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 4: Risk Management Applications, Instrument by Instrument

Chapter 27:
American Fixed Income Options

In Chapters 21 through 26, we have emphasized the valuation of securities where the option embedded in the security was exercisable at only one date. These European options, as we have shown, generally have explicit analytical solutions in the Vasicek family of term structure models. As we saw in Chapter 3, however, one factor interest rate models are very inaccurate models of interest rate movements. In this chapter, we continue to emphasize numerical solutions using the example of the 3 factor Heath Jarrow and Morton model from Chapter 9. In previous chapters, we have frequently remarked that an HJM Monte Carlo simulation was the preferred valuation methodology, with an HJM "bushy tree" used only for expositional purposes. In this chapter, the situation is reversed. Using the bushy tree approach is the most accurate technique for valuing American fixed income options. Monte Carlo simulation, even using the HJM approach, is one notch lower in accuracy for reasons we explain. Please note, however, that, even in the case of a one factor term structure model, numerical methods are necessary for the accurate valuation of American fixed income options. This is even more true when the issuer of the security has a positive probability of default. The interaction of the security's payoffs with the credit risk of the issuer has interesting and important implications.

The major advances in computer science since the first volume of this book was published have dramatically relaxed the computational constraints that formerly motivated the use of closed form solutions instead of highly accurate multi-factor simulations or bushy trees. The increasing use of 64-bit operating systems, with their huge expansion in available memory, makes the techniques in this chapter available at the transaction level for the entire enterprise.

We note also that the Jarrow-Merton put option, the key integrated risk measure that we focus on in this volume, can be thought of as an American option. The bank deposit insurance provided by the Federal Deposit Insurance Corporation in the United States is effectively an American put option on the value of deposits provided to depositors and paid for by the banks themselves (if priced correctly) and taxpayers (if priced incorrectly). Bank deposit insurance is an American option because the bank can default at any time, not just one point in time like a European option.

Much of the early work on lattice or bushy tree valuations of American options was done using Hull and White's [1990,1993, 1994] popular trinomial lattice. The lattice was initially developed for 1 factor term structure models and then expanded to the two factor case. As we saw in Chapter 3, between 5 and 10 factors are needed to accurate model even a relatively transparent yield curve like that for U.S. Treasury securities. A one or two factor model is simply not a good approximation to actual yield curve movements, and a lattice based on these assumptions will not capture some of the yield curve movements (like yield curve twists) that can have a big impact on the exercise of an American option like the prepayment option on a mortgage in the United States and Japan.

The reason for devoting the bulk of this and the two following chapters to American fixed income options is that almost all financial institutions have the majority of their balance sheets devoted to American fixed income options, explicitly or implicitly. A brief list of examples shows how critical this topic is. Typical American options on financial institution balance sheets include

- The right to terminate a life insurance policy in return for receipt of its surrender value
- The right to resign as the customer of an investment management firm
- The right to prepay a mortgage loan
- The right to withdraw as a bank from the consumer deposit gathering business
- The right to withdraw almost any consumer bank deposit at any time either at par or upon the payment of an early withdrawal penalty
- The right of a corporate borrower to declare bankruptcy
- The right of a financial institution to pay dividends
- The right to exercise a foreign exchange option
- The right to exercise a standard swaption contract

The list of examples is almost endless. Needless to say, the topic at hand is a critical one.

An Overview of Numerical Techniques for Fixed Income Option Valuation

Financial market participants generally use one of six approaches to the valuation of American fixed income options:

- Analytical approximations
- Monte carlo simulation
- Bushy trees
- Finite difference methods
- Binomial lattices (a special case of the bushy tree, because the tree "recombines")
- Trinomial lattices (another special case of bushy tree)

Analytical approximations aren't sufficiently accurate for us to devote time to here. We will discuss the alternatives to the bushy tree after first using the 3 factor Heath Jarrow and Morton framework from Chapter 9 on two examples of fixed income options, a callable fixed rate bond and a prepayable amortizing loan. There is a wide variety of American fixed income options in everyday finance, but loans and bonds that are callable or prepayable (the terms have the equivalent economic meaning) are particularly numerous.

An Example of Valuation of a Callable Bond with a 3 Factor HJM Bushy Tree

In Chapter 9, we analyzed the value of a bond with $100 in principal outstanding that pays 3% interest annually and matures at time T=4. The cash flows on the bond and its valuation in Chapter 9, 104.7071, are reproduced here in Figure 1. The valuation is calculated by multiply the cash flows in the cash flow table times the corresponding probability weighted discount factors that were displayed in Chapter 9:

Figure 1

	Maturity of Cash Flow Received				
	Current Time				
Row	0	1	2	3	4
1	0.0000	3.0000	3.0000	3.0000	103.0000
2		3.0000	3.0000	3.0000	103.0000
3		3.0000	3.0000	3.0000	103.0000
4		3.0000	3.0000	3.0000	103.0000
5			3.0000	3.0000	103.0000
6			3.0000	3.0000	103.0000
7			3.0000	3.0000	103.0000
8			3.0000	3.0000	103.0000
9			3.0000	3.0000	103.0000
10			3.0000	3.0000	103.0000
11			3.0000	3.0000	103.0000
12			3.0000	3.0000	103.0000
13			3.0000	3.0000	103.0000
14			3.0000	3.0000	103.0000
15			3.0000	3.0000	103.0000
16			3.0000	3.0000	103.0000
17				3.0000	103.0000
18				3.0000	103.0000
19				3.0000	103.0000
20				3.0000	103.0000
21				3.0000	103.0000
22				3.0000	103.0000
23				3.0000	103.0000
24				3.0000	103.0000
25				3.0000	103.0000
26				3.0000	103.0000
27				3.0000	103.0000
28				3.0000	103.0000
29				3.0000	103.0000
30				3.0000	103.0000
31				3.0000	103.0000
32				3.0000	103.0000
33				3.0000	103.0000
34				3.0000	103.0000
35				3.0000	103.0000
36				3.0000	103.0000
37				3.0000	103.0000
38				3.0000	103.0000
39				3.0000	103.0000
40				3.0000	103.0000
41				3.0000	103.0000
42				3.0000	103.0000
43				3.0000	103.0000
44				3.0000	103.0000
45				3.0000	103.0000
46				3.0000	103.0000
47				3.0000	103.0000
48				3.0000	103.0000
49				3.0000	103.0000
50				3.0000	103.0000
51				3.0000	103.0000
52				3.0000	103.0000
53				3.0000	103.0000
54				3.0000	103.0000
55				3.0000	103.0000
56				3.0000	103.0000
57				3.0000	103.0000
58				3.0000	103.0000
59				3.0000	103.0000
60				3.0000	103.0000
61				3.0000	103.0000
62				3.0000	103.0000
63				3.0000	103.0000
64				3.0000	103.0000
Risk Neutral Value =					104.7071

In Chapter 14, we assumed that we could observe prices of callable fixed rate coupon bonds, where the call option was European in nature: the bonds could be prepaid or "called" at only one point in time. In this example, we assume that the bonds could be called at "any time." In the context of our 4 step bushy tree from Chapter 9, "any time" effectively means the bonds can be called at time 0, time 1, time 2 or time 3, prior to maturity at time 4. If we need to evaluate the call option's impact more granularly, we would simply increase the number and shorten the length of the time steps in the bushy tree. Please note that a "Bermudan option" is an option that can be exercised at many points in time but not at "any time." In the context of this example, one could imagine that a bond could be called at time 2 or time 3 but not at time 0 or time 1. This is a special case of a Bermudan option. The bond would be described as "four years, non-call two." We assume that the call option can be exercised at any of the nodes of the lattice except the maturity date (because calling at that time is identical with normal principal and interest payments at maturity). Bermudan option analysis is identical except that the option to call or prepay is "turned off" at some points in time.

How would a rational corporate treasurer or bond investor value the probability that this 3% 4 year bond would be pre-paid (we'll use the term "called" when talking about bonds, "pre-paid" when talking about loans)? These probabilities need to be assessed in order to appropriately populate the cash flow table for Heath Jarrow and Morton valuation. The value of the bond at time zero, and the decision about whether to call the bond at time zero, depends on what the bond issuer would do at the 4 nodes at time T=1. At each of the four nodes at T=1, there are a set of 4 more nodes at time T=2, and the time 1 behavior depends on what the bond issuer would do at time 2. Finally, at each of the 16 nodes at time T=2, there are 4 nodes at time T=3 and the time 2 behavior depends on the actual decisions that the bond issuer would make at time T=3, where there are 64 nodes. This is a dynamic programming problem that we need to analyze in a recursive fashion. We start at the end of the bond's life and set our call strategy for the 64 nodes at time T=3. Having done that, we step backwards in time to the 16 nodes at time T=2. Because we know what the time T=3 actions are, we can deduce what to do on the 16 nodes at time T=2. We then repeat the process and step backwards in time to time T=1 (4 nodes) and time T=0 (1 node). This gives us the cash flows at each node. We insert these cash flows in the HJM cash flow table, multiply by the appropriate probability-weighted discount factors, and get the risk-neutral value. We do that now.

First we analyze whether or not to call the bonds at time T=3 at each of the 64 nodes. We assume the decision is made immediately after time T=3 interest of $3 is paid. The analysis is simple. We multiple the zero coupon bond price for a bond maturing at T=4 as of time 3 for that state st times the cash flow of $103 to get the net present value of the "no call" strategy, $P(3,4,s_t)103$. If this present value is more than 100, we call the bonds now and pay $100 in principal at time T=3 in addition to the $3 in interest paid 1 second earlier. If the present value is $100 or less, we leave the bonds outstanding and pay $103 at time T=4. Figure 2 analyzes the call or no call decision at each node at time T=3:

Figure 2

	Maturity of Cash Flow Received				
	Current Time				
Row	0	1	2	3	Action?
1	0.0000	0.0000	0.0000	97.0685	No Call
2		0.0000	0.0000	97.5686	No Call
3		0.0000	0.0000	97.9357	No Call
4		0.0000	0.0000	95.3373	No Call
5			0.0000	99.8991	No Call
6			0.0000	100.3984	Call
7			0.0000	100.9516	Call
8			0.0000	98.9787	No Call
9			0.0000	97.8892	No Call
10			0.0000	101.0018	Call
11			0.0000	99.8258	No Call
12			0.0000	97.6288	No Call
13			0.0000	97.6996	No Call
14			0.0000	98.2030	No Call
15			0.0000	98.5724	No Call
16			0.0000	95.9571	No Call
17				98.8942	No Call
18				102.0387	Call
19				100.8506	Call
20				98.6311	No Call
21				100.7792	Call
22				102.4334	Call
23				102.2163	Call
24				101.7783	Call
25				100.3673	Call
26				103.3860	Call
27				102.7336	Call
28				100.8145	Call
29				101.2022	Call
30				101.7080	Call
31				102.2685	Call
32				100.2698	Call
33				100.7035	Call
34				101.2068	Call
35				101.7645	Call
36				99.7757	No Call
37				100.7035	Call
38				101.2068	Call
39				101.7645	Call
40				99.7757	No Call
41				99.3869	No Call
42				102.3762	Call
43				101.7301	Call
44				99.8297	No Call
45				98.5580	No Call
46				101.6918	Call
47				100.5077	Call
48				98.2958	No Call
49				99.5157	No Call
50				102.6800	Call
51				101.4844	Call
52				99.2509	No Call
53				100.4064	Call
54				100.9082	Call
55				101.4642	Call
56				99.4813	No Call
57				100.4064	Call
58				100.9082	Call
59				101.4642	Call
60				99.4813	No Call
61				98.5589	No Call
62				100.7152	Call
63				100.2675	Call
64				98.0738	No Call

Based on these rational call strategies, the cash flows on a bond that could be called only at time T=3 are given in Figure 3. The value of this bond, a "4 year non-call 3" bond, would be 104.2480, less than the non-call version of the bond:

Figure 3

	Maturity of Cash Flow Received				
	Current Time				
Row	0	1	2	3	4
1	0.0000	3.0000	3.0000	3.0000	103.0000
2		3.0000	3.0000	3.0000	103.0000
3		3.0000	3.0000	3.0000	103.0000
4		3.0000	3.0000	3.0000	103.0000
5			3.0000	3.0000	103.0000
6			3.0000	103.0000	0.0000
7			3.0000	103.0000	0.0000
8			3.0000	3.0000	103.0000
9			3.0000	3.0000	103.0000
10			3.0000	103.0000	0.0000
11			3.0000	3.0000	103.0000
12			3.0000	3.0000	103.0000
13			3.0000	3.0000	103.0000
14			3.0000	3.0000	103.0000
15			3.0000	3.0000	103.0000
16			3.0000	3.0000	103.0000
17				3.0000	103.0000
18				103.0000	0.0000
19				103.0000	0.0000
20				3.0000	103.0000
21				103.0000	0.0000
22				103.0000	0.0000
23				103.0000	0.0000
24				103.0000	0.0000
25				103.0000	0.0000
26				103.0000	0.0000
27				103.0000	0.0000
28				103.0000	0.0000
29				103.0000	0.0000
30				103.0000	0.0000
31				103.0000	0.0000
32				103.0000	0.0000
33				103.0000	0.0000
34				103.0000	0.0000
35				103.0000	0.0000
36				3.0000	103.0000
37				103.0000	0.0000
38				103.0000	0.0000
39				103.0000	0.0000
40				3.0000	103.0000
41				3.0000	103.0000
42				103.0000	0.0000
43				103.0000	0.0000
44				3.0000	103.0000
45				3.0000	103.0000
46				103.0000	0.0000
47				103.0000	0.0000
48				3.0000	103.0000
49				3.0000	103.0000
50				103.0000	0.0000
51				103.0000	0.0000
52				3.0000	103.0000
53				103.0000	0.0000
54				103.0000	0.0000
55				103.0000	0.0000
56				3.0000	103.0000
57				103.0000	0.0000
58				103.0000	0.0000
59				103.0000	0.0000
60				3.0000	103.0000
61				3.0000	103.0000
62				103.0000	0.0000
63				103.0000	0.0000
64				3.0000	103.0000
Risk Neutral Value =					104.2480

Having done the call analysis for time T=3, we now step backwards in time to the 16 nodes at time T=2. At each of those 16 nodes, we look at the 64 nodes at time T=3 and study the net present value of the "no call" strategy as of time T=2. The net present value as of time T=2 is a function of the discount factor, shift probability and time T=3 value if the bonds are not called at time T=2. Recall the shift probabilities from Chapter 9, which we reproduce here as Figure 4:

Figure 4

		Shift Probabilities			
		Current Time			
Scenario	Row Numb	0	1	2	3
S-1, S-1, S-1	1	100%	12.50%	12.50%	12.50%
S-1, S-1, S-2	2		12.50%	12.50%	12.50%
S-1, S-1, S-3	3		25.00%	25.00%	25.00%
S-1, S-1, S-4	4		50.00%	50.00%	50.00%
S-1, S-2, S-1	5			12.50%	12.50%
S-1, S-2, S-2	6			12.50%	12.50%
S-1, S-2, S-3	7			25.00%	25.00%
S-1, S-2, S-4	8			50.00%	50.00%
S-1, S-3, S-1	9			12.50%	12.50%
S-1, S-3, S-2	10			12.50%	12.50%
S-1, S-3, S-3	11			25.00%	25.00%
S-1, S-3, S-4	12			50.00%	50.00%
S-1, S-4, S-1	13			12.50%	12.50%
S-1, S-4, S-2	14			12.50%	12.50%
S-1, S-4, S-3	15			25.00%	25.00%
S-1, S-4, S-4	16			50.00%	50.00%
S-2, S-1, S-1	17				12.50%
S-2, S-1, S-2	18				12.50%
S-2, S-1, S-3	19				25.00%
S-2, S-1, S-4	20				50.00%
S-2, S-2, S-1	21				12.50%
S-2, S-2, S-2	22				12.50%
S-2, S-2, S-3	23				25.00%
S-2, S-2, S-4	24				50.00%
S-2, S-3, S-1	25				12.50%
S-2, S-3, S-2	26				12.50%
S-2, S-3, S-3	27				25.00%
S-2, S-3, S-4	28				50.00%
S-2, S-4, S-1	29				12.50%
S-2, S-4, S-2	30				12.50%
S-2, S-4, S-3	31				25.00%
S-2, S-4, S-4	32				50.00%
S-3, S-1, S-1	33				12.50%
S-3, S-1, S-2	34				12.50%
S-3, S-1, S-3	35				25.00%
S-3, S-1, S-4	36				50.00%
S-3, S-2, S-1	37				12.50%
S-3, S-2, S-2	38				12.50%
S-3, S-2, S-3	39				25.00%
S-3, S-2, S-4	40				50.00%
S-3, S-3, S-1	41				12.50%
S-3, S-3, S-2	42				12.50%
S-3, S-3, S-3	43				25.00%
S-3, S-3, S-4	44				50.00%
S-3, S-4, S-1	45				12.50%
S-3, S-4, S-2	46				12.50%
S-3, S-4, S-3	47				25.00%
S-3, S-4, S-4	48				50.00%
S-4, S-1, S-1	49				12.50%
S-4, S-1, S-2	50				12.50%
S-4, S-1, S-3	51				25.00%
S-4, S-1, S-4	52				50.00%
S-4, S-2, S-1	53				12.50%
S-4, S-2, S-2	54				12.50%
S-4, S-2, S-3	55				25.00%
S-4, S-2, S-4	56				50.00%
S-4, S-3, S-1	57				12.50%
S-4, S-3, S-2	58				12.50%
S-4, S-3, S-3	59				25.00%
S-4, S-3, S-4	60				50.00%
S-4, S-4, S-1	61				12.50%
S-4, S-4, S-2	62				12.50%
S-4, S-4, S-3	63				25.00%
S-4, S-4, S-4	64				50.00%

The discount factors that are relevant are the 1 period zero coupon bond prices as of time T=2 for maturity at T=3 for each of the 16 nodes or states, $P(2,3,s_t)$. Those discount factors are shown in column x of Figure 5, which we derived in Chapter 9:

Figure 5

	Three Year Zero Coupon Bond Price		
	Current Time		
Row Number	0	1	2
1	0.9619	0.9425	0.9557
2		0.9816	0.9861
3		0.9692	0.9746
4		0.9640	0.9531
5			0.9780
6			0.9941
7			0.9919
8			0.9877
9			0.9811
10			0.9860
11			0.9915
12			0.9721
13			0.9784
14			0.9833
15			0.9887
16			0.9694

If the bonds are called at time T=3, the time T=3 value is the payoff of the bonds ($100) plus the time T=3 interest payment ($3), for a total of $103. If the bonds are NOT called at time T=3, the time T=3 value is the sum of the $3 interest payment made at time T=3 plus the time T=3 risk-neutral value of the final payment of $103. We gave these risk neutral values as of time T=3 in Figure 2.

For each of the 16 nodes at time T=2, we now calculate the risk-neutral value of the "no call" strategy at time T=2. This analysis is identical to the valuation of a 1 period zero coupon bond that could have alternative pay-offs, depending on the state. Figure 6 shows the call/no call analysis for each of the 16 nodes. If the net present value of not calling at time T=2 is greater than 100, the bonds will be called. If the net present value of not calling at time T=2 is less than or equal to 100, the bonds will not be called:

Figure 6

Time 2 Node	Time 2 Value	Action?	Discount Factor	Time 3 Node	Probability	Time 3 Value
1	95.0720	No Call	0.9557	1	0.1250	100.0685
			0.9557	2	0.1250	100.5686
			0.9557	3	0.2500	100.9357
			0.9557	4	0.5000	98.3373
2	101.0477	Call	0.9861	5	0.1250	102.8991
			0.9861	6	0.1250	103.0000
			0.9861	7	0.2500	103.0000
			0.9861	8	0.5000	101.9787
3	98.9261	No Call	0.9746	9	0.1250	100.8892
			0.9746	10	0.1250	103.0000
			0.9746	11	0.2500	102.8258
			0.9746	12	0.5000	100.6288
4	95.4169	No Call	0.9531	13	0.1250	100.6996
			0.9531	14	0.1250	101.2030
			0.9531	15	0.2500	101.5724
			0.9531	16	0.5000	98.9571
5	99.9293	No Call	0.9780	17	0.1250	101.8942
			0.9780	18	0.1250	103.0000
			0.9780	19	0.2500	103.0000
			0.9780	20	0.5000	101.6311
6	102.3874	Call	0.9941	21	0.1250	103.0000
			0.9941	22	0.1250	103.0000
			0.9941	23	0.2500	103.0000
			0.9941	24	0.5000	103.0000
7	102.1703	Call	0.9919	25	0.1250	103.0000
			0.9919	26	0.1250	103.0000
			0.9919	27	0.2500	103.0000
			0.9919	28	0.5000	103.0000
8	101.7325	Call	0.9877	29	0.1250	103.0000
			0.9877	30	0.1250	103.0000
			0.9877	31	0.2500	103.0000
			0.9877	32	0.5000	103.0000
9	100.9479	Call	0.9811	33	0.1250	103.0000
			0.9811	34	0.1250	103.0000
			0.9811	35	0.2500	103.0000
			0.9811	36	0.5000	102.7757
10	101.4524	Call	0.9860	37	0.1250	103.0000
			0.9860	38	0.1250	103.0000
			0.9860	39	0.2500	103.0000
			0.9860	40	0.5000	102.7757
11	101.9623	Call	0.9915	41	0.1250	102.3869
			0.9915	42	0.1250	103.0000
			0.9915	43	0.2500	103.0000
			0.9915	44	0.5000	102.8297
12	99.1233	No Call	0.9721	45	0.1250	101.5580
			0.9721	46	0.1250	103.0000
			0.9721	47	0.2500	103.0000
			0.9721	48	0.5000	101.2958
13	100.3484	Call	0.9784	49	0.1250	102.5157
			0.9784	50	0.1250	103.0000
			0.9784	51	0.2500	103.0000
			0.9784	52	0.5000	102.2509
14	101.0227	Call	0.9833	53	0.1250	103.0000
			0.9833	54	0.1250	103.0000
			0.9833	55	0.2500	103.0000
			0.9833	56	0.5000	102.4813
15	101.5794	Call	0.9887	57	0.1250	103.0000
			0.9887	58	0.1250	103.0000
			0.9887	59	0.2500	103.0000
			0.9887	60	0.5000	102.4813
16	98.7374	No Call	0.9694	61	0.1250	101.5589
			0.9694	62	0.1250	103.0000
			0.9694	63	0.2500	103.0000
			0.9694	64	0.5000	101.0738

Take the first node at time T=2 in Figure 6. Using the evolution of states from Chapter 9, the first node is a "Shift 1" state to get to time T=1 followed by another "Shift 1" to get to time T=2. Each of the first four rows in the columns labeled "probability" and "Time 3 value" represent the shift values and Time 3 values (derived as above) for Shifts 1, 2, 3, and 4 to reach time T=3. We multiply the probability of each Time 3 value by the probability of the state and then discount by the appropriate risk neutral discount factor (in the state of node 1 at time T=2, this discount factor is 0.9557) to get the time 2 value of 95.0720. Since this time 2 value is 100 or less, the decision at node 1 for time T=2 is "no call." The same analysis is repeated for all 16 nodes.

We now revise the cash flow table to recognize rational call behavior at either time T=2 or time T=3. If the bonds are called at node k at time T=2, all states which stem from node k at times T=3 and T=4 will have cash flows of zero. The cash flow at node k at time T=3 will have a cash flow equal to the sum of the \$3 interest plus the repayment of the \$100 principal, for a total of \$103. The revised cash flow table is shown in Figure 7. When these cash flows are multiplied times the appropriate probability-weighted discount factors, we get a risk-neutral value for this "4 year non-call 2" bond equal to 103.5473.

Figure 7

	Maturity of Cash Flow Received				
	Current Time				
Row	0	1	2	3	4
1	0.0000	3.0000	3.0000	3.0000	103.0000
2		3.0000	103.0000	3.0000	103.0000
3		3.0000	3.0000	3.0000	103.0000
4		3.0000	3.0000	3.0000	103.0000
5			3.0000	0.0000	0.0000
6			103.0000	0.0000	0.0000
7			103.0000	0.0000	0.0000
8			103.0000	0.0000	0.0000
9			103.0000	3.0000	103.0000
10			103.0000	103.0000	0.0000
11			103.0000	3.0000	103.0000
12			3.0000	3.0000	103.0000
13			103.0000	3.0000	103.0000
14			103.0000	3.0000	103.0000
15			103.0000	3.0000	103.0000
16			3.0000	3.0000	103.0000
17				3.0000	103.0000
18				103.0000	0.0000
19				103.0000	0.0000
20				3.0000	103.0000
21				0.0000	0.0000
22				0.0000	0.0000
23				0.0000	0.0000
24				0.0000	0.0000
25				0.0000	0.0000
26				0.0000	0.0000
27				0.0000	0.0000
28				0.0000	0.0000
29				0.0000	0.0000
30				0.0000	0.0000
31				0.0000	0.0000
32				0.0000	0.0000
33				0.0000	0.0000
34				0.0000	0.0000
35				0.0000	0.0000
36				0.0000	0.0000
37				0.0000	0.0000
38				0.0000	0.0000
39				0.0000	0.0000
40				0.0000	0.0000
41				0.0000	0.0000
42				0.0000	0.0000
43				0.0000	0.0000
44				0.0000	0.0000
45				3.0000	103.0000
46				103.0000	0.0000
47				103.0000	0.0000
48				3.0000	103.0000
49				0.0000	0.0000
50				0.0000	0.0000
51				0.0000	0.0000
52				0.0000	0.0000
53				0.0000	0.0000
54				0.0000	0.0000
55				0.0000	0.0000
56				0.0000	0.0000
57				0.0000	0.0000
58				0.0000	0.0000
59				0.0000	0.0000
60				0.0000	0.0000
61				3.0000	103.0000
62				103.0000	0.0000
63				103.0000	0.0000
64				3.0000	103.0000
Risk Neutral Value =					103.5473

We now step back one more time step to determine our optimal call/no call strategy from the perspective of the four nodes at time T=1. The analysis is identical to the exercise that we went through for the 16 nodes at time T=2. The probabilities for Shifts 1, 2, 3 and 4 are unchanged. The discount factors for a one period zero coupon bond as of time T=1 to mature at time T=2 depends on the state, $P(1,2,s_t)$. The values for these 1 period zero coupon bonds were given in Chapter 9. We reproduce them here in Figure 8:

Figure 8

	Two Year Zero Coupon Bond Price		
	Current Time		
Row Number	0	1	2
1	0.9841	0.9788	1
2		0.9931	1
3		0.9891	1
4		0.9866	1

We now calculate the probability weighted Time 1 values of the possible outcomes of the time T=2 nodes. Those calculations parallel the calculation of the Time 2 values. If the time 1 risk-neutral value of "no call" is greater than 100, we call the bonds. If the risk-neutral value of "no call" is 100 or less, the decision is "no call." Figure 9 shows the call/no call strategy at time T=1:

Figure 9

Time 1 Node	Time 1 Value	Action?	Discount Factor	Time 2 Node	Probability	Time 2 Value
1	97.7083	No Call	0.9788	1	0.1250	98.0720
			0.9788	2	0.1250	103.0000
			0.9788	3	0.2500	101.9261
			0.9788	4	0.5000	98.4169
2	102.2843	Call	0.9931	5	0.1250	102.9293
			0.9931	6	0.1250	103.0000
			0.9931	7	0.2500	103.0000
			0.9931	8	0.5000	103.0000
3	101.4485	Call	0.9891	9	0.1250	103.0000
			0.9891	10	0.1250	103.0000
			0.9891	11	0.2500	103.0000
			0.9891	12	0.5000	102.1233
4	100.9941	Call	0.9866	13	0.1250	103.0000
			0.9866	14	0.1250	103.0000
			0.9866	15	0.2500	103.0000
			0.9866	16	0.5000	101.7374

The results show that the bonds are called in 3 out of the 4 nodes at time T=1. We now revise our "call adjusted" cash flow table so that there are no cash flows on any branches of the bushy tree that follow a "call" decision at time T=1. The results of this adjustment are shown in Figure 10. The risk-neutral valuation for this "4 year non-call 1" bond is 102.4060.

Figure 10

	Maturity of Cash Flow Received				
	Current Time				
Row	0	1	2	3	4
1	0.0000	3.0000	3.0000	3.0000	103.0000
2		103.0000	103.0000	3.0000	103.0000
3		103.0000	3.0000	3.0000	103.0000
4		103.0000	3.0000	3.0000	103.0000
5			0.0000	0.0000	0.0000
6			0.0000	0.0000	0.0000
7			0.0000	0.0000	0.0000
8			0.0000	0.0000	0.0000
9			0.0000	3.0000	103.0000
10			0.0000	103.0000	0.0000
11			0.0000	3.0000	103.0000
12			0.0000	3.0000	103.0000
13			0.0000	3.0000	103.0000
14			0.0000	3.0000	103.0000
15			0.0000	3.0000	103.0000
16			0.0000	3.0000	103.0000
17				0.0000	0.0000
18				0.0000	0.0000
19				0.0000	0.0000
20				0.0000	0.0000
21				0.0000	0.0000
22				0.0000	0.0000
23				0.0000	0.0000
24				0.0000	0.0000
25				0.0000	0.0000
26				0.0000	0.0000
27				0.0000	0.0000
28				0.0000	0.0000
29				0.0000	0.0000
30				0.0000	0.0000
31				0.0000	0.0000
32				0.0000	0.0000
33				0.0000	0.0000
34				0.0000	0.0000
35				0.0000	0.0000
36				0.0000	0.0000
37				0.0000	0.0000
38				0.0000	0.0000
39				0.0000	0.0000
40				0.0000	0.0000
41				0.0000	0.0000
42				0.0000	0.0000
43				0.0000	0.0000
44				0.0000	0.0000
45				0.0000	0.0000
46				0.0000	0.0000
47				0.0000	0.0000
48				0.0000	0.0000
49				0.0000	0.0000
50				0.0000	0.0000
51				0.0000	0.0000
52				0.0000	0.0000
53				0.0000	0.0000
54				0.0000	0.0000
55				0.0000	0.0000
56				0.0000	0.0000
57				0.0000	0.0000
58				0.0000	0.0000
59				0.0000	0.0000
60				0.0000	0.0000
61				0.0000	0.0000
62				0.0000	0.0000
63				0.0000	0.0000
64				0.0000	0.0000
Risk Neutral Value =					102.4060

Now, what happens if we step back one more period and allow the bond to be called at time 0 (which would be the case if the bond had been issued some periods before)? Of course, the answer is obvious; since the risk-neutral value if not called is 102.4060, it's cheaper to pay the bonds off now at 100. Using the format of our prior analysis, we replicate this pricing and reach the "call" conclusion:

Figure 11

Time 0 Node	Time 0 Value	Action?	Discount Factor	Time 1 Node	Probability	Time 1 Value
1	102.4060	Call	0.9970	1	0.1250	100.7083
			0.9970	2	0.1250	103.0000
			0.9970	3	0.2500	103.0000
			0.9970	4	0.5000	103.0000

The cash flow table for this strategy is not very interesting. At time zero, the bonds pay 100 and all other cash flows are zero. Risk-neutral value is of course 100. While this may not be visually interesting, it's analytically interesting because bonds get called all the time and we have just demonstrated the analysis one would go through to make an accurate decision in this regard using the 3 factor Heath, Jarrow and Morton model of Chapter 9.

What is the Par Coupon on a Callable Bond?

Given that a 3% callable bond is "in the money" on the option to call the bond, one might ask this question: what coupon level produces a mark to market value of par on the assumption that the bonds are a new issue and will not be called at time zero? The answer can be obtained by using the "solver" function in common spreadsheet software to iterate the coupon and option-adjusted valuation in the Heath Jarrow and Morton framework. The option-adjusted par coupon is 1.8647%. The cash flows that result from the optimal call strategy are given in Figure 12. All nodes which show a cash flow equal to 101.8647 represent a "call" decision. All nodes which show a cash flow equal to 1.8647 represent a "no call" decision.

Figure 12

	Maturity of Cash Flow Received				
	Current Time				
Row	0	1	2	3	4
1	0.0000	1.8647	1.8647	1.8647	101.8647
2		101.8647	1.8647	1.8647	101.8647
3		1.8647	1.8647	1.8647	101.8647
4		1.8647	1.8647	1.8647	101.8647
5			0.0000	1.8647	101.8647
6			0.0000	1.8647	101.8647
7			0.0000	1.8647	101.8647
8			0.0000	1.8647	101.8647
9			1.8647	1.8647	101.8647
10			1.8647	1.8647	101.8647
11			101.8647	1.8647	101.8647
12			1.8647	1.8647	101.8647
13			1.8647	1.8647	101.8647
14			1.8647	1.8647	101.8647
15			1.8647	1.8647	101.8647
16			1.8647	1.8647	101.8647
17				0.0000	0.0000
18				0.0000	0.0000
19				0.0000	0.0000
20				0.0000	0.0000
21				0.0000	0.0000
22				0.0000	0.0000
23				0.0000	0.0000
24				0.0000	0.0000
25				0.0000	0.0000
26				0.0000	0.0000
27				0.0000	0.0000
28				0.0000	0.0000
29				0.0000	0.0000
30				0.0000	0.0000
31				0.0000	0.0000
32				0.0000	0.0000
33				1.8647	101.8647
34				101.8647	0.0000
35				101.8647	0.0000
36				1.8647	101.8647
37				1.8647	101.8647
38				101.8647	0.0000
39				101.8647	0.0000
40				1.8647	101.8647
41				0.0000	0.0000
42				0.0000	0.0000
43				0.0000	0.0000
44				0.0000	0.0000
45				1.8647	101.8647
46				101.8647	0.0000
47				1.8647	101.8647
48				1.8647	101.8647
49				1.8647	101.8647
50				101.8647	0.0000
51				101.8647	0.0000
52				1.8647	101.8647
53				1.8647	101.8647
54				1.8647	101.8647
55				101.8647	0.0000
56				1.8647	101.8647
57				1.8647	101.8647
58				1.8647	101.8647
59				101.8647	0.0000
60				1.8647	101.8647
61				1.8647	101.8647
62				1.8647	101.8647
63				1.8647	101.8647
64				1.8647	101.8647
Risk Neutral Value =					100.0000

This calculation has an enormous set of applications above and beyond the pricing of new issue bonds or swaptions at "fair value." It is used for transfer pricing and performance measurement for assets or liabilities which themselves have an embedded call option. All of the real world instruments listed in the introduction to this chapter require a similar kind of American fixed income options analysis. Because of the power of the Heath Jarrow and Morton framework, our probability-weighted discount factors do not change just because the instrument involves an American fixed income option. Only the cash flows used to populate the cash flow table change.

An Example of Valuation of a Rationally Pre-paid Amortizing Loan

By transaction count, a huge proportion of the assets held by major financial institutions around the world that involve an American fixed income option are consumer loans of some sort. Home mortgage loans, second trust deed/home equity loans, and auto loans commonly have an amortizing structure. We demonstrate the slight differences in the cash flow tables that result from an amortizing structure in the next two chapters. The process of valuation, however, is nearly identical to the one used in this chapter. We now turn to alternative techniques for valuation American fixed income options.

Monte Carlo Simulation

Monte Carlo simulation is justifiably popular in financial markets and in this volume, but it must be used with great care. Adams and van Deventer [1993] have discussed the reasons for such caution in detail. In general, if highly accurate closed form/analytical solutions exist, Monte Carlo simulation is not the fastest way to get an accurate answer. In the case of American fixed income options, where an accurate analytical solution is not available for a large number of interest rate risk factors, Monte Carlo simulation has to be used with care. We explain why in this section.

In general, however, Monte carlo simulation's limitations can be summarized as follows:

- As Hull notes, "one limitation of the monte carlo simulation approach is that it can be used only for European-style derivative securities."[1] To correctly value an American option, one measures value (as we did above) by working backwards from maturity to calculate value, assuming at each decision point that the option holder does the rational thing. Monte Carlo works by projecting one interest rate "path" at a time, so there is not enough information at any point on that path to correctly analyze whether or not an American option holder should prepay. For this reason, monte carlo simulation almost always requires the user to specify a decision rule regarding what the holder of the option should do in any given interest rate scenario, often in the form of a prepayment table or prepayment function in the case of a prepayment option. We analyze the implications of these approaches, which are very common, in the next two chapters. The use of a prepayment table or function is putting the cart before the horse, since the user has to guess how the option will be exercised before the user knows what the option is worth. Rather than going through this error-filled exercise, we think most users would get better results by just guessing the value of the option

[1] See Hull [1993], page 334.

directly as we have done so far in this chapter in the Heath Jarrow and Morton framework!

- o The calculation speed of Monte Carlo techniques is slow for problems where there is a small number of random variables. One example would be the case of a single factor term structure model with a few time steps. A bushy tree or Hull and White lattice would normally give an accurate and fast answer in this (MODEL RISK ALERT) unrealistic circumstance. Monte Carlo does have a speed advantage for problems with a large number of random variables, and it is this case which is typical for a complete analysis of integrated interest rate risk and credit risk. In that circumstance, the Heath Jarrow and Morton framework needs to be used to generate the Monte Carlo scenarios.
- o Monte carlo simulation by definition does not use all possible scenarios for valuation. There are an infinite number of interest rate scenarios, and Monte Carlo simulation, due to its speed problems, inevitably requires the user to use too few "scenarios" in the interests of time. In Chapters 36 to 41, we devote some time to this issue of the number of scenarios necessary to achieve statistical significance of a risk measure like value at risk. Some horrible mistakes have been made by some large institutions that we discuss at length in later chapters.
- o As a result of using less than all possible scenarios, Monte Carlo simulations have sampling error, which results from "throwing the dice" too few times. Many users of Monte Carlo simulation are under the mistaken impression that the beautiful probability distribution that is displayed on their computer screen reflects the true uncertainty about the value of a security as measured by Monte Carlo. Nothing could be farther from the truth. If you call a major securities firm to get a bid on a mortgage security, you get a bid in the form of one number, not a probability distribution. The beautiful probability distribution reflects the inaccuracy or sampling error of the technique itself, and this kind of graph reflects a weakness of Monte Carlo, not a strength. This kind of sampling error can lead to very serious problems when calculating hedge amounts. In most of the circumstances we discuss in the following chapters, there is no realistic alternative to a Monte Carlo simulation and sampling error, like the warning label on prescription medications, is a risk that needs to be dealt with on the basis of complete information.
- o To reduce sampling error to a meaningful level, a sophisticated analyst can calculate the number of simulations required to an answer with sampling error small enough to allow decision-making to be based on the Monte Carlo results.[2]

[2] Sophisticated users of the enterprise risk management system constructed by the authors and described on www.kamakuraco.com routinely do simulations with hundreds of thousands or millions of scenarios.

- Sampling error becomes even more important when basing hedges on the results of Monte Carlo simulation. The authors feel that the primary purpose of the calculations in this book are to define action, rather than simply to describe the amount of risk an institution currently has. Knowing how much risk you have without knowing what to do about it is nearly useless from a credit risk, market risk, liquidity risk and an interest rate risk perspective. Consider a hedger who values his portfolio using only 200 simulations via Monte Carlo (one of the 3 largest banks in North America has used fewer scenarios than this as recently as 2010 for enterprise wide risk management). The result shows a value of 100 and a sampling error standard deviation of, say, 2. This means that there is roughly a 65 percent probability that the true value lies between 98 and 102. In order to determine the proper amount of the hedge, the analysts shifts rates up by 10 basis points and repeats the analysis, getting a value of 99 and a sampling error standard deviation of 2 again. The analyst concludes that the "delta" of his portfolio is 100-99=1 and wants to base his hedge on this result. This is fine, as far as it goes, but the delta has sampling error also. The sampling error on the delta is a function of the sampling error on the two simulation runs, and it is calculated as follows:

$$\sigma_{hedge\ \Delta} = \sqrt{\sigma^2_{run\ 1} + \sigma^2_{run\ 2}}$$

  In the example given, the sampling error of the hedge delta of 1.00 works out to 2.828. What does it mean for the precision of the hedge? It means that there is a 36.2 percent chance that the hedge amount is not only the wrong magnitude *but the wrong SIGN*! This is a career-ending type of error that has happened often enough on Wall Street that it's become a familiar story.[3]

- The "delta" from a Monte Carlo simulation can only be derived from doing the calculation twice (or preferably three times for securities with high convexity), further aggravating its speed disadvantages. For simulations with a large number of macro factors driving total risk, there is a delta for each risk factor as well.

- Monte Carlo simulations must be done on the basis of the "risk neutral" distributions of all random variables, an adjustment which many users fail to make, for accurate valuation.

On the plus side, monte carlo simulation has a number of advantages which should not be overlooked:

[3] A major New York bank reported a loss of more than $100 million after discovering the Monte Carlo simulation routine it was using to value its portfolio was producing a "gamma" (or second derivative) with the wrong sign; the loss was the magnitude of the mark-to-market error discovered after using a more sophisticated technique to value the same portfolio.

- It is sometimes the only alternative where the cash flow is path dependent, as imprecise as it might be in that case. A bushy tree can also handle path dependence accurately and should not be overlooked as an alternative. Longstaff and Schwartz ("Valuing American Options by Simulation: A Simple Least-Squares Approach," Review of Financial Studies, 2001, Volume 14, No. 1, pp. 113-147) make a creative suggestion on how to enhance the accuracy of Monte Carlo simulation when a bushy tree approach is not available.

- Monte Carlo has speed advantages for problems with a large number of variables. Generally, this means three or more variables. Recent regulatory stress tests employ 20-40 macro-economic variables so Monte Carlo is almost an automatic choice for stress testing. This is typically the case for the integrated interest rate risk, market risk, liquidity risk and credit risk analysis that is the central focus of this book.

- Monte Carlo is relatively simple to implement, although a shockingly small number of enterprise risk management software firms have both built such a system and successfully commercialized it.

What are our conclusions about Monte Carlo simulation? First of all, it is a tool that all users should have access to. Risk management software packages which don't clearly display the sampling error of both value and hedges derived from Monte Carlo simulation calculations should be used with extreme caution. With this caveat, Monte Carlo simulation is essential to the complete integrated analysis of credit risk, market risk, liquidity risk and interest rate risk. Finally, Monte Carlo needs to be supplemented with the other techniques in this chapter in order to be used with efficiency and accuracy. For example, in a multiperiod Monte Carlo simulation that requires valuation at each point in time, closed form solutions can be used for this valuation when they are available, and the other techniques discussed below can be used when they are not available. The authors believe this "hybrid monte carlo" approach is firmly established as the best practice approach for integrated risk management, subject to the caveats above.

Finite difference methods

Finite difference methods are more complex but more general solution methods commonly used in engineering applications. Finite difference methods provide a direct general solution of the partial differential equation (that defines the price of a callable security), unlike lattice methods, which model the evolution of the random variables. These methods fall into two main groups, *explicit* and *implicit* finite difference methods. The explicit finite difference methods are equivalent to the lattice methods. Implicit methods are more robust in the types of problems that they can handle, but implicit methods are computationally more difficult. Both methods usually use a grid-based calculation method, rather than the lattice approach, to arrive at numerical solutions.

Finite difference methods can be used to solve a wide range of derivative product valuation problems, and they are not restricted to a small number of stochastic processes (i.e. normal or lognormal) that describe how the random variable moves. These methods can be used to value both American and European style options. Like lattice methods, the valuation by finite difference methods is performed by stepping backwards through time. This is a one stage process under the finite difference method; there is no need to model the evolution of the random variable before the valuation can be performed.

The authors believe that the finite difference method, in its grid rather than a lattice implementation, is a very useful tool for most institutions but its use to date in the financial services industry has been on single instrument valuation in a derivatives trading context, rather than a tool that has been successfully embedded in an enterprise-wide risk management system.

Binomial Lattices

A binomial lattice (sometimes called a binomial tree) is a discrete time model for describing the movement of a random variable whose movements at each node on the tree can be reduced to an up or down movement with a known probability. Unlike the one factor bushy tree used in Chapter 6 in the Heath Jarrow and Morton context, the binomial lattice is usually constructed by constraining the nature of movements of the underlying (usually single) random variable. The model is usually specified so that an upward movement followed by a downward movement gives the same value as a downward movement followed by an upward movement; this means there will be three possible values of the random variable at the end of the second time interval and k + 1 possible values at the end of time interval k. This recombining nature of the lattice is what makes it different from the bushy trees used in Chapters 6 through 9. In the most realistic risk applications, constraining random variable movements to recombine reduces accuracy substantially.

Jarrow [1996] uses the binomial tree concept for modeling interest rate movements, and Jarrow and Turnbull [1996] show how the same kind of tree (with appropriate modifications) can be used for modeling stock price movements. The binomial lattice can be used to value American style options as well as European style options.

Trinomial Lattices

In a number of papers with important theoretical and practical implications, Hull and White [1990,1993, 1994] developed the trinomial lattice valuation technique which has quickly established itself as a popular valuation technique, not only for the Vasicek model and its extended version (the Hull and White model), but also for a number of other single factor term structure models which are Markov in nature. A trinomial lattice (sometimes called a trinomial tree) is a discrete time model for describing the movement of a random variable whose movements at each node on the tree can be reduced to one of three possibilities: an up movement, a down movement or no change. The model is specified so that an upward movement followed by a downward movement and a downward movement followed by an upward movement give the same value as two movements where no change occurs; this means there will be three possible values of the random variable at the end of the first time interval and k + 2 possible values at the end of time interval k. This recombining feature maximizes the efficiency of the calculation.

The additional outcome at each node provides an additional degree of freedom which enhances the power of the model; trinomial lattices can be used to model almost any Markov stochastic process, including ones in which the parameters are functions of time and the random variable itself. In modeling interest rate-related derivatives, it allows rates to be modeled so that the current term structure of interest rates and term structure of rate volatility are matched exactly by the modeling process. This offsets the more complicated calculations of the values at the ends of each branch. The trinomial lattice can be used to value American as well as European style options.

In recent years, the Hull and White lattice approach has been overshadowed by the more general Heath Jarrow and Morton approach, with deals with a large number of interest rate risk factors much more easily as we have demonstrated in Chapters 6 through 9.

Heath Jarrow and Morton Valuation of American Fixed Income Options When Default Risk is Present

How does the possibility of default affect the call option on an ABC Company 10 year bond that is callable at any time after five years? The credit model and the option are not separable because they impact each other in the following ways:

- The value of the call option is reduced by the possibility that ABC company defaults before it is able to exercise the call option
- The value of the call option is reduced by the possibility that ABC's default probability is so high when exercising the call is rational that it cannot raise the cash to retire the bonds by calling them
- The potential losses from default are reduced by the possibility that the bonds have been called before default can occur.

This complex interaction is typical of multi-payment callable instruments that range from home mortgages to auto loans to bonds issued by corporations and governmental agencies. How is the HJM bushy tree or Monte Carlo modified to handle this interaction, particularly in the key case (our constant focus) where interest rates and other macro-economic factors are drivers of default probability levels? We discussed the appropriate adjustments in Chapters 17, 19 and 20 but we review them again here.
Jarrow and Turnbull [1995] describe this process in detail. At each node on the bushy tree, the default probability of the counterparty will take on a different value because the default probabilities vary by time and level of interest rates. Also at each node on the lattice, the company can either default or not default. If the company is in default, the payoff on the security is its defaulted value (in the Jarrow-Turnbull [1995] case, this is principal times recovery rate; in the Jarrow [1999, 2001] and Duffie-Singleton [1999] models, this is the recovery rate times the value of the security just an instant before default).

Jarrow [1999] goes on to show in detail how to construct lattice and bushy tree valuation techniques for the Jarrow model described in Chapter 16 where the default intensity is driven by both interest rates and other macro-economic factors.

In this chapter, we have assumed that the American option is exercised "rationally," with no considerations other than the obvious financial considerations driving exercise. We now turn to "irrational" exercise of American options and discuss the implications for integrated enterprise risk management.