

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 11:

Interest rate mismatching and hedging

In Chapters 1 through 4, we introduced the concept of the interest rate risk “safety zone.” The “safety zone” is that set of portfolio strategies such that it is impossible for the institution to fail to achieve its mission. In banking, for example, a new bank with \$10 billion in capital might fund \$90 billion in fixed rate loans (with no credit risk) for 5 years with \$90 billion in 5 year fixed rate borrowings. The \$10 billion remaining from the capital raise could be invested in any maturity from 1 day to 30 years without changing the default risk of the organization, provided that the assets concerned have no credit risk. This set of potential “zero failure” strategies on the ten billion of capital constitutes the “safety zone” for this new bank. Within this safety zone, shareholders of the bank should be indifferent as to what degree of interest rate risk is taken because the shareholder can take offsetting positions for the shareholder’s own account to achieve the interest rate the shareholder (not the bank) prefers. The assessment of the safety zone gets more complex when we allow for credit risk, an important topic in this book.

We will return to this safety zone concept frequently as we discuss interest rate risk and the implications of interest rate risk for the total risk, including credit risk, of the institution. In this chapter, we focus on the strategy and tactics of changing the interest rate risk of a financial institution from both a political and economic point of view, keeping in mind the lessons of the interest rate simulation Chapters 6 through 10. It is these more modern techniques which will allow actual implementation of the Merton-Jarrow approach to risk management, using the put option concept as a comprehensive measure of risk for the institution. We compare them to legacy approaches to interest rate risk management in Chapters 12 and 13 after setting the stage in this chapter. Our focus in this chapter is quite simple:

when should interest rate risk be hedged? When is it not necessary to hedge interest rate risk? We start by reviewing the kinds of analytical issues and political issues that frequently arise when making a decision to change the interest rate risk of an institution. By “institution,” we mean any legal entity which has interest rate risk, including (but not limited to) pension funds, mutual funds, hedge funds, banks, insurance firms, university endowments, credit unions, special purpose vehicles and so on.

Political Factions in Interest Rate Risk Management

In Chapters 2 and 3, we talked about how performance measurement systems differ by type of institution. These institutional differences have a big impact on the politics of measuring and changing the interest rate risk of a financial institution.

Pension Fund Considerations

As we discussed in Chapters 2 and 3, a pension fund is typically divided between the asset side of the organization and the liability side of the organization. The liability side of a pension fund has a heavy actuarial component, and expert actuaries lay out the most likely benefits that must be paid to the beneficiaries of the pension fund. The Board of Directors of the organization then establishes the “best mix” of various asset classes that the fund will invest in. Pension funds which consist solely of “defined contribution” plans instead of “defined benefit” plans have the easiest task, because in the former case the beneficiary of the pension selects the fund’s assets. In the defined benefit case, however, the asset classes selected include

- **Fixed income securities**, which are designed to provide a steady amount of cash flow toward the pension obligations. The manager of this sub portfolio is typically measured against a naïve benchmark like the Barclays (pre-2008, “Lehman Brothers”) Government Bond index.
- **Equity securities**, which are designed to provide (but sometimes don’t, as in the post-bubble Japan case) a higher long run expected return than fixed income securities while at the same time partially capturing any benefits from an inflationary environment to help offset a rise in pension obligations that inflation would trigger. The manager of such a fund would be measured versus an index like the DAX, Topix or S&P 500 equity index.
- **Commodities and real estate**, which are almost exclusively designed to capture returns from inflation to offset an increase in pension benefits that occur in an inflationary environment.

The amount of risk that the fund has is analyzed on two levels. The first level is at the macro level, where the decision is made about what the proper share of the total portfolio is that should go to these three components. The second level is to analyze the level of risk of the sub portfolio versus the naïve benchmark that is selected as the performance benchmark for the sub portfolio. The total amount of risk taken by the organization is largely determined by the first of these two decisions. The Jarrow-Merton “put option” framework offers considerable benefits in quantifying the many factors that can cause a mismatch in returns on the asset side and obligations on the liability side (inflation, interest rate movements, changes in the life expectancy of the pension beneficiaries, movements in home prices, building prices, commodity prices, and so on). This is an extremely important issue that most sophisticated pension funds are just beginning to quantify. There will be enormous progress on this issue over the next decade that will involve simultaneous analysis of

mark to market technology and multi-period simulation, techniques which we discuss in the rest of this book. What does the Jarrow-Merton “put option” measure in the pension fund case? The expected risk neutral present value of the cost of “bailing out” the pension fund when it does not have sufficient funds to meet its promised “defined benefit” obligations.

Things are more manageable at the sub portfolio level. Consider the fixed income sub portfolio. The manager of the fixed income sub portfolio will have “tracking error” versus the Barclays government bond index in two dimensions. One contribution to tracking error comes from the potential for default if the manager is buying fixed income securities with credit risk. The second contribution to tracking error comes from the mismatch in the maturity structure of cash flows in his actual portfolio versus the cash flow associated with the components of the Barclays government bond index.

The politics of what fixed income portfolio tracking error to accept is more straightforward in pension funds than it is in most institutions because of the relatively small number of people involved. Normally the Board of Directors will specify the maximum tracking error versus the benchmark index and the head of fixed income (or a small committee) will decide where the portfolio should fall within that tracking error benchmark. This will be determined in many cases by advantages in execution rather than by outright interest rate “bets.”

How about the risk-return trade-off at other institutions?

Life Insurance Companies and Property and Casualty Insurance Companies

Changes in the level of interest rate risk in a life insurance company share many of the characteristics of the pension fund decision process. The liability side is analyzed by actuarial staff and the head of the investment department, supervised by the Board of Directors, is given specific performance benchmarks whose performance he must exceed while keeping tracking error within pre-set limits. Most of the politics of risk taking in an insurance company is focused on the credit risk of the benchmark fixed income portfolio and who bears credit losses—if the actuaries ask for a BBB-rated return but the investment department suffers any related credit losses, the incentive system is flawed.

One other complexity is found in the life insurance business which pension funds don’t face so directly. Insurance companies must report earnings quarterly or semi-annually, at least to regulators and often to outside shareholders. The level of scrutiny is much higher for life insurance firms than it is for pension funds except in the worst of circumstances, as in the credit crisis of 2007-2009.

In that sense, the insurance company risk-return trade-off involves both a mark-to-market orientation (how does the value of the Jarrow-Merton put option on our portfolio drop if we take less risk?) and a net income orientation. We turn to that in our discussion of the banking industry in the next section.

Property and casualty insurance companies have still another added complexity—the payoffs on the liability side are much more volatile from year to year because many important classes of property

and casualty insurance losses are much more volatile than payouts on life insurance policies. This brings into play many of the liquidity risk considerations that we discuss in Chapter 37.

State of the art risk management at insurance companies is moving rapidly toward an enterprise-wide view of risk, in which both sides of the balance sheet are managed jointly. This is an essential move toward best practice.

Commercial Banks

The politics of interest rate risk-return strategy in commercial banking is the most complex because all of the complications mentioned above (with the exception of property and casualty insurance losses) are present and the number of players in the decision is very large. A partial list includes the following parties to the interest rate risk and return strategy decision:

- Board of Directors of the Bank
- Chief Executive Officer
- Chief Financial Officer
- Chief Risk Officer
- Managing Committee of the Bank (Asset and Liability Management Committee or “ALCO”)
- Heads of Major Business Units
- Day-to-day Manager of the Transfer Pricing Book which includes all internal funds transfers among business units
- Manager of the Trading Floor
- Bank Regulators

The politics is particularly complex because (unlike pension funds, for example) many of the decision makers are decision makers not because of their expertise in interest rate risk but because of the size of the businesses they run. Someone who is running a network of 3,000 retail branches is too busy to be an expert in interest rate risk strategy, but they are often still a voter in major interest rate risk decisions.

In Chapter 2, we introduced the impact that a “common practice” transfer pricing system has on bank interest rate risk strategy. We continue that discussion later in this book when we discuss performance measurement in more detail:

- Day to day interest rate risk mismatches are moved via the transfer pricing system to a central “transfer pricing book” where the manager has interest rate risk limits consistent with his seniority and the limits he would have if he were trading bonds for the bank
- Huge macro mismatches are the responsibility of the ALCO and must be approved by the Board
- Basis risk from non-market sensitive pricing indexes (like the prime rate in many countries or the average cost of funds of the 12th District savings and loan associations in the United States, once an extremely important pricing index for home mortgages) is born by the heads of line units

Within this general structure, there are a number of political games and tactics that exist to advance the career aspirations or personal wealth of one manager instead of the aspirations of the shareholders. These are too numerous (and too humorous) to cover in depth, so we'll just hit the highlights:

- “I can predict interest rates, why can't you?” This strategy is the oldest in the book. A member of senior management claims he can predict interest rates and urges the transfer pricing book management to take on a huge interest rate risk position. He does. If rates move as predicted, the senior manager takes the credit. If rates move the opposite direction, the transfer pricing book manager takes the blame. This is combated with the limits structure on the transfer pricing book and a separate “book” for the ALCO committee where every position is carefully recorded and the advocates noted by name.
- “My assets are liquid so give me a short term interest rate.” This is the second oldest ploy in the book, usually argued by the head of the trading floor or the head of a portfolio that can be securitized. It was an argument that almost certainly contributed to the failures of Bear Stearns and Lehman Brothers in 2008. The advocates of this position want to avoid paying a matching maturity cost of funds, even on a floating rate basis, as we discuss in detail in Chapter 37 on liquidity risk. This political strategy is combated by saying “no” and following a more modern approach to transfer pricing which we discuss in Chapters 36 to 41. On the argument that “my assets are liquid and we can securitize” we note that a sharp rise in interest rates puts the mark to market value of a fixed rate loan portfolio below par value—very few institutions will securitize if it means recognizing this mark to market loss. The credit crisis of 2007-2009 made it very clear that securitization markets can completely disappear if the integrity of a market (like the subprime home loan market in the U.S.) is suddenly called into question.
- “I want the long run historical weighted average cost of funds, not the marginal cost of funds.” If rates are rising, this is a common request by the asset side of the organization. If rates are falling, the liability side of the organization wants the long run weighted average cost of funds. This is why the marginal cost of funds system has to be very comprehensive or this kind of petty gaming of the transfer pricing system is endless.

If the performance measurement system is designed correctly, politics (for the most part) falls by the wayside and the relevant decision makers can focus on the right risk-return trade-offs.

It is to that task we now turn.

Making a Decision on Interest Rate Risk and Return: The Safety Zone

There are two environments in which management has to make a decision about whether to change the interest rate risk of the institution.

- Inside the “safety zone”
- Outside the “safety zone”

If the institution is still inside the safety zone which we discussed above, then shareholders don't care (as a group) what the interest rate risk position is. They don't care because they can form a portfolio of the institution's common stock and fixed income securities for their own account that exactly offsets any decisions made by the management of the financial institution and gets the combined “pseudo financial institution” exactly where the individual shareholder would like it to be from an interest rate risk perspective. If the financial institution is inside the safety zone, we can focus on the pure economics

of changing the interest rate risk position.

If the financial institution is outside the safety zone, shareholders are being forced to suffer some of the expected costs of the potential bankruptcy of the institution. The interest rate risk-return estimation should include the reduction of these expected losses from bankruptcy as part of the calculated benefits of reducing risk. As this book was written, the legal costs and advisory costs associated with the bankruptcy of Lehman Brothers are still being tallied, but they are certain to be in the billions of dollars. An advocate of an interest rate risk strategy that makes the likelihood of incurring these costs higher should include these potential bankruptcy costs in the pros and cons of the strategy suggested.

Obvious Interest Rate Risk Decisions

Some decisions about interest rate risk have obvious “yes or no” decisions, while others are more complex. Some of the obvious decisions are listed here:

1. Is the “put option” that eliminates the interest rate risk of the institution in the Jarrow-Merton sense from Chapter 1 so expensive that the financial institution can’t afford to buy it? If that is the case, the interest rate risk of the institution is way too large and has to be reduced by other means, like a capital issue.
2. If a change in the cash flow structure from an interest rate-related transaction smoothes net income but leaves the Jarrow-Merton “put option” value unchanged, then the transaction is worth doing. This has an “information” benefit to shareholders who might otherwise misperceive the risk of the institution based on short run changes in net income
3. If the change in cash flow structure from an interest rate-related transaction smoothes net income but leaves the bank well within the safety zone, then it is probably worth doing

Assessing the Risk and Return Trade-offs from a Change in Interest Rate Risk

For all other changes in the financial institution’s interest rate risk and return policy, there are a number of tools for quantifying the benefits of a change in the interest rate risk position of the institution. The two most important components of the benefits of a change in the interest rate risk of an institution consist of the following:

- The change in cost of the Jarrow-Merton put option which would be necessary to completely eliminate the financial institution’s interest rate risk. Most of the rest of this book is devoted to valuing this put option, using the tools of Chapters 6 through 10 and the tools of later chapters on credit and market risk.
- The change in the expected costs of potential bankruptcy of the institution, which we discussed above.

These two issues are critical to deciding when to hedge and when not to hedge. If the financial institution is in deep inside the safety zone, we don’t need to spend much time on this analysis. If the bank is near the boundaries of the safety zone or outside it, this analysis becomes very important.

In the next chapter, we introduce traditional or “legacy” tools for supplementing the Jarrow-Merton put option analysis and analysis of the expected costs of bankruptcy. In the view of the authors, the worst of these legacy tools have been replaced by the Heath Jarrow and Morton techniques of

Chapters 6 through 10. The best of the legacy tools can be shown to be special cases of the HJM approach, but, in almost every case, they are too simple to be used accurately in practice. We explain why in the next chapter.