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 5: Portfolio Strategy and Risk Management

Chapter 38:
Performance measurement: Plus alpha versus transfer pricing

In Chapter 2 of this book, we introduced the common practice of performance measurement in large financial institutions. There are a number of key points about the way in which large financial institutions manage risk:

- Commercial banks use a mix of financial accounting and market-oriented performance measurements.
- 'Transfer pricing' in commercial banks is a financial accounting-driven performance measure that has the virtue of moving interest rate risk into one central 'transfer pricing book'. The system has a 40-year history, but it mistakenly implies that each new transaction should be judged versus the average marginal risk of the bank (via the bank's marginal cost of funds) instead of the marginal risk of the asset being considered. The fact that a traditional transfer pricing system would reject the purchase of risk-free U.S. Treasuries shows this process is flawed at the transaction level, since all banks in the U.S. own at least some U.S. Treasuries
- Insurance companies are moving to adopt the transfer pricing technology as a way of finally integrating risk assessment of insurance assets and liabilities
- Fund managers, pension funds, and insurance companies (on the asset side of the balance sheet) almost all base performance measurement on a market price basis. A good performer is one who can generate 'plus alpha' performance over and above a predefined benchmark with

equivalent risk. Banks (until recently) have almost never used this approach outside of their own fund management business, even in their own trading activities.

- Banks allocate capital to judge the performance of major business units, and almost no other financial institutions do. The capital allocation is done so that each business unit can be judged using the financial accounting measures that would be applied to the business unit if it were a stand-alone financial institution.

Our purpose in this chapter is to talk about the implications of the Jarrow-Merton put option as a potential guide to performance measurement. Matten [1996] presents an excellent summary of current common practice in commercial bank capital allocation that will be of interest to readers who would like to supplement the material in Chapter 2 before plunging into the rest of this chapter.

Transaction Level Performance Measurement versus
Portfolio Level Performance Measurement

There is a central tenet of performance management that we need to keep in mind for the rest of this chapter:

> *A performance measurement system which sends incorrect signals at the transaction level will be inaccurate at the portfolio level as well, since the portfolio is the sum of its individual transactions.*

In Chapters 2 and in van Deventer and Imai [2003], 'buy low/sell high' was outlined as the ultimate transaction level guide to asset and liability selection. If an asset worth 102 can be purchased at a cost of 100, we should do it. If another asset worth 101 can be purchased at a cost of 100, we should buy that too. What could be simpler than that? This basic performance measurement system has been at the heart of capitalism for thousands of years, and our objective in this chapter is to strip back some of the institutional 'baggage' of financial institutions' performance measurement systems to reveal performance as clearly as these examples.

In Chapters 19 through 35, we showed that every single type of asset and liability issued by major financial institutions can be valued precisely both before the asset or liability is added to the balance sheet and every instant thereafter. We now know enough to do the simple assessment in the previous paragraph. Are there any real or imagined complexities that may complicate the buy low/sell high transaction level performance assessment? There are only a few and we mention them briefly in this section:

Costs of origination *may have both a large fixed cost component and a marginal cost component which impact our assessment of buy low/sell high.*

A classic example is the consumer deposit gathering franchise of a commercial bank. On the transaction level, a 90-day consumer certificate of deposit may cost 25 basis points less than a wholesale certificate of deposit with the identical maturity. Based on this information alone, the first pass buy low/sell high decision is 'buy'—originate the retail CD. The costs of origination and servicing need to be included in the analysis for a complete answer. What if there was a $1 per month fixed cost of servicing the consumer CD that would not impact the wholesale CD? We just need to add this to the effective cost of the retail CD to have a more accurate answer. More realistically, the head of the retail banking division should be constantly re-evaluating this calculation: I can sell the building housing my River City branch for $500,000 today. Or I can keep it, pay staff expenses of $450,000 per year, and issue a large number of consumer CDs. Which is better? This is still a buy low/sell high calculation that requires no capital allocation to do correctly. It is a question that should be continually asked. A related question is whether the 25 basis point margin on the consumer CD is the margin that produces maximum present value (and marginal revenue = marginal cost from basic economic theory is relevant here). There are other calculations that come into play that we mention below.

Future business may depend on doing current business*, with an actuarially-likely arrival of good buy low/sell high business going forward.*

The example above is a good example to illustrate another complexity in assessing performance of any business. The River City branch's existence may give us the 'real option' to do lots of retail CD business going forward. This retail business depends both on our pricing and our own default probability as we saw in Chapter 37 on liquidity risk. The probability of a client rolling over one profitable retail CD has characteristics that are very similar to that of the reduced form credit models we discussed in Chapter 16, the credit default swaps we discussed in Chapter 20, and the insurance policies we valued in Chapter 35. Doing a really good buy low/sell high decision analysis depends on taking these facts into account in a careful way. It has nothing to do with capital allocation—it's a straightforward net present value calculation using the analytics presented in Chapters 19 to 35.

Some ***costs may be shared*** *with other products*.

What if the River City branch also makes profitable small business loans (valued correctly using the technology of Chapter 19 and 20), and that there is a fixed cost of monitoring them that is in part embedded in the $450,000 staff costs of the River City branch? Again, assessing whether to close the branch or keep it and optimize pricing is a risk-adjusted present value calculation like those in Chapters 19 through 35—shared fixed costs simply broaden the boundaries of the calculation that has to be done and the transactions that have to be included. Capital allocation is again irrelevant. What is relevant is the available market return that has the same risk as keeping the River City branch open.

Cash might be tight.

Van Deventer and Imai [2003] discuss how the buy low/sell high transaction level performance measurement strategy has to be tempered by the financial institution's ability to buy low. Chapter 37 discusses how, as a financial institution's risk rises, the supply of funds to the financial institution can fall or reach zero. At such a high level of risk, the financial institution has lost its real option to buy low/sell high. This is one of the reasons why financial institutions typically find that their optimal credit risk from a shareholder value-added perspective is consistent with default probabilities in a low range, a lesson they learned the hard way in the credit crisis of 2007 to 2010. We deal with this issue in more detail in the next chapter. This "own default risk" issue is why the buy low/sell high performance measurement test outlined above is necessary, but not sufficient to result in the origination of an asset or liability. The financial institution's own default risk impacts its ability to exercise its option to buy low/sell high.

Transaction Level versus Portfolio Level Performance Measurement

The examples above show that, even at the transaction level, portfolio considerations can be important. The decision to issue retail CDs depends on the total volume that can be done (the portfolio of CDs), not just the pricing on one CD. The decision on whether to keep or sell the branch depends on transaction level benefits of the retail CD and small business loan portfolio transactions at the branch. Finally, the decision to buy low/sell high can depend on the entire existing portfolio of assets and liabilities, since they determine the current credit risk of the institution.

Plus Alpha Benchmark Performance versus Transfer Pricing

In Chapter 2, we discussed both the financial accounting-oriented transfer pricing system used by commercial banks and the market value-based benchmark performance system used by almost every other type of financial institution. The authors believe that commercial banks would benefit from adopting the benchmark-based performance measurement approach, while being careful to avoid the pitfalls of such an approach as outlined in Chapter 2.

Why Default Risk Is Critical in Performance Measurement of Equity Portfolios

The risk management techniques in common use for equity portfolio management include Nobel prize winning concepts like the capital asset pricing model, arbitrage pricing theory, and the efficient frontier concept. The list of researchers behind these concepts reads like a who's who of finance: William Sharpe, Harry Markowitz, John V. Lintner, Jr., Jan Mossin, Steven Ross, and many other key contributors. Like any theory, however, the CAPM, APT and efficient portfolio concept are based on the same flaw that has doomed simple implementations of value at risk and the copula approach for CDO valuation: a normal distribution of equity returns is at the heart of CAPM, APT, and the efficient frontier. Default risk is ignored. The current credit crisis makes a more sophisticated and realistic assessment of risk in equity portfolios a mission critical function. We talk about how to do that in this section.

In our discussion of value at risk in Chapter 36, we observed that the assumption that equity returns are normally distributed with their historical mean and standard deviation is so highly stylized that it makes default disappear like magic. In Bear Stearns' case, for example, the historical mean and variance of monthly returns from January 1990 onward implied that the probability of a -100% return in a given month was 0 to six decimal places. The fact that this -100% return did occur is not a "Black Swan" a la Nassim Taleb, it's a very predictable result that is a measurable probability for any public firm. The decision to ignore default in measuring risk and return in equities is a simplification that has proven extraordinarily costly to many in the 2007 to 2010 credit crisis. Still we hear things like this all the time from equity portfolio managers:

"Why do I need default probabilities? I manage an equity portfolio, not a fixed income portfolio."

On hearing this, Professor John Y. Campbell at Harvard University laughingly responded, "Don't they know that equity is the most highly subordinated liability on the balance sheet of a public firm?" He's exactly right. A more subtle argument for ignoring default in equity portfolio management is embedded in this comment:

"I don't need default probabilities to manage my equity portfolio. My benchmark is the S&P 500 and no firm in the S&P 500 has ever defaulted."

It turns out, thanks to the web site of Standard & Poor's, that this comment is dramatically untrue. Remember Dana Corporation? Calpine? Delphi? Winn-Dixie? All of these firms failed while they were a component of the S&P 500. There is a more important phenomenon, however, that is very much credit driven that affects equity managers. Consider firms like Washington Mutual, Wachovia, Merrill Lynch, General Growth Properties and Lehman Brothers. All of them were thrown out of the S&P 500 as their credit quality deteriorated. These changes in the composition in the index happen "after hours." The stock price of the firm dropped from the index will be down 10-20% at the next day's open, and the stock price of the firm replacing them in the index will be up 10-15% on the open. Even if the equity portfolio manager is running a perfect replica of the S&P 500, the manager will lose the equivalent of 20% (best case) to 35% on one of the 500 elements of the index because of this phenomenon. This is a serious "negative alpha" phenomenon. Changes in index composition are similar to what the always insightful Michael Lewis pointed out in his baseball book Money Ball (2004): walking is a skill, not an accident. Similarly, changes in index composition are not accidents either. They are highly predictable, and the default probabilities of the companies in the index have a very high degree of correlation with their probability of being dropped from the index.

Just as important to equity managers is the dis-connect between the credit default swap market and equity returns. The CDS market is focused on senior unsecured debt. What if Citigroup is bailed out continuously but not nationalized? Most observers expect that equity holders will lose everything (by being completely diluted) and debt holders will lose nothing. Only default probabilities that predict the probability of the company failing (i.e. the shareholders being wiped out) capture this impact correctly. The CDS quotes for this example, Citi, would be much lower than the correct default probabilities because they reflect anticipation of 100% recovery like that which has accompanied the creeping nationalization of AIG.

"Plus Alpha" Performance Measurement in Insurance and Banking

With the proper adjustments discussed in the prior section, the comparison of actual mark to market performance relative to a naïve index of the same risk is a highly useful and very simple concept. The only trick in implementing this approach in the commercial banking and insurance arena is the need to create benchmarks. For example, as van Deventer and Imai [2003] suggest, the head of retail banking who owns the three-year auto loan portfolio should be judged by the market value of a U.S. Treasury index that replicates the exact cash flow timing and amount of each auto loan on a full option-adjusted basis. We show in Chapter 40 that this is a modest computer science effort that is already commercially available and it is merely a question of institutional will to move in this direction.

The manager of the three-year auto loan portfolio is a 'good performer' if over time the mark-to-market value of the auto loans, after expenses and after actual credit losses, is a larger number than the mark-to-market value of the artificial 'Treasury auto loan' index.

The manager of the auto loan portfolio will use transaction level analytics to decide which loans to originate and which loans to pass on. An important part of this relates directly to the reduced form modeling technology we discussed in Chapters 16 and 17:

- The **default probability** of each auto loan borrower has to be correctly assessed using state of the art logistic regression such as that discussed in the public firm default probability context in Chapter 16.
- The proper **credit spread** for the auto loan has to be assessed in light of expenses of origination and monitoring and in conjunction with the macro factors driving default. In addition, the probability of prepayment has to be taken into account as discussed in Chapters 27 to 29.
- Once 'break even' pricing is set for each borrower, the bank can **accept or reject the loan**. Skillful loan selection creates plus alpha performance at the portfolio level by creating value one loan at a time.

This is true buy low/sell high implementation, consistent with state of the art credit risk and interest rate risk management. A benchmark performance that is plus alpha will show management exactly how much value has been created by this business above and beyond what would have been achieved by a matched maturity investment in U.S. Treasury securities with identical payment dates, payment amounts and options characteristics. Of course the more traditional transfer pricing methodology used in commercial banking can be used to show this as well, but the transfer pricing yield curve will be the U.S. Treasury yield curve, not the marginal cost of funds yield curve of the bank. In addition, the transfer pricing rate will include a premium that reflects the prepayment option held by the auto loan borrower.

Decomposing the Reasons for Plus or Minus Alpha in a Fixed Income Portfolio

One of the most interesting set of cultural differences we have come across in finance is the difference between investment managers and bankers. One could write 100 books and 2 million jokes about those differences, but this section has a more modest ambition: to show how to improve fixed income performance attribution in both investment management and in banking by combining best practice from both types of institutions. This section is an introduction to that topic.

We can briefly summarize the key causes of the cultural differences between fixed income investment managers and bankers in a few bullet points:

- **Nature of assets managed**: The overwhelming majority of the assets held by fixed income fund managers is publicly traded. The overwhelming majority of assets held by banks, at least by number of transactions, is not publicly traded.
- **Nature of performance measurement**: In fixed income fund management, performance measurement is based on market valuation and both absolute returns and returns versus a benchmark portfolio. In spite of more than 40 years of debate in banking, that industry remains largely driven by financial accounting-based performance measurement, not market value based performance measurement
- **Granularity of performance measurement**: The CFA Institute, formerly known as the Association for Investment Management and Research, has long recommended calculating performance using daily time periods. Outside of the trading floor, banks are largely driven by monthly or quarterly time frames.
- **Volume of transactions**: Most investment managers would have no more than 5,000 or 10,000 positions, and many would have far less. One of the world's 5 largest banks has more than 700 million distinct assets and liabilities. This creates a much bigger information technology footprint in banking by necessity.
- **Staff count**: Staffing in investment management is lean and mean, because very little retail business is done, at least relative to banking. The largest banks in North America, by contrast, have a staff count in the hundreds of thousands.
- **Analytical rigor**: There are brilliant people in both banking and fixed income investment management who are well beyond best practice in their analysis. That being said, it is much more tempting to manage risk with a spreadsheet or risk analysis that is not much more powerful than a spreadsheet when the portfolio has a relatively small transaction count, something very common in the hedge fund industry in particular.

We recently came across a European paper on fixed income performance attribution that relied heavily on standard yield to maturity calculations and duration as the basis for improved fixed income performance attribution. The purpose of this section, in an introductory way, is to explain why it is particularly important in fixed income portfolio attribution to apply much more accurate and modern techniques than the duration concept, because the errors embedded in the yield to maturity concept and duration (as we discussed in Chapter 4) are large and can result in wildly inaccurate assessments of fixed income performance attribution:

- Yield to maturity discounts cash flows at all payment dates at the same interest rate
- If the same issuer has two bonds outstanding with different yields to maturity, interest payments on the same dates will be discounted at different rates for each bond
- Common spread calculations are usually the simple difference between the yields to maturity on two bonds with similar but not identical maturity dates

Now we can see the benefits of cross-pollenization of risk management between bankers and fixed come fund managers:

- We take the **mark to market orientation** of fund managers
- We take the exact day count "**matched maturity transfer pricing**" concept from banking
- We take the **yield curve smoothing** concepts that are essential to transfer pricing in banking
- We take the macro factor driven **stress testing approach** from banking and bank regulators that has come out of the 2007 to 2010 credit crisis
- We then use this technology to analyze the reasons for mark to market **performance changes** in a way already familiar to fixed income fund managers but with greater accuracy.

We know that changes in a large number of macro-economic factors affect performance of fixed income portfolios. A small sample of them is listed here:

- Level and shape of the risk free yield curve
- Home prices
- Commercial real estate prices
- Oil prices
- Other commodity prices
- Foreign exchange rates
- Market volatility
- Unemployment rates
- Changes in real gross domestic product
- Government surpluses or deficits

Best practice risk managers are increasingly measuring the impact of each of these factors on mark to market performance by tracing their impact on credit spreads, prepayment and default risk of borrowers from retail to small business to major corporations to sovereigns. In the example of this section, we keep it simple.

- We analyze the performance of one bond
- We assume that we only need to analyze 3 drivers of performance: changes in the risk free yield curve and the credit spread for the bond issuer, along with the passage of time.

- We look from the perspective of July 15, 2010 and ask the question, "how much did these risk factor movements contribute to the gain or loss in the value of our bond since April 15, 2010?" We could just as easily do the analysis for a portfolio versus a benchmark. We keep it very simple here for expository purposes.

A more complete analysis has a much longer list of drivers of performance that includes all of the macro factors above plus many more.

A Worked Example of Modern Fixed Income Performance Attribution

We now go through a worked example of a more modern approach to fixed income performance attribution with one bond issued by "ABC Company" and three drivers of performance: the risk free yield curve, the spread and the passage of time.

We assume the ABC bond has a par value of 1000, a 10% coupon, semiannual coupon payment dates of June 30 and December 31, and a maturity on June 30, 2019. We observe in the market place these prices:

April 15, 2010: $1532.60 (net present value=price plus accrued interest)

July 15, 2010: $1580.18

We ask this question: What factors have caused the net present value of the bond to change from $1532.60 to $1580.18? How much was due to the passage of time, how much to changes in the risk free rate, and how much to spread changes?

We answer this question by using the maximum smoothness forward rate approach to yield curve smoothing outlined in Chapter 5. We apply these techniques to the risk free yield curve and to the ABC bonds outstanding to derive a "risk free" and "risky" yield curve on April 15, 2010 and July 15, 2010. As a substitute for a risky yield curve derived from all ABC bonds outstanding, we assume that, by coincidence, the yield curve for ABC is identical to the US dollar libor-swap curve.

For example, here are the zero coupon bond yields for the risk-free and risky (ABC curve or swap curve) on April 15, 2010 as reported in Kamakura's "Friday Forecast" available on www.kamakuraco.com :

Figure 1

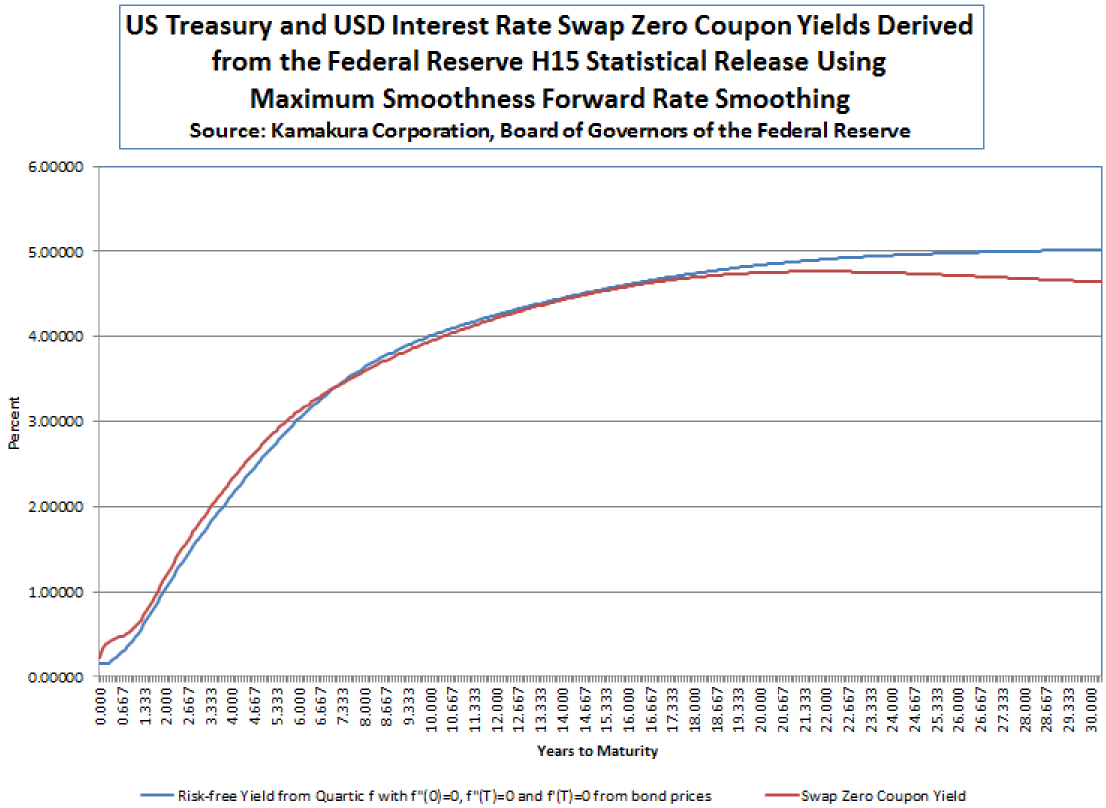

Note that, on this day, the risk-free yield curve's zero coupon yields were actually higher than the ABC Company/swap curve zero yields at the longer maturities. This is an increasingly common phenomenon that recognizes that the risk free or sovereign yield curve is not, in fact, risk free.

We also take the zero coupon yields for the risk free and risky ABC/swap yield curve on July 15, 2010:

Figure 2

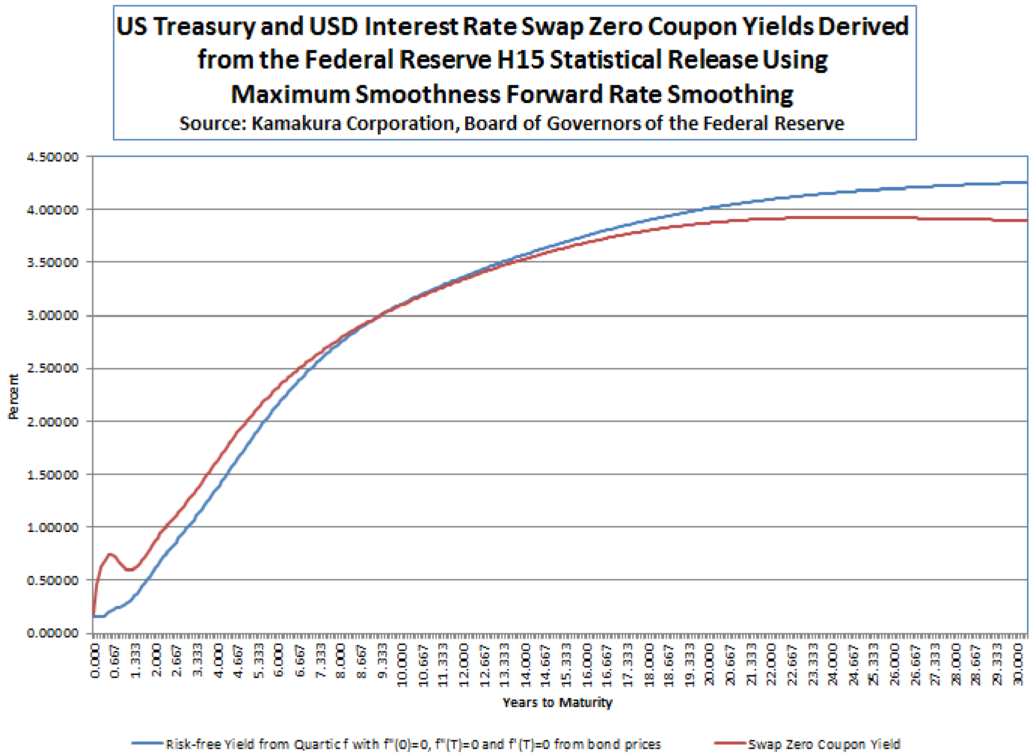

Over this three month period, we can see that the risk free zero coupon yield curve has fallen substantially:

Figure 3

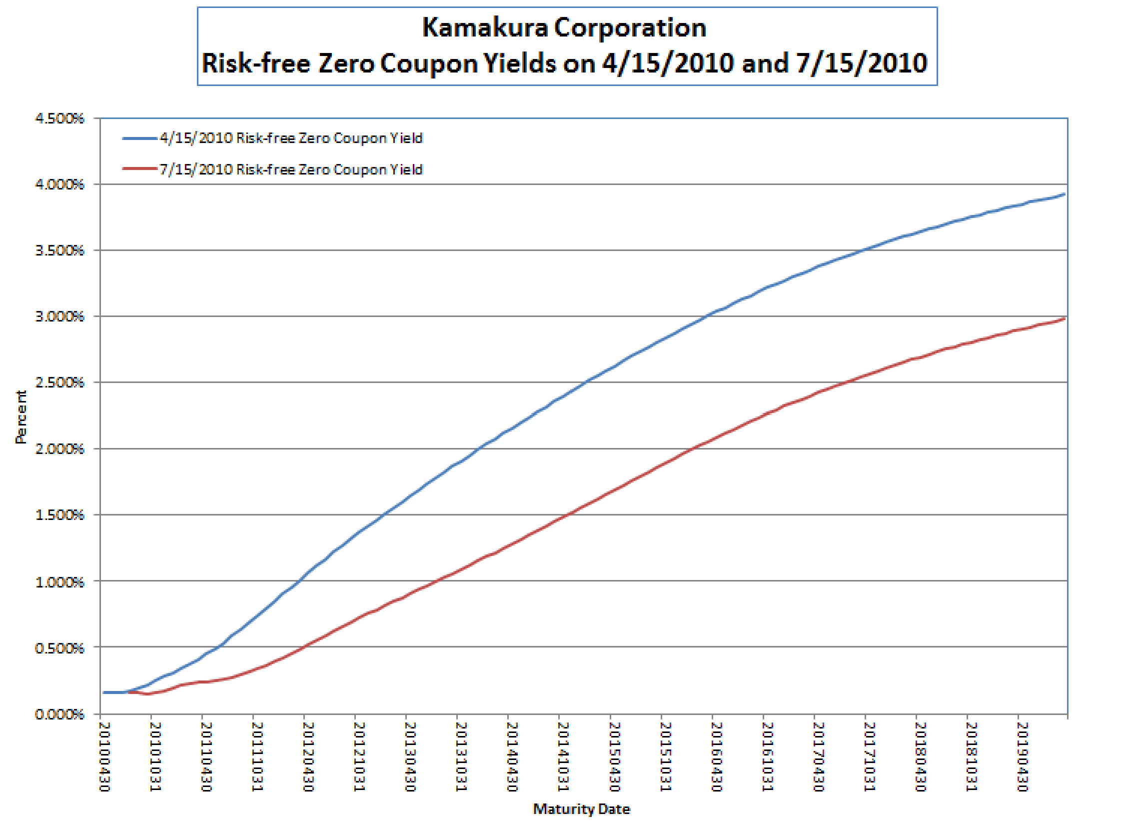

Risky zero coupon yields, with the exception of the short maturities, have also moved downward, as this graph shows:

Figure 4

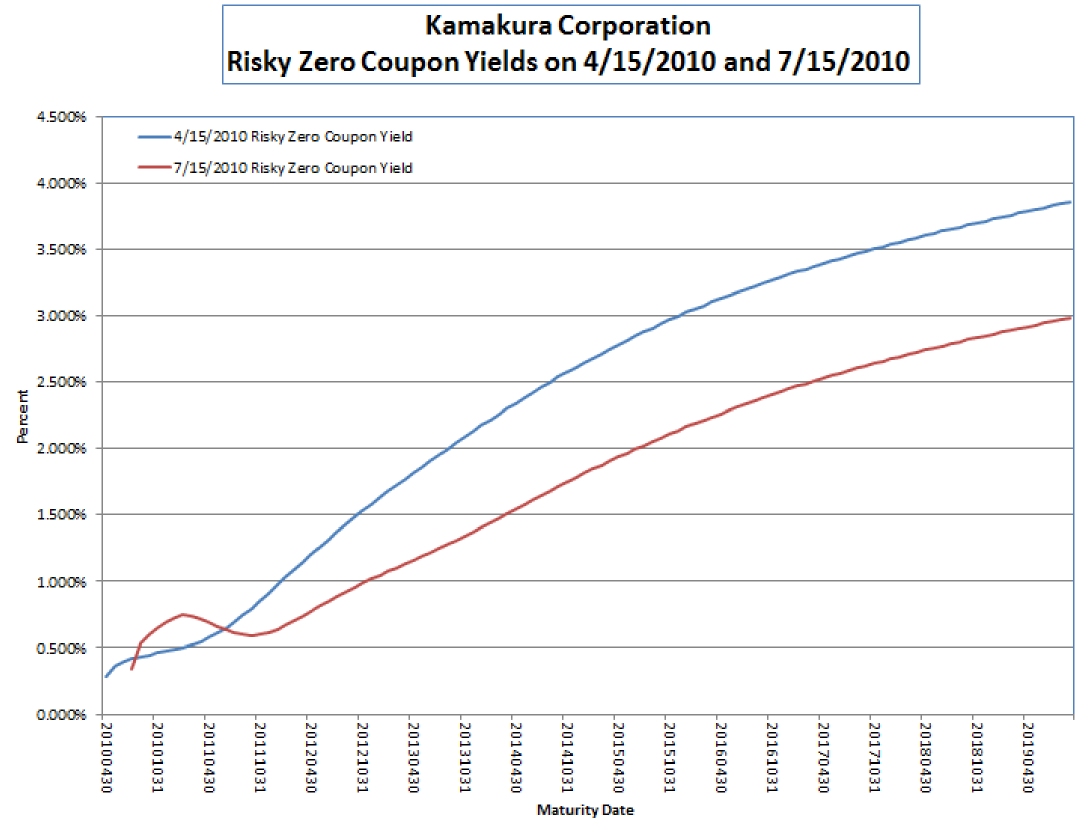

The zero coupon credit spread for ABC company/swap curve have moved as shown below:

Figure 5

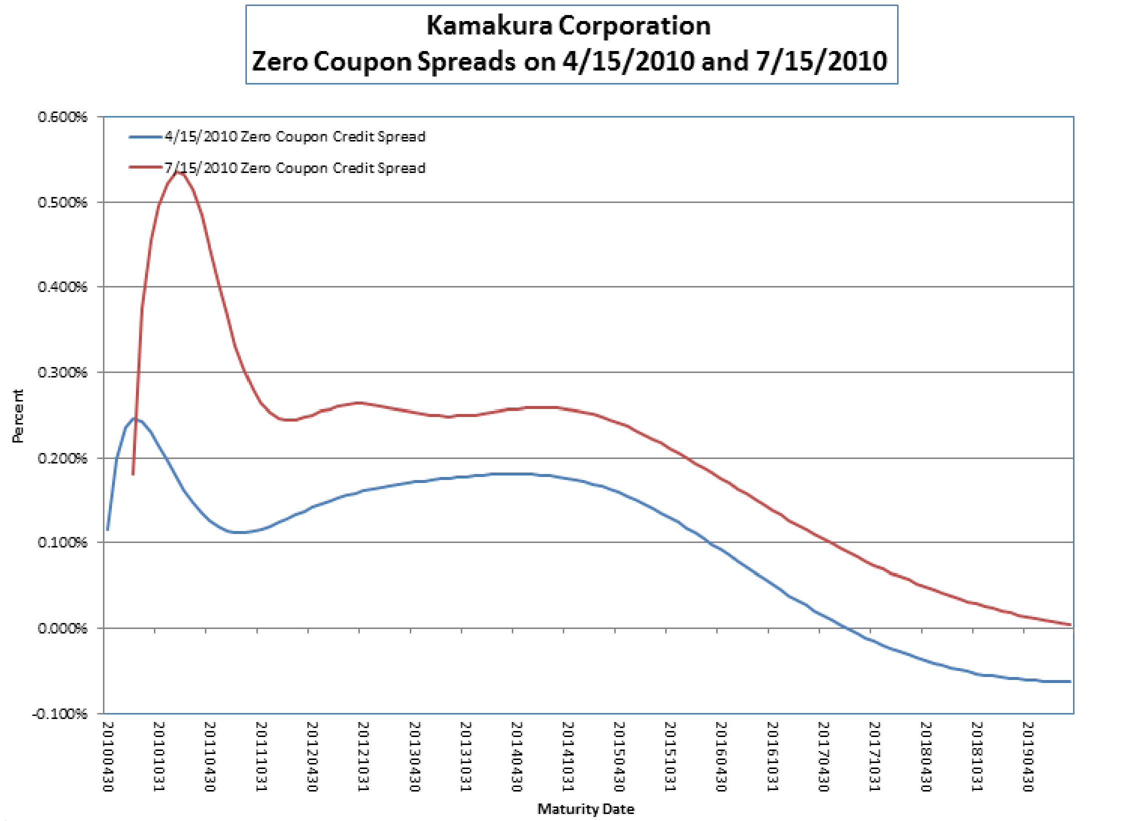

Many researchers have found that it is common for credit spreads to rise when risk-free yields fall, and that is consistent with yield curve movements between these two dates. We extract the following zero coupon bond prices from both yield curves on both dates and apply it for performance attribution to the cash flows on the ABC bond:

Figure 6

Actual Zero Coupon Yields and Zero Coupon Credit Spreads									
Date	4/15/2010 Risk-free Zero Coupon Yield	4/15/2010 Risky Zero Coupon Yield	4/15/2010 Zero Coupon Credit Spread	7/15/2010 Risk-free Zero Coupon Yield	7/15/2010 Risky Zero Coupon Yield	7/15/2010 Zero Coupon Credit Spread	Change in Risk-free Zero Yields	Change in Risky Zero Yields	Change in Zero Spreads
20100430	0.163%	0.279%	0.116%						
20100531	0.160%	0.358%	0.198%						
20100630	0.159%	0.394%	0.235%						
20101231	0.310%	0.488%	0.178%	0.193%	0.729%	0.535%	-0.025%	0.280%	0.305%
20110630	0.537%	0.651%	0.114%	0.264%	0.631%	0.367%	-0.147%	0.085%	0.232%
20111231	0.850%	0.974%	0.124%	0.394%	0.641%	0.247%	-0.293%	-0.159%	0.134%
20120630	1.171%	1.321%	0.149%	0.593%	0.850%	0.257%	-0.420%	-0.300%	0.120%
20121231	1.468%	1.633%	0.165%	0.787%	1.048%	0.261%	-0.536%	-0.434%	0.102%
20130630	1.737%	1.910%	0.174%	0.967%	1.217%	0.250%	-0.637%	-0.557%	0.080%
20131231	1.995%	2.175%	0.180%	1.155%	1.407%	0.252%	-0.712%	-0.637%	0.075%
20140630	2.239%	2.420%	0.180%	1.352%	1.611%	0.259%	-0.766%	-0.688%	0.078%
20141231	2.476%	2.648%	0.172%	1.560%	1.813%	0.253%	-0.800%	-0.724%	0.076%
20150630	2.697%	2.848%	0.150%	1.763%	1.995%	0.232%	-0.825%	-0.755%	0.070%
20151231	2.908%	3.025%	0.117%	1.963%	2.163%	0.200%	-0.841%	-0.776%	0.065%
20160630	3.099%	3.177%	0.079%	2.149%	2.312%	0.163%	-0.856%	-0.791%	0.065%
20161231	3.274%	3.312%	0.039%	2.321%	2.448%	0.127%	-0.867%	-0.799%	0.068%
20170630	3.427%	3.431%	0.004%	2.474%	2.568%	0.093%	-0.878%	-0.805%	0.073%
20171231	3.564%	3.539%	-0.024%	2.613%	2.677%	0.065%	-0.885%	-0.809%	0.076%
20180630	3.682%	3.638%	-0.044%	2.733%	2.775%	0.041%	-0.891%	-0.815%	0.077%
20181231	3.788%	3.731%	-0.056%	2.842%	2.864%	0.023%	-0.895%	-0.821%	0.074%
20190630	3.879%	3.817%	-0.062%	2.936%	2.945%	0.009%	-0.898%	-0.829%	0.069%

We want to understand the drivers of changes in bond value so we analyze these specific combinations of changes in the key 'macro factors' that we have chosen to analyze:

Figure 7

Scenarios for Fixed Income Performance Measurement

	Scenario Description				
Attribution Driver	**Actual 4/15/2010**	**Actual 7/15/2010**	**Risk-free Shift Only**	**Spread Shift Only**	**Date Shift Only**
Risk-free Yields	4/15/2010 Risk-free	7/15/2010 Risk-free	7/15/2010 Risk-free	4/15/2010 Risk-free	4/15/2010 Risk-free
Risky Spreads	4/15/2010 Spreads	7/15/2010 Spreads	4/15/2010 Spreads	7/15/2010 Spreads	4/15/2010 Spreads
Cash Flow Dates	4/15/2010 Dates	7/15/2010 Dates	4/15/2010 Dates	4/15/2010 Dates	7/15/2010 Dates
Gross Bond Value	1532.60	1580.18	1623.37	1525.47	1499.56
Less Value of 6/30/Coupon	-49.96	0.00	-49.96	-49.96	0.00
Net Bond Value	1482.64	1580.18	1573.41	1475.51	1499.56
Dollar Change in Bond Value		97.55	90.78	-7.12	16.92
% Change vs Actual 4/15/2010	Not Applicable	6.579%	6.123%	-0.480%	1.141%

The first two scenarios use the real inputs for April 15 and July 15, 2010. The next three scenarios shift one of our three key drivers of returns, one at a time. The "risk-free shift only" scenario can be explained as follows: "Assume we are at April 15, 2010 and the risk free curve took July 15, 2010 levels but the credit spread was the April 15 spread." The "spread shift only" scenario means "Assume we are at April 15, 2010 and risk free curve is the April 15 curve, but the spread is the July 15 spread." In the final scenario, the "date shift only" scenario applies the April 15 risk-free curve and spread on the July 15, 2010 value date, when all cash flows are 3 months closer than they were on April 15.

In each case, the ABC bond's present value is the sum of the relevant zero coupon bond times the cash flow on that particular date. Those zero coupon bonds are shown here:

Figure 8

	Discount Factors Used				
Date	**Actual 4/15/2010**	**Actual 7/15/2010**	**Risk-free Shift Only**	**Spread Shift Only**	**Date Shift Only**
4/30/10	0.999885	0.000000	0.999885	0.999858	0.000000
5/31/10	0.999549	0.000000	0.999551	0.999326	0.000000
6/30/10	0.999180	0.000000	0.999196	0.998725	0.000000
12/31/10	0.996530	0.996632	0.997060	0.994357	0.997924
6/30/11	0.992169	0.993972	0.994839	0.990181	0.994783
12/31/11	0.983462	0.990664	0.989529	0.981383	0.988359
6/30/12	0.971222	0.983458	0.981552	0.968789	0.977682
12/31/12	0.956632	0.974488	0.972106	0.954309	0.964121
6/30/13	0.940508	0.964612	0.961174	0.938243	0.948833
12/31/13	0.922381	0.952409	0.948210	0.919769	0.931598
6/30/14	0.903134	0.938171	0.933437	0.900173	0.912958
12/31/14	0.882634	0.922227	0.917207	0.879599	0.892903
6/30/15	0.862090	0.905735	0.900325	0.859106	0.872436
12/31/15	0.841234	0.888488	0.883094	0.838142	0.851572
6/30/16	0.820835	0.871166	0.865981	0.817459	0.830999
12/31/16	0.800504	0.853554	0.848926	0.796693	0.810561
6/30/17	0.780775	0.836253	0.832022	0.776583	0.790663
12/31/17	0.760962	0.818773	0.815034	0.756466	0.770766
6/30/18	0.741694	0.801735	0.798145	0.737103	0.751354
12/31/18	0.722329	0.784601	0.781114	0.717813	0.731909
6/30/19	0.703509	0.767940	0.764216	0.699266	0.712941

When we apply these scenarios, we get the following components of performance:

Figure 9

Attribution Driver	Scenario Description: Actual 4/15/2010	Actual 7/15/2010	Risk-free Shift Only	Spread Shift Only	Date Shift Only	Impact of All Factors' Joint Movement	Total
Risk-free Yields	4/15/2010 Risk-free	7/15/2010 Risk-free	7/15/2010 Risk-free	4/15/2010 Risk-free	4/15/2010 Risk-free		
Risky Spreads	4/15/2010 Spreads	7/15/2010 Spreads	4/15/2010 Spreads	7/15/2010 Spreads	4/15/2010 Spreads		
Cash Flow Dates	4/15/2010 Dates	7/15/2010 Dates	4/15/2010 Dates	4/15/2010 Dates	7/15/2010 Dates		
Gross Bond Value	1532.60	1580.18	1623.37	1525.47	1499.56		
Less Value of 6/30/Coupon	-49.96	0.00	-49.96	-49.96	0.00		
Net Bond Value	1482.64	1580.18	1573.41	1475.51	1499.56		
Dollar Change in Bond Value		97.55	90.78	-7.12	16.92		
% Change vs Actual 4/15/2010	Not Applicable	6.579%	6.123%	-0.480%	1.141%	-0.204%	6.579%

Note that we first deduct from the actual April 15, 2010 "gross bond value" or net present value the present value of the June 30, 2010 coupon payment of $50. This adjustment is necessary to insure we are comparing the value of two securities (ABC bond on April 15 and July 15) with the same attributes. This adjustment would not be necessary if we were using the daily return intervals that the CFA Institute recommends, since the coupon payment would just be an element of daily total return. In this case, we deduct the present value of that first coupon rather than speculate on reinvestment returns from June 30 to July 15.

The analysis shows that, after adjusting for the first coupon, our ABC bond returned a total of 6.579% from April 15, 2010 to July 15, 2010. The downward shift in the risk-free curve contributed 6.123% to this total return. The widening of credit spreads by itself reduced return by 0.480%. The passage of time, which moved all cash flows 3 months closer, contributed 1.141% to total return. A final component, the impact of all factors moving together, reduced total return by 0.204%.

In this fashion, we have a complete and accurate understanding of the contribution of the risk factors we have chosen to analyze to total absolute return. Return relative to a benchmark index, or tracking error, is analyzed in the same way. We have avoided the serious performance attribution errors that can arise from the simplifying assumptions embedded in yield to maturity and simple credit spread calculations by using modern yield curve and credit spread smoothing on an exact day count basis. The process is simple and straightforward when using modern enterprise risk management software.

The Jarrow-Merton Put Option and Capital

We now turn back to more traditional measures of performance, using financial accounting as a basis, which have a long tradition in banking and (more recently) in insurance. In Chapters 36 and 37, we discussed the practical calculation of the Jarrow-Merton put option that is the comprehensive measure of integrated credit risk, market risk, liquidity risk and interest rate risk that we first discussed in Chapter 2. In Chapters 19 through 35, we have illustrated how we can value the assets and liabilities of any financial institution using the simulation techniques of Chapter 36. We can use this calculation as a basis for capital allocation to the various business units in a financial institution just we can for the institution as a whole.

Using the Jarrow-Merton Put Option for Capital Allocation: An Introduction

As we noted in Chapter 2, most of the business units of major financial institutions have either assets only or liabilities only. They rarely have both assets and liabilities at the same time. When this is the case, there is no need to allocate capital because the benchmark performance system we discuss above can be used. Nonetheless, many financial institutions allocate capital to achieve multiple objectives:

- They want all business units, after the allocation of debt and equity, to have the same relative degree of risk.

- They then feel comfortable, using either historical market returns or financial accounting returns, in expressing performance as the ratio of historical return to allocated capital.

As we noted in Chapter 2, this approach implicitly says that 'good' business units have good returns on risk-adjusted capital and should get more emphasis. This portfolio measure of performance is thought of by many as a substitute for the buy low/sell high approach discussed in the introduction to this chapter.

Using the Jarrow-Merton Put Option Concept for Capital Allocation
In this section, we show how to use the Jarrow-Merton put option technology for capital allocation. We start with a one-period example and then generalize to a multi-period example that is consistent with the Loss Distribution Model of Jarrow et al announced by the Federal Deposit Insurance Corporation on December 10 2003.

Capital Allocation for a Single Period
In Chapter 36, we outlined a careful process for simulating the impact of macro-economic factors on the default probabilities and legal payment obligations of all counterparties from retail borrowers to sovereigns. We will use the identical process in this section. If we set a single time horizon as relevant for capital allocation, the process works like this:

1. Choose the time horizon *T* (one year is a common choice).
2. Obtain from market sources the default probability for the institution. This default probability can be derived in any method with which the institution is comfortable as discussed in Chapters 16 and 17. Obtain the business unit's complete transaction-level listing of assets and liabilities.
3. Simulate the macro factors over *N* scenarios as of time *T*.
4. Calculate the default probabilities for all counterparties.
5. Calculate *N* values for the value of business unit assets less the amount of liabilities
6. Assume the default probability for the business unit with no Jarrow-Merton put option is d^*. We need to solve for the value of the Jarrow Merton put option, which reduces d^* to the institution's own default probability *d*.
7. Examine the distribution of negative net worths from the simulation.
8. Determine the amount *K* of a digital put option payable on default such that the percentage of negative net worth scenarios after payment of *K* is reduced from d^* to *d*. For example, if *d* is 1%, but 200 of the 10,000 scenarios produce negative net worth on a mark-to-market basis *(d^*=2%)*, we need to find *K* such that 100 of the negative scenarios would be erased. If the 100th largest loss is \$72.27, then the digital option would pay \$72.27 on default, reducing d^* to 1%= *d*.
9. Calculate the value of the Jarrow-Merton put option which pays *K* (\$72.27) upon default (i.e. negative mark-to-market before *K* is paid). Let us call the value of this put option *X*. This is the price of portfolio insurance for a default probability of *d*.
10. Adjust the capital and the assets of the business unit for the value of the put option. Before the put option, the capital of the business unit was *V-B* and its capital ratio was *(V-B)/V*. After the put option, capital will be increased by *X* and assets will be increased by *X*. Assets will be *V+X* and capital will be *V-B+X*. The ratio of capital to assets will be *(V-B+X)/(V+X)*.
11. This process equalizes the default probabilities of every business unit at the institutional level and explicitly incorporates the risk of each individual asset or liability on the books of the business unit.

There are other choices for the terms of the pay-off on the Jarrow-Merton put option. The most obvious one is one that equalizes the 'expected loss' of each business unit. Even if default probabilities are equalized, the expected loss can differ and vice versa. This is matter for institutional strategy.

In any event, the technology for the Jarrow-Merton put option in the first 37 chapters of this book allows us to explicitly price the value of portfolio insurance (both credit risk and interest rate risk) on all assets and liabilities of each business unit so that the total default risk of each business unit after portfolio insurance is equalized.

Extending the Jarrow-Merton Capital Allocation to a Multi-Period Framework

While most financial institutions allocate capital with a one-year time horizon in mind, the credit crisis of 2007 to 2010 makes it clear that a longer time horizon may be more appropriate. The fact that the credit default swap market has crystallized around a five-year maturity is indirect confirmation that the market believes that a longer time horizon has the most value from a credit insurance policy point of view.

For this reason, it may be appropriate to analyze the Jarrow-Merton put option as a multi-period payout option that insures the survival over a longer time horizon just like the Loss Distribution Model of the FDIC (December 10 2003, available on www.fdic.gov) does.

This changes the valuation of the Jarrow-Merton put option from the previous section in only modest ways:

- Scenario simulation has to deal with reinvestment of cash flows over multiple periods for correct assessment of the probability of business unit 'default'.
- The 'default' of a business unit can happen at multiple points in time and multiple payments to the business unit may be necessary to insure survival over the entire time horizon. This is highly likely given the high correlation in defaults that we have discussed throughout this book.
- This will raise the cost of the Jarrow-Merton put option for the business unit, all other things being equal.

Summing Up

As noted in Chapters 2, backward looking measures of asset performance have minimal use as predictors of future performance. For that reason, the authors prefer forward looking measures of performance (buy low/sell high) for asset selection at both the transaction level and the portfolio level.

Nonetheless, the Jarrow-Merton put option concept can be applied neatly for capital allocation for those institutions that have the need to apply the capital allocation concept. We explicitly price the portfolio insurance (integrating credit risk, market risk, liquidity risk and interest rate risk) for each business unit such that we equalize the (a) probability of default or (b) expected loss for each business unit. The adjusted asset totals and capital totals for each business unit reflect the incremental cost of this credit insurance, the value of the Jarrow-Merton put option.

We turn next to a related topic, the safety and soundness of our own financial institution.