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 1: Risk Management: Definitions and Objectives

Chapter 2:
Risk, return, performance measurement, and capital regulation

In Chapter 1, we defined risk management as the discipline designed to help management understand the relative risks and returns from different strategies, both at the portfolio level (the perspective of the Chief Executive Officer) and at the transaction level (the perspective of the trader, portfolio strategist, or lending officer). In this chapter, we focus on the nature of risk and return in a practical sense, not a philosophical sense. The practical definition of risk and return reflect the differences in perception of different parts of the financial world. The definitions also introduce a number of biases into perceived risk and return that can be very dangerous. We will point out these dangers in this chapter and then spend the rest of the book talking about how to avoid these pitfalls in the conventional wisdom.

There is an on-going debate in the financial services industry as to how risk and return should be measured. Fund managers and the asset side of pension funds are almost always on a mark to market basis. Insurance companies and commercial banks use a mix of mark to market risk/return management and a variation on financial accounting called "transfer pricing" which separates the interest rate risk portion of net income from the non-rate related components. We introduce the "common practice" of transfer pricing and its history in the last half of this chapter. We discuss the "best practice" of transfer pricing and performance measurement in the remainder of this book, especially in Chapter 38.

Practical Quantification of Risk

In December, 2002 at a major risk management conference in Geneva, Nobel Prize winner Robert Merton told a story about the equity salesman from a major securities firm who was trying to convince a portfolio manager to switch the emphasis of his fund from fixed income securities to equities. Merton quoted the salesman as saying "Over the 40 year time horizon that your fund has, you have a 99% probability that an all equity portfolio will return more than a portfolio of bonds. There's no way that you'll have less money in 40 years than the current par value of the bonds you own now." Merton said the portfolio manager thought for a minute and then said to the salesman, "Well, if you're right, I'd like to buy a put option from your firm to put the equity portfolio you want me to create back to your firm in 40 years at my current portfolio value. Since your argument is so persuasive, I am sure that put option will be very cheap." Naturally, he never heard from the salesman again.

The salesman Merton spoke about was using a tried and true technique, talking about expected returns and ignoring risk. This was the sales approach used by Wall Street to foist low quality, overrated collateralized debt obligations (CDOs) on hapless investors in the lead-up to the 2006-2010 credit crisis. In the Merton story, however, the portfolio manager replied in a way that reflects the highest standard of risk management theory and practice: he asked the price of a contract that insures perfectly[1] against the risk that concerns him. This approach is being pursued very aggressively by practitioners and academics who feel, like we do, that many traditional measures of risk are much too simple. Robert Jarrow and Amiyatosh Purnanandam, in "Generalized Coherent Risk Measures: The Firm's Perspective" (*Finance Research Letters*, 2005), argue the virtues of put premiums like the one requested by the portfolio manager as a particularly high quality measure of risk.[2]

Other measures of risk have understated true risk in a more subtle way. Some of the understatement is an accidental reflection of how much more we understand now about risk than we understood when the risk measure was introduced. Sometimes, risk measures which understate risk have a longer life than they deserve when some parties have a vested interest in perpetuating the understatement.

We turn to various risk measures to point out these potential pit-falls in the next section.

Perils and Pitfalls in the Measurement of Risk: The Impact of Selection Bias

A major bank with a significant portfolio of loans in the Asia Pacific region once retained one of the authors as a consultant in assessing the risk of the portfolio. As part of the assignment, the Chief Financial Officer of the bank requested a historical value at risk analysis of the bank's 100 largest counterparties in Asia. We of course did the work he requested but provided the analysis in the context of a letter which explained the nature of the results. The work was done, we explained, using industry standards and "best practices" of value at risk analysis but the 99th percentile "worst case" we described in the results section was in fact the "best case worst case" that the bank would experience.

[1] Assuming of course that the securities firm providing the insurance is riskless.
[2] Working paper, Cornell University and Kamakura Corporation, February 2001.

"How does that help me with bank regulators?" the Chief Financial Officer asked, not fully appreciating that we were trying to save his job. We explained the bias in value at risk patiently to him as follows:

1. You have selected 100 of your current counterparties in Asia
2. By definition, they are your counterparties now because they have not defaulted.
3. You asked us to do a historical analysis of value at risk using the bonds of these counterparties over the last 12 months
4. We have done that and produced a 99th percentile "worst case" for you as you requested BUT
5. It is the best case worst case scenario because we are doing a historical analysis on 100 names over a period in which none of them defaulted. The very nature of the risk measure and the nature of the counterparties guarantee that this seriously underestimates the credit risk in this portfolio. If the Chief Financial Officer wanted a more accurate assessment of the risk, the test should have either been based on 100 counterparties randomly chosen as of a year before (and indeed some of them subsequently defaulted) or we should have done a forward looking simulation that allows for default on a name by name basis.

He conceded the point and hedged himself with the Board, avoiding the career-ending scenario if a risk that turned out to be much worse than the advertised "worst case."

We now turn to some of the most common examples of this selection bias in measuring risk.

Biases in Return versus a Relative Benchmark

In the fund management and pension fund business, most portfolio managers are judged by measures of their "plus alpha" or risk adjusted return versus a predefined benchmark. An example might be a fund manager who managers a "large cap" portfolios of common stock and is given the Standard & Poor's 500 stock index as his benchmark index. This fund manager is told that he must keep tracking error of less than x% versus the index and that he is expected to earn a return in excess of the return on the S&P 500. This is a very, very common kind of risk and return measurement system on the "buy side" of financial markets.

What is wrong with this as a risk and return measurement system? Everything that was wrong in the story in the previous section.[3] We can see this by taking apart this risk management system piece by piece.

First, how is tracking error measured? There are many variations but in most cases a regression is run on two data series: the stock price returns for ABC Company and the returns on the S&P 500 index. By definition, tracking error will only be known for ABC Company if that stock price series exists (i.e. if the company has not been delisted from the exchange due to bankruptcy). Our risk measure for ABC Company is a "no bankruptcy risk" risk measure. Similarly, what about the bankruptcy risk of the S&P 500? Generally speaking, as a company's credit quality (and or market capitalization) declines, they are dropped from the S&P 500. This happens after the close of

[3] The authors would like to thank Leo de Bever, Barbara Zvan, Francois Gingras, and Barbara Chun of Ontario Teachers Pension Plan for their insights and inspiration of this section of the book.

business, and on the subsequent opening, the price of the stock drops considerably in a way that is unhedgable unless one predicts the change in S&P 500 composition. The S&P index is not affected by this drop but the holder of the common stock of the company dropped from the index will suffer a large loss. The composition of the S&P 500 has included some firms that went bankrupt while they were in the index. We suspect, however, that the index is ultimately restated to offset the impact of these failures.

Therefore, the most common performance measure used in the equity markets ignores credit risk almost completely, even though on average between 1.00% and 1.5% of listed American companies have defaulted over the 1963-2012 time period. We are substantially understating the risk component of this traditional performance measure.

How does the risk of bankruptcy affect our fund manager if the risk occurs? It depends on whether

A. ABC company is in the manager's portfolio or not
B. ABC company is in the S&P 500 or not
C. ABC goes bankrupt while it is in the index or is dropped from the S&P 500 for credit quality reasons before going bankrupt

If ABC Company is dropped from the index, the manager will have a decline of as much as 20% or more when the "drop" is announced, and yet there will be no corresponding change in the index. The manager will have a big "negative alpha" because he couldn't get out of ABC Company until after the announcement that ABC Company is not in the index any longer. If ABC in fact does go bankrupt while still in the index, the drop in stock price could be as much as 100% but at least the index will also be affected. We strongly recommend that managers with an equity index as a benchmark independently confirm how a bankruptcy affects the index.

If the manager owns a stock which goes bankrupt and it is not in the index, he has a negative alpha versus the index that tracking error completely ignores. Even if the company is merely dropped from the index, the manager will suffer a negative 20% or so stock price decline that is not offset by the movement of the index and which is not reflected in measured "tracking error." This same kind of problem affects fixed income performance measurement versus a benchmark like the former Lehman Brothers Government Bond Index (now renamed after Barclays Capital), but in more subtle ways.

We turn to another common measure of risk with similar pitfalls.

Historical Value at Risk: Selection Bias Again

There are as many approaches to value at risk as there are risk management analysts, but in general value at risk calculations fall into three basic categories:

A. A "historical" value at risk which measures dollar price changes for specific securities over a period in history
B. A variance/covariance approach which measures the variances and covariances of past returns and implies (from this implicit assumption of normally distributed returns) the Nth percentile worst case

C. A forward looking simulation of returns over the desired time horizon

The comments of this section apply only to methods A and B.

What are the concerns of these calculations which have been much used and much discussed over the past decade? Credit risk is completely ignored. We have the selection bias discussed earlier in this chapter, because (by definition) none of the current securities holdings would include the securities of a bankrupt firm.
Still, there are two other serious problems with these two approaches to value at risk which are just as serious:

1. They ignore all cash flows between time zero and the value at risk horizon (10 days, one month, three months, one year) since these approaches are effectively a single period analysis which implicitly assumes you own the same things at the end of the period that you owned at time zero
2. They cannot answer the question posed by the portfolio manager in the story by Robert Merton: what is the hedge and how much does it cost?

Both of these drawbacks seriously affect value at risk as a risk measure. In the end, on a fully credit adjusted basis, we need to know what the hedge is and how much it costs as Robert Merton and Robert Jarrow advise.

Monte-Carlo Based Value at Risk

Many thoughtful risk managers have sought to bring credit risk into traditional historical or variance-covariance value at risk by either simulating default/no default during the VAR time period or by using a "transition matrix" approach.

As an "interim fix" this approach is a step forward but again some key problems remain that understate the risk remain:

- Common macro-economic factors which drive correlated default are not specified. This will generally result in a significant underestimate of the "fat tails" from credit losses. Van Deventer and Imai [2003] discuss the impact of macro factors in great detail and we turn to them later in this book. The obvious lesson of the collapse of the Japanese bubble and the 2006-2010 credit crisis is that macro factors affect everything. A value at risk calculation prior to 2006 that did not simulate home price movements accurately would have been grossly misleading.
- The default probabilities are implicitly held constant over the time period. In reality, these default probabilities risk and fall through the credit cycle and holding them constant again understates the level of credit risk implicit in the portfolio
- The timing of defaults during the period is ignored, when in fact the exact timing is very important (see the next section for the reasons)
- Interim cash flows are again ignored (unless it is a multiperiod simulation)
- What's the hedge? Again, we can't answer the portfolio manager's question in the Robert Merton story.

Expected Losses on Tranches of Collateralized Debt Obligations
In the last decade, the market for credit default swaps and collateralized debt obligations has exploded and then collapsed completely. Ratings have been assigned and CDO tranches have been marketed on the basis of the "expected loss" on the CDO tranche, and the, with amazing speed, marked down to

"junk" or default.
The focus up through 2007 was purely on expected losses and ignored the fat tails that are typical of the worst part of the credit cycle. The index summarizes some of the tuition payments made by large international banks and securities firms in this regard.
As always, to understand the risk completely, we need the full probability distribution of the losses and most important, we need to be able to stress test it in order to calculate the hedge.
We turn now to the other side of the risk-return horizon.

Measuring Return: Market versus Accounting Returns
The nature of risk measurement is complex enough, as we have seen in previous sections, but there is more complexity involved: the definition of return. There are a few key determinants of what approach is used to return measurement:

- The nature of the institution involved
- The nature of the regulatory agencies involved
- The availability of observable market prices
- The time horizon of greatest concern

For managers of common stock portfolios, the decisions are pretty straightforward. The market prices of the common stock are observable, regulation is market-value based, and there is an observable daily benchmark. Moreover, industry expert bodies like the CFA Institute (formerly the Association for Investment Management and Research) have instituted a step by step approach to daily performance benchmarking.

In other industries, things are more complex. How do you deal with life insurance policy liabilities and an asset mix of equities, bonds, real estate and commodities? What would you do if your liabilities were pension fund obligations? What if they were passbook savings?

The result is often a system of splitting risk and return using a technique called "transfer pricing," which in this context has a much different meaning than its meaning to tax experts. We turn now to an introduction of how commercial banks have established this practice and what the current state of the art is. We discuss the mechanics of transfer pricing in Chapter 38 in light of the analytical tools we discuss in the first 37 chapters of this book.

Introduction to Transfer Pricing:

Extracting Interest Rate Risk in a Financial Accounting Context

The subject of transfer pricing represents one of the greatest differences in asset and liability management practice between large banks internationally and between banks and the fund managers mentioned in the previous section. In the United States and Australia, the so-called "matched maturity transfer pricing system" is accepted as an absolute necessity for the successful management of a large financial institution. In Japan, by way of contrast, the first bank to use multiple rate transfer pricing adopted a very simple system more than 20 years after Bank of America initiated a much more sophisticated system in 1973. As of this writing, transfer pricing has become an accepted banking discipline in a large number of countries, but it is still rare in the insurance industry.

The rational for "transfer pricing" as an alternative to the investment management business' reliance on a mark to market approach is a simple one. With employee counts of tens of thousands of people and hundreds of business units which invest in assets with no observable price, it is literally impossible to

manage a commercial bank on a mark to market basis at the working level. Transfer pricing is the trick or device by which the rules of the game are changed so that (a) actions are consistent with a mark to market approach but (b) the vehicle by which the rules of the game are conveyed are in financial accounting terms.

We start by presenting a selective history of the transfer pricing discipline in the United States. In Chapter 38, we provide a detailed description of the current “common practice” in transfer pricing, with examples of its practical use.

Bank of America, 1973-1979

Interest rate risk management has advanced in bursts of energy that usually follows a period of extremely high interest rates in the U.S. and other markets. Interest rate risk has, up until recently, been the main focus of transfer pricing. In the early 1970s, the entire home mortgage market in the United States was made up of fixed rate loans. The floating rate mortgage had yet to be introduced. Interest rates were beginning to be deregulated, and the first futures contracts (on U.S. Treasury bills) had not yet been launched. In this environment, a rise in market interest rates created a double crisis at banks with a large amount of retail business: consumer deposits flowed out of banks (which had legal caps on deposit rates at that time) into unregulated instruments like Treasury bills or newly introduced money market funds. Simultaneously, the banks suffered from negative spreads on their large portfolios of fixed rate loans.

A spike in interest rates beginning in 1969 triggered a serious reexamination of financial management practice at the largest bank in the United States (at the time), the Bank of America. A team of executives that included C. Baumhefner, Leland Prussia, and William "Mack" Terry recognized that changing interest rates were making it impossible for existing management "accounting" systems to correct allocate profitability among the bank's business units and to set responsibility for managing interest rate risk. Until that time, the bank had been using a single internal transfer pricing rate that was applied to the difference between assets and liabilities in each business unit to allocate some form of interest expense or interest income in a way that the units' assets and liabilities were equal. This rate, know as the “pool” rate system for transfer pricing, was a short-term interest rate at Bank of America calculated as a weighted average of its issuance costs on certificates of deposit.

The liabilities of such a single rate “pool rate” system are outlined in later chapters of this book. Given the sharp changes in rates that were beginning to look normal in the United States, senior management at the Bank of America came to the conclusion that the existing system at the bank was dangerous from an interest rate risk point of view and made proper strategic decision making impossible. Strategic decision making was handicapped because the single rate transfer pricing system made it impossible to know the true risk-adjusted profitability of each business unit and line of business. The numbers reported by the existing system mixed results due to business judgment ("skill") and results due to interest rate mismatches ("luck," which could be either good or bad).

Although the bank had a very sizable accounting function, the responsibility for designing the new transfer pricing system was given to the brilliant young head of the Financial Analysis and Planning Department of the bank, MIT graduate Wm. Mack Terry. This division of labor at the bank was a beginning of a trend in U.S. banks that resulted in the total separation of the preparation of official financial reports, a routine and repetitive task that required a high degree of precision and a low level of

imagination, from management information (not managerial "accounting") on how and why the bank was making its money. Mack Terry, who reported directly to Chief Financial Officer Lee Prussia, took on the new task with such vigor that most knowledgeable bankers of that era acknowledge Mack Terry as the father of the new discipline of "matched maturity transfer pricing."

Mack Terry and the bank were faced by constraints that continue to plague large banks today:

- The bank was unable to get timely information from loan and deposit mainframe computer application systems regarding the origination, maturity, payment schedules, and maturities of existing assets and liabilities
- The bank had neither the time nor the financial resources to develop a new transfer pricing system on a main frame computer.

Personal computers, needless to say, were not yet available as management tools.

The danger of future rate changes and incorrect strategic decisions was so great that bank management decided to make a number of crude but relatively accurate assumptions in order to get a "quick and dirty" answer to the bank's financial information needs. The fundamental principle developed by the Financial Analysis and Planning Department was the "matched maturity" principle. A three year fixed rate loan to finance the purchase of an automobile would be charged a three year fixed rate. A thirty year fixed rate mortgage loan would be charged the cost of fixed rate 30-year money. At other U.S. banks, it was more common to use the bank's average cost of funds or its marginal cost of new three month certificates of deposit. The average cost of funds method had obvious flaws, particularly in a rising rate environment, and it ultimately contributed to the effective bankruptcy of Franklin National Bank and many U.S. savings and loan associations. We discuss how these transfer pricing rates are calculated in later chapters.

The Financial Analysis and Planning ("FAP") Department team agreed that the bank's "marginal cost of funds" represented the correct yield curve to use for determining these matched maturity transfer pricing rates. The bankers recognized that the bank had the capability to raise small amounts of money at lower rates but that the true "marginal" cost of funds to the bank was the rate that would be paid on a large amount, say $100 million, in the open market. This choice was adopted without much controversy, at least compared to the controversy of the "matched maturity" concept itself. In more recent years, many banks have used the London Interbank Offered Rate and the interest rate swap curve as a proxy for their marginal costs of funds curve. It is clear from the perspective of today, however, that the Libor-swap curve is a poor substitute for an accurate marginal cost of funds curve.

One of the first tasks of the FAP team was to estimate how the internal profitability of each unit would change if this new system were adopted. It was quickly determined that there would be massive "reallocations" of profit and that many line managers would be very upset with the reallocation of profit away from their units. In spite of the controversy that management knew would occur, management decided that better strategic decision making was much more important than avoiding the short run displeasure of half of the management team (the other half were the ones who received increased profit allocations).

Once a firm decision was made to go ahead with the new system, implementation decisions had to be dealt with in great detail. The systems "queue" was so long at the bank that it was impossible to keep track of transfer pricing rates on a loan by loan or deposit by deposit basis. The first unit of measurement, then, was to be the "portfolio" consisting of otherwise identical loans which differed only in maturity and rate, not in credit risk or other terms. In addition, the bank (to its great embarrassment) did not have good enough data on either its own loans or its historical cost of funds to reconstruct what historical transfer pricing rates would have been for the older loans and deposits making up most of the bank's balance sheet. Estimates would have to do.

These estimates were at the heart of the largest of many internal political controversies about the matched maturity transfer pricing system. Much of the fixed rate real estate portfolio was already "under water"--losing money--by the time the transfer pricing system was undergoing revision. If the current marginal cost of funds were applied at the start of the transfer pricing system to the mortgage portfolio, the profitability of the portfolio would have been essentially zero or negative. On the other hand, Mack Terry and his team recognized that in reality this interest rate risk had gone unhedged and that the bank had lost its interest rate bet. Someone would have to "book" the loss on the older mortgage loans.

Ultimately, the bank's asset and liability management committee was assigned a "funding" book that contained all of the interest rate mismatches at the bank. The funding book was the unit that bought and sold all funds transfer priced to business units. At the initiation of the matched maturity system, this funding book was charged the difference between the historical marginal cost of funds that would have been necessary to fund the mortgage portfolio on a matched maturity basis and the current marginal cost of funds. This "dead weight loss" of past management actions was appropriately assigned to senior management itself.

Because of the lack of data, the bank's early implementation of the matched maturity system was based on the use of moving average matched maturity cost of funds figures for each portfolio. While this approximation was a crude one, it allowed a speedy implementation that ultimately made the bank's later troubles less severe than they otherwise would have been.

Finally, the system was rolled out for implementation with a major educational campaign aimed at convincing lending officers and branch managers of the now well accepted discipline of "spread pricing," pricing all new business at a spread above the marginal cost of matched maturity money[4]. Controversy was expected, and expectations were met. A large number of line managers either failed to understand the system or didn't like it because their reported profits declined. Soon after the announcement of the system, Mack Terry began receiving anonymous "hate mail" in the inter-office mail system from disgruntled line managers.

[4] From a 2012 perspective, this emphasis on the marginal cost of funds as a basis for pricing is showing its age. We discuss problems with this concept in later chapters as part of the continual battle between those with a financial accounting basis ("how can we make money if we don't charge at least as much of our cost of funds") and those with a market orientation ("this credit is priced 10 basis points over market for AA-rated companies and the fact that we are a triple BBB bank is irrelevant"").

Nonetheless, senior management fully backed the discipline of the new system, and for that the Bank of America receives full credit as the originator of the matched maturity transfer pricing concept. In an ironic footnote, the bank suffered heavily from an interest rate mismatch as interest rates skyrocketed in the 1979-1980 time period. The transfer pricing system made it clear that senior management was to blame since the asset and liability management committee had consciously decided not to hedge most of the interest rate risk embedded in the bank's portfolio.

First Interstate, 1982-1987

In the years to follow, a number of banks adopted the matched maturity transfer pricing system. The oral history of U.S. banking often ranks Continental Illinois (later acquired by the Bank of America) as the second bank to move to a matched maturity transfer pricing system, not long after the Bank of America implemented the idea. At most banks, however, the idea was still quite new and the early 1980s were largely consumed with recovering from the interest rate-related and credit risk-related problems of the 1979-1981 period. By the mid-1980s, however, banks had recovered enough from these crises to turn back to the task of improving management practices in order to face the more competitive environment that full deregulation of interest rates had created. First Interstate Bancorp, which at the time was the 7th largest bank holding company in the United States, approached the transfer pricing problem in a manner similar to that of many large banking companies in the mid-1980s.

First Interstate's organization was much more complex than Bank of America's since the company operated 13 banks in 9 western states, all of which had separate treasury functions, separate management, separate government regulation, and separate systems[5]. In addition, the bank holding company legal entity First Interstate Bancorp had a number of "non-bank" subsidiaries that required funding at the holding company ("parent") level. Most U.S. bank holding companies consisted of a small parent company whose dominant subsidiary was a lead bank that typically made up 90 percent of the total assets of the company in consolidation. The lead banks were almost always considered a stronger credit risk than the parent companies because of the existence of Federal deposit insurance at the bank level but not the parent level and because of the richness of funding sources available to banks compared to bank holding companies.

In the First Interstate case, things were more complex. The lead bank, First Interstate Bank of California, represented only 40 percent of the assets of the holding company and therefore its credit was generally felt by market participants to be weaker than that of the parent. Moreover, the First Interstate Banks did not compare funding needs, and as a result, it was often said that a New York bank could buy overnight funding from one First Interstate Bank and sell it to another First Interstate Bank at a good profit. The transfer pricing system at First Interstate had to cure this problem as well as addressing the correct allocation of profits and interest rate risk as in the Bank of America case.

Management took a two-fold approach to the problem. At the holding company level, the corporate treasury unit began "making markets" to all bank and non-bank units within the company. Because First Interstate was much more decentralized than the Bank of America, funds transfers between the holding company and subsidiaries were voluntary transfers of funds, not mandatory. In addition, since each unit

[5] The multi-state nature of commercial banking is commonplace in the U.S. and most other countries in 2012 but it was a significant barrier to implementation of comprehensive risk management in the mid-1980s, just as national barriers are today.

was a separate legal entity, a transfer pricing transaction was accompanied by the actual movement of cash from one bank account to another.

The holding company transfer pricing system began in early 1984 under the auspices of the holding company's funding department. The department agreed to buy or sell funds at its marginal cost at rates which varied by maturity from 30 days to 10 years. No offers to buy or sell funds were to be refused under this system. The transfer pricing system, since it was voluntary, was received without controversy and actually generated a high degree of enthusiasm among line units. For the first time, the units had a firm "cost of funds" quotation that was guaranteed to be available and which could be used to price new business in line units on a no-interest rate risk basis. Demand from line units was very strong, and the parent company became an active issuer of bonds and commercial paper to support the strong funding demand from both bank and non-bank subsidiaries.

Among bank subsidiaries, the "on-demand" transfer pricing system had the effect of equalizing the cost of funds across subsidiary banks. High cost of funds banks immediately found it cheaper to borrow at the lower rates offered by the parent company. Regulatory restrictions kept the parent company from borrowing from subsidiary banks, but there was an equivalent transaction that achieved the same objective. After the transfer pricing system had been in operation for some months, the parent company had acquired a substantial portfolio of certificates of deposit of subsidiary banks. These certificates of deposit could be sold to other banks within the system. By selling individual bank certificates of deposit to other First Interstate banks, the holding company reduced the size of this portfolio and effectively "borrowed" at its marginal cost of funds, the yield it attached to the certificates that it sold.

The transfer pricing system at the holding company did have implementation problems of a sort. Generally, rates were set at the beginning of a business day and held constant for the entire day. The parent company soon noticed that borrowings from affiliates, and one subsidiary in particular, would increase when open market rates rose late in the day, allowing subsidiaries to "arbitrage" the parent company treasury staff by borrowing at the lower rate transfer price set earlier in the day. This got to be a big enough problem that rates for all borrowings above $10 million were priced in "real time." Most large banks use a similar "real time" quotation system now for pricing large corporate borrowings. What will surprise many bankers, however, is that "in-house" transactions ultimately also have to be priced in real time in many cases.

At the same time that the parent company was implementing this system, the parent company's asset and liability management department and the financial staff of First Interstate Bank of California ("FICAL") began to design a Bank of America-style matched maturity transfer pricing system for FICAL. Personal computer technology at the time did not permit PCs to be used as the platform, so the company undertook a very ambitious mainframe development effort. After a complex design phase and a development effort that cost close to $10 million and two years of effort, the system was successfully put into to action with considerably less controversy than in the Bank of America case. Such a system today would cost much, much less from a third party software vendor.

The biggest practical problem to arise in the FICAL transfer pricing system was a subtle one with significant political and economic implications. During the design phase of the system, the question was raised about how to handle the pre-payment of fixed rate loans in the transfer pricing system. The financial management team at First Interstate dreaded the though of explaining "option-adjusted"

transfer pricing rates to line managers and decided to do the following: internal transfer pricing would be done on a "non-prepayable" basis. If the underlying asset were prepaid, then the treasury unit at the center of the transfer pricing system would simply charge a "mark-to-market" prepayment penalty to the line unit, allowing it to extinguish its borrowings at the same time that the asset was prepaid. This simple system is standard practice in many countries, including the retail mortgage market in Australia.

This decision led to unforeseen consequences. During the latter half of the 1980s and continuing into the 1990s, interest rates declined in the United States and were accompanied by massive prepayments of fixed rate debt of all kinds. As a result, the FICAL transfer pricing system's reported profits for line units were soon dominated by huge mark-to-market penalties that line managers didn't have control over and generally didn't understand.

Bank management quickly moved to a transfer pricing system that allowed for "costless prepayment" by incorporating the cost of a prepayment option in transfer prices. This trend is firmly established as standard practice in most banks today. With this political and organizational background on transfer pricing in mind, we now put it in the context of performance measurement and capital regulation. We return to the mechanics of transfer pricing in Chapter 38.

Performance Measurement and Capital Regulation

Our primary purpose in this chapter and subsequent chapters is to present “common practice” for performance measurement in various wings of the financial services business and to contrast the approaches used by different institutions. When we see institution A and institution B managing similar risks but using different approaches to the measurement and management of risk, we will carefully note the differences and seek an explanation. Often, institutional barriers to change delay the synthesis and common practices that one would expect from two different groups of smart people managing the same kind of risk.

This difference is particularly stark in the role of capital in risk management and performance measurement. Commercial banks are extensively focused on capital-based risk measures, and almost no other types of financial services businesses are. Why? We begin to answer this question with some additional perspectives on the measurement and management of risk. We then talk about managing risk and strategy in financial institutions, business line by business line, and how capital comes into play in this analysis. We then turn to the history of capital-based risk regulations in the commercial banking business and discuss its pros and cons.

Perspectives on Measuring Risk: One Source of Risk or Many Sources of Risk?

Risk management has been an endless quest for a single number which best quantifies the risk of a financial institution or an individual business unit. In Chapter 1, we discussed how Robert Merton and Robert Jarrow have suggested that the value of a put option that fully insures the risk is the current “best practice” in this regard. This risk measure also provides a concrete answer to the question which best reveals the strengths or weaknesses of a risk management approach: “What is the hedge?”

Over the last 50 years, the evolution of risk management technology toward the Merton and Jarrow solution has gone through two basic steps:

1. Assume there is a single source of risk and a risk measurement statistic consistent with that source of risk
2. Recognize that in reality there are multiple sources of risk and revise the measure so that we have an integrated measure of all risks and sources of risk

Interest Rate Risk Management Evolution

In Chapters 3 to 14, we review traditional and modern tools of interest rate risk management. The first sophisticated interest rate risk management tool was the Macauley [1938] duration concept. In its traditional implementation, managers who own fixed income assets shift a yield curve up or down by a fixed percentage (often one percent) and measure the percentage change in the value of the portfolio. This measure shifts yield curves at all maturities by the same amount, essentially assuming that there is only one type of interest rate risk, parallel movements in the yield curve. As we show in Chapter 3, that is not the case in reality and therefore the duration concept measures some but not all interest rate risks.

The next step forward in sophistication was to recognize that yield curves in fact move in non-parallel ways. The term structure models that we discuss in Chapters 6-14 allow analysts to move one or more key term structure model parameters (typically including the short-term rate of interest) and see the response of the full yield curve. This provides a more sophisticated measure of duration that allows for non-parallel shifts of the yield curve.

Fixed income managers in both commercial banks and other financial services companies recognize, however, that both of these approaches are abstractions from reality. As we show in Chapter 3, there are multiple factors driving any yield curve, whether it is a credit-risk-free government yield curve or a yield curve that embeds the default risk and potential losses of a defaultable counterparty.

In recognition of these N factors driving the yield curve, fixed income managers have overlaid practical supplements on the theoretical models discussed above:

1. Interest rate sensitivity "gaps" (see legacy approaches to interest rate risk in Chapter 12) which show the mismatches between the maturities of assets and liabilities (if any) period by period. These gaps provide visibility to managers that allows them to implicitly reduce the risk that interest rates in gap period K move in a way different from what the assumed duration or term structure model specifies.
2. Multiperiod simulation of interest rate cash flows using both single scenario interest rate shifts and true monte-carlo simulation of interest rates. Managers can measure the volatility of cash flows and financial accounting income period by period, again to supplement the interest rate models which may understate the true complexity of interest rate risk movements

What is the equivalent of the Merton and Jarrow "put option" in the interest rate risk context? It is the value of a put option to sell the entire portfolio of the financial institution's assets and liabilities at a fixed price at a specific point in time. From a pension fund perspective, a more complex string of options which guarantee the pension fund's ability to provide cash flow of X(t) in each of N periods to meet obligations to pensioners would be necessary.

We return to this discussion in more detail in Chapters 36 to 41.

Equity Risk Management Evolution

Risk management in the equity markets began in a similar fashion with the powerful and simple insights of the capital asset pricing model. This model initially suggested that all common stocks had their returns driven by a single common factor, the return on the market as a whole, and an idiosyncratic component of risk unique to each stock. By proper diversification, the idiosyncratic component of risk could be diversified away. The single risk factor, the exposure to changes in return on the market, was measured by the beta of both individual common stocks and on the portfolio held by the equity manager.

As in the interest rate case, risk managers quickly seized on the insights of the model and then generalized it to recognize there are multiple drivers of risk and the multi-period nature of these risk drivers. Earlier in this chapter, we introduced the need to add credit risk to the multiperiod equity risk models which reflect current "common practice."

Option Risk Management Evolution

The Black-Scholes option model brought great hope to options risk managers in much the same way that the capital asset pricing model did. Options portfolio managers (and managers of broader portfolios of instruments with options characteristics) quickly adopted the "delta" of an options position, the equivalent position in the underlying stock, as an excellent single measure of the risk of an options position.

With experience, however, analysts realized that there was in fact a "volatility smile." The volatility smile was the graph which showed that the Black-Scholes options model implied a different level of volatility at different strike prices with the same maturity. Analysts have tended to "bend" the Black-Scholes options modeling by using different volatilities at different strike prices rather than looking for a more elegant explanation for the deviation of options prices from the levels predicted by the Black-Scholes model. In effect, analysts are using the Black-Scholes model as an N factor model, not a single factor model with the stock price as the risk driver with a constant volatility. See Jarrow [March 15, 2011] on the dangers of such an ad hoc approach.

Credit Risk Management Evolution

The integration of credit risk, market risk, asset and liability management, liquidity risk management, and performance measurement is one of the central themes of this book. Analysts have moved from traditional credit analysis on a company by company basis to more modern technology over the last 25 years. The more traditional credit analysis was a single company analysis that involved financial ratios and occasionally simulation to estimate a default probability or rating that summarizes the risk level of each company, one by one.

Now the kind of quantitative credit models that we discuss in Chapters 15 to 18 allow us to recognize the M common macro-economic risk factors that cause correlated defaults among a portfolio of companies. These insights allow financial institutions to move to a macro-hedging program for a large portfolio of credits for the first time.

It also provides a framework for the quantification of risk as Robert Merton and Robert Jarrow have suggested, with a put option on the credit portfolio.

We now turn to the management process in commercial banks, insurance companies, and pension funds and analyze how and why capital came to play a role in commercial banking risk management and strategy but not in other institutions.

Managing Risk and Strategy, Business by Business

One of the major differences between financial institutions is the number of people who are involved in asset generation. In a life insurance company, a vast majority of the people are involved in the generation, servicing and pricing of life insurance policies. On the asset side, investment activities are much more "wholesale" and the number of people involved per one million dollars of assets is much less. In an asset management company, the total staff count is also generally less than 1,000 people. The same is true for pension funds. Bank of America, by contrast, had a staff of 288,000 at the end of 2011. For very large commercial banks, the number of people involved is very large because almost all banks these days are heavily involved in the retail generation of assets and liabilities.

There is another important difference between commercial banks and the rest of the financial community. For asset managers, pension funds, and insurance companies, most of the people managing the assets of the company work in one central location. In a commercial bank, there can be thousands of branches involved in asset generation across a broad swath of geography.

For that reason, in what follows we will largely discuss management of risk and business strategy in a banking context because most of what happens in other financial institutions is a "special case" of the banking business, including even the actuarial nature of liabilities of all of these types of financial institutions.

Risk and Strategy Management in a Complex Financial Institution

From the introduction to this chapter, we know that in the banking business various parts of the institution are managed by one or more bases from both a risk and return point of view:

- A mark to market orientation for the institution as a whole, the trading floor, and the transfer pricing book
- A financial accounting point of view for both the institution as a whole and almost all non-trading business units

Moreover, as we say in the introduction to this chapter, multiple risk measures are applied depending on the business unit involved:

- **Complex interest rate risk measures** are applied to the institution as a whole, the trading floor, and the transfer pricing book
- **Less sophisticated interest rate risk measures** are usually sufficient for most business units because the transfer pricing process in most institutions removes the interest rate risk from line units. Depending on the financial institution, the basis risk (a lack of synchronicity between the asset floating rate pricing index and market rates) usually stays with the business unit it the head of that business has the responsibility for selecting that pricing index for the bank
- **Yes/no credit risk measures** are applied at the "point of sale" in asset generating locations, often controlled by a central location
- **Portfolio credit risk measures** are applied by the business unit responsible for maintaining or hedging credit risk quality after the generation of assets
- **Mark risk and credit risk measures** are almost applied together on the trading floor because so much of the mark risk taking business has taken the form of credit risk arbitrage, even in the interest rate swap market. The percentage of derivatives contracts where at least one side is held by one of the top ten dealers is astronomical, especially as the bank merger trend continues in the United States.
- **Return measures are stated either on a market-based total return basis or a net income basis,** usually after interest rate risk is removed via the transfer pricing system

Figure 1

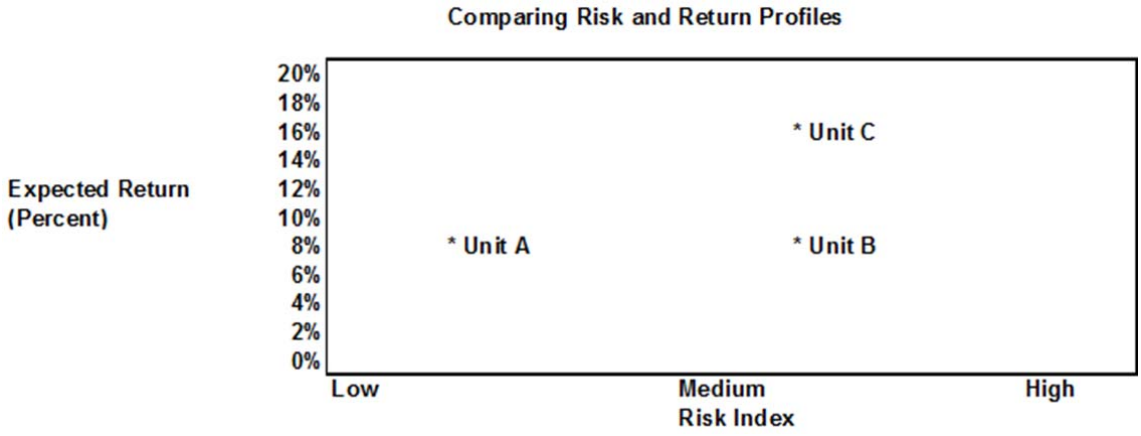

We defined risk management in Chapter 1 as the discipline by which management is made aware of the risk and returns of alternative strategies at both the transaction level and the portfolio level. Once we have coherent measures of risk and return, management often runs across situations like the graph presented above, where the risk and returns of three alternative business units A, B, and C are plotted. Unit C is clearly superior to Unit B, because it has a higher expected return than B. Similar, Unit A is superior to Unit B because Unit A has the same expected return but less risk than Unit B. It is much harder to compare Units A and C. How to we select business units which are "over achievers" and those which are "under achievers"? This is one of the keys to success in financial institutions management.[6] The Sharpe ratio is commonly cited as a tool in this regard, but there is no simple answer to this question. We hope to provide extensive tools to answer this question in the remainder of the book. There are a number of hallmarks to best practice in this regard:

- A tool which cannot provide useful information about which assets are the best to add at the **transaction level** cannot be useful at the portfolio level, since the portfolio is just a sum of individual transactions
- **Buy low/sell high** contains a lot of valuable wisdom, as simple as it seems. If the financial institution is offered the right to buy asset 1, with a market value of 100, at a price of 98, it should generally say yes.
- **If the bank has to choose** between asset 1 and asset 2, which has a market value of 100 and can be purchased for a price of 96, the bank should buy asset 2. If the bank is solvent and its risk is not adversely affected at the portfolio level, the bank should buy both. The market price summarizes risk and return into one number.
- Risk-return tools should provide value from **both a risk management perspective and the perspective of shareholder value creation**. These tools are not separate and distinct—they are at the heart of good management.

What causes financial institutions to fail?

[6] See Chapter 1 of *Financial Risk Management in Banking* by Dennis Uyemura and Donald R. van Deventer for more on this chart.

In deciding at the transaction and portfolio levels which assets to select for a financial institution as a whole, the criteria for selection are not independent of which assets and liabilities are already on the books of the financial institution.[7]

The dominant reason for financing institutions to fail in the 21st century is credit risk, now that the interest rate risk tools and techniques discussed in the next few chapters are used with care. Even before the obvious lessons of the 2006-2010 credit crisis, this had been confirmed by many studies of bank failures in the United States, most importantly the Financial Institutions Monitoring System introduced by the Financial Institutions Monitoring System implemented by the Board of Governors of the Federal Reserve in the United States in the mid-1990s.[8] The effort to better predict financial institutions failure and the subsequent losses to deposit insurance funds gained new momentum with the announcement of its new Loss Distribution Model by the Federal Deposit Insurance Corporation on December 10, 2003.[9] The FDIC Loss Distribution Model, with our colleague Prof. Robert A. Jarrow as lead author, correctly predicted that the FDIC insurance fund was grossly underfunded a full 3 years before the advent of the 2006-2010 credit crisis. We discuss the loss distribution model in great detail in later chapters.

Given the importance of credit risk in the failure of financial institutions, an integrated treatment of credit risk, market risk, asset and liability management, and performance measurement is critical or we may totally miss the boat from a risk perspective. In part for this reason, capital has become a critical component of both regulatory and management supervision of the diverse business units involved in a complex financial institution.

The Role of Capital in Risk Management and Business Strategy

Given the diverse business units in a large, complex financial institution, lots of complexities come into play. First of all, some types of financial institutions have a very large institutional divide between the asset side of the organization and the liability side. Pension funds, life insurance companies, property and casualty insurance companies, and asset management companies are almost always clearly divided between the two halves of this balance sheet. This means that most managers either oversee assets or liabilities but almost never both. The author would argue that this split is unfortunate from a "best practice" risk management point of view since many risk factors affect both sides of the balance sheet. Nonetheless this political reality simplifies managerial reporting.

In commercial banks, senior management is faced with three types of business units:

- Those who own only assets (specialized lending offices)
- Those who own only liabilities (the unit issuing wholesale certificates of deposit)
- Those who own both assets and liabilities (retail branches, business banking offices, trading floors doing repurchase agreements, etc.)

[7] Donald R. van Deventer and Kenji Imai discuss this issue extensively in *Credit Risk Models and the Basel Accords* [2003] in Chapters 1 and 2.
[8] Ibid, Chapter 9
[9] A copy of the FDIC Loss Distribution Model is available on the research page on www.kamakuraco.com .

The largest number of business units falls into the last category, where the business units have both assets and liabilities. Management has two challenges:

1. Put the business units on a common basis
2. Adjust differences between units so that an apples to apples performance comparison can be done

Implicit in these two tasks are the assumptions that a common basis is possible and that backwards looking market data or backwards looking financial statements are a good indicator of future performance. The latter assumption is more tenuous than the first.[10]

There are two steps in accomplishing these two tasks:

A. Specify the method by which business units with different risks can be compared to each other (especially in the base of a comparison like Unit A and Unit C graphed above)
B. Do the calculations that make this comparison possible

The "common practice" in this regard is what we focus on in this section. In the last section of this chapter, we talk about the pitfalls of the common practice and spend the last chapters of this book refining these calculations using the tools that we develop in the rest of the book.

- Step 1: **Choose the basis for the calculations** (market value basis or financial accounting basis). As mentioned earlier in this chapter, most commercial banks use a mix of these approaches but choose a financial accounting basis for the vast majority of business units. We assume that choice is made here.
- Step 2: **Select the adjustments to nominal net interest income** for the business unit that must be made to allow for units of different risk to be compared
- Step 3: **Select a common basis for comparing business units** of different levels of next income to better understand relative risk adjusted performance.

For a pension fund, this process is pretty easy and straightforward. The basis is the market-value-based return on the asset portfolio being analyzed. There are no adjustments to this basis except to subtract the relative performance of the benchmark index by which performance is judged. A check is also done to ensure that the risk (usually tracking error) of the funds invested were within an acceptable tolerance around the risk of the benchmark.
For a bank, this process is much more convoluted. Some of the complexity is productive and some is done to insure maximum employment in the performance measurement unit.

Here is a sample of the kinds of calculations usually made on a sample portfolio of external transactions like this:

[10] Ibid., Chapter 9.

Assets

3 year fixed rate auto loans	100

Liabilities

90-day certificates of deposit	120

All assets are funded with matched maturity borrowings from the transfer pricing center as discussed earlier in this chapter. Transfer prices are assigned using the techniques of the following chapters. All funds from the certificates of deposits gathered are sold to the transfer pricing center. "Capital" is then assigned depending on the nature of the assets and the nature of liabilities. Proceeds of the capital are sold back to the transfer pricing center with a maturity to match the period over which the amount of capital is held constant. The result is a financial-accounting-based management information system that would look something like this:

Table 1

Risk-Adjusted Return on Capital

Table 1
Risk-Adjusted Return on Capital

Assets	Balance Sheet Amount	Yield (Percent)	Annual Income
3-year fixed rate auto loans	100.0	6.76	6.76
90-day loans to the transfer pricing center	120.0	1.50	1.80
Capital-related loans to the transfer pricing center	7.7	1.50	0.12
Total assets	227.7		8.68
Liabilities			
90-day certificates of deposit	120.0	1.37	1.64
3-year borrowings from transfer pricing center	100.0	5.18	5.18
Total Liabilities	220.0		6.82
Capital			
Capital assigned for 90-day certificates of deposit	1.2		
Capital assigned for 3-year auto loans	6.5		
Total capital assigned	7.7		
Total liabilities and capital	227.7		1.86
Risk-adjusted return on auto loans			24.31%
Risk-adjusted return on deposits			13.00%
Total risk-adjusted return			24.05%

The portfolio of three year auto loans yields 6.76%. The exact-day count matched maturity transfer pricing system assigns $100 in borrowings to fund the three year auto loans on a matched maturity basis. These borrowings have a yield of 5.18%. We discuss how to calculate a transfer price for auto loans like these in Chapters 4 and 5. In a similar way, the 90-day funds raised by the business unit from issuing certificates of deposit are sold on a matched maturity basis to the transfer pricing center at 1.50%, a premium of 0.13% over the 1.37% cost of the certificates of the deposit to the business unit. This premium is a function of the marginal cost of issuing large-lot 90 day certificates of deposit, not an arbitrary premium.

The next step is where there is a considerable lack of consensus among financial institutions for the reasons we outline in the next section. Most leading international banks assign “risk adjusted capital” to each asset and (less often) each liability. The rationale for this assignment of capital varies widely from bank to bank but they include the following common themes:

1. The total amount of capital assigned, when added across business units, should sum to the total amount of capital in the organization. The "total amount of capital in the organization" in some banks means financial accounting capital and in other banks it means the market value of the bank's common stock.
2. The amount of capital assigned to each asset depends on the risk of the asset class. Van Deventer and Imai [2003] show how the Shimko-Tejima-van Deventer [1993] model of risky debt can be used to do this while considering both the interest rate risk and the credit risk of the asset category.
3. The amount of capital assigned to each liability depends on its maturity and the "liquidity" protection the deposits provide. This is perhaps the most controversial step in the calculation
4. Risk-adjustments come in two places. Interest rate risk is removed from the business unit by the assignment of matched maturity transfer pricing assets and liabilities to each transaction. Credit risk is reflected in the risk-weighted capital assignment.
5. The result is a "risk-adjusted return on capital." For the auto loan portion of the portfolio, the interest-rate-risk-adjusted net income is the $6.76 in income on the auto loans less the $5.18 cost of matched maturity funding on that portfolio, for an interest-rate-risk-adjusted net income of $1.58. This is a 24.31% return on the $6.50 in risk-adjusted capital assigned to the auto loan portfolio. Note that the interest rate adjustment comes in the numerator of the return calculation (since net income is net of interest rate risk) and the credit risk adjustment comes in the denominator via the risk-adjusted capital assigned.
6. In a similar way we can calculate a 13% risk-adjusted return on the deposit portfolio and a 24.05% risk-adjusted return for the business unit as a whole.

For a comprehensive survey of this risk-adjusted capital allocation system, interested readers should see Matten [1996].

Common practice in the banking industry says that, armed with these figures, we can do the following:

1. We can clearly judge which business unit creates the most shareholder value
2. We can rank the business units from best to worst
3. We can decide which business units should be expanded (receive more capital) and which business units should be discontinued
4. We can carry this managerial process to the individual transaction level and set correct pricing for every asset and liability for a clear "accept/reject" criterion. More generally, this risk-adjusted capital allocation system correctly calculates the true value/market price of each asset and liability and the bank should buy assets which have an offered price (not yield) below this in the market place and issue liabilities that have a market price (not yield)

If this is true, we have a powerful tool on our hands. If this is more than a touch of overselling the concept, we need to be careful how we use it. We turn to those issues next.

Capital-Based Risk Management in Banking Today: Pros and Cons

The previous section outlines the "common practice" for risk-adjusted asset and liability selection in the commercial banking business. We discussed the same process in the pension fund industry earlier in this chapter:

1. Select a benchmark index that best matches the nature of the assets being analyzed
2. This benchmark should be a naïve strategy that any institution could achieve with a minimum of internal staff and analysis (such as buy every stock in the S&P 500 or buy every U.S. Treasury bond outstanding)
3. Measure the "tracking error" of the asset class being studied to insure that the asset class is within the limits for tracking error versus the benchmark
4. A good manager who has assets whose total return in excess of the benchmark return is considered "plus alpha," someone who has generated a positive risk adjusted return.

This approach is starkly different than the approach taken by the banking industry:

a. It is based on market returns, not financial accounting returns
b. Risk adjustment is reflected in the returns of the benchmark and the limits on tracking error
c. There is no "transfer pricing" of matched maturity funds
d. There is no capital allocation
e. There is no risk adjusted return on capital calculation
f. There is also no pretense that the methodology provides a tool for pricing the assets in question, only that skill of the manager selecting the assets can be measured

While we discussed the concerns about this calculation above, it has a simplicity and elegance and lack of arbitrariness that distinguishes it as a managerial tool.

Why did the banking industry turn to such a convoluted calculation instead?

- The banking industry has many more people to manage
- The banking industry has many classes of assets and liabilities that don't have observable market prices
- The banking industry feels it cannot manage its huge number of business units on a mark to market basis
- The banking industry has many business units which have both assets and liabilities so the concept of "capital" is a natural one
- In many business units, they originate more liabilities than assets so there has to be some "adjustment" to a normal basis
- Many of the regulatory restrictions on banks revolve around capital

There are a few more important points about the common practice in banking risk-adjusted capital that need to be mentioned. Van Deventer and Imai [2003] in Chapters 9 and 10 go into these points at length:

- The risk-adjusted return on capital is backward looking, not forward looking. It has long been known that past returns on a security are fully reflected in its market price and therefore provide no guidance on the level of future returns

- The risk-adjusted return on capital assigns capital in a well-intended way unsupported by a sophisticated theory of capital structure (although the approach van Deventer and Imai [2003] recommend is a step in this direction).
- No theory of asset valuation in financial economics yet incorporates the capital of the potential buyer in the valuation formula

Are these concerns just academic concerns or do they indicate severe problems with the approach taken in the banking industry?
A few examples prove that the problems are severe:

1. The common practice in banking industry risk-adjusted return on capital indicates a negative risk-adjusted return on holding government securities because they have a negative spread versus the marginal cost of funds for the bank. This signal says no banks should own risk-free government securities, but in fact almost all banks do. Therefore the decision rule is wrong. Banks seek to overcome this obvious "miss" in the rule with ad hoc adjustments
2. Anecdotal evidence from many banks who have carried this analysis to the transaction level has accumulated that the prices indicated are consistently "off the market" which is no surprise, since the transactions are not benchmarked in current market prices
3. If this method were superior to the approach used by fund managers, then fund managers would use it

Calculation of "shareholder value added," which measures true cash flow rather than financial net income, is a partial step in the right direction. It falls short of the mark in applying a minimum hurdle rate for acceptable return which is not consistent with actual returns in the market (in contrast to the approach that pension funds take). For more on the conventional implementation of shareholder value-added, see Chapter 2 of Uyemura and van Deventer [1993].

As a result of these concerns about risk-adjusted capital allocation, the authors recommend that the risk-adjusted capital allocation approach be restricted to the largest consolidated business units in the banking business (for example, wholesale corporate lending, investment banking, retail banking, and so on) and not be carried down to the smaller business unit level or the transaction level. They don't work on a micro-level and their errors at the macro level are obvious enough that senior management can bear the problems in mind when making decisions.

The methodology is not sufficiently accurate to provide a high quality guide to asset and liability selection. We discuss advanced methodologies that provide better guidance to asset and liability management selection in Chapters 36 to 41. We now turn to regulatory views of capital management.

History of Capital-Based Regulations in Commercial Banking

Capital allocation by banks and regulators has had a complex chicken and egg relationship that has obscured bank management's ability to adopt the correct risk adjusted strategies for asset selection. When regulatory constraints are imposed that are not based on market values and market returns, bank management teams are forced into regulatory "arbitrage" that stems from their conflicting dual obligation to do (a) what regulators insist they do and at the same time (b) seek to maximize risk-adjusted returns to shareholders over the long run.

As Uyemura and van Deventer [1993] outline, the regulatory process in the United States has historically had a tri-partite structure:

1. System liquidity provided by the Board of Governors of the Federal Reserve System
2. Interest rate-related regulations, the now defunct Regulation Q
3. Deposit insurance managed by the Federal Deposit Insurance Corporation

The role of the Federal Reserve in providing both institution-specific liquidity when necessary (for example in the wake of the September 11, 2001 attacks on New York City) and systemic liquidity during the 2006 to 2010 credit crisis is well known. We discuss the credit crisis liquidity support by the Federal Reserve for specific institutions in Chapter 37.

What is less well-know is why the United States regulators maintained a limit on the maximum deposit interest rate that any commercial bank could pay, known as Regulation Q. For many years the maximum rate that could be paid on savings deposits with balances of $100,000 or less were at levels of 5.00-5.25%. The theory was that this limit would prevent aggressive risk-seeking banks from bidding up deposit rates to expand their activities rapidly.

The emergence of "money market funds" which pooled the funds of small depositors so they could earn the higher returns paid on deposits of more than $100,000 spelled the end of Regulation Q in the United States, which required regulators to take a more sophisticated approach to risk management. Ironically, the low rates that have prevailed since the 2006-2010 credit crisis and the "breaking of the buck" (the $1.00 net asset value of each dollar provided to a money fund) by the Reserve Fund (Bloomberg, September 16, 2008) in the wake of the bankruptcy of Lehman Brothers have impacted the growth of money market funds in a major way.

In 1980, the passage of the Financial Deregulation and Monetary Control Act ended Regulation Q and led financial regulators in the United States to turn to capital-based regulations to insure the safety and soundness of financial institutions. The regulators' first attempt in this regard was to define the concept of "primary capital." The regulatory definition of capital was firmly rooted in financial accounting-based capital, not market values. They included in "primary capital" common equity and retained earnings, perpetual preferred stock (an instrument which until then had rarely been issued), the reserve for loan losses (even though this amount was usually set by commercial banks to be consistent with expected losses), and so-called mandatory convertible notes, a specific type of debt instrument which became the primary vehicle for arbitrage of the primary capital concept by commercial banks.

This regulation triggered a number of actions by bank management which led regulators to conclude that the primary capital calculations were misguided and that the actions they engendered increased, rather than decreased, the risk in the banking system:

- Banks reduced their assets subject to the capital regulations by selling their safest, most liquid assets (government securities), increasing liquidity risk at many banks

- Banks sold their head office buildings to book an accounting "gain" that increased capital ratios but as a result added to the fixed costs that banks would have to incur on an on-going basis.
- Banks began to engage in an increasing number of "off balance sheet transactions," obscuring the accuracy of financial statements on which regulations were based

Bank failures dramatically increased throughout the 1980s so it is hard to argue that the concept of primary capital was effective in reducing the risk in the banking system. The primary capital concept was ultimately scrapped in favor of a system that was both more international in scope and more sophisticated.

In 1986 the Board of Governors of the Federal Reserve System and the Bank of England announced the concept of risk-based capital. The risk-based capital concept was specifically designed to address some of the major deficiencies of the primary capital concept:

1. Make a distinction among asset classes by riskiness
2. Incorporate off-balance sheet transactions to reduce regulatory capital arbitrage
3. Adopt common standards internationally for a fair international competition in financial services

The risk-based capital regulations had two basic components:

a. A revised definition of regulatory capital
b. An introduction of the "risk weighted asset concept"

Risk-weighted assets are a variation of the bank risk-adjusted capital concept outlined in the previous section except that the riskiness of an asset class is reflected in "risk weighted assets" associated with the financial accounting assets for that class, rather than the "risk-adjusted capital" assigned to the asset class. Low risk assets were assigned risk weighted assets that were a small fraction of their financial accounting amounts or market values. Higher risk assets were assigned a risk-weighted asset amount closer to or equal to their financial accounting values. This change in capital regulations was designed to encourage banks to retain liquid assets on their balance sheets by reducing the capital ratio penalty of the primary capital concept. These changes were ultimately labeled "Basel I" after the headquarters city of the Bank for International Settlements which coordinated the international negotiations for the capital accords.

Bankers and regulators quickly noted that the regulations were too simple (and therefore inaccurate) in two respects. First, they almost completely ignored interest rate risk. Second, the risk-weighting schemes of Basel I were so few in number than finer gradations of credit risk were ignored, again penalizing asset classes with a risk level lower than average in their Basel I risk category.

The Board of Governors of the Federal Reserve in the United States moved to address the first point with a proposal that would link interest rate risk to capital levels. This proposal made use of

Macauley's duration concept [1938} that we discussed in Chapter 1 and which we review in detail in later chapters. This proposal was never adopted as the mass of banks in the United States were incapable of understanding the basics of duration, or at least they claimed as much. Cynics responded that it was harder to get a drivers license than to be a bank CEO, where no license is required that certifies a basic understanding of how to run a bank. The Fed, even today, is very nervous about safety and soundness regulations which exceed a minimum level of complexity and that has affected U.S. participation in "Basel II," to which we now turn.

The Basel Committee on Banking Supervision began its deliberations on how to address the shortcomings of Basel I in the late 1990s, and it has issued a steady stream of pronouncements and proposals beginning in 2001. In June, 2004, the Basel II guidelines were announced that rely on a combination of capital requirements, supervisory review and market discipline to reign in bank risk taking. Van Deventer and Imai [2003] point out the biggest weakness of the Basel II approach: capital ratios, however derived, are weak predictors of the safety and soundness of financial institutions and they significantly under perform models like the recently announced Loss Distribution Model of the Federal Deposit Insurance Corporation (see the December 10, 2003 press release of the FDIC available on www.kamakuraco.com or on the website of the FDIC).

The Basel III provisions were announced by the Basel Committee on Banking Supervision in series of pronouncements that now total over 1,000 pages in 2010 and 2011. The proposals add liquidity ratios to the Basel II approach, but with much less insight and usefulness than the best practice approaches we discuss in Chapter 37 on liquidity risk. The Basel III regulations add leverage restrictions to their inverse, the capital ratios of Basel II. In this edition of Advanced Financial Risk Management, we have been candid in our de-emphasis of discussions about legacy approaches to risk management that have failed the test of time. The credit crisis of 2006-2010 resulted in the failure or nationalization of the biggest banks in Europe and the United States. Neither the Basel II regulations nor the proposals for Basel III would have prevented these failures. Indeed, the implication of the Basel II and Basel III regulations, that government securities are the only truly safe asset, is now a source of amusement in the wake of ISDA credit events in the credit default swap market for Ecuador and Greece, with an artful dodge of an ISDA default for Ireland on a technicality. Government financial difficulties continue at this writing in the Eurozone.

We understand that the Basel Committee on Banking Supervision was charged with an impossible task: forging an international standard for banking risk regulation among a group of countries none of which had yet succeeded in defining such a standard in their own financial systems. Rather than review the lowest common denominator proposals that emerged, we focus in the remainder of this book on best practice and emerging best practice. Readers interested in the evolution of the Basel Accords are encouraged to review the proposals in their original form on www.big.org.

We begin our journey through best practice risk management with the very foundation of financial institutions risk management, interest rate risk.