Kamakura Corporation

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

Advanced Financial Risk Management:

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

April 2012

Donald R. van Deventer
Kenji Imai
Mark Mesler

Part 2: Risk Management Techniques for Interest Rate Analytics

Chapter 10:
Valuation, liquidity and net income

Chapters 6 through 9 introduced modern multi-factor interest rate models of Heath, Jarrow and Morton. In this chapter, we introduce some of the key issues that we need to deal with for practical application of the HJM concepts. We will spend the rest of the book addressing these issues in great detail, but we start here with an overview of some of the most important topics.

How many risk factors are necessary to accurately model movements in the risk free yield curve?

In Chapters 6 and 7, we implemented a one factor model of yield curve movements in the HJM framework. We extended that to a two factor model in Chapter 8 and a three factor model in Chapter 9. This begs the question, "how many risk factors are necessary to accurately model the movements of the forward rates (zero yields, and zero prices) of the risk free curve?" There are two answers to this question. The most important answer is the perspective of a sophisticated analyst whose primary concern is accurately reporting the risk that his or her institution faces. The second most important perspective is the requirements of a sophisticated regulator. We address these both in turn.

We start with a perspective of a sophisticated risk manager whose primary concern is accuracy in modeling. There are two conditions for determining the risk number of risk factors driving the yield curve:

1. **Necessary condition**: all securities whose value is tied to the risk free yield curve and its volatility should be priced accurately if a market price is visible
2. **Sufficient condition:** providing condition 1 is met, the residual correlation of changes in forward rates 1 through n should be statistically indistinguishable from zero.

Note that the legacy Nelson-Siegel yield curve fitting technique fails even the necessary condition in general because it lacks sufficient parameterization to price all input securities accurately, violating the no arbitrage condition. We can illustrate these conditions for more sophisticated yield curve fitting techniques using data from Chapters 8 and 9. In Chapter 9, we faced this question: "Do we need to add a third factor to the two risk factors which we used in Chapter 8?" In Chapter 8, we denoted the change in the 1 year U.S. Treasury spot rate as risk factor 1 and the residual risk factor was denoted "everything else." We fitted econometric relationships that sought to explain the movements of the 1 period forward rates maturing at times 2 through 10. We took the residuals from those relationships and analyzed their correlation with each other:

Figure 1

Correlation of Residuals from Regression 1 on U.S. Treasury Spot and Forward Rates, 1962-2011
Source: Kamakura Corporation and Federal Reserve

	1 Year U.S. Treasury Forward Rate Maturing in Year:							
Maturity in Year:	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
2	1							
3	0.92016	1						
4	0.84091	0.93618	1					
5	0.80762	0.82393	0.94274	1				
6	0.79099	0.80811	0.87585	0.95039	1			
7	0.7531	0.80368	0.82023	0.85957	0.96701	1		
8	0.73779	0.80699	0.83312	0.83894	0.91046	0.9611	1	
9	0.6775	0.7417	0.80867	0.79686	0.78653	0.81691	0.93905	1

The lowest correlation of the remaining variation in forward rate changes after accounting for risk factor 1 was 67.75%. Even though a 1 factor model could meet the necessary condition above for no arbitrage (as we saw in Chapter 6), it fails the second condition that the correlation of the remaining variation in forward rate changes be indistinguishable from zero. What if we add a second variable called "everything else"? That variable would meet the sufficient condition only if the correlations remaining after risk factor 1 were all either 100% or -100%, indicating all remaining random shifts were perfectly correlated with the "all other" risk factor, either positively or negatively. That is not the case here.

For that reason, we decided to add a third factor in Chapter 9 and we formally denoted the second risk factor as the one period change in the 10 year forward rate. After we take the residuals remaining after risk factor one and try to explain them with our new risk factor two, we have another correlation matrix remaining:

Figure 2

Correlation of Residuals from Regression 2 on U.S. Treasury Spot and Forward Rates, 1962-2011
Source: Kamakura Corporation and Federal Reserve

	1 Year U.S. Treasury Forward Rate Maturing in Year:							
Maturity in Year:	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
2	1.000000							
3	0.887509	1.000000						
4	0.771819	0.906347	1.000000					
5	0.720069	0.726640	0.901236	1.000000				
6	0.700086	0.711934	0.803607	0.924316	1.000000			
7	0.644186	0.704231	0.710768	0.776320	0.949579	1.000000		
8	0.607691	0.695730	0.687453	0.705748	0.871682	0.960999	1.000000	
9	0.511367	0.592157	0.632456	0.619087	0.674017	0.738484	0.882839	1.000000

The use of the second risk factor has reduced the lowest correlation of residual risk from 67.75% to 51.14%. That represents series progress, but adding a third factor "everything else" will not eliminate the risk because the remaining correlations after the two designated risk factors are extracted are not all either 100% or -100%. A sophisticated analyst knows from this that we must add more factors. How do we know how many factors are necessary? One method is to keep incrementally adding risk factors until the residual correlations are either all 100% or -100% (in which case the last factor will be "everything else") or we add factors until the remaining residual risk has correlation that is statistically indistinguishable from zero. Dickler, Jarrow, and van Deventer estimate that six to ten factors are needed to realistically model forward rate movements of the U.S. Treasury curve from 1962 to 2011. Numerous academic authors have investigated the same issue and conclude that at least three factors, and probably more, are needed. For a video of the daily movements of the forward rate curve, see www.kamakuraco.com or youtube.com, searching on "Kamakura."

What is the regulatory perspective on the problem? Perhaps the most insightful comment comes from the Basel Committee on Banking Supervision. In its December 31, 2010 *Revisions to the Basel II Market Risk Framework*, the Committee states its requirements clearly on page 12:

> "For material exposures to interest rate movements in the major currencies and markets, banks must model the yield curve using a minimum of six risk factors."

Revisiting the Phrase "No Arbitrage"

Condition 1 above sets as a minimum standard for the number (and magnitude) of the risk factors the existence of "no arbitrage." This phrase often has a different meaning in an academic context than it does in practical application, but the true meaning should be the same as stated in Condition 1: all securities whose pricing is fully determined by the risk free yield curve and its assumed volatility should be priced accurately if market prices are visible. If that is not the case, there is one of three types of errors: (a) the number of risk factors is too few, (b) the magnitude and term structure of at least one risk factor is not correct, or (c) there is an error in the software implementation or market data inputs. One common market data input "error" might be simply the bid offered spread on a given bond, the most benign type of value discrepancy that will not go away. A more refined statement of "priced accurately" is "priced such that the valuation is at or between the bid and offered prices" for the security.

In the Heath, Jarrow and Morton context, it was assumed that the market for zero coupon bonds is complete, perfectly competitive, and that one can go long or short zero coupon bonds in any amount costlessly. Because if this assumption, it is common but wrong for an analyst to conclude that his HJM implementation for the U.S. Treasury curve is complete if we accurately value the zero coupon bonds used as input as we saw in Chapters 6 through 9. This conclusion would be wrong if there were other securities tied to the risk free curve and its volatility that are mispriced after implementation. Consider the 3% coupon bond that we valued in Chapters 6 through 9. What if this bond had been callable, and the theoretical pricing were not correct, even though another non-callable bond with the same terms is priced correctly? This means that the number of risk factors or their magnitude remains incorrect. The solution is to add ALL securities with observable prices to assure that this source of error does not come about.

We discuss this in most detail in Chapter 14, estimating the parameters of interest rate models. It is most complex in modeling a yield curve tied to London Interbank Offered Rates and the interest rate swap curve because of many complications: (a) the large number of instruments tied to this curve, (b) the fact that this curve is not risk free, (c) the fact that the credit risk of Eurodollar deposits (for which Libor is quoted) and the credit risk of the swap curve are not the same, and (d) the mounting evidence and legal proceedings on the potential manipulation of Libor. These issues get special attention in Chapter 17.

Valuation, Liquidity Risk and Net Income

A complete analysis of risk of a large portfolio or a large financial institution usually includes a number of calculations, just a few of which are listed here:

a. Valuation of all assets and liabilities, both at time 0 and at user-specified dates in the future
b. Gross and net cash received and cash paid, for liquidity risk assessment, for any set of user-defined time intervals
c. Financial accounting (GAAP or IFRS basis) net income, interest income, interest expense and (often) non-interest income and expense for any set of user-defined time periods

What extensions of the HJM examples in Chapters 6 through 9 are needed to achieve these objectives? There are three major extensions that are necessary:

1. A distinction between risk neutral probabilities and empirical probabilities that a given interest rate scenario will come about.
2. Monte Carlo simulation, not just bushy tree simulation, particularly to ensure that a large number of interest rate scenarios can be generated even in near term time intervals.
3. An implementation of GAAP and IFRS accounting accruals.
4. Expansion of analysis to include "insurance events" like default, prepayment, mortality, property and casualty risks, and a large array of macro-economic risk factors beyond the risk free yield curve.

The third extension is very important, hard to do well, and best covered in other volumes and in the user documentation for a state of the art enterprise risk management system. We will not address it in this book. The fourth extension is at the heart of the rest of this book. We turn now, however, to the important distinction between risk neutral and empirical default probabilities.

Risk Neutral and Empirical Probabilities of Interest Rate Movements

In Chapters 6 through 9, we have followed Jarrow (2002) and showed how "risk neutral valuation" allows us to accurately value zero coupon bonds (and other securities, as we see in the following chapters) using pseudo probabilities which, without loss of generality, can be selected for the user's convenience. In Chapters 6 and 7, the probabilities of an up and down shift were set at 1/2. In Chapter 8, with two risk factors, the probabilities of an up shift, mid shift and down shift were set at 1/4, 1/4 and 1/2. In Chapter 9, with three factors, the probabilities for shifts 1, 2, 3 and 4 were set at 1/8, 1/8, 1/4 and 1/2. All of these choices (and many others) result in accurate valuation provided that the formulas for the shifted forward rates and zero coupon bond prices are adjusted correctly. We need to ensure that the expected value and variance of the changes in forward returns fit these two constraints and that the drift in forward rates μ (adjusted for the choice of pseudo probabilities and the number of risk factors) is consistent with the formula in Chapter 9:

$$E\left[\frac{\log F(t+\Delta,T)}{\Delta}-\frac{\log F(t,T)}{\Delta}\right]=\mu(t,T)\Delta$$

$$Var_t\left[\frac{\log F(t+\Delta,T)}{\Delta}-\frac{\log F(t,T)}{\Delta}\right]=\left[\sum_{i=1}^{n}\sigma_i\,(t,T)^2\right]\Delta$$

The reader can confirm that these relationships hold for the three factor example in Chapter 9, for example, by taking the expected value of the four shifted forward rate values, weighted by the probabilities 1/8, 1/8, 1/4, and 1/2. The variance can be confirmed in the same way. These are the key relationships that one uses to parameterize a Monte Carlo simulation using the HJM approach.

What if an analyst is interested in using the actual probabilities of forward rate shifts (also called the empirical or statistical probabilities) instead of the pseudo probabilities (also called risk-neutral probabilities). Jarrow shows that these empirical probabilities are straightforward shifts of the pseudo probabilities. The amount of the shift is a function of the risk aversion embedded in the term structure of risk-free interest rates. The formulas are given in Jarrow (2002), Chapter 15, page 280.

When would the empirical probabilities be used instead (or in addition) to the pseudo probabilities? Whenever the analyst is doing a simulation of actual cash flows and probabilities, actual net income and actual probabilities, or the actual "self-assessment" of the institution's probability of default. For a one factor interest rate model, the expected value and variance above would be shifted in a simple way using the time dependent function $\theta(t)$:

$$E\left[\frac{\log F(t+\Delta,T)}{\Delta}-\frac{\log F(t,T)}{\Delta}\right]=[\mu(t,T)-\theta(t)\sigma(t,T)]\Delta$$

$$Var_t\left[\frac{\log F(t+\Delta,T)}{\Delta}-\frac{\log F(t,T)}{\Delta}\right]=\sigma(t,T)^2\Delta-\theta(t)^2\sigma(t,T)^2\Delta^2$$

Instead of the pseudo probability of an upshift, which was set at 1/2 in Chapter 6 and 7, the empirical probability of an upshift would be set at

$$q_t(s_t=up)=\frac{1}{2}+\frac{1}{2}\theta(t,s_t)\sqrt{\Delta}$$

See Jarrow (2002), Chapter 15, page 279 for details. Within a net income simulation, the empirical probabilities would be used to generate an empirical distribution of net income in each user-defined period. Part of the net income calculation, under GAAP accounting for example, might involve the amortization of a premium or discount of a security purchased at a market value different from par value. This amortization is required under GAAP accounting to be updated as the security's value changes. Valuation of this security would be done using the pseudo probabilities that are associated with the empirical probabilities. This is a feature one can find in a best practice enterprise-wide risk management system.

How is the function $\theta(t)$ determined? It is done using a historical study like Dickler, Jarrow and van Deventer (2011) in which forward rates maturing at time T*, observed at various points in time t, are compared statistically with the actual spot interest rates which came about for maturity at time T*. The graph below, for example, shows the convergence of all 1 year U.S. Treasury forward rates (out to 9 years forward) to the actual 1 year U.S. Treasury spot rate for 24 dates on which the one year spot Treasury rate was 0.27%.

Figure 3

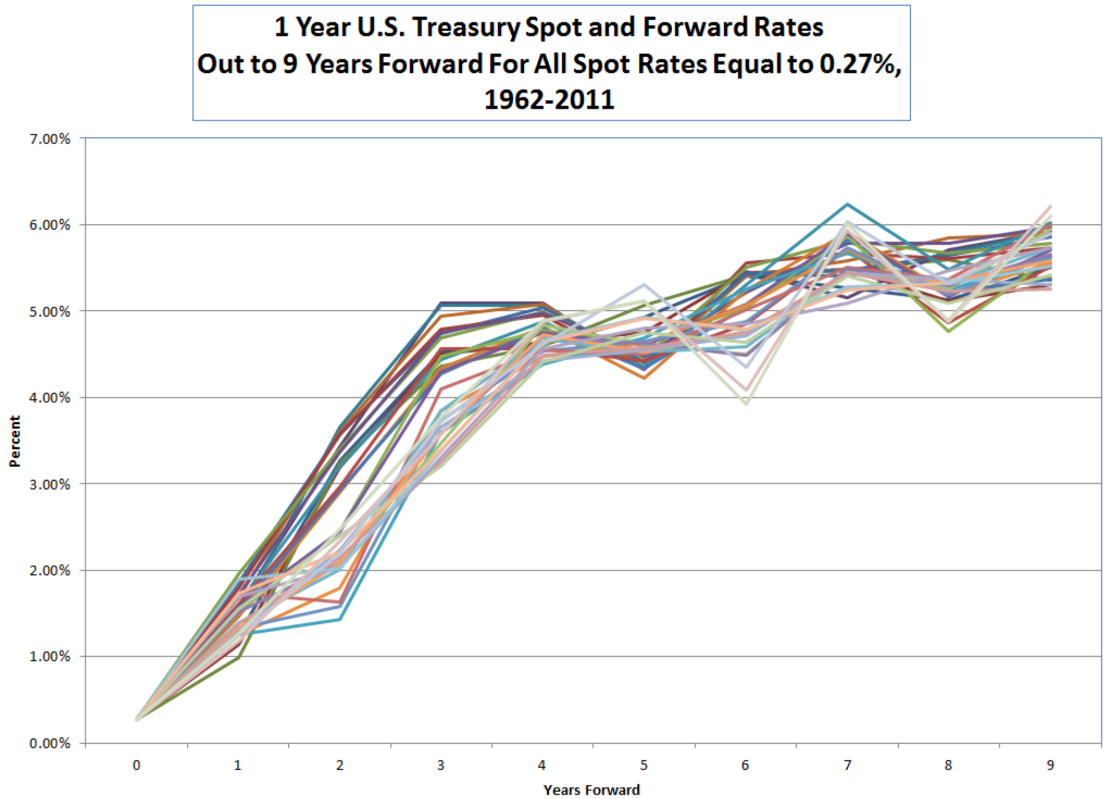

Monte Carlo Simulation Using Heath Jarrow and Morton Modeling

How does one use Monte Carlo simulation in the Heath Jarrow and Morton framework? For reasons outlined below, in most cases, Monte Carlo simulation will be preferred to the busy tree approach we illustrated in Chapters 6 through 9. The process follows the "pseudo code" here for M scenarios and N periods into the future, using the example of the three factor model in Chapter 9:

001 Load relevant zero coupon bond prices
002 Load assumptions for risk factors 1, 2, and 3
003 Set ScenarioCounter =1
004 Set RemainingPeriodCounter=1
010 Get relevant zero coupon bond prices, forward rates and sigmas for new current time t. Remember that in Chapters 7, 8 and 9 the sigmas chosen were a function of the spot rate at the current time t.
011 If RemainingPeriodCounter<=N, take one random draw each for risk factors 1, 2 and 3 else go to 100
012 Shift all forward rates and zero coupon bond prices based on the new shift amounts, which will be a randomly drawn number of sigmas up or down for all three risk factors
013 Store results for valuation and other calculations
014 RemainingPeriodCounter=RemainingPeriodCounter+1
015 Go to 010

100 If ScenarioCounter<M, ScenarioCounter=ScenarioCounter+1 else 500
500 Do valuations
501 Do net income simulations with default "turned on"
502 Do liquidity risk analysis with default "turned on"
503 Do value at risk with default "turned off"
504 Do credit adjusted value at risk with default "turned on"
998 Generate reports
999 End

This approach is a very straightforward modification of what we did in Chapters 6 through 9. That being said, many analysts "get it wrong," even with a good understanding of Chapters 6 through 9 and this Monte Carlo process. In the next section, we discuss some common pitfalls in interest rate risk analysis.

Common Pitfalls in Interest Rate Risk Management

The most common pitfalls in interest rate risk analysis are very fundamental mistakes, but they are surprisingly common

a. Using a number of risk factors or volatility assumptions that are unrealistic, violating conditions 1, 2 or both. The most common mistake is the most serious one: using one or two factor interest rate models in a world where both realism and regulators demand 6-10 factors.
b. Selecting interest rate paths from a bushy tree or interest rate lattice as a substitute for Monte Carlo, without realizing the huge errors that result from this detour from best practice.

We discuss each of these problems in turn.

Pitfalls in the Use of One Factor Term Structure Models

One factor models of the term structure of interest rates were developed to provide insights into the valuation of fixed income options. Our implementation of the one factor Heath Jarrow and Morton model in Chapters 6 and 7 is a good example of a modern implementation. Since the one factor models' early development from 1977 to 1993, they have been consistently misapplied in asset and liability (interest rate risk) management. Many analysts, choosing one of the common one factor models, create random interest rate scenarios to evaluate the interest rate risk of the firm. This section uses 50 years of Dickler, Jarrow and van Deventer's U.S. Treasury yield curve information to prove that the major implications of one factor term structure models are inconsistent with historical data. A multi-factor approach is necessary to create realistic interest rate simulations, as we have emphasized in Chapters 8 and 9.

In this section, we take the same approach as Jarrow, van Deventer and Wang (2003) in their paper "A Robust Test of Merton's Structural Model for Credit Risk." In that paper, the authors explained how, for any choice of model parameters, the Merton model of risky debt implies that stock prices would rise and credit spreads would fall if the assets of the firm were to rise in value. The opposite movements would occur if the assets of the firm were to fall in value. Jarrow, van Deventer, and Wang used a wide array of bond data and proved that this implication of the Merton model was strongly rejected by the data at extremely high levels of statistical significance.

We now take a similar approach to one factor term structure models. Speaking broadly, the original Macaulay (1938) duration concept can be thought of as the first one factor term structure model. Robert Merton (1970) developed a similar term structure model in a continuous time framework. Oldrich Vasicek (1977) brought a greater realism to one factor term structure modeling by introducing the concept of mean reversion to create interest rate cycles. Cox, Ingersoll, and Ross (1985) pioneered the affine term structure model, where the magnitude of random shocks to the short term interest rate varies with the level of interest rates. Ho and Lee (1986) and Hull and White (1993) developed one factor term structure models which could be fit exactly to observable yields. Black, Derman and Toy (1990) and Black and Karasinski (1991) focused on term structure models based on the log of the short rate of interest, insuring that rates would not be negative. We will cover the basics of these legacy models briefly in Chapter 13.

All of these models are special cases of the general Heath, Jarrow and Morton (1990, 1990, and 1992) framework for N-factor term structure models that we have illustrated in Chapters 6 through 9. For two excellent summaries of the literature on term structure models, please see Duffie and Kan (1994) and Jarrow (2009).

Among one factor term structure models, for any set of parameters, there is almost always one characteristic in common. The majority of one factor models make all interest rates positively correlated (which implies the time-dependent interest rate volatility sigma(T) is greater than 0). We impose this condition on a one factor model. The non-random time-dependent drift(T) in rates is small and, if risk premia are consistent and positive (for bonds) across all maturities T, then the drift will be positive as in the Heath Jarrow and Morton (1992) drift condition which we labeled μ in Chapters 6 through 9. With more than 200 business days per year, the daily magnitude of the interest rate drift is almost always between plus one and minus one basis point. Hence, one might want to exclude from consideration yield curve changes that are small, so the effect of the drift is removed. We ignore the impact of drift in what follows.

Given these restrictions, for a random shift upward in the single risk factor (usually the short rate of interest in most of the legacy one factor models), the zero coupon yields at all maturities will rise. For a random shift downward in the single factor, all zero coupon yields will fall. If the short rate is unchanged, there will be no random shock to any of the other zero coupon yields either.

Following Jarrow, van Deventer and Wang, we ask this simple question: "What percent of the time is it true that all interest rates either rise together, fall together, or stay unchanged?" If this implication of one factor term structure models is consistent with actual yield curve movements, we cannot reject their use. If this implication of one factor term structure models is inconsistent with actual yield movements, we reject their use in interest rate risk management and asset and liability management.

Results from the U.S. Treasury Market, 1962-2011

In this section we use the same data as Dicker, Jarrow and van Deventer (2011): U.S. Treasury yields reported by the Federal Reserve in its H15 statistical release from January 2, 1962 to August 22, 2011. During this period, the Federal Reserve changed the maturities on which it reported data frequently:

Analysis of Federal Reserve H15 Release Data Regimes

Start Date	End Date	Data Regime	Number of Observations
1/2/1962	6/30/1969	1, 3, 5 and 10 years	1,870
7/1/1969	5/28/1976	1, 3, 5, 7 and 10 years	1,722
6/1/1976	2/14/1977	1, 2, 3, 5, 7 and 10 years	178
2/15/1977	12/31/1981	1, 2, 3, 5, 7, 10 and 30 years	1,213
1/4/1982	9/30/1993	3 and 6 months with 1, 2, 3, 5, 7, 10 and 30 years	2,935
10/1/1993	7/30/2001	3 and 6 months with 1, 2, 3, 5, 7, 10, 20 and 30 years	1,960
7/31/2001	2/15/2002	1, 3 and 6 months with 1, 2, 3, 5, 7, 10, 20 and 30 years	135
2/19/2002	2/8/2006	1, 3 and 6 months with 1, 2, 3, 5, 7, 10, and 20 years	994
2/9/2006	8/22/2011	Second era: 1, 3 and 6 months with 1, 2, 3, 5, 7, 10, 20 and 30 years	1,388
		Total	12,395

We calculate the daily changes in yields over this period, eliminating the first daily change after a change in data regime, leaving 12,386 data points for analysis. We calculate daily yield changes on three bases:

- Using the yields from the Federal Reserve with no other analysis
- Using monthly forward rates
- Using monthly zero coupon bond yields, expressed on a continuous compounding basis

Monthly forward rates and zero coupon yields were calculated using the maximum smoothness forward rate approach that we outlined in Chapter 5. We then asked the question posed above: "What percentage of the time were yield curve movements consistent with the implications of one factor term structure models?" Yield curve movements were considered consistent with one factor term structure models in three cases:

- Negative shift, in which interest rates at all maturities declined
- Positive shift, in which interest rates at all maturities increased
- Zero shift, in which interest rates at all maturities remained unchanged

By definition, any other yield curve movement consists of a combination of positive, negative and zero shifts at various maturities on that day. We label this a yield curve "twist," something that cannot be explained by one factor term structure models.

The results are striking. Using only the input data from the H15 statistical release, with no yield curve smoothing, 7,716 of the 12,386 days of data showed yield curve twists, 62.3% of all observations. This is a stunningly high level of inconsistency with the assumptions of one factor term structure models.

Summary of U.S. Treasury Yield Curve Shifts
Type of Daily Rate Shifts
January 2, 1962 to August 22, 2011

Type of Shift	U.S. Treasury Input Data	Monthly Zero Coupon Bond Yields	Monthly Forward Rates
Negative shift	2,276	1,486	297
Positive shift	2,306	1,496	327
Zero shift	88	88	88
Subtotal	4,670	3,070	712
Twist	7,716	9,316	11,674
Grand Total	12,386	12,386	12,386

For monthly zero coupon bond yields, yield curve twists occurred on 75.2% of the 12,386 days in the sample. When forward rates were examined, yield curve twists prevailed on 94.3% of the days in the sample. One factor term structure models would have produced results that are consistent with actual forward rates only 5.7% of the 1962-2011 period.

Summary of U.S. Treasury Yield Curve Shifts Percentage Distribution of Daily Rate Shifts January 2, 1962 to August 22, 2011			
Type of Shift	**U.S. Treasury Input Data**	**Monthly Zero Coupon Bond Yields**	**Monthly Forward Rates**
Negative shift	18.4%	12.0%	2.4%
Positive shift	18.6%	12.1%	2.6%
Zero shift	0.7%	0.7%	0.7%
Subtotal	37.7%	24.8%	5.7%
Twist	62.3%	75.2%	94.3%
Grand Total	100.0%	100.0%	100.0%

Using the approach of Jarrow, van Deventer and Wang (2003), we can now pose this question: is the consistency ratio we have derived different to a statistically significant degree from the 100% ratio that would prevail if the hypothesis of a one factor term structure is true?

We use the fact that the measured consistency ratio p is a binomial variable (true or not true, consistent or not consistent) with a standard deviation s as follows:

$$s = \sqrt{\frac{p(1-p)}{n}}$$

where p is the measured consistency ratio and n is the sample size, 12,386 observations. The chart below shows the standard deviations for each of the three data classes and the number of standard deviations between the measured consistency ratio p and the ratio (100%) that would prevail if the one factor term structure model hypothesis were true.

Consistency of U.S. Treasury Yield Daily Shifts with One Factor Models of the Term Structure of Interest Rates January 2, 1962 to August 22, 2011			
Yield Curve Analysis	U.S. Treasury Input Data	Monthly Zero Coupon Bond Yields	Monthly Forward Rates
Consistency with 1 Factor Models	37.7%	24.8%	5.7%
Standard Deviation of Consistency %	0.4%	0.4%	0.2%
Number of Standard Deviations from 100%	143.1	193.9	450.6

We reject the hypothesis that the consistency ratio is 100%, the level that would prevail if one factor term structure models were an accurate description of yield curve movements. The hypothesis that one factor term structure models are "true" is rejected with an extremely high degree of statistical significance, by more than 100 standard deviations in all three cases.

There are a number of very serious errors that can result from an interest rate risk and asset and liability management process that relies solely on the assumption that one factor term structure models are an accurate description of potential yield movements:

a. Measured interest rate risk will be incorrect, and the degree of the error will not be known. Using the data above, 62.3% of the time actual yield curves will show a twist, but the modeled yield curves will never show a twist.

b. Hedging using the duration or one factor term structure model approach assumes that interest rate risk of one position (or portfolio) can be completely eliminated with the proper short position in an instrument with a different maturity. The duration/one factor term structure model approach assumes that if interest rates on the first position rise, interest rates will rise on the second position as well so "going short" is the right hedging direction. We discuss this in Chapter 11. The data above shows that on 62.3% of the days from 1962 to 2011, this "same direction" assumption was potentially false (some maturities will show same direction changes and some will show opposite direction changes) and the hedge could actually ADD to risk, not reduce risk.

c. All estimates of prepayments and interest rate-driven defaults will be measured inaccurately

d. Economic capital will be measured inaccurately

e. Liquidity risk will be measured inadequately

f. Non-maturity deposit levels will be projected inaccurately

This is an extremely serious list of deficiencies. The remedy is to move to the recommended 6-10 factors. Academic assumptions about the stochastic processes used to model random interest rates have been too simple to be realistic. Jarrow (2009) notes,

> "Which forward rate curve evolutions (HJM volatility specifications) fit markets best? The literature, for analytic convenience, has favored the affine class but with disappointing results. More general evolutions, but with more complex computational demands, need to be studied. How many factors are needed in the term structure evolution? One or two factors are commonly used, but the evidence suggests three or four are needed to accurately price exotic interest rate derivatives."

This view is shared by the Basel Committee of Banking Supervision, which we quoted above as requiring six factors for market risk analysis. A move to two risk factors doesn't begin to address the problem, as we explained earlier in this chapter. We now turn to another common pitfall in implementing interest rate risk models.

Common Pitfalls in Asset and Liability Management

Interpolating Monte Carlo Results, or How to Prove Augusta National is Not a Golf Course

The authors would like to thank colleagues Dr. Predrag Miocinovic and Dr. Alexandre Telnov for their contributions to this section.

With the advances in software design and computer speeds of recent years, risk managers now have very powerful techniques for Monte Carlo simulation at their disposal. That is the reason why the multi-factor interest rate framework of Heath, Jarrow and Morton is so useful and important. That being said, we find many shortcuts have been taken that can provide very misleading estimates of risk levels. This section explains how the misuse of Monte Carlo interest rate simulation can prove that Augusta National is not a golf course and that the Grand Canyon does not exist.

A common approach to risk measurement can be described by this example, which was explained in a public forum by one of the largest banks in the world:

- We have three factors driving the yield curve: level, shift and bend
- We chose "no arbitrage values" for each of these factors and choose 7 values for each factor, the mean and plus or minus 1, 2 and 3 standard deviations from the mean
- We then calculate portfolio values for all 7 x 7 x 7 = 343 combinations of these risk factors
- We calculate our nth percentile value at risk from the smoothed surface of values that results from these 343 portfolio values, using 1 million draws from this smoothed surface to get the 99.97th percentile of our risk.

What's wrong with this approach? It certainly seems reasonable at first glance and could potentially save a lot of computing time versus a large Monte Carlo simulation like the one outlined earlier in this chapter. Let's look at it out of the finance context. We are posed this problem:

Our golf ball is placed in the location marked by the white square in the lower left-hand corner of this aerial photo of a golf course in Augusta, Georgia. However, we are blindfolded and cannot see. We are told that the ball is either on a green at Augusta National Golf Course or that it is not, and our task is to determine which case is the true one.

We calculate the standard deviation of the length of the first put made by every golfer on every hole at last year's Masters golf tournament. From the location of our golf ball, we then get the latitude and longitude of 49 points: our current golf ball location plus all other locations that fall on a grid defined by 1, 2, or 3 standard deviations (of putting length) North-South or East-West with our golf ball in the middle. The grid looks like this:

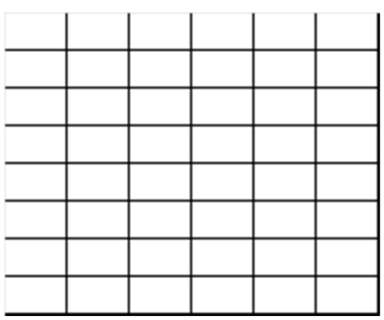

At each point where the North-South and East-West lines intersect, we are given very precise global positioning satellite estimates of the elevation of the ground at that point. Upon reviewing the 49 values, we smooth the surface connecting them and we find that the 99.97^{th} percentile difference in elevation on this smoothed surface is only 2 inches. Moreover, we find that the greatest difference in elevation among any of the points on the grid is only slightly more than two inches. We know by the rules of golf that a golf hole must be at least 4 inches in depth. We reject the hypothesis that we are on a green. ***But we are completely wrong!***

This silly example shows the single biggest pitfall of interpolating between Monte Carlo values using some subset of a true Monte Carlo simulation. For the same reasons, using the Heath Jarrow and Morton bushy tree is also not "best practice" when Monte Carlo is available, with the exception of selected securities valuations which we discuss later in this book. In the Masters example, we have assumed that the surface connecting the grid points is smooth, but the very nature of the problem contradicts this assumption. The golf green is dramatically "unsmooth" where there is a hole in the green. Our design of the Monte Carlo guaranteed a near certainty that our analysis of "on a green"/"not on a green" would be no better than flipping a coin.

Similarly, we could have been placed blindfolded on a spot somewhere in the United States and told either we are within 100 miles of the Grand Canyon or we are not. We again form a grid of 9 points all of which are plus or minus 100 miles North-South or East-West of where we were placed blind folded, as shown below:

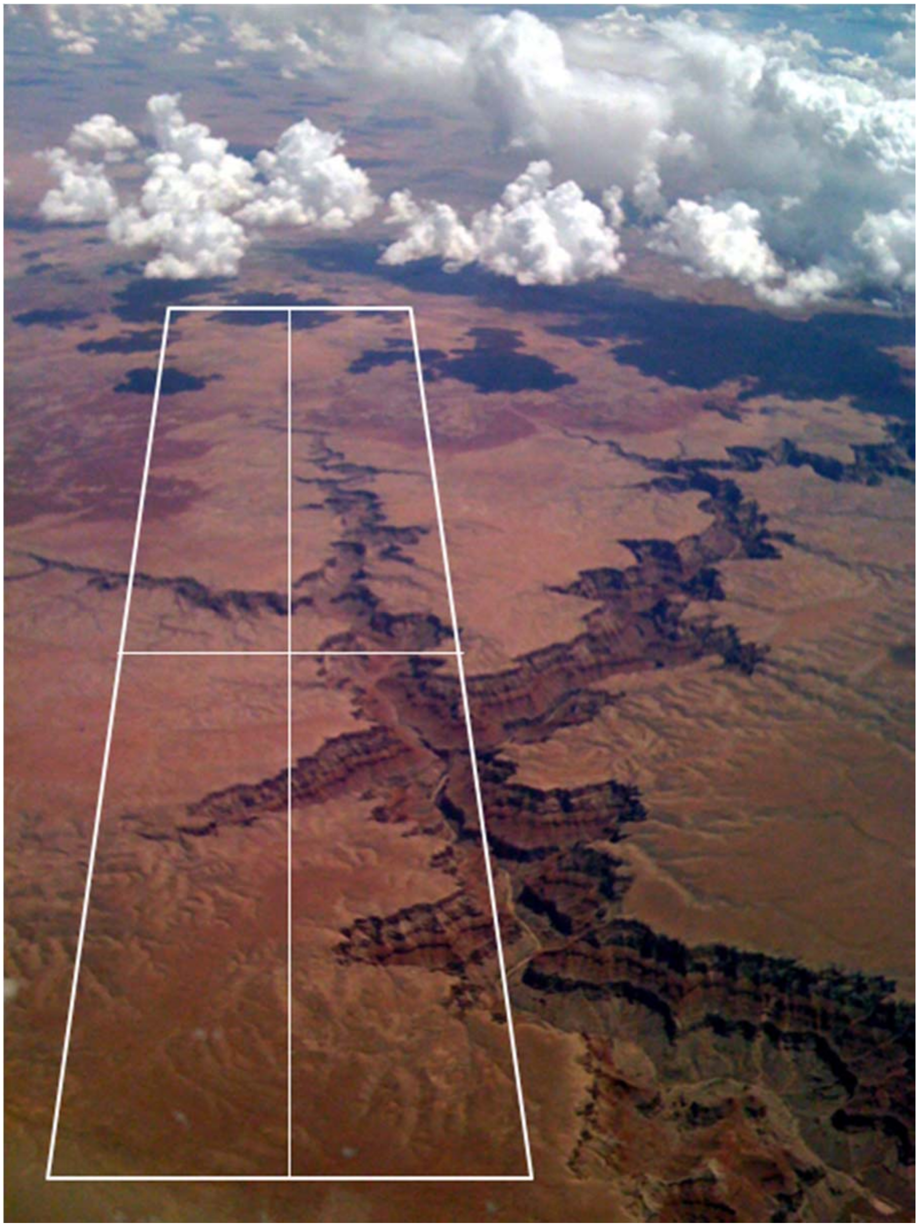

Again, we take the GPS elevation at these 9 grid points, smooth the surface, and take the 99.97th percentile elevation difference on the smoothed surface. We find that our elevation difference is only 100 meters at the 99.97th percentile. We reject the hypothesis that we are within 100 miles of the Grand Canyon since we know the average depth of the Grand Canyon is about 1,600 meters. Again, we are grossly wrong in another silly example, because our assumption that intervening values between our grid points are connected by a smooth surface is wrong.

Summarizing the Problems with Interpolated Monte Carlo Simulation for Risk Analysis

For relatively simple, low dimensional systems that arise in physics and engineering, where correlations are well-controlled even in the tails, a reliable and believable description of the asymptotic behavior of the global probability density function may indeed be possible, using a phrase like “smooth” and multi-factor smoothing in an extension of Chapter 5. In high dimensional cases, such as in the risk management of financial institutions, things are much more complex with respect to the global probability density function of values, cash flows and net incomes. Nothing illustrates that more clearly than the $2 billion losses announced by JP Morgan Chase on May 8, 2012. The simple examples above illustrate two problems that can be extremely serious problems if risk is measured by interpolating between Monte Carlo valuations on a grid or small subset of possible outcomes (like the branches of a Heath Jarrow and Morton bushy tree or a Hull and White interest rate lattice):

- **Smoothness is an assumption, not a fact**, as our simple examples show. Interpolating yield curves using a smoothness criterion is reasonable for good economic reasons, but as we show below, it is highly likely that portfolio values and cash flows in financial institutions have sharp discontinuities.
- **The Nth percentile “worst case” derived from the simulation is not relative to the true “worst case,” it is relative to the worst case on the grid** defined by the analyst. In our Grand Canyon example, the “worst case” elevation differential was 100 meters, but the true “worst case” elevation difference was the true depth of the Grand Canyon, 1600 meters.

Almost the entire balance sheet of financial institutions is filled with caps, floors, prepayment risk and defaults. In the JPMorgan 2012 case, the balance sheet was filled with huge credit default swap positions taken by the trader called the “London Whale.” One cannot assume that the law of large numbers will make these 0/1 changes in values and cash flows vary in a smooth way as the number of transactions rises. We can show that with a simple example. Assume it’s February 2000. 3 month US dollar Libor is 6.00%. The five year monthly standard deviation of 3 month US dollar Libor from March 1995 to February 2000 was 32 basis points. We have only 3 transactions on our balance sheet:

- The first digital option pays us $1 if 3 month Libor is 6.25% or higher in 18 months
- The second digital option pays us $1 if 3 month Libor is 6.05% or less in 18 months
- The third digital option requires us to pay $100 if 3 month Libor is below 4.00% in 18 months

We set up our grid of 7 values for 3 month Libor in 18 months centered around the current level of Libor (6.00%) and plus and minus 1, 2, and 3 standard deviations. We do valuations at these Libor levels in 18 months:

- Scenario 1: 6.96% Libor
- Scenario 2: 6.64% Libor
- Scenario 3: 6.32% Libor
- Scenario 4: 6.00% Libor
- Scenario 5: 5.68% Libor
- Scenario 6: 5.36% Libor
- Scenario 7: 5.04% Libor

Our valuations of our portfolio are $1.00 in every scenario:

- Scenario 1: 6.96% Libor means we earn $1 on digital option 1
- Scenario 2: 6.64% Libor means we earn $1 on digital option 1
- Scenario 3: 6.32% Libor means we earn $1 on digital option 1
- Scenario 4: 6.00% Libor means we earn $1 on digital option 2
- Scenario 5: 5.68% Libor means we earn $1 on digital option 2
- Scenario 6: 5.36% Libor means we earn $1 on digital option 2
- Scenario 7: 5.04% Libor means we earn $1 on digital option 2

Typical interpolation of these Monte Carlo results would smooth these calculated values. Ardent interpolation users then draw from this smoothed surface (we've heard up to 1 million draws from this interpolated surface). In this case, all values on the grid are 1, the smoothed surface is 1 everywhere and we announce proudly "With 99.97% confidence we have no risk because we earn $1 in every scenario."

WRONG! We earn nothing in the range between 6.05% and 6.25%. We completely missed this because we didn't sample in that range and assumed cash flows were smooth, but they weren't! Even worse, the actual value of 3 month Libor in August 2001 was 3.47%, and we had to pay out $100 on digital option 3!

This is a simple example of a very real problem with valuing from a small number of grid points instead of doing a true Monte Carlo simulation. One of the authors remembers a scary period in the mid-1980s when typical interest rate risk simulations never simulated interest rates high enough to detect the huge balance of adjustable rate mortgages in bank portfolios with interest rate caps at 12.75% and 13.00%. The young analysts, with no knowledge of history and no examination of the portfolio at a detailed level, did not detect this risk, just as our simple example above missed medium sized risk within the grid and huge risk outside the grid.

This graph of the 10 year and 5 year moving average standard deviations of month-end 3 month Libor shows how easy it is to design "grid widths" that end up being horribly wrong.

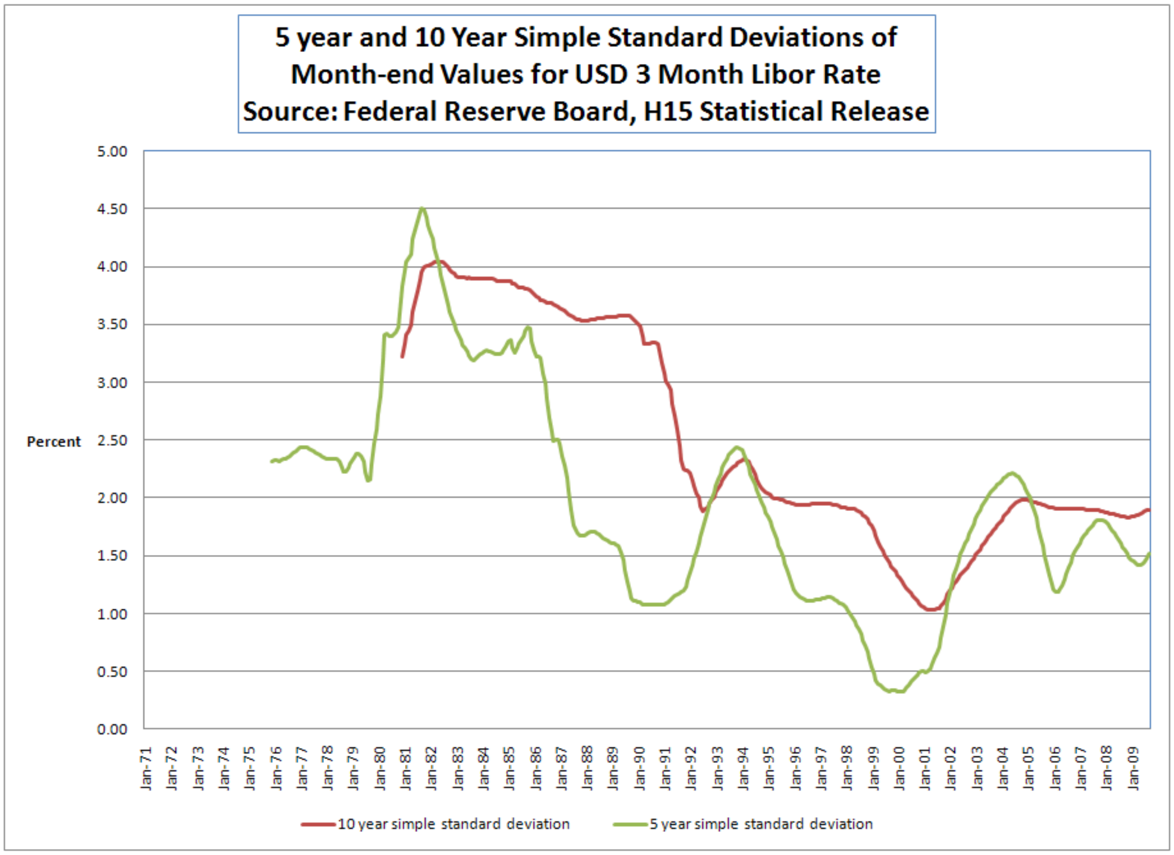

Take the 32 basis point standard deviation of 3 month Libor over the five years ended in February 2000. The standard deviation of 3 month Libor for the 5 years ended in September 1981 was 450 basis points!

Going back to our August National example, if we don't know the true width of the hole on a golf course, we at least know that 49 grid points is not enough. We do 1,000 or 10,000 or 100,000 random north-south/east-west coordinates looking for an elevation change of at least 4 inches (the depth of a golf hole). The degree of confidence we seek determines the number of random coordinates we draw. If we know by the rules of golf that the whole is exactly 4.25 inches wide, we can be much more refined in our search techniques. The trouble in risk management, however, is that we don't know if we're looking for a golf hole (a few mortgage defaults in 2005) or the Grand Canyon (massive mortgage defaults from the huge drop in U.S. and U.K. home prices in 2007-2009).

Given the current state of risk systems technology, there is no reason not to do a true Monte Carlo simulation that defines the amount of risk embedded in the portfolio with any reasonable desired level of precision. The short cut of interpolating between a small subsample of scenarios is extremely dangerous. The full power of N-factor interest rate modeling using the Heath Jarrow and Morton approach deserves a full Monte Carlo simulation.

We now turn to interest rate mismatching and hedging in a modern framework.