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

Advanced Financial Risk Management:

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

April 2012

Donald R. van Deventer
Kenji Imai
Mark Mesler

Part 5: Portfolio Strategy and Risk Management

Chapter 36:
Value at risk and risk management objectives revisited at the portfolio and company level

In Chapter 1, we defined risk management in a practical way:

> ***Risk management is the discipline that clearly shows management the risks and returns of every major strategic decision at both the institutional level and the transaction level. Moreover, the risk management discipline shows how to change strategy in order to bring the risk return trade-off into line with the best long and short-term interests of the institution.***

In Chapter 1, we noted that this definition of risk management includes within it the overlapping and inseparable sub-disciplines such as:

- Credit risk
- Market risk
- Asset and liability management
- Liquidity risk
- Capital allocation
- Regulatory capital calculations
- Operational risk
- Performance measurement
- Transfer pricing

and many other sub-disciplines. Our primary focus in this book is to show how to execute the practice of risk management in a way that is fully integrated and makes no distinction between these sub disciplines. We do that in this chapter.

The Jarrow-Merton Put Option as a Measure of Total Risk: An Example

Value at risk as a risk measure has dominated the last 20 years of risk analytics, especially for traded instruments. The authors grant that this focus was well-intended. Alas, as the large JPMorgan credit trading losses of 2012 and the listing of large institution losses in the Introduction make clear, the VAR concept hasn't worked well in preventing losses and failures at the institutions that have relied on it. In the rest of this chapter, we draw a clear distinction between VAR and best practice, state of the art risk management.

Unlike VAR, the Jarrow-Merton put option as a measure of total risk is an observable, tradable number. Figure 1 shows put option prices on the common stock of Citigroup prevailing on July 25, 2012.

Figure 1

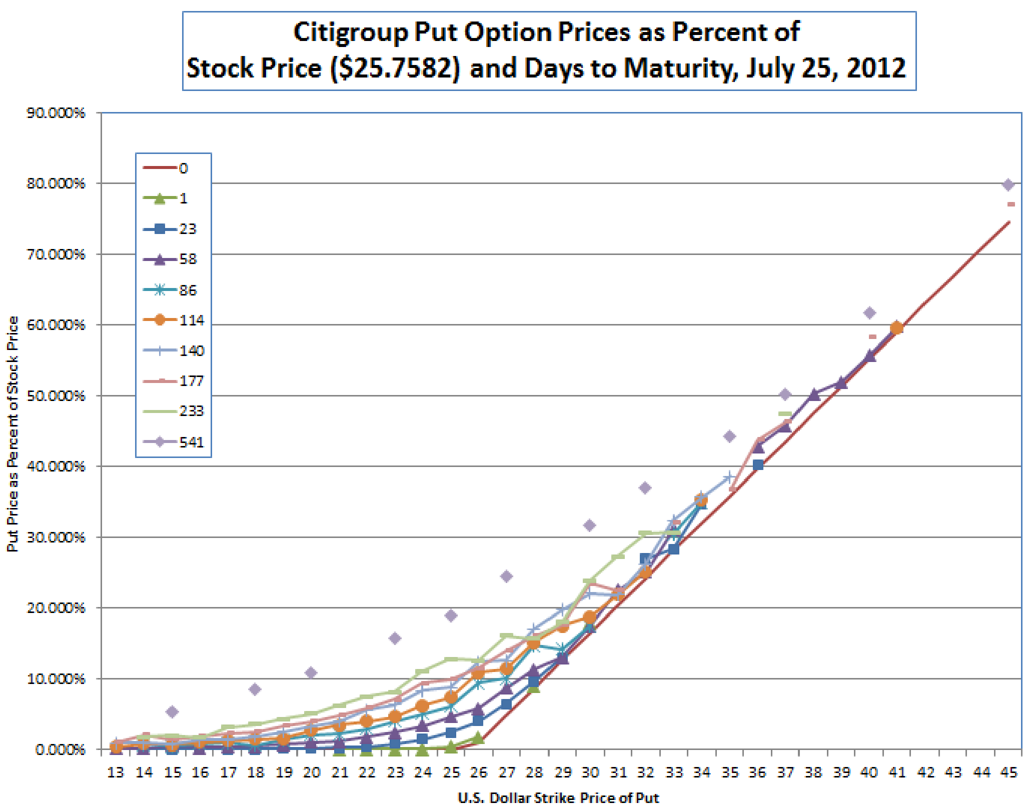

This data is available, often in real time, on websites accessible even to retail investors, let alone CEOs of major financial institutions. The chart makes a number of things much more clear that the kind of value at risk calculations outlined by Jorion's [third edition, 2006] classic overview of the subject:

- All assets and liabilities contribute to total risk, not just traded assets
- Total risk is what matters, not just a narrow definition of market risk
- Risk has a term structure and the price of protection against total risk changes as the time horizon expands beyond the myopic 10 day to 30 day focus of most VAR analysis

The Jarrow-Merton put option allows a risk analyst to definitively answer practice questions on capital adequacy. Figure 2, the data behind the graph in Figure 1, gives us the ability to deal with a number of questions very specifically:

Figure 2

Citigroup
Put Option (Last) Price as Percent of Stock Price
July 25, 2012
Source: Yahoo Finance and Kamakura Corporation
Stock Price 25.7582

	Days to Maturity of Option Contract									
	In the Money Amount									
Strike Price	0	1	23	58	86	114	140	177	233	541
13	0.000%		0.078%	0.194%		0.427%	0.893%	0.971%		
14	0.000%		0.078%	0.194%	0.660%	0.776%	0.893%	2.096%	1.708%	
15	0.000%		0.039%	0.621%	0.505%	0.699%	0.854%	1.475%	1.980%	5.358%
16	0.000%		0.078%	0.272%	0.776%	0.971%	1.359%	1.863%	1.631%	
17	0.000%		0.078%	0.388%	0.893%	1.281%	1.320%	2.291%	3.183%	
18	0.000%		0.078%	0.466%	0.621%	1.398%	1.902%	2.485%	3.572%	8.619%
19	0.000%		0.116%	0.660%	1.320%	1.592%	2.485%	3.339%	4.309%	
20	0.000%		0.155%	0.971%	2.058%	2.679%	3.222%	3.999%	5.047%	10.948%
21	0.000%	0.039%	0.233%	1.126%	2.213%	3.416%	3.999%	4.814%	6.250%	
22	0.000%	0.039%	0.427%	1.747%	2.912%	3.999%	5.590%	5.862%	7.532%	
23	0.000%	0.039%	0.776%	2.446%	3.999%	4.620%	6.289%	7.105%	8.153%	15.723%
24	0.000%	0.078%	1.398%	3.339%	5.047%	6.212%	8.386%	9.434%	11.103%	
25	0.000%	0.349%	2.407%	4.620%	6.095%	7.454%	8.813%	9.977%	12.811%	19.023%
26	0.939%	1.747%	4.038%	5.823%	9.473%	10.870%	12.423%	11.530%	12.617%	
27	4.821%		6.483%	8.696%	10.094%	11.491%	12.617%	13.976%	16.111%	24.458%
28	8.703%	8.852%	9.589%	11.336%	14.753%	15.141%	17.082%	16.111%	15.723%	
29	12.586%		13.006%	13.006%	14.170%	17.470%	19.800%	17.664%	18.053%	
30	16.468%	17.858%		17.470%	17.470%	18.829%	22.129%	23.488%	23.876%	31.835%
31	20.350%			22.517%		21.935%	21.935%	22.517%	27.331%	
32	24.232%		26.982%	25.041%		25.235%	26.205%		30.631%	37.076%
33	28.115%		28.340%	31.058%	30.476%		32.417%	32.223%	30.670%	
34	31.997%		34.746%		34.940%	35.329%	35.717%			
35	35.879%						38.628%	36.843%		44.258%
36	39.761%		40.181%	42.899%				43.870%		
37	43.644%			45.811%				46.393%	47.596%	50.314%
38	47.526%			50.314%						
39	51.408%			52.061%						
40	55.290%			55.710%				58.428%		61.806%
41	59.173%			59.787%		59.593%				
42	63.055%									
43	66.937%									
44	70.819%									
45	74.702%							77.140%		79.858%

If dynamic new Citigroup Chairman of the Board Michael O'Neill asked this question: "How much additional capital does Citigroup need to ensure that our stock price is at least $25 in January 2013?" Figure 2 contains the answer. On July 25, Citigroup's stock was trading at $25.7582. A put option with a strike price of $25 and a maturity of January 18, 2013 (177 days from July 25, 2012) costs $2.57 per share or 9.977% of the current Citigroup stock price. What trade would Citigroup need to do to buy this insurance? We note that this trade can't be done in practice for legal reasons but it does give us the market price of total risk protection. The hypothetical trade is this, assuming no change in Citigroup stock price if a new equity issue is done:

1. Issue N new shares at $25.7582 per share, bringing the total number of shares from the current number of X to X+N
2. Buy put options with a total cost of $2.57 per share (9.977% of share price) on X+N shares for a total cost of $2.57 (X+N)
3. On January 18, 2013, if Citigroup's share price is below $25, exercise the puts

With the current number of shares outstanding at Citigroup at about 2.93 billion shares, puts on those shares would be \$2.57 (2.93 billion) = \$7.5 billion. We then solve for the number of new shares N such that the amount of the new issue is \$7.5 billion plus \$2.57N (since we need puts on the new shares too):

$$25.7582\,N = 2.57(2.93\ billion) + 2.57N$$

The number of shares that need to be issued to protect a stock price level of \$25 per share on January 18, 2013 is N=325 million. If Citigroup's total risk were lower, N would be a lot lower number and the put price of \$2.57 would be much lower.

Value at risk provides next to no information on this practical question.

A Four Question Pass Fail Test for Financial Institutions CEOs and Boards of Directors

Above and beyond the practical question posed above, the four most important risk management questions that a financial institutions CEO and Board should be able to answer are easy to summarize in a pass/fail test.

Question 1: What happens to the market capitalization and net income of the firm if any of these risk factors change: home prices, foreign exchange rates, commercial real estate prices, stock index levels, interest rates, commodity prices?

If the answer is "I don't know," the institution fails the test.

Question 2: Using an insider's knowledge of the assets and liabilities of the firm, both "on balance sheet" and "off balance sheet," what is the best estimate, monthly for the next ten years, of the probability that the firm will fail in each of these 120 monthly periods?

If the answer is "I don't know," the institution fails the test.

Question 3: Using only information available to an outsider, what is the best estimate of the probability of the failure of the firm in both the short run and the long run?

If the answer is "I don't know," the institution fails the test.

Question 4: If the firm is able to answer Questions 1, 2, and 3, what hedging position is necessary to insure that the macro factor sensitivity of the firm and default probability of the firm reach the target levels set by the Board of Directors?

If the answer is "I don't know," the institution fails the test.

Why Do These Four Questions Matter?

When one looks at the ex-post analysis of why institutions like Merrill Lynch, UBS and Citigroup got into such trouble during the current crisis, the feedback from the Board and CEO levels is amazingly consistent. We summarized key quotations in this regard in Chapter 1. In short, management admitted that they could not answer these questions and that is why the institutions failed. The purpose of this book is to provide a framework that allows an automated fully integrated highly accurate answer to these questions that is consistent with the observable Jarrow-Merton put prices outlined above.

But aren't there more than 4 questions one should be able to answer? Of course. Here are some other questions that financial executives using best practice risk management can answer and should answer.

An Alphabet of 26 Extra Credit Questions

A modern approach to risk management can provide answers to a longer list of additional questions than these four key questions posed above. To illustrate the power of a multi-period credit adjusted approach to risk management, we pose additional "extra credit" questions that a financial institution can answer with a high degree of accuracy using a best practice approach.

A. What is the probability that the firm will have to increase its reserve for loan losses by more than X per quarter, quarterly for the next 10 years?
B. Given that macro factors drive actual losses, what pattern for the allowance for loan losses gives the most stable net income for the institution?
C. To what extent does lending in commercial real estate provide diversification versus the bank's mortgage lending businesses?
D. What interest rate risk posture provides the most effective offset of credit losses given the links between interest rates and defaults?
E. Given cyclical credit losses, what timing of common stock offerings (a) meets capital needs through the cycle and (b) results in the highest average common stock price on new issues?
F. How do regulatory capital ratios like the Basel II risk measures vary over the business cycle and to what extent are they correlated with the firm's actual risk of default?
G. What 10 macro factors represent the highest threat to the safety and soundness of the institution, from highest to lowest?
H. What risk limits should the firm have in place with respect to these 10 macro factors?
I. What hedging instruments are available to mitigate the risk of these macro factors if their impact on the institution's default risk grows too high?
J. How should these macro factor risk exposures be displayed and reported in the institution's annual report and regulatory filings like the 10-k report of the U.S. Securities and Exchange Commission?
K. How should compensation systems for senior management be modified in light of the risk that the institution has from its macro factor exposure?
L. How should the institution replace external and internal ratings, which obscure macro factor links to default risk, with a more transparent "credit risk CAT scan" of risk?
M. Looking backwards in time, which macro factors had the largest historical impact on the firm's market capitalization and default risk?
N. How do macro factor movements affect the supply of liabilities (alleged "core" deposits, certificates of deposit, commercial paper, bonds, derivatives) to the firm?
O. What percentage decline can be expected in the supply of liabilities to the firm if there is an adverse movement of 10%, one by one, in the 10 most important macro factors affecting the institution?
P. What strategies can be put in place to mitigate this macro factor driven credit risk and liquidity risk?
Q. Which class of business counterparties has the greatest sensitivity to each of these macro factors, and how should pricing of credits to these counterparties reflect the risk?
R. Is the institution as diversified as it could be with respect to this macro factor risk? If not, how could diversification is improved?
S. How is operational risk linked to movements in the same top 10 macro factors?
T. To what extent do legacy silo risk systems and ratings systems obscure the answers to the questions above?

U. Can the systems infrastructure be improved so that an integrated risk assessment of macro factor sensitivity is both more accurate and more efficiently produced?
V. Can the management team and board be strengthened so there is an enhanced understanding of macro factor driven risk across the organization?
W. Has the firm worked closely with auditors to ensure that the auditors have a clear understanding of macro-factor driven risk?
X. Is the firm's self-assessment of its own risk level consistent with the pricing on put options on its common stock and credit default swaps?
Y. If the answer is no, what is the reason for the inconsistencies?
Z. If the market prices of puts and CDS are overly pessimistic, what investor education campaign is necessary to more effectively communicate the true risk of the institution?

A firm whose management team can answer this alphabet of questions truly represents best practice in risk management and shareholder value creation.

Is your value at risk from "Value at Risk"?

How does a firm's ability to answer these critical 4 + 26 risk management questions differ if value at risk is the primary risk measure? We summarize the well-known problems of value at risk in this section.

VAR has as many incarnations as there are risk managers. In the mid-1990s, J.P. Morgan made the decision to publicize the technique and to spin off an affiliate that would resell the VAR technology and related information. Hundreds of financial institutions adopted the VAR technology as a supplement for or replacement of traditional trading floor-based risk limits like "PVBP" (present value of a basis point move in interest rates) or "delta" type measures with respect to other risk factors, like stock prices and foreign exchange rates. Philippe Jorion [2006] has done a masterful job in his book at documenting various ways to implement the VAR concept. The purpose of this section is much simpler--to answer the question "Does it work?" Could a sophisticated user of the VAR concept have anticipated the 2007 to 2010 credit crisis and taken action early on to mitigate the impact of the crisis? The answer is...maybe. That's not a sufficiently strong endorsement to justify the considerable sums that many institutions have invested in legacy VAR technology. Why is the answer so lukewarm? We turn now to the reasons.

First of all, there are three principal variations of VAR technology, and they have different strengths and weaknesses:

A. Historical VAR, which estimates the nth percentile in potential losses based on historical returns on the relevant traded assets.

B. Variance-Covariance VAR, which takes the historical volatility and correlation of returns on traded assets to calculate the nth percentile in potential losses

C. Monte carlo VAR, which simulates the environment forward, rather than using historical data like A and B.

Most of the initial publicity around the VAR concept focused on a 10-day time horizon for the calculation, since it was advertised as an index of short term trading risk for assets traded in an active market. Most implementations assume that the portfolio of assets and liabilities that one is analyzing stays constant over the one period horizon (with length 10 days).

How well did these techniques serve their users in the credit crisis? On January 28, 2008, a story on bloomberg.com entitled "Death of VAR Evoked as Risk Taking Vim Meets Taleb's Black Swan" reported that Merrill Lynch's historical VAR measure was $92 million at a time that CDO losses had already reached $18 billion at Merrill. With all due respect to Nassim Taleb, this was no "black swan." The error in the VAR measure at Merrill was due to simplifying assumptions in the calculation that were simply not true, leading to incorrect risk assessment. With respect to historical VAR, here are the major problems that would lead one to underestimate risk by a factor of 200 like Merrill did:

Problem 1: Future returns almost never equal historical returns. Home prices dropped, mortgage default rates sky rocketed, and the historical returns (which are typically measured over a period of less than a year) did not adequately reflect the possibility of these events

Problem 2: Much of the losses in the 2007 to 2010 crisis have stemmed from assets that either always have been or are currently "non traded": subprime mortgage loans, Alt-A loans, and structured products using them as reference collateral. They were simply "left out' of the calculation

Problem 3: The balance sheet you end with is not the balance sheet you start with, even though that's the assumption made by most VAR users. A veteran risk manager at one of the largest life insurance firms in North America once asked, "This VAR concept doesn't make any sense. If we extend the time horizon to a year, a huge percentage of our assets would have matured by then. How could anyone assume a constant balance sheet for VAR?"

Most of these problems apply to variance-covariance VAR as well, with an additional monkey wrench thrown into the gears:

Problem 4: Returns on most securities are NOT normally distributed. If one assumed that the returns on Bear Stearns common stock were normally distributed at their historical mean and variance, one would have implied a probability of failure (a monthly return of -100%) of 0.000000. For CDO tranches (see Chapter 20) with narrow bands of losses, the outcome is almost exclusively "all or nothing," since it becomes highly unlikely for the tranche to suffer partial losses when the size of the tranche is a small percentage of total notional principal of the collateral underlying the CDO. The net impact of the normality assumption is a dramatic underestimation of risk.

When one turns to Monte Carlo VAR, still another problem can be important:

Problem 5: Returns in period N are not independent of the returns in period N-1, even though this assumption of independence is very, very common in financial theory and practice. Take the Case-Shiller home price index from www.sandp.com and test the hypothesis that home price returns are independent of returns on prior periods to see for yourself--the returns are very, very highly cyclical and dependent on past returns.

Can VAR be saved? The answer is "yes" if the analyst is very careful and has a powerful enterprise wide software package available. Here are the steps that one needs to take to get maximum accuracy from the VAR way of thinking:

Fix Number 1: Do the VAR calculation on a **multi-period basis** that allows for a dynamic balance sheet, with cash flow reinvested, new assets and liabilities generated, and embedded options exercised

Fix Number 2: **Allow for defaults to occur**, since they are the biggest source of "non-normality" in the real world. Kamakura Risk Manager uses default probabilities from the Kamakura Risk Information Services default service (or any other source) for this purpose

Fix Number 3: Recognize that the drivers of risk, typically **macro economic factors**, can be cyclical with probability distributions in period N that depend on prior periods.

Fix Number 4: Look at the **entire distribution of value**, not just one percentile level, at multiple points in time

Fix Number 5: **Stress test VAR and mark to market value with respect to all relevant macro factors** (like home prices) to avoid being lulled into complacency. This requires explicit links between the macro factors, default probabilities, and credit spreads.

With these fixes, a multi-period default-adjusted VAR calculation can add some insight that creates incremental risk-adjusted shareholder value added. We now turn to some additional important differences between VAR and the Jarrow-Merton put option concept.

VAR versus the Put Option for Capital Allocation

What are the real differences between using the VAR approach to capital allocation and use the put option approach? This section answers that question.

A worked example shows the differences very clearly. Let's assume that there are only 100 scenarios that can impact our firm ABC Company. In best practice net present value terms, in 93 scenarios there is no value destruction and capital is unchanged. In seven scenarios, numbered 94 to 100, we have these losses:

Scenario	Net Present Value of Losses
94	1
95	3
96	4
97	6
98	12
99	17
100	27

We can then easily determine that the average loss over all 100 scenarios is 0.70. The 99th percentile VAR is 17, and the 95th percentile VAR is 3. We need only $3 of capital to survive 95% of the time, but we need $17 to survive 99% of the time. This is simple enough. How does the put option approach compare?

Let's make it very clear that, when we are talking about a put option, we could equivalently call it an insurance policy that reimburses us for any value destruction in our current portfolio. Puts come at various strike prices and insurance policies come with deductibles and caps. The table below shows the payments that an insurance company would make to us in each of the loss scenarios for four policies: three policies where the losses are capped (at 5, 10 and 20) and one policy with no caps:

Scenario	Net Present Value of Losses	Losses Capped at 5	Losses Capped at 10	Losses Capped at 20	No Cap on Losses
94	1	1	1	1	1
95	3	3	3	3	3
96	4	4	4	4	4
97	6	5	6	6	6
98	12	5	10	12	12
99	17	5	10	17	17
100	27	5	10	20	27

How much would each of these put options or insurance policies cost? Since our losses are stated in net present value terms, in a competitive market with perfect competition the put/policy should be priced at its expected value over the 100 scenarios. For the policy capped at a maximum loss of 5, the cost is 0.28 (the expected losses of 1+ 3 + 4 + 5 + 5 + 5 + 5 = 28/100 trails). For the 10 dollar cap, the cost is 0.44. For the 20 loss cap, the cost is 0.63. For a policy that reimburses us for all losses, the cost of the put is exactly the expected loss on the portfolio: 0.70. How much capital would be needed if we were comfortable with insurance only up to 20 in losses? We issue additional capital of 0.63, and we use the 0.63 to buy the insurance policy.

Why are the VAR and Put Approaches So Different: Self Insurance versus Third Party Insurance

How do we explain the fact that the 99th percentile VAR capital needed is 17 and the put option with a cap at 20, where the company also loses money in only 1 of the 100 scenarios, costs only 0.63, not 17? The VAR capital calculation assumes that you are "self insuring" the company against losses. Even though 17 or more in losses happens in only 2 of the 100 scenarios, you issue 17 in capital, which turns out to be unnecessary 98 times out of 100.

The Jarrow-Merton put premium or insurance cost of protecting your portfolio of losses up to $20 reflects not only the amount of the losses, but also their probability of occurrence. The insurance provider can diversify, and a competitive market place will ensure that the benefits of diversification are passed along to the clients of the insurance provider. In a nutshell, the VAR approach assumes you are self insured and can't diversify. The put option approach assumes you buy protection from a third party who is diversified. It's all very simple.

As we noted above, when it comes to capital allocation using the put option concept, we don't even have to make the calculation for bank holding companies that are listed on stock exchanges in the United States. Put options on their common stock will be publically observable, and, for any "deductible" (strike price); there will be a different amount of capital required for that level of protection as we saw in Figures 1 and 2 for Citigroup.

Another approach to protect one's capital has tremendous moral hazard for your counterparty—you can buy credit default swaps on yourself. Rumor has it that AIG was "sophisticated" enough to be willing to do this…But what if Nassim Taleb is right and the CDS market is the equivalent of buying marine insurance from another passenger on the Titanic? The put/insurance policy can still be valued and purchased. We do a "credit risk cat scan" that looks through all the assets and liabilities of the firm and all of the counterparties to see how they are affected by key macro factors like interest rates, home prices, foreign exchange rates, and stock indices. This is exactly what we describe below. We then buy the appropriate insurance protection on these macro factors. Almost all of them have exchange traded futures and (in some cases) options, so we don't have a Titanic problem.

In short, the difference between a VAR approach to capital needs and a put option or insurance approach to capital needs is the difference between self-insurance or insurance in an efficient market that recognizes diversification. In most cases, the latter assumption will be a much more accurate indicator of how much capital protection is necessary.

As we explained above, a put option or insurance policy that insures the firm against the same loss profile is the probability weighted present value of all scenarios. Given identical loss projections, this normally means that the "true" VAR will be much higher than the put price (unless the losses are the same in all scenarios) because the put is probability weighted and the VAR is not. Then what about this quote from www.bloomberg.com on January 28, 2008?: "Merrill's highest one-day value at risk in the third quarter was $92 million, indicating that the firm's maximum expected cost during the 63-trading day period would be $5.8 billion. In fact, the firm wrote down $8.4 billion from the value of collateralized debt obligations, subprime mortgages and leveraged finance commitments, 45 percent more than the worst- case scenario." How could that happen?

What was reported in the Bloomberg story was "false" VAR, a gross underestimate of risk that has a host of problems that we listed above in the section "Is Your Value at Risk from Value-at-Risk?" The Bloomberg story goes on to state that most of the large U.S. securities firms were using "false" VAR based on 1-4 years of historical data. Given the history of U.S. home prices, a 4 year historical "false" VAR would have predicted minimal losses (as it did) because over the 4 year period ending December 31, 2007, home prices were strongly up. Even though home prices peaked in Los Angeles in September 2006, they did not begin to have a powerful impact on mortgage defaults until the second half of 2007. "False" VAR is based on an average of history. "True" VAR, which is what we assumed in our example above, is a complete and accurate listing of all possible future outcomes and their probabilities, Nobel laureate Prof. Kenneth Arrow's famous "state space." If we base "true" VAR on that complete listing of outcomes and probabilities, the put price or cost of insuring risks via a third party than can diversify will almost always be less than the high percentile (say 99th) "true" VAR number. That "true" VAR number is the amount of self-insurance you need to survive with 99 percent probability.

Wasn't the miss with respect to VAR just a "Black Swan" as Nassim Taleb argued? No, it was just bad math. It was an assumption that the world was flat. As we explained above, historical VAR contains an implicit assumption that neither Lehman nor Bear Stearns could fail, because (based on historical equity returns) the implied probability of -100% stock returns within a month for each of them was 0.000000%. Similarly, if the monthly changes in home prices have been strongly positive over the measurement period, the historical VAR approach will implicitly assume away the probability of a decline. The "true" VAR calculation and a rationally priced put will consider all possible future scenarios, not just those that have actually come about in the last 12 or 48 months.

Calculating the Jarrow-Merton Put Option Value and Answering the Key 4 + 26 Questions

'What is the hedge?' In other chapters in this book, we posed this question as the best single sentence test of risk management technology. For that reason, the Jarrow-Merton put option is a very powerful concept because the put option they propose, if properly structured, *is* the hedge. 'What is the equivalent of the Merton and Jarrow put option in the interest rate risk context?' It is the value of an option to buy the entire portfolio of the financial institution's assets and liabilities at a fixed price at a specific point in time.

Here are some examples of the Jarrow-Merton put option as a practical risk management concept:

- Instead of the 10-day **value at risk** of a trading portfolio, what is the value of a 10-day put option on my current portfolio with an exercise price equal to the portfolio's current market value? The price of the put option will increase sharply with the risk of my portfolio, and the put option's price will reflect all possible losses and their probability, not just the 99th percentile loss as is traditional in value at risk analysis
- Instead of **stress testing the 12-month net income** of the financial institution to see if net income will go below $100 million for the year, what is the price of a put option in month 12

which will produce a gain in net income exactly equal to the short-fall of net income versus the $100 million target? The more interest rate risk in the balance sheet of the financial institution, the more expensive this put option will be. The put option will reflect all levels of net income shortfall and their probability, not just the shortfalls detected by specific stress tests

- Instead of the **Basel II, Basel III or Solvency II risk-weighted capital ratio** for the institution, what is the price of the put option that insures solvency of the institution in one year's time? This put option measures all potential losses embedded in the financial institution's balance sheet and their probability of occurrence, including both interest rate risk and credit risk, as we discuss at the end of this chapter.
- Instead of '**expected losses' on a collateralized debt obligation tranche's B tranche**, what is the price of a put option on the value of the tranche at par at maturity? This put option reflects all losses on the tranche, not just the average loss, along with their probability of occurrence.
- Instead of '**expected losses on the Bank Insurance Fund'** in the U.S., the Federal Deposit Insurance Corporation has valued the put option of retail bank deposits at their par value as discussed in the FDIC's loss distribution model announced on December 10 2003.

The Jarrow-Merton put option concept helps to reconcile what many regard as conflicting objectives to be managed from a risk management perspective:

- **Net interest income (or net income),** a multi-period financial accounting-based figure that includes the influence of both instruments that the financial institution owns today and those that it will own in the future.
- **Market value of portfolio equity**, the market value of the assets a financial institution owns today less the market value of its liabilities.
- **Market-based equity ratio**, which is the ratio of the mark-to-market value of the equity of the portfolio ('market value of portfolio equity' in bank jargon) divided by the market value of assets. This is most closely related to the capital ratio formulas of the primary capital era, Basel I, Basel II, Basel III and Solvency II.
- **Default probability of the institution**, which is another strong candidate as a single measure of risk.

We next discuss how the valuation of these options can be done using the technology in the first 38 chapters of this book.

Valuing and Simulating the Jarrow-Merton Put Option

How do we value the Jarrow-Merton put option? In some of the cases mentioned above, we already have known valuation approaches from Chapters 20 through 35 using an N-factor Heath Jarrow and Morton interest rate model. In those chapters, we used the 3 factor HJM model from Chapter 9 for exposition purposes. For the full on and off-balance sheet portfolio of a financial institution, we simply repeat the process with finer time granularity and more transactions. This is what computers are for. We need to calculate the 'risk neutral' expected values of the cash flows that would be paid on the Jarrow-Merton put option that we are analyzing.
We outlined the steps in this simulation in Chapter 19 in the context of a portfolio of risky bonds:
For the reduced form model, the steps are as follows:

1. Simulate the risk-free term structure using an N-factor HJM approach

2. Choose a formula and the risk-drivers for the liquidity component of credit spread for all of the relevant asset classes as discussed in Chapter 17.
3. Choose a formula and the risk-drivers for the default intensity process in the Jarrow model as discussed in Chapter 16 for all asset classes.
4. Simulate the random values of the drivers of liquidity risk and the default intensity for *M* time periods over *N* scenarios, consistent with the risk-free term structure.
5. Calculate the default intensity and the liquidity component at each of the *M* time steps and *N* scenarios for each counterparty, from retail to small business to corporate to sovereign. Note that these will be changing randomly over time because interest rates and other key macro factors (like home prices) are changing as well. This is essential to capturing cyclicality in traded asset prices and defaults.
6. Apply the Jarrow model as we have done in Chapters 16 and 17 to get the zero-coupon bond prices for each maturity for each transaction for each counterparty.
7. Calculate the dates and cash flow amounts for the Jarrow-Merton 'put option' structure that is most relevant.
8. Calculate the put option's value by discounting by the appropriate risk-free zero-coupon bond price for each of the *M* time steps and *N* scenarios.

This is a general valuation procedure that we can apply for all instruments discussed in Chapters 20 through 26.

What's the Hedge?

Let's say management has become comfortable with the Jarrow-Merton put option as a comprehensive measure of risk. How do we address the key test of risk management technology: 'What's the hedge?'

If we have done a comprehensive job of fitting our default probability models as in Chapters 16-17 and a thorough job of fitting our interest rate models as discussed in Chapters 5 through 10, we can stress test the put option value that we derived in the previous section. We pick a macroeconomic risk factor to stress test. If the financial institution is a lender in the U.S., the S&P500 is a commonly used macro factor that can be proven to be a statistically significant driver of default probabilities for a large range of counterparties. We can calculate what position in S&P 500 futures is needed to control changes in the Jarrow-Merton put option as a measure of risk in exactly the same way van Deventer and Imai [2003] discuss macro hedging of a portfolio of risky credits:

1. Select the base case value *X* for the S&P 500 and all other macro factors.
2. Value the appropriate Jarrow-Merton put option using the approach of the previous section.
3. Select the stress test value *Y* for the S&P500 and use all of the same macro factor values as in step 1. The same monte carlo 'seed value' as step 1 also has to be used.
4. Get the stress-tested value of the Jarrow-Merton put option.
5. The change in the Jarrow-Merton put option is the Value in Case 1 – Value in Case 2.
6. The change in the S&P 500 is X-Y.
7. The proper hedge depends on what type of S&P500 hedging instruments are being used (futures contracts or not, maturity of futures contracts, etc.). The number of contracts to be used for the hedge is the number such that the change in the value of the futures contracts if the S&P500 moves from *X* to *Y* will exactly offset the change in the value of the Jarrow-Merton put option from its Case 1 value to its Case 2 value.

Using this approach, we always know the proper integrated hedge of interest rate risk and credit risk. Our objectives have been achieved.

Liquidity, Performance, Capital Allocation and 'Own Default' Risk

Using the Jarrow-Merton put option approach that we have built in this chapter on the foundation of Chapters 1 through 35, we can now definitively address key risk management issues like those in our 4 + 26 questions above:

> What is the liquidity risk of my organization?
>
> What amount of capital should I have in the institution as a whole?
>
> What amount of capital should I have in each business unit?
>
> What is the performance of each business unit?

We turn to that task in the next three chapters.