis: Equity Invested in Overnight Funds

Balance Sheet Category	1 Day	2-7 Days	8-30 Days	2 Months	3 Months	4 Months	5 Months	6 Months	12 Months	2 years	3-5 years	Over 5	Other	Total
Assets														
Overnight Funds	200													200
Commercial Loan								800						800
	200	0	0	0	0	0	0	800	0	0	0	0	0	1000
Liabilities														
Certificates of Deposit								800						800
	0	0	0	0	0	0	0	800	0	0	0	0	0	800
Interest Rate Sensitivity Gap	**200**	0	0	0	0	0	0	**0**	0	0	0	0	0	**200**
Cumulative Rate Sensivity Gap	200	200	200	200	200	200	200	200	200	200	200	200	200	

This financial institution is "asset sensitive" because it has more short term "asset" exposure than it does liabilities. Without doing any multiperiod simulation, we know from this simple table that net income will rise sharply if interest rates rise.

What if equity instead were invested in long term securities? Then the interest rate sensitivity gap would look like this:

Standard Interest Rate Sensitivity Gap Analysis: Long Term Bonds

Balance Sheet Category	1 Day	2-7 Days	8-30 Days	2 Months	3 Months	4 Months	5 Months	6 Months	12 Months	2 years	3-5 years	Over 5	Other	Total
Assets														
Long Term Bonds												200		200
Commercial Loan								800						800
	0	0	0	0	0	0	0	800	0	0	0	200	0	1000
Liabilities														
Certificates of Deposit								800						800
	0	0	0	0	0	0	0	800	0	0	0	0	0	800
Interest Rate Sensitivity Gap	**0**	0	0	0	0	0	0	**0**	0	0	0	200	0	**200**
Cumulative Rate Sensivity Gap	0	0	0	0	0	0	0	0	0	0	0	200	200	

The net income of the financial institution will not change in response to short run changes in interest rates because the capital of the institution is invested in long term securities.

The Safety Zone

All of the interest rate risk positions of the financial institution we have discussed so far are in the "safety zone." The financial institution can never go bankrupt from interest rate risk alone because all short term liabilities are invested in matched maturity assets. The capital of the bank can be invested at any maturity. What does this imply for the interest rate sensitivity of the financial institution's common stock? As we discussed in Chapter 4, it moves in three different ways for these three balance sheets:

Case 1: The common stock trades like $200 in six month instruments, since capital is invested in the six-month reset loan, plus the present value of the constant spread on $800 of the loan

Case 2: The common stock trades like $200 in overnight instruments, since capital is invested in the overnight market, plus the present value of the constant spread on $800 of the loan

Case 3: The common stock trades like $200 in long term bonds, since capital is invested in the long term bond market, plus the present value of the constant spread on $800 of the loan

Since all of these investments are "at market" at the time they are done, they neither create nor destroy value. They are all in the safety zone, but the financial institution's common stock will respond differently in the three scenarios to changes in interest rates. In no case, however, can the financial institution go bankrupt from interest rate risk so no hedge is necessary.

What if the management thinks shareholders want stable net income? Then they should invest in long term bonds instead of the overnight market, but either position is in the safety zone and shareholders should be indifferent—as discussed in Chapters 4 and 5, the shareholders can convert the $200 from overnight money to long term bonds or vice versa by doing transactions for their own account in combination with their holdings of the financial institution's common stock.

What's Wrong with Gap Analysis?

What's wrong with gap analysis? As an expositional tool for communicating with non-specialists, nothing is wrong with it.

As a decision-making tool, though, it has serious flaws that led of the Office of the Comptroller of the Currency of the United States to stop requiring gap analysis to be reported on the Quarterly Report of Condition for commercial banks in the United States.

MODEL RISK ALERT: We can summarize some of the problems with gap analysis briefly:

- It ignores the coupons and cash flows on all assets and liabilities.
- Gap analysis typically uses accounting values of assets and liabilities, not market values. Two bonds with accounting values of $100 but market values of $120 and $80 respectively are both put in the gap chart at $100.
- The time buckets typically used are very focused on short term maturities and tend to be very crude for the long term interest rate risk, where the real big risk to the institution often lies
- Gap analysis ignores options—An 8% 30 year mortgage may have an effective interest rate sensitivity of 30 years when market interest rates are 12% but an effective interest rate sensitivity of 30 days when interest rates are at 3%
- Gap analysis can't deal with other instruments where payments are random, like property and casualty insurance policies, savings deposits, demand deposits, charge cards, and so on
- Some interest rates have a lagged adjustment to changes in open market rates, like the prime rate in the United States and Japan and home mortgage indexes in Japan and Australia
- Many balance sheet items (charge card balances, deposit balances) have a strong seasonal variation—they flow out and flow bank in again. Do you classify them by the first outflow date and ignore the seasonal inflow?
- Assets with varying credit risks are mixed and visibility on credit risk is completely blurred across the balance sheet

Sophisticated risk managers use detailed transaction level, exact day count interest rate risk analysis like that in the remainder of this book. They use gap charts to communicate their conclusions to senior management, if the senior managers are non-specialists.

Legacy Rate Risk Tools: Multi-period Simulation

Simulation analysis was the next tool to which major financial institutions turned once progress in computer science allowed financial institutions to take a more detailed look at interest rate risk. In the 1970s, most of the focus was on net income for a 1-2 year horizon. This focus was understandable but fatal for thousands of U.S. financial institutions who had interest rate risk of thirty year fixed rate mortgages (the market standard at the time) which one year net income simulation ignores after the first year. Fortunately, we have much more sophisticated tools available, which we outline in the rest of this book. Nonetheless, it is important to review the kinds of considerations that played an important role in the simulation exercises of major financial institutions over the last four decades. We outlined one of those pitfalls in Chapter 11, the predilection of some bankers to select interest rates from an interest rate lattice or bushy tree like those studied in Chapters 6 through 9. In this section, we are focused on other areas of model risk that stem from a focus on net income simulation.

Key Assumptions in Simulation

Simulation can be done in many ways:

1. Simulate N discrete movements in interest rates with the current balance sheet held constant and produce market values and net income. This is a variation on the assumption behind traditional value at risk, a risk calculation which critics have assaulted with full force in the wake of the $2 billion in “hedge” losses announced by VAR-sponsor JPMorgan Chase on May 8, 2012. MODEL RISK ALERT: Assets and liabilities mature, prepay, and default. They are not constant.
2. Simulate the current balance sheet as if it were constant and as if interest rates remained at current levels. MODEL RISK ALERT: Assuming interest rates do not change implies a riskless arbitrage versus the current yield curve.
3. Move yield curves through multiple specified changes over the forecasting horizon with specific new assets and liabilities added. MODEL RISK ALERT: While specific scenario analysis is a good supplement to other risk techniques and a frequent requirement of regulators, it is safe on a stand-alone basis only for expositional purposes.
4. Randomly move interest rates over the forecast horizon and dynamically simulate which assets and liabilities will be added as interest rates changes. MODEL RISK ALERT: The simulation of interest rates that has typically been done by major financial institutions has not been consistent with no arbitrage. The Heath Jarrow and Morton no arbitrage constraints of Chapters 6 through 9 need to be imposed on any Monte Carlo simulation of rates.

Most financial institutions employ some combination of these techniques, but they do so in the more modern way that we discuss in later chapters.

Even in the more modern style, however, simulation involves a number of key assumptions:

- **Movements in Balance Sheet Amounts**: How will seasonal variation in deposit balances and commercial loans be modeled? How will cash flow thrown off from coupon payments be invested? Will these decisions be independent of interest rate levels or shall they be set as a function of interest rates? MODEL RISK ALERT: These projections are no more accurate than any other forecast made by humans. One needs to avoid hedging assets and liabilities that one does not have now or that one does not have a legal obligation to originate later.
- **Maturity of New Assets and Liabilities**: What will be the maturity structure and options characteristics of new assets and liabilities? For example, the mix of home mortgage loans between fixed rate loans and floating rate loans changes substantially over the interest rate cycle. How should this be modeled?
- **Administered Rate Analysis**: How should interest rates which respond with a lag to changes in open market rates be modeled? MODEL RISK ALERT: Many of these “open market rates” have been historically tied to the London Interbank Offered Rate in major currencies. We discuss lawsuits filed in 2011 and revised in 2012 alleging manipulation of Libor rates by market participants in Chapter 17.

- **New Business**: How should "new business" (as opposed to cash flow from existing assets "rolling over") be modeled?

Data Aggregation in Simulation Modeling

One of the issues in traditional simulation was the degree to which individual transactions should be aggregated for analysis. Many analysts even today think that aggregation is normally a harmless convenience, but more thoughtful risk managers are horrified at the thought. It is done only as a necessary evil which produces some offsetting benefit, like allowing a larger number of Monte Carlo scenarios or faster processing times. MODEL RISK ALERT: Except in the most extreme circumstances, like Chinese banks with total assets and liabilities consisting of more than 700 million separate transactions, a systems need for data aggregation is an indication that the relevant risk system is very out of date. In the 21st century with 64 bit operating systems, data aggregation serves no purpose except to boost employment at the financial institution's information technology department.

In an earlier era with limited computer speeds and limited memory, data summarization and how to do it was a major issue. In a more modern era, no one would "average" a $100,000 4% coupon mortgage with a $100,000 6% mortgage and say they were the effective equivalent of $200,000 in a 5% mortgage. The entire discipline of market risk is built around the constant search for more accuracy (and slightly different, higher precision), but in an earlier era in asset and liability management this kind of approximation was necessary under the 80/20 rule we introduced earlier in this chapter. Now that the 98/2 rule is relevant, data "stratification" is a euphemism for MODEL RISK.

Constraining the Model

Assets and liabilities of a financial institution can't be accurately be modeled independently of each other. This is a truism, but, particularly in the insurance industry, the "Chinese Wall" between the liability side and asset side of insurance firms is often a serious impediment to integrated risk management. If the objective of the modeling exercise is a detailed estimate of statutory (insurance industry) or general accepted accounting principles net interest income, then of course assets and liabilities plus accounting capital have to be equal. On a cash flow basis, the basis preferred by the authors, total cash in must equal total cash out. This is not a trivial exercise in simulation, as the amount of certificates of deposit that a bank, for example, must issue depends on cash thrown off or consumed by every other asset and liability on the balance sheet. This is determined by insuring that cash in equals cash out, and cash in's "plug" item is the amount of new certificates of deposit or some other residual funding source. Similar concepts apply in pension funds, endowments, mutual funds, hedge funds, property and casualty insurance, and life insurance.

Modeling the Maturity Structure of a Class of Assets

In the early days of simulation modeling, the maturity structure of a pool of assets (like home mortgages) was not visible to the analyst because computer systems were too antiquated to provide that information on a timely basis to management. The team led by Wm. Mack Terry at Bank of America (see Chapters 2 and 38) in the 1970s for example was not able to get the exact maturity dates and payment schedules of each mortgage loan at the bank. They had to estimate a "tractor" type roll-off of old mortgages and their replacement by new mortgages (see Chapter 7 in Uyemura and van Deventer [1993] for more on this approach).

Fortunately modern financial systems now should be based on exact day count and exact maturity dates at the individual transaction level. This is actually less work than summarizing the data in a way that doesn't distort the results.

Periodicity of the Analysis

In an earlier era, the periods of the simulation were often "hard coded" to be of equal length. For example, periods could be monthly but the differences in lengths of the month (from 28 days to 31 days) were often ignored. Now analysts specify the exact start date and end date of each period, which can differ in length from month to month. Modern risk analysis is firmly rooted (for the most part) in exact day count technology. Intra-day and near real-time analysis is increasingly standard.

Exceptions to the Exact Day Count Trend

An exception to the trend toward exact day count precision of simulation is the brief risk management detour of "value at risk" analysis, which we discussed earlier in this chapter and which we describe fairly harshly in Chapter 36. MODEL RISK ALERT: This single period analysis effectively ignores all cash flows (and their exact day count) between the date of the analysis and the value at risk horizon. A similar set of assumptions is commonly applied to get quick (and very dirty) answers on instruments like collateralized debt obligations. These simplistic approaches are being rapidly replaced by more sophisticated simulation like that discussed in Chapter 20 on collateralized debt obligations, which many claim was the "financial instrument of mass destruction" at the heart of the 2007-2009 credit crisis.

Legacy Rate Risk Tools: Duration and Convexity

The use of the duration has gone through a curious evolution. For most of the period since 1938, duration was under-utilized, for nearly inexplicable reasons, as an enterprise risk management tool. For the last 15 years or so, the opposite is true. Duration has become "the conventional wisdom" and now, for equally inexplicable reasons, it is over utilized as a risk management tool in spite of enormous progress in interest rate risk management like the Heath Jarrow and Morton models of Chapters 6 through 9.

As we discussed in Chapter 1, the use of duration as a tool for total balance sheet risk management met with more than five decades of resistance, in spite of the fact that duration and its variants have been standard techniques on trading floors for most of that time. The duration concepts are really the foundation upon which all further advances in interest rate analytics are based. The main model risk now is clinging to duration long after the progress triggered by duration has resulted in far better techniques. On the other hand, from a total balance sheet management perspective, there are still many management teams at very large financial institutions who sincerely believe that net income simulation is a solid basis for risk management and that interest rate risk analytics used on the trading floor either don't apply or are too hard to understand. While many of these CEOs were fired in the aftermath of the 2007-2009 credit crisis, this will be a persistent risk management issue for decades.

In this chapter, we discuss the duration concepts in view of these two constituencies—one group who views duration as "too simple" and which has moved on to more advanced techniques, and another group which views duration as "too hard" to apply. As we discussed in Chapter 1, a lot of this difference in view is due to gaps between the generations and differences in educational experience. As one of the authors is the "older generation," he advocates the latter explanation!

Even back in 1978, Jonathan Ingersoll, Jeffrey Skelton, and Roman Weil published an appreciation of Frederick R. Macaulay's 1938 work on duration entitled "Duration Forty Years Later." In that work, the authors pointed out a number of key aspects of duration and potential pitfalls of duration that we want to emphasize in this chapter. In what follows, we introduce the concepts of duration and convexity as they have traditionally been used for "bond math." We then relate them to Macaulay's 1938 work. "Standard practice" in financial markets has gradually drifted away from Macaulay's original concept of duration. Since Ingersoll, Skelton, and Weil published their article, the "standard practice" has forked into two streams, the traditional adherents of "bond math" and the term structure model school of thought, which we outlined in Chapters 6 through 10 and which we review from a historical perspective in Chapter 13. The purpose of this chapter is to reemphasize the points made about duration by Ingersoll, Skelton and Weil and to highlight the errors that can result from deviating from Macaulay's original formulation of duration. Most of the vocabulary of financial markets still stems from the use of the traditional fixed income mathematics used in this chapter, but the underlying meanings have evolved as financial instruments and financial mathematics have become more complex.

Macaulay's Duration: The Original Formula

As in Chapter 4, we use the basic valuation formulas for calculating the present value of a bond.

We define the price of a zero coupon bond with maturity of t years as P(t). We let its continuously compounded yield be y(t), and we label the cash flow in period t as X(t). As we discussed in Chapter 4, the present value of a series of n cash flows is

$$\begin{aligned} \textit{present value} &= \sum_{i=1}^{n} P(t_i)X(t_i) \\ &= \sum_{i=1}^{n} e^{-y(t_i)t_i} X(t_i) \end{aligned}$$

The last line substitutes the relationship between y(t) and P(t) when yields y(t) are continuously compounded. Macaulay investigated the change in present value as each yield y(t) makes a parallel shift of amount x (lower case), so the new yield at any maturity y(t)* = y(t) + x. The change in present value that results is

$$\frac{\partial\ present\ value}{\partial x} = \sum_{i=1}^{n} -t_i e^{-y(t_i)t_i} X(t_i)$$

$$= \sum_{i=1}^{n} -t_i P(t_i) X(t_i)$$

In earlier chapters, particularly Chapters 9 and 10, we reviewed the Dickler, Jarrow and van Deventer (2011) result that 94% of forward rate curve movements are inconsistent with the parallel shift assumption in the U.S. Treasury market on a daily basis from 1962 to 2011. For the time being, however, let's assume that a parallel shift assumption is good enough that we need to at least measure our exposure to this kind of risk. MODEL RISK ALERT: The parallel shift assumption is not good enough for interest rate risk management, and we use it in the remainder of this chapter to explain to a modern reader how duration adherents think about risk.

Macaulay defined duration as the percentage change (expressed as a positive number, which requires changing the sign in the equation above) in present value that results from this parallel shift in rates:

$$Duration = -\frac{\frac{\partial present\ value}{\partial x}}{present\ value}$$

$$= \frac{\sum_{i=1}^{n} t_i P(t_i) X(t_i)}{\sum_{i=1}^{n} P(t_i) X(t_i)}$$

From this formula, one can see the reason that duration is often called the "present value-weighted average time to maturity" of a given security. The time to maturity of each cash flow t_j is weighted by the share of the present value of the cash flow at that time in total present value.

In the case of a bond with n coupon payments of C dollars (note that C is not the annual percentage interest payment unless payments are made only once per year) and a principal amount of 100, this duration formula can be rewritten as

$$Duration = -\frac{\frac{\partial present\ value}{\partial x}}{present\ value}$$

$$= \frac{C\left[\sum_{i=1}^{n} t_i P(t_i)\right] + t_n P(t_n)100}{C\left[\sum_{i=1}^{n} P(t_i)\right] + P(t_n)100}$$

This form of duration, the form its inventor intended, has come to be known as Fisher-Weil duration. All discounting in the present value calculations is done at a different continuous yields to maturity $y(t_i)$ for each maturity. We know from Chapter 5 that we can get these continuous yields to maturity from the various yield curve smoothing techniques in that chapter.

Using Duration for Hedging

How is the duration concept used for hedging? Let's examine the case where a huge financial institution has its entire balance sheet in one unit of a security with present value B_1. Let's also assume that the financial institution is uncomfortable with its risk level. One thing the financial institution could do is to try to buy the Jarrow-Merton put option that we discussed in Chapter 1. This put option would provide the financial institution with price insurance against the decline in value of this security. The valuation of this put option should be consistent with the valuation methods used for interest rate options in Chapters 6 through 9 and reviewed in more detail in Chapter 24. Let's assume that the financial institution wants to hedge in a more traditional way. Let's assume that the financial institution wants to form a no risk portfolio that includes the hedging security (which has present value B_2). What amount of this second security should the financial institution use to hedge? Let the "zero risk" amount of the hedging security be w. Then the total value of this portfolio W is

$$W = B_1 + w B_2$$

We want the change in the value of this hedged portfolio to be zero for infinitely small parallel shifts in the term structure of interest rates. We can call the amount of this shift x. For the hedge to be successful, we must have

$$\frac{\partial W}{\partial x} = \frac{\partial B_1}{\partial x} + w\frac{\partial B_2}{\partial x} = 0$$

From this equation, we can calculate the right amount of B_2 to hold for a perfect hedge as

$$w = -\frac{\frac{\partial B_1}{\partial x}}{\frac{\partial B_2}{\partial x}}$$

Notice that this formula doesn't directly include duration. We can incorporate the duration calculation by modifying the original equation for a perfect hedge as follows:

$$\frac{\partial W}{\partial x} = B_1 \frac{\frac{\partial B_1}{\partial x}}{B_1} + w\, B_2 \frac{\frac{\partial B_2}{\partial x}}{B_2}$$
$$= - B_1 \; Duration[\, B_1] - wB_2 \; Duration[\, B_2] = 0$$

Therefore the hedge amount w can be rewritten

$$w = \frac{- B_1 \; Duration[\, B_1]}{B_2 \; Duration[\, B_2]}$$

Using either formula for w, we can establish the "correct" hedge for the financial institution’s portfolio. While we have talked about hedging as if the financial institution only owned one security, this traditional analysis applies equally well to a financial institution which owns millions of transactions whose risk is captured by the artificial “security” we have described above. MODEL RISK ALERT: This hedging technique has the following characteristics:

- It is correct only for infinitely small parallel shifts in interest rates
- It has to be "rebalanced" like a Black-Scholes options delta hedge whenever
    1. rates change
    2. time passes

The hedge will not be correct if there are large parallel jumps in interest rates, nor will it be correct for non-parallel shifts in interest rates. In other words, the basic theory of duration doesn’t result in a “buy and hold” hedge. This insight is perhaps the model’s greatest contribution to the world of finance. At the same time, it leaves many management teams looking for the insurance policy against risk that they can buy once and don’t have to deal with again for many years. That insurance policy is the Jarrow-Merton put option.

Comparing a Duration Hedge with Hedging in the Heath Jarrow and Morton Framework

How does this traditional hedge compare with hedging in the Heath Jarrow and Morton framework of Chapters 6 and 7 (with 1 risk factor driving interest rates), Chapter 8 (2 risk factors driving rates), or Chapter 9 (3 risk factors driving rates)? We summarize the differences here:

Assumption about Yield Curve Shifts

Duration: Parallel

HJM: No restrictions on nature of yield curve movements, which are implied by the number and nature of the risk factors

Size of Yield Curve Movements

Duration: Assumed to be infinitely small

HJM: Can be implemented in either discrete form (Chapters 6 through 9) or in continuous time form with infinitely small yield shifts

Number of Instruments for "Perfect" Hedge Implied by Model

Duration: 1 hedging instrument, and any maturity can be made into a perfect hedge by holding in the correct amount

HJM 1 Factor Model: 1 hedging instrument and "cash"

HJM 2 Factor Model: 2 hedging instruments and "cash"

HJM 3 Factor Model: 3 hedging instruments and "cash"

HJM n Factor Model: n hedging instruments and "cash"

Size of Time Step over Which Interest Rates Change

Duration: Implicitly assumed to be zero

HJM: Equal to length of period in discrete implementation or infinitely small in the continuous time case

In the Heath Jarrow and Morton discrete implementations, the analysis is simple. Since the HJM examples are built on the assumption of no arbitrage, a riskless hedged position will earn the one period risk free rate prevailing at time zero. The return on $1 invested in "cash," the riskless one period bond, earns 1/P(0,1), where P(0,1) is the time 0 value of a zero coupon bond maturing at time 1. Let's assume that the instrument being hedged is bond B. We write the value of bond B as B(t,j) where t is the then-current time and j is the relevant scenario that prevails. Let's assume we are at time 0 in Chapter 9 where 3 risk factors are driving interest rates. At time 1, there are four nodes on the bushy tree, denoting scenarios 1 (shift 1), 2 (shift 2), 3 (shift 3) and 4 (shift 4). The hedging instruments are a position of w_1 dollars invested in cash (the riskless bond maturing at time 1), w_2 dollars invested in the zero coupon bond maturing at time 2, w_3 dollars invested in the zero coupon bond maturing at time 3, and w_4 dollars invested in the zero coupon bond maturing at time 4. We know that the initial investment in bond B and the hedging instruments is

Time zero value $Z(0)=B(0,*)+w_1+w_2+w_3+w_4$

If the value of a zero coupon bond at time 1 maturing at time T in scenario j is P(1,T,j), then the time one value of our hedged portfolio in scenario j will be

$$Z(1,j)= B(1,j)+w_1[P(1,1,j)/P(0,1,*)]+w_2[P(1,2,j)/P(0,2,*)]+w_3[P(1,3,j)/P(0,3,*)]+w_4[P(1,4,j)/P(0,4,*)]=Z(0)/P(0,1,*)$$

The last equality is the "no arbitrage" condition that insures that a hedged portfolio earns the time zero risk free rate. Since there are 4 unknowns and we have 4 scenarios, we can solve for the perfect hedge in the 3 factor HJM case in Chapter 9. In the 2 factor HJM case, we have 3 unknowns for the 3 scenarios. In the 1 factor HJM case, we have 2 unknowns for the 2 scenarios.

MODEL RISK ALERT: Note that the HJM framework allows us to measure the error in a duration hedge precisely. Since only one instrument is used in a duration hedge, it will be a perfect hedge in (at most) only one of the 4 shifts in a 3 factor Heath Jarrow and Morton model. We can calculate the hedging error from a duration strategy by valuing the hedged portfolio just as we valued simpler instruments in Chapters 6 through 9.

Note also that the HJM hedge needs to be rebalanced with each forward movement in time. In general, the hedge amounts will be different at different nodes on the bushy tree, even at the same current time t. This is a common phenomenon in finance, like the continuous delta hedging that is necessary when using the Black-Scholes options model.

Duration: The Traditional Market Convention

For many years, the conventional wisdom has drifted over time from Macaulay's original formulation of the duration concept. With the explosion in interest rate derivatives over the last two decades, the original formulation and its more modern relatives dominate among younger and more analytical interest rate risk experts. The "evolved" Macaulay duration is passing back out of fashion because it relies on old fashioned "bond math" rather than the Fisher-Weil continuous compounding and the related concepts of Chapter 4. The reason for the change to "bond math" from Macaulay's original formulation is a simple one: until relatively recently, it has been difficult for market participants to easily take the P(t) zero coupon bond values from the current yield curve using the techniques that were outlined in Chapter 5. Now that yield curve smoothing has become easier, this problem has declined in importance, but the market's simpler formulation persists in many areas of major financial institutions. The best way to explain the "bond math" market convention for duration is to review the formula for yield to maturity.

The Formula for Yield to Maturity

For a settlement date (for either a secondary market bond purchase or a new bond issue) that falls exactly one (MODEL RISK ALERT) equal-length "period" from the next coupon date, the formula for yield to maturity can be written as follows for a bond which matures in n periods and pays a fixed dollar coupon amount C m times per year as shown in Chapter 4:

$$price + accrued\,\text{int}\,erest = present\ value = C\left[\sum_{i=1}^{n} P(t_i)\right] + P(t_n) * Principal$$
$$where\ P(t_i) = \frac{1}{\left(1+\frac{y}{m}\right)^{i}}$$

Yield to maturity is the value of y that makes this equation true. This relationship is the present value equation for a fixed coupon bond, with the addition that the discount factors are all calculated by dividing one dollar by the future value of one dollar invested at y for i equal length periods with interest compounded m times per year. In other words, y is the internal rate of return on this bond. Note also that the present value of the bond, using these discount factors, and the "price" of the bond are equal since there is no accrued interest.

Yield to Maturity for Long or Short First Coupon Payment Periods

For most outstanding bonds or securities, as we saw in Chapter 4, the number of periods remaining usually contains a short (or occasionally a long) first period. The length of the short first period is equal to the number of days from settlement to the next coupon, divided by the total number of days from the last coupon (or hypothetical last coupon, if the bond is a new issue with a short coupon date) to the next coupon. The exact method of counting days for purpose of this calculation depends on the interest accrual method associated with the bond (examples are actual/actual, the National Association of Securities Dealers 30/360 method in the United States, or the Euro market implementation of the 30/360 accrual method). We call the fraction of remaining days (as appropriately calculated) to days between coupons (as appropriately calculated) x. Then the yield to maturity can be written

$$present\ value = C\sum_{i=1}^{n}\left[P(t_i)\right] + P(t_n) * Principal$$
$$where\ P(t_i) = \frac{1}{\left(1+\frac{y}{m}\right)^{i-1+x}}$$

As in the case of even first coupon periods, y is the yield to maturity that makes this relationship true, the internal rate of return on the bond. Note that when x is 1, this reduces to the formula above for even length first periods. We can solve for y in simple spreadsheet software in a number of ways, including standard yield to maturity/internal rate of return formulas in the spreadsheet or more generally using the “solver” function.

Applying the Yield to Maturity Formula to Duration

The conventional definition of duration is calculated by applying the implications of the yield to maturity formula to the original Macaulay duration formula, which is given below:

$$True\ Duration = -\frac{\frac{\partial present\ value}{\partial x}}{present\ value}$$

$$= \frac{\sum_{i=1}^{n} t_i P(t_i) X(t_i)}{\sum_{i=1}^{n} P(t_i) X(t_i)}$$

We can rewrite this formula for a conventional bond by substituting directly into this formula on the assumptions that market interest rates are equal for all maturities at the yield to maturity y and that interest rates are compounded at discrete, instead of continuous, intervals. MODEL RISK ALERT: This assumption is false except by coincidence. It assumes that the yield curve is flat. MODEL RISK ALERT: Instead of allowing for time periods of different length as we did with the Macaulay formulation (say semi-annual periods of 181 and 184 days, for example), this formulation assumes all periods except the first or last period periods have equal length. This variation of the Macaulay formula results in the following changes:

- o the discount factors $P(t_i)$, taken from the smoothed yield curve by Macaulay, are replaced by a simpler formulation which uses the same interest rate (the yield to maturity) at each maturity:

$$P(t_i) = \frac{1}{\left(1 + \frac{y}{m}\right)^{i-1+x}}$$

- o the time t_i is directly calculated as

$$t_i = \frac{i-1+x}{m}$$

Using this formulation, the "conventional" definition of duration for a bond with dollar coupon C and principal of 100 as

$$Conventional\ Duration = \frac{C\left[\sum_{i=1}^{n} t_i P(t_i)\right] + t_n P(t_n)100}{present\ value}$$

$$= \frac{C\left[\sum_{i=1}^{n} \frac{i-1+x}{m}\left(\frac{1}{\left(1+\frac{y}{m}\right)^{i-1+x}}\right)\right] + \frac{n-1+x}{m}\left(\frac{1}{\left(1+\frac{y}{m}\right)^{n-1+x}}\right)100}{present\ value}$$

This is just a discrete time transformation of Macaulay's formula with the additional false assumptions of a flat yield curve and equal length periods (other than the first or last periods). As shown below, however, this formula doesn't measure the percentage change in present value for a small change in the discrete yield to maturity, as shown in the next section.

Modified Duration

In the case of Macaulay duration, the formula was based on the calculation of

$$Duration = -\frac{\frac{\partial present\ value}{\partial x}}{present\ value}$$

$$= \frac{\sum_{i=1}^{n} t_i P(t_i) X(t_i)}{\sum_{i=1}^{n} P(t_i) X(t_i)}$$

The second line of the equation immediately above results from the fact that the yield to maturity of a zero coupon bond with maturity t_i, written as $y(t_i)$, is continuously compounded. The conventional measure of duration, given in the previous section, is not consistent with Macaulay's derivation of duration when interest payments are discrete instead of continuous. What happens to the percentage change in present value when the discrete yield to maturity (instead of the continuous yield to maturity analyzed by Macaulay) shifts from y to y+z for infinitely small changes in z? Using the yield to maturity formula as a starting point and differentiating with respect to z, we get the following formula, which is called modified duration:

$$Modified\ Duration = \frac{C\left[\sum_{i=1}^{n}\frac{i-1+x}{m}\left(\frac{1}{\left(1+\frac{y}{m}\right)^{i+x}}\right)\right]+\frac{n-1+x}{m}\left(\frac{1}{\left(1+\frac{y}{m}\right)^{n+x}}\right)100}{present\ value}$$

$$= \frac{Conventional\ Duration}{1+\frac{y}{m}}$$

This modified duration measures the same thing that Macaulay's original formula was intended to measure: the percentage change in present value for an infinitely small change in yields. MODEL RISK ALERT: Unfortunately the measurement is inaccurate because the assumptions underlying the model are false. The Macaulay formulation and the conventional duration are different for the following reasons:

- The Macaulay measure (most often called Fisher-Weil duration) is based on continuous compounding and correctly measures the percentage change in present value for small parallel changes in these continuous yields. The conventional duration measure, a literal translation to discrete compounding of the Macaulay formula with the additional implied assumption of a flat yield curve, does NOT measure the percentage change in price for small changes in discrete yield to maturity. The percentage change in present value is measured by conventional modified duration. When using continuous compounding, “duration” and “modified duration” are the same.

- The Macaulay measure uses a different yield for each payment date instead of using the yield to maturity as the appropriate discount rate for each payment date.

The impact of these differences in assumptions on hedging are outlined below.

The Perfect Duration Hedge:

The Difference between the Original Macaulay and Conventional Durations

Both the original Macaulay (Fisher-Weil) and conventional formulation of duration are intended to measure the percentage change in "value" for a parallel shift in yields. MODEL RISK ALERT: We remind the reader that yield curve twists on daily movements in the U.S. Treasury curve occurred 94% of the time between 1962 and 2011. Of the remaining shifts, although rates moved up or down together, a small fraction were truly parallel shifts in equal amounts at each maturity. Using duration, the percentage change is measured as the percentage change in present value; it is important to note that it is *NOT* the percentage change in *price* (present value - accrued interest) that is measured. The difference between a hedge calculated based on the percentage change in present value (the correct way) versus a hedge based on the percentage change in price (the incorrect way) can be significant if the bond to be hedged and the instrument used as the hedging instrument have different amounts of accrued interest. The impact of this “accrued interest” effect varies with the level of interest rates and the amount of accrued interest on the underlying instrument and the instrument used as a hedge. If payment dates on the two bonds are different and the relative amounts of accrued interest are different, the hedge ratio error can exceed five percent.

The second source of error when using modified duration to calculate hedge ratios is the error induced by using the yield to maturity as the discount rate at every payment date instead of using a different discount rate for every maturity. The hedging error (relative to Fisher-Weil duration) from this source can easily vary from -2 percent to plus 2 percent as a linear yield curve takes on different slopes. The hedge ratio error is obviously most significant when the yield curve has a significant slope.

The third source of error is the dominant yield curve movement: yield curve twists instead of parallel shifts. If one is hedging a five year bond with a three year bond and one yield rises while the other yield falls, one can lose money on both sides of the “hedge.” The only way to avoid this phenomenon is to AVOID duration models and to use an HJM framework with an accurate number of underlying risk factors.

Convexity and Its Uses

The word "convexity" can be heard every day on trading floors around the world, and yet it's a controversial subject. It is usually ignored by leading financial academics, and it's an important topic to two diverse groups of financial market participants: adherents to the historical "yield-to-maturity" bond mathematics that predates the derivatives world, and modern "rocket scientists" who must translate the latest in derivatives mathematics and term structure modeling to hedges that work. This section is an attempt to bridge the gap between these constituencies.

Convexity: A General Definition

The duration concept, reduced to its most basic nature, involves the ratio of the first derivative of a valuation formula f to the underlying security's value:

$$duration = \frac{-f(y)}{f(y)}$$

The variable y could be any stochastic variable which determines present value according to the function f. It may be yield-to-maturity, it may be the short rate in a 1 factor HJM term structure model, etc. It is closely related to the delta hedge concept which has become a standard hedging tool for options related securities. The underlying idea is exactly the same as that discussed at the beginning of this chapter: for infinitely small changes in the random variable, the first derivative (whether called "duration" or "delta") results in a perfect hedge if the model incorporated in the function f is "true." So far in our discussion, the constituencies mentioned above are all in agreement.

As we showed in the section above, the duration and modified duration concepts are *not* "true" whenever the yield curve is not flat. The academic community and the derivatives hedges agree on this point and neither group would be satisfied with a hedge ratio based on the discrete yield to maturity formulas presented in this chapter, despite their long and glorious place in the history of fixed income markets. All constituencies agree, however, that when the random variable in a valuation formula "jumps" by more than an infinitely small amount, a hedging challenge results.

Leading academics solve this problem by either assuming it away or by breaking a complex security into more liquid "primitive" securities and assembling a perfect hedge that exactly replicates the value of the complex security. This is exactly what we did in the case of the 3 factor HJM model hedge of our bond B above. In the case of a non-callable bond, we know that the perfect hedge (perhaps one of many perfect hedges, depending on the number of risk factors driving rates) is the portfolio of zero coupon bonds that exactly offsets the cash flows of the underlying bond portfolio. Traditional practitioners of "bond math" and derivatives experts often are forced to deal with (a) the fact that no simple cost-efficient perfect hedge exists or (b) the fact that required rebalancing of a hedge can only be done at discrete, rather than continuous, intervals.

The fundamental essence of the convexity concept involves the second derivative of a security's value with respect to a random input variable y. It is usually expressed as a ratio to the value of the security, as calculated using the formula f:

$$convexity = \frac{f''(y)}{f(y)}$$

It is very closely linked with the "gamma" calculation (the second derivative of an option's value with respect to the price of the underlying security).

When a security's value makes a discrete jump because of a discrete jump in the input variable y, the new value can be calculated to any desired degree of precision using a Taylor expansion, which gives the value of the security at the new input variable level, say y+z, relative to the old value at input variable level y:

$$f(y+z) = f(y) + f'(y)z + f''(y)\frac{z^2}{2!} + \ldots + f^{[n]}(y)\frac{z^n}{n!} + \ldots$$

Rearranging these terms and dividing by the security's original value f(y) gives us the percentage change in value that results from the shift z:

$$\begin{aligned}\frac{f(y+z) - f(y)}{f(y)} &= \frac{f'(y)}{f(y)}z + \frac{f''(y)}{f(y)}\frac{z^2}{2!} + \ldots + \frac{f^{[n]}(y)}{f(y)}\frac{z^2}{n!} + \ldots \\ &= (duration)z + (convexity)\frac{z^2}{2!} + error\end{aligned}$$

The percentage change can be expressed in terms of duration and convexity and an error term. Note that "modified duration" should replace "duration" in this expression if the random factor is the discrete yield to maturity. We will see in Chapter 13 that this Taylor expansion is very closely related to Ito's lemma, the heart of the stochastic mathematics we introduce in Chapter 13. These techniques, combined with the yield curve smoothing technology of Chapter 5, form the basis for an integrated analysis of interest rate risk, market risk, liquidity risk and credit risk which is booming in popularity today. The Heath Jarrow and Morton framework of Chapters 6 through 9 is the fundamental foundation for this integrated risk framework.

Convexity for the Present Value Formula

For the case of the present value of a security with cash flows of X_i per period, payments m times per year, and a short first coupon of x periods, the convexity formula can be derived by differentiating the present value formula twice with respect to the yield to maturity y and then dividing by present value:

$$Convexity = \frac{\left[\sum_{i=1}^{n} X(t_i)\left(\frac{-(i-1+x)}{m}\right)\left(\frac{-(i-1+x)-1}{m}\right)\left(\frac{1}{1+\frac{y}{m}}\right)^{-(i-1+x)-2}\right]}{Present\ Value}$$

When present value is calculated using the continuous time approach,

$$Present\ Value = \sum_{i=1}^{n} e^{-t_i y(t_i)} X(t_i)$$

then the continuous compounding or Fisher-Weil version of convexity is

$$Convexity = \frac{\left[\sum_{i=1}^{n} t_i^2 P(t_i) X(t_i)\right]}{Present\ Value}$$

Hedging Implications of the Convexity Concept

To the extent that rebalancing a hedge is costly or to the extent that the underlying random variable jumps rather than moving in a continuous way, the continuous rebalancing of a duration (or delta) hedge will be either expensive or impossible. In either case, the "best hedge" is the hedge that comes closest to a "buy and hold" hedge, a hedge that can be put in place and left there with no need to rebalance until the portfolio being hedged has matured or been sold. A hedge which matches both the duration (or modified duration in the case of a hedger using a present value based valuation formula) and the convexity of the underlying portfolio will come closer to this objective than most simple duration hedges. MODEL RISK ALERT: Even with a convexity adjustment, however, the implicit assumption is still that interest rate changes are driven only by one parallel yield curve shift factor. Since this assumption is dramatically false, hedging accuracy will be low except by coincidence. Consider the case of a financial institution with a portfolio that has a value P and two component securities (say mortgages and auto loans) with values S_1 and S_2. Denoting duration by D and convexity by C, the best (simple) hedge ratios w_1 and w_2 for a hedge using both securities S_1 and S_2 can be obtained by solving the same kind of equations we used above in the duration case. We need to solve either

$$P' = w_1 S'_1 + w_2 S'_2$$
$$P'' = w_1 S''_1 + w_2 S''_2$$

or

$$P\, D_P = w_1 S_1 D_1 + w_2 S_2 D_2$$
$$P\, C_P = w_1 S_1 C_1 + w_2 S_2 C_2$$

for the hedge ratios w_1 and w_2.

Hedging a Bank Balance Sheet, Hedging an Insurance Company's Equity, or Hedging a Pension Fund's Assets and Liabilities: Conclusion

The concepts we have illustrated in this chapter are generally analyzed using a very simple "position" of a trader or a financial institution—there was generally only one security in the portfolio and only one hedging instrument. In practical application, traders have 10, 50, 100, or 100,000 positions and financial institutions have 50,000, 1 million or 700 million positions. The process of measuring the change in the value of the position (some of which are assets and some of which are liabilities) is identical whether there is one position or millions. We illustrate the calculation of these interest rate "deltas" or durations in the next few chapters, but we do it with emphasis on the multi-factor HJM approach. The calculation of the perfect hedge is also the same from a total balance sheet perspective with one modest variation. As we saw in Chapters 1 through 5, the interest rate risk of the proceeds of the financial institution's equity is irrelevant to shareholders as long as all of the other assets are funded in such a way that there is a constant hedged cash flow from them. If this assumption is correct, then the financial institution cannot go bankrupt from interest rate risk and the financial institution is in the "safety zone." Hedging of the interest rate risk of an institution that is already in the safety zone, regardless of the hedging technique (HJM or duration) used, is simply wasteful because it creates no shareholder value.

Over the last seven decades, the market's conventional definition of duration has drifted from Macaulay's original concept to a measure that appears similar but which can be very misleading if used inappropriately. Fortunately, the trends in interest rate risk management are now swinging forcefully back to the original Macaulay concept and its implementation in a multi-factor Heath Jarrow and Morton interest rate framework. As Ingersoll, Skelton and Weil point out, the Macaulay formulation offers considerable value as a guide to the correct hedge if one takes a back to basics approach: the measure should focus on the percentage change in present value, not price, and the discount rate at each maturity should reflect its actual continuous yield to maturity as obtained using the techniques in Chapter 5 and driven by multiple interest rate factors as we did in Chapters 8 and 9. Assuming the yield to maturity is the same for each cash flow or coupon payment on a bond, even when the payment dates are different, can result in serious hedging errors. The convexity approach provides a useful guide to maximizing hedging efficiency, but the term convexity has outlasted the simple convexity formula commonly used for bonds. The formula itself has been supplanted by the Heath Jarrow and Morton hedging techniques which we illustrated above.

How do we implement these concepts for a bond option? How do we calculate them for a life insurance policy or a first to default swap? How do we calculate them for a mortgage which prepays? How do we calculate them for a savings deposit without an explicit maturity? That is a task to which we will turn after a little more foundation in interest rate risk analytics.