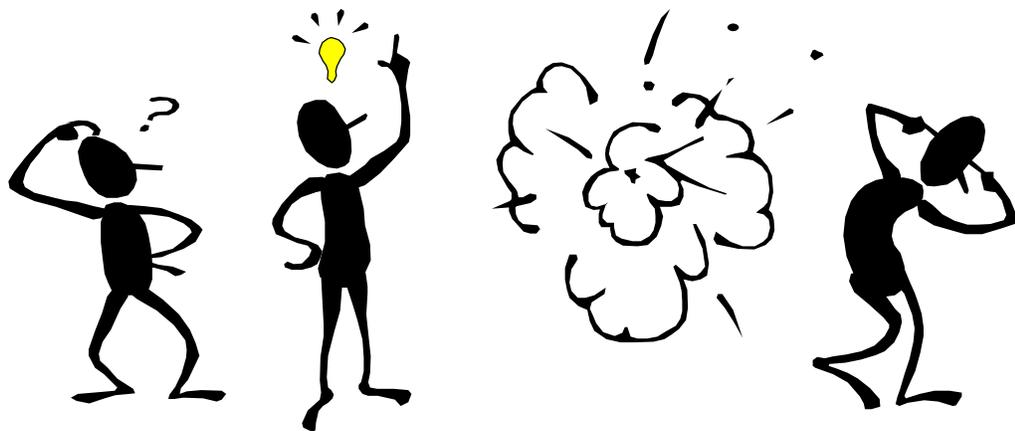


Risk Analysis Techniques

1998 GARP FRM Exam Review Class Notes





Risk Analysis Techniques

GARP FRM Exam Review Class Notes

Ian Hawkins September 1998

Table of Contents

Introduction	3
VaR Assumptions	3
Delta-Normal Methodology	5
Delta-Gamma Methodology	6
Historical Simulation.....	7
Stress Testing.....	8
Monte Carlo Simulation.....	8
Raroc	9
Model Risk	10
Implementation Strategy	13
Warning Signs	14
Conclusion.....	15
Other Sources.....	15
Answers to Sample Exam Questions	16
Answers to Risktek Risk Olympics™ Questions	16

Introduction¹

The purpose of this class is to provide an idiosyncratic review of the techniques for risk analysis that a risk management professional should be familiar with. This document contains a large number of references and you should spend some time tracking a few of them down, particularly in areas where you feel less comfortable with your own experience or the class content. There is no guarantee that the information presented here is either correct or what the examiners will be questioning you on.

Let's assume that our overall goal is to create a quantitative measure of risk that can be applied to the business unit we are responsible for. Sitting on our hands doing nothing is not an option. We need a measure of risk that can be applied at all levels of an organization either to an isolated business unit or in aggregate to make decisions about the level of risk being assumed in those business units and whether it is justified by the potential returns. A sensible objective is to conform to (no more than) industry best practice at a reasonable cost. The standard industry approaches set out below are a starting point. I will defer most of the debate as to why and whither VaR to the end of the session.

This class material can be organized into three areas: the syllabus topics, some additional practical topics I think you will find of interest, and sample exam questions and answers.

Let's begin with the syllabus. The structure of the syllabus follows Thomas Wilson's chapter in the Handbook of Risk Management and Analysis². Duffie and Pan provide another good review³. While the sheer volume of material can be overwhelming the RiskMetricsTM technical document and the quarterly RiskMetrics monitors⁴ are excellent and well worth whatever time you can spend with them. For a gentler read try Linsmeier and Pearson⁵, or, if you have an expense account, Gumerlock, Litterman et al⁶.

VaR Assumptions

The Value at Risk of a portfolio is defined as the portfolio's maximum expected loss from an adverse market move, within a specified confidence interval, over a defined time horizon. There is a considerable amount of research directed at removing the qualifier "within a specified confidence interval"^{7,8,9,10}, but this paper will stick with current market practice.

¹ Thanks to Lev Borodovsky, Randi Hawkins, Yong Li, Christophe Rouvinez, Rob Samuel and Paul Vogt for encouragement, helpful comments and/or reviewing earlier drafts. Please send any questions or comments to IanHawkins@aol.com © Ian Hawkins 1997-8

² Calculating Risk Capital, Thomas Wilson, in the Handbook of Risk Management and Analysis. Carol Alexander (ed) Wiley 1996 ISBN 0-471-95309-1

³ An Overview of Value at Risk, Darrell Duffie and Jun Pan, Journal of Derivatives, Spring 1997, pp7-49

⁴ <http://www.jpmorgan.com/RiskManagement/RiskMetrics/pubs.html>

⁵ Risk Measurement: An Introduction to Value at Risk, Thomas Linsmeier and Neil Pearson, University of Illinois at Urbana-Campaign, July 1996 at

<http://econwpa.wustl.edu/eprints/fin/papers/9609/9609004.abs>

⁶ The Practice of Risk Management, Robert Gumerlock, Robert Litterman et al, Euromoney Books, 1998 ISBN 1 85564 627 7

⁷ Thinking coherently, Phillipe Artzner et al, Risk V10, #11, November 1997

⁸ Expected Maximum Loss of Financial Returns, Emmanuel Acar and David Prieul, Derivatives Week, September 22, 1997

⁹ Living On The Edge, Paul Embrechts et al, Risk, January 1998

¹⁰ History Repeating, Alexander McNeil, Risk, January 1998.

Implementing VaR for any reasonable sized organization is a heroic undertaking that requires both heroic assumptions and heroic compromises to succeed. We start with the most common assumptions about the portfolio and the markets behind VaR measurements and discuss how we can make VaR a robust tool .

Describing the portfolio

The first assumption made is that the portfolio does not change over the VaR time horizon. It is hard to imagine any trading organization for which this could possibly be true. VaR is usually measured against closing positions for a one-day to two-week horizon. We know that overnight position limits are smaller than intra-day limits – so what happens if the crash hits in the middle of the day when you are halfway through hedging a large deal? Even though your position is a lot larger than the closing position you are probably going to do something about it a lot sooner than the VaR measurement horizon.

Second, we assume the portfolio can be summarized by its sensitivities with respect to a small number of risk factors. Do we have a suitable set of factors for the type of portfolio under consideration? A controller once asked me about the risk of a floor trading operation that she was responsible for overseeing. The positions showed essentially flat greeks in each contract month. Either the traders were only taking intra-day positions or they were running strike spreads that did not show on the report. Not surprisingly it was the latter. While option strike risk has gained heightened interest post NatWest, most VaR systems do not capture *change* in smile as a risk factor, even if smiles are used for revaluation. In fact it is usually easier to catalogue which risk factors are present rather than which are missing (bond option portfolios repo exposures, commodity portfolio contango risks, swap portfolio basis risk 1M vs. 3M vs. 6M vs 12M Libor, cash/futures divergence). VaR cannot replace the rich set of trading controls that most businesses accumulate over the years. Over-reliance on VaR is simply an invitation for traders to build up large positions that fall outside the capabilities of the implementation.

Third, we assume that the sensitivities can be captured by the first (and possibly second derivatives) with respect to the risk factors – often dropping any cross partial derivatives. Not surprisingly, Taylor series work well *only* for portfolios with sensitivity profiles that are close to linear (or possibly quadratic) forms¹¹, and work very poorly if, for example, the portfolio is short large amounts of very deep out of the money puts (Niederhoffer and UBS equity derivatives) that have essentially zero sensitivity to local movements in spot prices.

Describing the market

We begin by assuming that past market behavior can tell us something about the future. Second we have to decide how much of the past market behavior we wish to consider for our model. As we are interested in rare events, it might seem reasonable to constrain our market history to the rare events, but in most cases we use the complete history for a particular time frame, as it is difficult to form a statistically meaningful sample if we only study the rare events. Given a data set, we now have to propose a model for the market data innovations. Most analytic methods are based on a set of normally distributed risk factors, with independent increments and a stationary variance-covariance matrix. Third and finally, we have to estimate parameters for the model, and then assume those parameters can be applied to a forward-looking analysis of VaR.

Most of the research on non-parametric estimation of the process for the spot interest rate or the yield curve challenges all of these assumptions. Non-parametric estimation of a model means using a large amount of data to estimate the real-world probability distribution. I am

¹¹ Taylor, Black and Scholes: Series Approximations and Risk Management Pitfalls, Arturo Estrella, FRBNY Research Paper #9501

sure you can find similar references for other markets. Ait-Sahalia¹² and Wilmot et al.¹³ both reject the family of one-factor models in common use and propose models that are significantly different and more complicated. Ait-Sahalia actually finds that interest rates do not follow a process that is either a diffusion or Markov¹⁴.

One ray of hope is a recent paper by Pritsker¹⁵ that suggests that the earlier tests may be flawed when applied in a time series context. However he also implies that estimation of the true process for rates is even more complicated than work that is already inaccessible to most practitioners.

Robust VaR

Just how robust is VaR? In most financial applications we choose fairly simple models and then abuse the input data external to the model to accommodate the market. We also build a set of rules about when the model output is likely to be invalid. VaR is no different. Consider the Black-Scholes analogy: one way we abuse the model is by varying the volatility according to the strike. We then add a rule to not sell very low delta options at the model value because even with a steep volatility smile you just can't get the model to charge enough to make it worth your while. A second Black-Scholes analogy is the modeling of stochastic volatility by averaging two Black-Scholes values (using market volatility +/- a perturbation).

Given the uncertainties in the input parameters (with respect to position, liquidation strategy/time horizon and market model) and the potential miss-specification of the model itself it seems reasonable to attempt to estimate the uncertainty in the VaR. This can either be done formally to be quoted whenever the VaR value is quoted – or informally to flag the VaR value because it is extremely sensitive to the input parameters or the model itself.

Consider a simple analysis of errors for single asset VaR. The VaR is given by confidence interval * risk factor_{STD} * risk factor_{POSITION}: if the position is off by 15% and the standard deviation is off by 10% then relative error of VaR is 15+10 = 25%! Note that this error estimate excludes the problems of the model itself.

This does not indicate that VaR is meaningless – just that we should exercise some caution in interpreting the values that our models produce. Now let's proceed to the methodologies.

Delta-Normal Methodology¹⁶

The standard RiskMetrics methodology measures positions by reducing all transactions to cash flow maps. The volatility of the returns of these cash flows is assumed to be normal i.e. the cash flows each follow a lognormal random walk. The change in the value of the cash flow is then approximated as the product of the cash flow and the return (i.e. using the first term of a Taylor series expansion of e^x).

¹² Testing Continuous Time Models of the Spot Rate, Yacine Ait-Sahalia, Review of Financial Studies 2, No 9, 1996, p385-426

¹³ Spot-on Modeling, Paul Wilmot et al., Risk, Vol 8, No 11, November 1995

¹⁴ Do Interest Rates Really Follow Continuous-Time Markov Diffusions?, Yacine Ait-Sahalia, Working Paper, Graduate School of Business, University of Chicago

¹⁵ Non-parametric Density Estimation and Tests of Continuous Time Interest Rate Models, Matt Pritsker, Federal Reserve Board of Governors Working Paper FEDS 1997-26 at <http://www.bog.frb.fed.us/pubs/feds/1997/199726/199726pap.pdf>

¹⁶ The accompanying spreadsheet has some simple numerical examples for the delta-normal and historical simulation methods.

Cash flow mapping can be quite laborious and does not extend beyond price and interest rate sensitivities. The Delta-Normal methodology is a slightly more general flavor of the standard RiskMetrics methodology, which considers risk factors rather than cash flow maps. The risk factors usually correspond to standard trading system sensitivity outputs (price risk, vega risk, yield curve risk). One benefit is a huge reduction in the size of the covariance matrices. Even if additional risks beyond price and interest rate are considered, you typically replace sixteen RiskMetrics maturities with no more than three yield curve factors (parallel, tilt and bend).

The risk factors are assumed to follow a multivariate normal distribution and are all first derivatives. Therefore the portfolio change in value is linear in the risk factors and the position in each factor and the matrix math looks identical to RiskMetrics even though the assumptions are rather different¹⁷.

Assuming that the sensitivity of a position can be captured entirely by first derivatives is quite crude. The following sections describe various ways to improve on this.

Delta-Gamma Methodology

There are two methodologies commonly described by the term delta-gamma. In both cases the portfolio sensitivity is described by first and second derivatives with respect to risk factors.

Tom Wilson works directly with normally distributed risk factors and a second order Taylor series expansion of the portfolio's change in value. He proposes three different solution techniques, two of which require numerical searches. The third method is an analytic solution that is relatively straightforward. The gamma of a set of N risk factors is an NxN matrix. The diagonal is composed of second derivatives – what most people understand by gamma.

The off diagonal or cross terms describe the sensitivities of the portfolio to joint changes in a pair of risk factors. For example a yield curve move together with a change in volatility. Tom orthogonalizes the risk factors. The transformed gamma matrix has no cross terms, so the worst case in each risk factor will also be the worst case risk for the portfolio. He then calculates an adjusted delta that gives the same worst case P/L for the market move corresponding to the confidence interval as the original delta as the original volatility, the original delta and the original gamma. Picture the adjusted delta as a straight line from the origin to the worst case P/L, where the straight line crosses the curve representing the actual portfolio P/L. Given this picture, we can infer that the VaR number is correct only for a specified confidence interval and cannot be re-scaled like a delta-normal VaR number.

An ad-hoc version of this approach can be applied to un-transformed risk factors – provided the cross terms in the gamma matrix are small. To make things even simpler you can require the systems generating delta information to do so by perturbing market rates by an amount close to the move implied by the confidence interval and feed this number into your delta-normal VaR.

RiskMetrics¹⁸ takes a very different approach to extending the delta-normal framework. The delta and gamma are used to calculate the first four moments of the portfolio's return distribution. A function of the normal distribution is chosen to match these moments. The percentile for the normal distribution can then be transformed to the percentile for the actual

¹⁷ As a reference for the variance of functions of random variables see "Introduction to Mathematical Statistics" Robert Hogg and Allen Craig, Macmillan, ISBN 0-02-355710-9, p 176ff

¹⁸ RiskMetrics Technical Document, Fourth Edition, p130-133 at <http://www.jpmorgan.com/RiskManagement/RiskMetrics/pubs.html>

return distribution. If this sounds very complicated, think of the way you calculate what a 3-std move in a log-normally distributed variable is worth. You multiply the volatility by 3 to get the change in the normal variable, and then multiply the spot price by e^{change} to get the upper bound and divide by $e^{-\text{change}}$ to get the lower bound. (Hull and White¹⁹ propose using the same approach for a slightly different problem.)

Now let's consider how we can address the distribution assumptions.

Historical Simulation

Historical simulation is the process of calculating P/L by applying a historic series of the changes in risk factors to the portfolio (with the sensitivity captured either using risk factors [as many terms as you like], a P/L spline or, much less often, a complete revaluation for each set of historic data. This approach addresses the problem of modeling the market if old data is "representative" and potentially also addresses the issue of using only a local measure of risk, depending on the implementation. The portfolio change in value is then tabulated and the loss percentile in question can simply be looked up.

While the problems of modeling and estimating parameters for the market are eliminated you are obviously sensitive to whether the historic time series captures the features of the market that you want to be represented – whether that is fat tails, skewness, non-stationary volatility or the presence of extreme events. Naturally absence of a historic time series for a risk factor you want to include in your analysis is a problem! For instance, OTC volatility time series are difficult to obtain (you usually have to go cap in hand to your option brokers) and entry into a new market for an instrument that has not been traded for long is a real problem.

The method is computer resource intensive compared to delta-normal and delta-gamma particularly in CPU and possibly also in the space required for storing all the historic data. However, note that the time series takes up less data than the covariance matrix if the number of risk factors is more than twice the number of observations in the sample²⁰.

Instead of just looking up the loss percentile in the table of the simulation results, the distribution of the portfolio change in value can be modeled and the loss inferred from the distribution's properties^{21,22}. This approach uses information from all the observations to make inference about the tails.

Finally, incremental VaR is a hazier concept in a historic simulation, as the days that contribute maximum loss for two positions may be different, and the VaR will change less smoothly than for an analytic model with any reasonably small data set. (You may see a similar effect in any model that accounts for non-linear portfolio behavior, as the maximum loss scenario may be quite different for an incremental change in the portfolio.)

From historical simulation of a set of market changes it is natural to move on to stress testing, which considers single historic events.

¹⁹ Value At Risk When Daily Changes In Market Variables Are Not Normally Distributed, John Hull and Alan White, *Journal of Derivatives*, Spring 1998

²⁰ A General Approach to calculating VaR without volatilities and correlations, Peter Benson and Peter Zangari, *RiskMetrics Monitor* Second Quarter 1997.

²¹ Streamlining the market risk measurement process, Peter Zangari, *RiskMetrics Monitor* First Quarter 1997

²² Improving Value-at-Risk Estimates by Combining Kernel Estimation With Historic Simulation, J Butler and B. Schachter, OCC, May 1996.

Stress Testing

Stress testing is the process of replaying the tape of past market events to see their effect on your current portfolio. The BIS²³ lists the 87 stock market crash, Sterling's exit from the ERM, and the 94 bond market crash as events whose impact should be studied. Other events worth looking at are the Asian contagion, the Mexican peso devaluation, the Bunker Hunt silver market squeeze, the collapse of copper market prices in the summer of 1996²⁴ and the collapse of tin prices after the demise of the ITC in 1985. Note that the BIS requires you to tailor scenarios to the bank's portfolio. I would add that you need to take a forward-looking approach to devising scenarios – if anything it is more important to spend time devising events that might happen rather than concentrating on those that already have happened.

In addition to understanding the size of the market moves that occur at times of stress it is also instructive to read broader descriptions of the events – particularly if you have not seen any major market moves yourself. The GAO description of the 87 crash²⁵ contains a wealth of information – two things I take from the report are the need for crisis management plans to be in place before the event happens and the fact that, while the exchanges performed well, the NYSE specialist system did not.

Fung and Hsieh²⁶ conclude that large movements in the level of interest rates are highly correlated with large movements in yield curve shape in contrast to the statistical behavior of the curve when considering all movements.

Just as a reminder when you are reviewing the results of your stress tests – it is imprudent to enter into any transaction whose payoff if triggered, however unlikely that trigger event might be, would significantly impact the viability of the business unit. One rule of thumb is to never commit more than 10% of your capital to any one bet or any one client. There is a conflict between the risk reducing effects of business and client diversification, and the desire of institutions to exploit core competencies, find niches, and expand client relationships.

Monte Carlo Simulation

Monte Carlo simulation uses a model fed by a set of random variables to generate risk factor innovations rather than historical data. Each simulation path provides all the market data required for revaluing the whole portfolio. The set of portfolio values can then be used to infer the VaR as described for historical simulation. Creation of a model for the joint evolution of all the risk factors that affect a bank's portfolio is a massive undertaking. This approach is also extremely computationally intensive and is almost certainly a hopeless task for any institution that does not already use similar technology in the front office.

While, in principle, Monte Carlo simulation can address both the simplifying assumptions in modeling the market and representing the portfolio it is naïve to expect that most implementations will actually achieve these goals. Monte Carlo is used much more

²³ Amendment to the Capital Accord to incorporate Market Risks, Basle Committee on Banking Supervision, January 1996 at <http://www.bis.org/publ/bcbs23.htm> and <http://www.bis.org/publ/bcbs24.htm>

²⁴ Copper and Culpability, Euromoney magazine July 1996 at <http://www.euromoney.com/contents/euromoney/em.96/em.9607/em96.07.4.html>

²⁵ Stock Market Crash of October 1987, GAO Preliminary Report to Congress, CCH Commodity Futures Law Reports Number 322 part II, February 1988.

²⁶ Global Yield Curve Event Risks, William Fung & David Hsieh, Journal of Fixed Income, Sep 96, p37-48

frequently as a research tool than as part of a production platform in financial applications, except possibly for MBS.

Performance measurement is a natural complement to risk management as the data needed are typically collected as part of the risk management function.

Raroc

Senior bank management are essentially deciding how to allocate capital among a portfolio of businesses, and they need a measure of performance that takes into account both the returns and the risk of a business activity to do it.

Markovitz²⁷ introduced the concept that investors should choose portfolios that offer the highest return for a given level of risk rather than just maximizing expected return. Implementing Markovitz' mean-variance analysis has many parallels with calculating VaR. Sharpe²⁸ developed a simpler index – originally intended for measuring the performance of mutual funds – equal to the incremental return over a benchmark divided by the standard deviation of the incremental returns. Although the Sharpe Ratio is a widely used benchmark, note John Bogle's (founder of the Vanguard mutual fund business) comment that the Sharpe ratio fails to capture how precious an additional 100bps of return is relative to an additional 100bps of risk, for a long term investor – his view is that risk is weighted too heavily in the Sharpe ratio²⁹.

Banks typically use Risk Adjusted Return on Capital (Raroc) to measure performance. Smithson³⁰ defines Raroc as adjusted net income/economic capital where net income is adjusted for the cost of economic capital. Smithson also highlights the different flavors of capital measure that should be used for different types of decisions. For allocation decisions the capital measure should reflect any potential diversification benefit offered by a business when placed in the bank portfolio whereas for performance measurement the capital measure should reflect the economic capital of the business as a stand alone unit.

Shimko³¹ relates Raroc, VaR and the Sharpe ratio, given the strong assumption that VaR corresponds to economic capital.

Traders have a put on the firm. Bonus pools are typically funded according to a set percentage of net income. The traders' income is a linear multiple of the firm's income, with a floor at their base salary. Given this payoff the way for the traders to earn the most income is to increase the variance of the P/L as much as possible (large negative returns will be absorbed by the firm). You may or may not believe that ethical considerations and the risk of getting fired temper this. Losing large amounts of money does not seem to be correlated with career failure. Asset managers have similar incentives³². In theory traders and asset managers should be compensated on the basis of a measure that takes returns and risk into account but in practice this is rare.

²⁷ Portfolio Selection, H. Markovitz, Journal of Finance, March 1952

²⁸ The Sharpe Ratio, William F Sharpe, Journal of Portfolio Management, Fall 1994 or <http://www-sharpe.stanford.edu/sr.htm>

²⁹ The Four Dimensions Of Investment Return, John Bogle, Speech to Institute for Private Investors Forum, May 21, 1998 at www.vanguard.com/educ/lib/bogle/dimensions.html

³⁰ Capital Budgeting, Charles Smithson, Risk, Vol 10, No 6, June 1997 p40-41

³¹ See Sharpe Or Be Flat, David Shimko, Risk, Vol 10 No 6, June 1997, p33

³² Investment Management Fees: Long-Run Incentives, Robert Ferguson and Dean Leistikow, Journal of Financial Engineering , Vol 6, No 1, p 1-30

Any modeling effort is susceptible to conceptual errors and errors in execution. The next sections consider what can go wrong in the modeling and implementation.

Model Risk³³

While the VaR methodologies implemented at most firms have many flaws, their simplicity is actually an asset that facilitates education of both senior and junior personnel in the organization. VaR is just the first step along the road. Creating the physical and intellectual infrastructure for firm wide quantitative risk management is a huge undertaking. Successful implementation of a simple VaR model is a considerable achievement that few institutions have accomplished in a robust fashion.

We have already discussed the assumptions behind VaR. As with any model you should understand the sensitivity of the model to its inputs. In a perfect world you would also have implemented more than one model and have reconciled the difference between their results. In practice this usually only happens as you refine your current model and understand the impact of each round of changes. Beder³⁴ shows a range of VaR calculations of 14 times for the same portfolio using a range of models – although the example is a little artificial as it includes two different time horizons. In a more recent regulatory survey of Australian banks, Gizycki and Hereford³⁵ report an even larger range (more than 21 times) of VaR values, though they note that “crude, but conservative” assumptions cause outliers at the high end of the range.

Note that most implementations study the *terminal probabilities* of events, not *barrier probabilities* i.e. the possibility of the event happening at *any* time over the next 24 hours rather than the probability of the event happening when observed after 24 hours have passed. Naturally, the probability of exceeding a certain loss level at any time over the next 24 hours is higher than the probability of exceeding a certain loss level at the end of 24 hours. This problem in handling time is similar to the problem of using a small number of terms in the Taylor series expansion of a portfolio’s P/L function. Both have the effect of masking large potential losses inside the measurement boundaries.

The regulatory multiplier^{36, 37} takes the VaR number you first thought of and multiplies it by at least three – and more if the regulator deems necessary. Even though this goes a long way to addressing the modeling uncertainties I would still not think of VaR as a measure of your downside on its own. Best practice requires that you establish market risk reserves³⁸ and model risk reserves³⁹. Model risk reserves should include coverage for potential losses that relate to risk factors that are not captured by the modeling process and/or the VaR process. Whether such reserves should be included in VaR is open to debate. Remember that VaR measures *uncertainty* in the portfolio P/L, and reserves are there to cover *potential* losses.

³³ Emanuel Derman’s articles are required reading: Model Risk, Risk, Vol 9, No 5, May 1996, p 34-37 and Valuing Models and Modeling Value, The Journal of Portfolio Management, Spring 1996, p 106-114

³⁴ VaR: Seductive but dangerous, Tanya Styblo Beder, Financial Analysts Journal Vol 51, no 5 (Sep/Oct 1995), p12-24 or <http://www.cmra.com/fajvar.pdf>

³⁵ Differences of Opinion, Marianne Gizycki and Neil Hereford, Asia Risk, August 1998, p42-47. I recommend reading this paper!

³⁶ Three Cheers, Gerhard Stahl, Risk, V10, #5, May 1997

³⁷ Market Risk Capital, Darryll Hendricks and Beverly Hirtle, Derivatives Week, April 6, 1998

³⁸ Derivatives: Practices and Principles, Global Derivatives Study Group, Group of Thirty, Washington, DC Recommendations 2 and 3

³⁹ Derivatives: The Realities of Marking to Model, Tanya Styblo Beder, Capital Market Risk Advisors at www.cmra.com/research.htm

Certain changes in the P/L or actual losses, even if not captured by the models used for revaluation, should be included in the mark to market as adjustments.

The process that a major derivatives loss typically follows is: realization of a major mark to market loss over an extended period of time (weeks to months); review of mark to market, model and position; revision of mark to market, model and/or risk appetite; liquidation. You lose a lot of money repeatedly (i.e. multiples of the VaR). Then you rethink what you have been doing and take a further hit as you change the model and/or change your mark to market. Then you pay to liquidate the position and the bill for the lot is more than the sum of the VaR and all the reserves (if you have them). Then you fire your traders and hire the guys who blew up just as badly at another institution.

Implementation Risk

Every VaR system carries implementation risk. Marshal and Siegel⁴⁰ study the range of vendor VaR estimates for the same set of positions and the same model (the original RiskMetrics model). The aggregate VaR estimates range from 2.9 to 5.4 around a mean of 4.2 million and for a linear portfolio and from 0.75 to 2.1 million around a mean of 1.1 million for the options portfolio. Note that this was after the vendors had feedback on the rest of the sample and the opportunity to revise their numbers. Buried in the endnotes is a comment that if different implementations exhibit systematic biases for particular instruments to a greater extent than different models integration may increase model and implementation risks, i.e. you are better off finding a way to add the VaRs from separate systems, each of which handles the distinctive features of a particular business well, rather than using a single system that handles all instruments poorly.

Of course every VaR system is also subject to implementation and model risk in the supporting systems that supply position and risk factor data over and above the VaR calculation itself.

In addition to specifying a model for the market we have to provide the model with parameters.

Parameter Estimation

Hendricks⁴¹ studies the use of equally and exponentially weighted estimators of variance for a number of different sample sizes in the delta normal and historical simulation methods. His results indicate that there is very little to choose between either the different estimators or the two VaR methods. The range of the scaling factors required to make the VaR achieve a given level of coverage (15%) is a lot smaller than the model risk results above. This is primarily because he studies simple, linear, foreign exchange portfolios and I interpret his results as an indication that all else held equal VaR results are not that sensitive to the choice of estimator.

Boudoukh et al⁴² study the efficiency of different weighting schemes for volatility estimation. The “winner” is non-parametric multivariate density estimation (MDE). MDE puts high weight on observations that occur under conditions similar to the current date. Naturally this requires an appropriate choice of state variables to describe the market conditions. For example the authors use yield curve level and slope when studying Treasury bill yield

⁴⁰ Value at Risk: Implementing a Risk Measurement Standard, Chris Marshall and Michael Siegel, The Journal of Derivatives, Spring 1997, p 91-111

⁴¹ Evaluation of VaR Models Using Historical Data, Darryll Hendricks, FRBNY Policy Review, April 1996 or http://www.ny.frb.org/rmaghome/econ_pol/496end.pdf

⁴² Investigation of a Class of Volatility Estimators, Jacob Boudoukh, Matthew Richardson and Robert Whitelaw, Journal of Derivatives, Spring 1997, p 63-71

volatility. MDE does not seem to represent a huge forecasting improvement given the increased complexity of the estimation method but it is interesting that one can formalize the concept of using only representative data for parameter estimation.

My own preference is for long run (one to five year) equally weighted estimators that are updated once a year. This eliminates change in VaR due to change in anything other than position, and also reflects my belief that VaR should not decay rapidly (a la RiskMetrics) if the market does little for a couple of weeks. <http://www.riskex.com/vargraph.html> shows the VaR for a portfolio that matches the FT-30 index. I see a repeated pattern of decaying VaR followed by a spike corresponding to an extreme move, followed by a jump up in VaR i.e. the VaR increases *after* the blowup, and the use of a short term estimator exacerbates the extent to which losses exceed VaR more than predicted by the model. Hoppe⁴³ provides evidence to the contrary.

Naïve use of a delta normal approach requires estimating and handling very large covariance matrices. Alexander⁴⁴ advocates a divide and conquer strategy – break down the risk factors into a sets of highly correlated factors. Then perform principal components analysis to create a set of orthogonal risk factors – then estimate variances of the orthogonal factors and covariances of the principal components. Alexander also concludes from back testing that there is little to choose between the regulatory 1 year equally weighted model and Garch(1,1) while the RiskMetrics estimator performs less well.

Having implemented a model we then have to measure how well it performs.

Back Testing

VaR measurement requires making statistical inference about rare events. Kupiec⁴⁵ describes the difficulty devising tests to check VaR models that will reliably reject models that underestimate VaR and at the same time not reject models that correctly estimate VaR. Even for large data windows there may not be enough observations make reliable statements about the tails of the portfolio returns distribution. See also Jackson et al⁴⁶.

Jorian and Taleb

The Siskel and Ebert of Risk Management... if you have not read the Derivatives Strategy articles by Taleb and Jorian then please make the effort to check them out⁴⁷. Taleb is an excellent writer and gives good copy. I may be reading too much in to a phrase, but it seems that much of Taleb's ire is directed at the risk management "profession" rather than their choice of tools. As Taleb probably has better quantitative skills, more market knowledge and more product knowledge than those who would monitor him we can all empathize.

It is not clear to me what if anything Taleb is proposing as an alternative for aggregation and comparison of risk capital. The main arguments in favor of VaR are the flaws in the institutional and regulatory metrics it replaces, or the void it fills. However, as I stated earlier,

⁴³ VaR and the Unreal world, Richard Hoppe, Risk, July 1998

⁴⁴ Splicing Methods for VaR, Carol Alexander, Derivatives Week, June 1997, p8-9 and On the Covariance Matrices Used in VaR Models, Alexander and Leigh, Journal of Derivatives, Spring 1997 p50-62

⁴⁵ Techniques for Verifying the Accuracy of Risk Measurement Models, Paul Kupiec, The Journal of Derivatives, Winter 1995, p 73-84

⁴⁶ Bank Capital and Value at Risk, Patricia Jackson, David Maude and Willian Perraudin, The Journal of Derivatives, Spring 1977, p 73-89

⁴⁷ The World According To Nassim Taleb, Joe Kolman, Derivatives Strategy, Dec/Jan 1997, p37-40 or <http://pw1.netcom.com/~ntaleb/index.html>

I see VaR as a complement to other sets of rules rather than a single grand unified solution to all our risk monitoring problems.

Implementation Strategy

I think the two most common mistakes made in project management are expansion of scope beyond the minimum necessary to get the job done and failure to kill off projects that are failing. Two thumping good reads on project management are Brooks⁴⁸ and Alexander⁴⁹.

My experience is that most projects overrun time and cost by a factor of at least two or three (and that is if they don't fail completely). If you have a project manager with a batting average above .400 they are good (note that I measure the average with respect to the original scope and plan). The only way to manage the problems inherent in building large systems is to break the problem up. Break the overall project up into intermediate deliverables (spreadsheet, then Access system, then Sybase/C++) and push as much functionality as possible into later phases. Break the problem up into modules with well-defined interfaces, even at the expense of overall functionality. Break up the functional requirements of different users into different systems. For example, while Lawrence⁵⁰ proposes integration of liquidation horizon and liquidation cost with the VaR analysis I think that it is more practical to separate the two. It is better to have humble goals and deliver.

Conversely, beware project managers that don't schedule major deliverables at least every six months and beware risk managers who want to build transaction warehouses (before exhausting all other possibilities). The archetypal train of events for a project that will go wrong is as follows:

- IT (note not trading or risk management) engages in study of global risk management systems. After a six-month world tour they deliver a hundred page report concluding that:
 - ◆ there is a great deal of redundancy in the current systems,
 - ◆ the current systems don't meet the organizations needs and
 - ◆ the organization needs a single global solution that will provide everyone with what they want at much less than the cost of upgrading what they have now.
- IT engages in a study of vendor systems. After a six-month world tour they deliver a hundred page report concluding that we need to build the system in-house. Given two years and \$5million they can do the job and save the organization money by eliminating all the redundant systems.
- Risk Management sign up for the show. (Though to be fair sometimes Risk Management drive the whole process.) Senior management feels powerless to resist after reading a devastating risk management G30 Principles and Practices audit – and after all the cost benefit analysis is persuasive. Functional specifications now explicitly include valuation and risk management of every transaction in the bank.
- A team of ten to fifteen consultants is hired at head office. Vast sums of money are spent on new hardware and development software. Resources to the legacy systems (i.e. the ones that are installed now and actually provide useful results) are cut and the first

⁴⁸ The Mythical Man-Month, Frederick Brooks, Addison-Wesley, Anniversary Edition, 1995, ISBN 0-201-83595-9

⁴⁹ The Oregon Experiment, Christopher Alexander et al., Oxford University Press, New York, 1975 ISBN 0-19-501824-9

⁵⁰ Working Liquidity into your VaR, Colin Lawrence, Derivatives Strategy, February 1997, p 45-47

employees quit. Risk Managers can only get new functionality by downloading information into spreadsheets and manipulating the data there.

- Six months into the project massive volumes of specifications are published, distributed and used as doorstops for conference rooms globally. The project plan has no deliverables until global rollout in two years time.
- The middle and end game variations are more varied but the result is usually a very large investment of time and money in return for very little. My view is that the only really effective control is a requirement that the revenue generating units sign off on any major expenditure that will be charged back to them, even though this does compromise the independence of reporting lines to some extent.

My specific recommendations with regard to VaR systems are:

- If in doubt buy it. Don't buy anything that has not been built yet – verify that everything you need on day one is in the vendor product now – otherwise you are just contracting out a build and that will be as painful as managing an in-house build (if not more so).
- If you buy it, structure payments as a lease, rather than a fee plus maintenance. Encourage the vendor to manage an ongoing relationship rather than just make the sale, complete the installation and move on to the next client.
- Leave the valuation functionality in the trading systems and export risk factors. Keep the VaR system as an aggregation/reporting engine. If you can't generate all the risk factors you need from the trading system that indicates a problem with the trading system that should be fixed. It is almost certainly cheaper to fix ALL the trading systems you use than to build a valuation engine that can handle ALL the transactions you have – and it is much more likely that you will have something to show for your efforts at the end of the day.
- Forget about real time VaR unless you have EVERYTHING that you want once a day.
- Remember that the implementation cost, first time through, is going to be comparable to the buy or build cost (either in employee heartache or real overtime/consulting dollars).

All systems and processes can be defeated by fraud – we can only do our best to ensure that fraud is discovered quickly and cannot be perpetrated easily or by individuals acting alone.

Warning Signs

Market losses are often compounded by the attempts of trading staff to conceal them. The first and last line of defense of any organization is the quality and integrity of their staff in both trading and risk management.

Be sensitive to incidents that cause you to doubt the integrity of a trader. If the new hire brought a briefcase full of software from his last employer he engaged in theft. Why would you think (s)he wouldn't steal from the new employer as well?

Be sensitive to changes in the behavior of trading staff – either hours or temperament. It can take a long time to rig positions or P/L, so if a trader suddenly starts working late after leaving each day at 5, you might want to take a closer look. Living with the strain of hiding a large loss can make a trader even crankier than usual! If the outburst sent you scurrying away then it was an effective tactic.

If you don't understand how a business makes its money, you can't monitor it. Traders should be able (and willing) to explain to a risk manager what they are doing to generate profits in a reasonably short amount of time.

If you are not comfortable with *any* information you discover don't ignore it. There is probably a reasonable explanation but make sure you obtain it (especially if you have a fiduciary responsibility to do so!). If you are unhappy with the responses you get from trading or trading management take the issue to your management.

Any student of the major debacles can point to the warning signs that were either ignored or not followed to an appropriate resolution along the way. Be particularly sensitive to large suspense items and large P/L adjustments particularly if they are carried for extended periods of time.

Look extremely hard at businesses that show steady profits with little P/L volatility – particularly if the position size is growing and the maturity of the portfolio is such that it does not “turn over” in a short time frame (say 6-12 months – for example most swaps portfolios don't). Senior management has a tendency to preach risk adjusted returns, but embrace high earnings of any quality.

When things go wrong a lot of people will lie about what happened. Ask the same questions of different people and note inconsistencies. (For example: Trading manager says he reviews all tickets every day, TA says he has only been getting the copies for the last week i.e. after the problem occurred.) Apologize up-front and warn those involved that you are probably going to have to go over the same ground many times.

Some losses are caused by stupidity. If you see hedging or accounting practice that appears wrong take the time to find out why it is correct or prove to yourself that it is wrong and fix it.

Learn from other people's mistakes - read the accounts of big losses.

Conclusion

The challenges facing risk managers today are to provide their services in a sustainable fashion and to use their authority in a responsible manner. The well-publicized industry losses and fear of their recurrence can be used to take control of an inappropriately large resource or to excessively restrict trading activities – either of which could be as damaging to the firm as the problems it is being protected against.

Other Sources

Most of you know about this one: <http://www.garp.com/>

Barry Schachter's VaR bibliography is the best single resource I have seen:
<http://pw2.netcom.com/~bschacht/varbiblio.html>

Capital Markets Risk Advisors has good overviews of policies and processes with regard to Risk Management: <http://www.cmra.com>

Contingency Analysis has a lot of relevant reference material and briefing papers:
<http://www.contingencyanalysis.com>

As general finance and economics references try www.finweb.com and
<http://econwpa.wustl.edu>

Many of the www references in the paper have related materials and you may want to browse around the sites in addition to the pulling the reference.

Answers to Sample Exam Questions

- 4) **Assuming random walk markets and normally distributed returns, if a one day VaR on an asset is 100,000 what would be the approximate VaR corresponding to holding this asset for a year?**
- a) 36,500,000
 - b) 24,500,000
 - c) 1,900,000
 - d) 1,600,000 >> the standard deviation over a time horizon t of a normally distributed variable with standard deviation of σ per unit time is given by $\sigma\sqrt{t}$. The number of trading days in the year is about 250 ($365 \times 5/7 - 10$ or so holidays). $\sqrt{250}=15.8$, which multiplied by 100,000 is about 1,600,000.
- 5) **What value is added by stress testing a portfolio as compared to just computing value at risk of the portfolio using the delta-nomal covariance approach?** VaR numbers are best as descriptions of the likely behavior of a portfolio under normal circumstances. Stress testing provides a sense of the sheer magnitude of potential losses caused by a large market move – VaR math simply does not capture “how bad it could possibly be”. Stress testing may also capture the impact of changes in the covariance matrix during a market break. For example when the stock market is collapsing ALL stocks go down and their correlation goes to 1.
- 6) **If a VaR on a portfolio is \$100,000 at 95% one tailed confidence level, one day holding period, how often should losses exceed \$100,000?**
- a) One in ten days
 - b) One in twenty days >> 5% of the time
 - c) Nineteen out of twenty days
 - d) Almost never
- 7) **Assume that losses in the previous question exceed 100,000 more often (than one in twenty days) what does that say about the VaR of the portfolio?** It may say nothing at all – we are attempting to infer information about the ex-ante VaR from the ex-post P/L sample. We may correctly infer that the VaR from our model is too low or we may incorrectly reject an accurate model (Type I error)⁵¹.

Answers to Risktek Risk Olympics™ Questions

<http://www.ozemail.com.au/~ristek/heat1.html> and <http://www.ozemail.com.au/~ristek/heat2.html>

Questions Copyright © 1996 Risk Technology Systems Pty Limited

⁵¹ Basle's Backtesting Framework, Tuna Amobi, Derivatives Week, September 2 1996 p5

- 1) Heat 1. Question 1 - the humble cash flow. A positive cash flow has a fixed face value. A sensitivity analysis is performed to measure the increase in its present value from a sudden 1% decrease in yields. The same analysis is performed on a second cash flow, which is identical in all respects, except it has a longer duration. Assume a flat yield curve at say, 10% semi-annual. The change in present value of the longer term cash flow:
- A. Is necessarily greater than for the shorter term cash flow.
 - B. Is necessarily less than for the shorter term cash flow.
 - C. May be greater or lesser than the shorter term cash flow. *Assuming continuous compounding the change in present value depends on the ratio of the product of the duration and the discount factor for each cash flow. You can construct either case.*
 - D. The present values actually decrease, not increase.
- 2) Question 2 - the age of VaR. A risk manager is presented with a Value at Risk (VaR) measure of a complex diversified derivatives portfolio. The VaR is a single statistic intended to express the maximum one day loss within a 99% confidence limit. The VaR indicates the potential for loss is too high. The risk manager must formulate directly an effective hedge or hedges to reduce the portfolio's market risk exposure.
- A. The risk manager can do this simply by knowing the VaR.
 - B. The asset types, not specific instrument details, in the portfolio must also be known.
 - C. The asset types, spot and forward market rates must also be known.
 - D. The risk manager cannot determine the required hedge(s) given only the VaR, the relevant asset types and market data. *The answer to this one depends on your definition of effective—I would only have taken B if the question used language like crude or approximate.*
- 3) Question 3 - exposed to what? A corporate enters into a long term FX forward to hedge a contingent liability. The hedge may need to be closed and settled at any time. The present value of the profit or loss on the hedge is, in general, exposed to:
- A. Changes in the spot (exchange) rate only.
 - B. Changes in the spot rate and the domestic interest rate only.
 - C. Changes in the spot rate and the interest rates of both currencies only.
 - D. Changes in the spot rate, both interest rates as well as time lapse. *The forward rate is a function of spot and both interest rates and time to maturity. Passage of time changes both the forward and the discounting of the PV back to today.*
 - E. No market rates once a reverse FX forward with equal face value and maturity is entered into. *Note that this is not true because an off market reverse even if it offsets the FX cash flow will still leave a cash flow/PV in the base currency that has interest rate exposure.*
- 4) Question 4 - is this real? A trader performs a stress test on an interest rate related derivatives portfolio using uniform yield shifts with time and all other economic parameters unchanged. The portfolio seems to profit when yields increase. It also appears to profit when yields decrease.
- A. This suggests the net delta is zero and gamma is zero or positive. *Yes, and see last question below.*

- B. This suggests the stress analysis software is defective - this result violates arbitrage theory.
- C. This suggests the portfolio will tend to gain value as time lapses.
- D. This effect can occur only if options are present in the portfolio. *No – you can be delta neutral to a parallel shift in the yield curve and long gamma from a bond position (substitute duration and convexity for delta and gamma).*
- 5) Question 5 - does it matter? An exporter can use either FX futures or forwards to hedge a long dated FX exposure. Both futures and forwards are available with expiries coinciding with the payment date. Domestic and foreign currency interest rates exceed 10%. To achieve near risk equivalence, the face value of the futures contracts:
- A. Should be greater than an equivalent forward contract.
- B. Should be less than a forward contract.
- C. Should be the same as a forward contract.
- D. May be greater or less than a forward, depending on which currency has the higher interest rate.
- 6) Question 6 - a fundamental thing The critical dimension, that is the fundamental unit of measure, of a simple interest rate is:
- A. A unit of the currency of the underlying instrument (e.g. dollars of US dollar based instruments).
- B. It has no dimension, interest rates are dimensionless numbers.
- C. $1/\text{Time}$. *Multiply rate by time to get return. Note that you have to multiply by notional to get dollars i.e. dollars are not part of the fundamental units.*
- D. Time.
- E. %
- 7) Question 7 - the business of banking. A banker uses interest rate derivatives to manage a fixed rate mortgage portfolio. The Bank's board views exposure management based on rate forecasts as speculative and prohibits it. A careful stress analysis indicates that the portfolio contains risk holes, it is exposed to yield curve changes. The risk manager then formulates and executes a hedge parcel that reduces significantly market risk exposure. With less market risk the expected or average profit to the bank:
- A. Increases.
- B. Decreases, but only by the amount of the hedge parcel transaction costs.
- C. Decreases significantly since less risk means less return.
- D. Stays the same.
- E. May increase or decrease. *Even if we assume efficient markets the market price of risk could be positive or negative so while the hedged portfolio should only earn the risk free rate the unhedged portfolio might earn more or less.*
- 8) Heat 2. Question 1 - Bonds and forwards. A three year 10% coupon bond (fixed interest rate instrument) is trading at par. A risk manager determines the sensitivity of the bond's price to a 1 basis

point change in various 90 day forwards. In general, which forward rate should the risk manager expect the bond's price to be most sensitive to?

- A. A near dated forward. *The near dated forward will affect all the cash flows of the bond and produce the greatest PV change.*
 - B. A forward with maturity close to the bond's duration.
 - C. A forward with maturity coinciding with the bond's maturity. *No but this would be the right answer for a spot rate.*
 - D. The manager is unlikely to know until a thorough stress analysis is performed.
- 9) Question 2 - The VaR thing. A risk manager is presented with a Value at Risk (VaR) measure for two derivatives portfolios, Port1 and Port2. The VaR is a single statistic intended to express the maximum one day loss within a 99% confidence limit. Both VaRs are within policy limits. But the VaR of Port1 is much less than the VaR of Port2. The risk manager must decide which portfolio is more exposed to a sudden large jump in market rates - outliers exceeding the 99% confidence limit.
- A. Port1 is definitely safer than Port2.
 - B. Port2 is definitely safer than Port1.
 - C. Port1 may be more exposed than Port2. *About all you can say from the info above – there is not reason portfolio 1 could not have a very large exposure for a move greater than that implied by 99% even if it has a small VaR.*
 - D. Neither portfolio can produce unexpected results.
- 10) Question 3 - FRA exposure. A risk manager analyses the present value sensitivity of a single 90 day borrowers FRA to isolated movements in various forward rates. A borrowers FRA is used to protect against rising forward interest rates. In present value terms, in general, the risk manager should expect the FRA:
- A. To always profit from a rising forward rate.
 - B. To incur a profit or loss from a rising forward rate. *Any FRA maps to a long and a short cash flow so you should be able to construct scenarios that show either profit or loss depending on the maturity of the forward rate relative to the FRA.*
 - C. To be insensitive to time lapse.
 - D. To be equivalent to a pure discount bond.
- 11) Question 4 - Foreign currency loans. A corporate enters into a long term foreign currency loan and immediately converts the proceeds to their domestic currency. No exchange rate hedge is used. At the time to repay the principal, the foreign currency required is purchased on the spot exchange market using domestic currency. The amount of domestic currency required to repay the principal may be:
- A. Up to 50% more than the amount of domestic currency originally received.
 - B. Up to 100% more than the amount originally received.
 - C. Up to 200% more than the amount originally received.
 - D. More than 200% of the amount originally received. *I don't see how the devaluation of the domestic currency can be bounded.*

12) Question 5 - Gamma zero? In Heat 1 the following question appeared: "A trader performs a stress test on an interest rate related derivatives portfolio using uniform yield shifts with time and all other economic parameters unchanged. The portfolio appears to profit when yields increase. It also appears to profit when yields decrease". The most popular answer which a minority of players selected was choice A "This suggests the net delta is zero and gamma is zero or positive". The Games Master also chose A. However, a number of astute people commented that the gamma could not be zero.

A. The Game producers made a slight error - the gamma cannot be zero.

B. The Game producers made a slight error - the gamma can also be negative.

C. The Game producers were careful to include all possibilities because some exotic instruments can produce this effect. *If the delta and the gamma were zero then you would have to look for the fourth derivative to get a symmetric profit. From a Taylor series perspective why not? The higher terms don't enter into the PDE because it only addresses infinitesimal changes in rates.*

D. The Game producers were careful to include all possibilities because this effect can occur even with rather elementary instruments such as simple cash flow portfolios.