Morgan Stanley

Buy Side Risk Management

Kenneth Winston

March, 2006

The key is not to predict the future, but to be prepared for it. (Pericles, 495-429BC)

Summary

There are many forms of currently active risk management. Even within the financial services industry, risk management means very different things in different areas. This article surveys current practices in one area of financial services, buy side risk management. This form of risk management has experienced rapid growth and increased interest recently. In part 1, we discuss differences between buy side and sell side risk management. In part 2, we look at some current practices in buy side risk management, where the gap between theory and practice can be large. In part 3, we note that long-short portfolios present special problems and present a framework for evaluating them.

Kenneth Winston
Chief Risk Officer
Morgan Stanley Investment Management
1221 Avenue of the Americas
New York, NY 10020
212.762.4502

1. Business models and risk management implications

1A. Financial services

We can put financial sector activities in one or more of the following categories:

- Services in which firm capital is put at risk;[1]
- Services consisting of decision-making on behalf of a client in order to put the client's capital at risk;
- Services that pool capital, such as banks and insurance companies;
- Services such as custody in which putting capital at risk is not the main activity.

For our purposes, the **sell side** consists of those firms (or those sections of large diversified firms) that effect transactions ordered by clients. These firms may put firm capital at risk to (1) facilitate the completion of such transactions; and/or (2) make a direct profit.

The **buy side** consists of firms that are hired by clients to exercise discretion as to how to put the client's capital at risk in order to attain a financial goal or reward. We include pension funds and endowments in this category since in effect the group that invests the money is servicing clients – the stakeholder(s) in the pool of money they invest.

Aggregation services such as banks and insurance companies collect large numbers of comparatively small pools of capital (insurance premiums, bank accounts) into small numbers of large pools of capital. They provide services that take advantage of the pooling effect.

The financial services industry is large and diverse, so there are many activities that have aspects of more than one of our categories. Concentrating on the first two, which center on what capital is put at risk, there is a spectrum of activities in which what we have called the sell side is at one extreme and what we have called the buy side is at another. The spectrum, however, is not uniformly populated, and most activities cluster close to one end or the other. Hedge funds are an interesting hybrid case, since hedge fund managers (who often come from sell side desks) trade in instruments and use techniques usually associated with the sell side, but they do so with client capital sometimes invested alongside their own capital.

[1] By virtue of being in business every company puts firm capital at risk – a restaurant, for example, can do well or poorly, and thereby gain or lose firm capital. We omit that natural aspect of business from our definitions.

1B. Risk management definition

Risk is imperfect knowledge about the outcomes of an action. This definition admits of the possibility that an action can be risky for one party but not risky for another. Most dictionary definitions of risk involve loss or harm, but if loss or harm is certain then managing it is either futile or straightforward. A vampire knows the sun will come up tomorrow and knows he will die if exposed to it, so he either gets back in his coffin before sunup or dies. It is only when there is imperfect knowledge that risk management becomes nontrivial.

Imperfect knowledge includes what Frank Knight [Knight, 1921] defined as risk: a set of outcomes with known probabilities not forming a point distribution. It also includes Knightian uncertainty. In Keynes's [Keynes, 1937] formulation:

> By `uncertain' knowledge, let me explain, I do not mean merely to distinguish what is known for certain from what is only probable. The game of roulette is not subject, in this sense, to uncertainty...The sense in which I am using the term is that in which the prospect of a European war is uncertain, or the price of copper and the rate of interest twenty years hence...About these matters there is no scientific basis on which to form any calculable probability whatever. We simply do not know.

We don't know what will happen on any particular turn of the roulette wheel, but we know enough about the distribution of outcomes so that variations in casino profits are affected mainly by how well they advertise, their amenities, their ambiance – and very little by unexpected events at the roulette wheel.

Knightian uncertainty is more imperfect than Knightian risk. The former can include a set of known outcomes with only partial knowledge of outcome probabilities, but it can also include the situation where we know one fact: the number of possible outcomes is greater than one. Thus we may not be able to quantify risk at all – we may just know that we don't know.

All forms of risk management start with analysis of the activity being managed. If we truly know only the minimum – there is more than one outcome - there is nothing we can do to manage the activity's risk. We don't necessarily need to be at the other extreme – Knightian risk where we know each outcome and an associated probability, but we need to have some information in order to get traction. We have to be in a middle ground between complete ignorance and complete certainty about the future.

Buy side risk management takes the analysis and turns it into actions that

- Profit from anticipated outcomes; and
- Limit the cost of unanticipated outcomes.

Sell side risk management takes the analysis and turns it into actions that

- Avoid uncertain outcomes; and
- Lock in a certain or near-certain spread.

The buy side's business model is to profit from successfully taking risk, while the sell side's business model is to profit from avoiding risk.[2]

1C. Sell-side Sharpe ratios

Many on the sell side object to this formulation. They protest that, rather than collecting a fee for interposing themselves between market participants and perhaps providing a little capital as a lubricant for the engines of commerce, they get paid for subjecting firm capital to variable outcomes. We are, most on the sell side would say, not mere intermediaries extracting a toll from our clients because we belong to a small cadre of firms that provide access to markets. Instead, they say, we are skilled risk takers who expose ourselves to danger in order to keep the capital markets functioning.

Is this true? A cynical theory is put forth by Eric Falkenstein [Falkenstein 2003]. He calls it "alpha deception:"

> When I was a risk manager at a large bank I remember many traders who portrayed themselves as speculators, price takers who made money using market savvy, as opposed to buying for a fraction less than they sold all day. In reality they were liquidity providers with privileged access to retail trading flow, but through bullying and their use of an imposing trade vernacular, they were able to propagate the belief they where big alpha alchemists worthy of hefty compensation... For many activities in large institutions there is no alpha and there is little conventional risk, just a predictable cash flow due to franchise value.

While it may be true that some individuals exaggerate their skill (as individuals do in any business), we do not ascribe this motive or model to the entire industry. We can, however, look at the ratio of reward to risk on the sell side. Consider Table 1C.1:

Table 1C.1: Sales&Trading Revenue (2005) vs. Trading VaR (99% 1-day)

Company	Sales&Trading Revenue	Trading VaR
Goldman Sachs	15,927	113
JP Morgan	9,041	93
Citigroup	12,673	90
UBS	12,009	85
Deutsche Bank	13,251	84
Morgan Stanley	11,586	81
Merrill Lynch	10,705	70
Credit Suisse	8,195	53
Lehman	9,807	48

Source: Company filings, Morgan Stanley compilation. Figures in $MM.

[2] As we noted at the end of section 1A, in reality such a pure distinction does not exist. This is an exaggeration in order to make a point.

These figures are derived from the purified sell side portions of large diversified financial firms. To translate these figures to an annualized Sharpe ratio, we put the Sales & Trading Revenues in the numerator. To obtain a standard deviation for the denominator, we observe that the inverse normal multiplier $N^{-1}(.99)=2.33$. However, to account for fat tails a rough rule of thumb is to change this to 3.5.[3] We therefore multiply the Trading VaR figures given in the table by the square root of 252 (to annualize from daily) and divide by 3.5 (to convert to standard deviation).

This gives us:

Figure 1C.1: Trading annualized Sharpe ratios from Table 1C.1

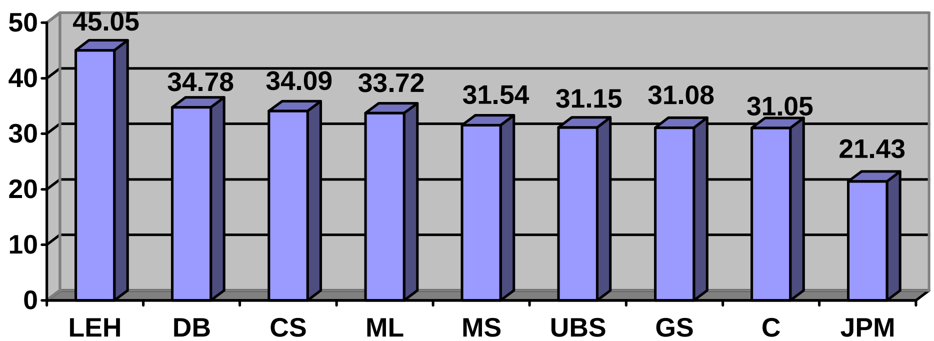

Sharpe ratios like this are not seen on the buy side. There are many methodological criticisms one can make here, including riding roughshod over issues like serial correlation and distributions violating the VaR=3.5 standard deviations heuristic. One obvious problem is that due to data availability we put revenues in the numerator rather than profits. On the other hand we know that trading profits are generally on the order of one third to one half of trading revenues, so the Sharpe ratios would still be very large.

A rough check can be obtained by looking at 20 quarters of PBT (profit before taxes) for a slightly different set of companies. These are full company PBT's, not just the sales & trading portion. For the numerator we take the arithmetic average PBT over the 20 quarters and multiply by 4 to annualize. For the denominator we take the standard deviation of the PBTs over the 20 quarters and multiply by 2 to annualize. That gives the following:

[3] [Litterman 1996] notes "Given the non-normality of returns that we find in most financial markets, we [Goldman Sachs] use as a rule of thumb the assumption that four-standard deviation events in financial markets happen approximately once per year." This is comparable to the tail-fattening assumption that 99% 1-day VaR represents 3.5 standard deviations. Over longer periods (as we will see below), leptokurtosis becomes less severe, so for example [deBever et. al. 2000] use 2.6 instead of 2.33 to fatten tails of monthly distributions. Given the magnitude of the Sharpe ratios noted here it makes no qualitative difference to our conclusions.

Figure 1C.2: PBT annualized Sharpe ratios

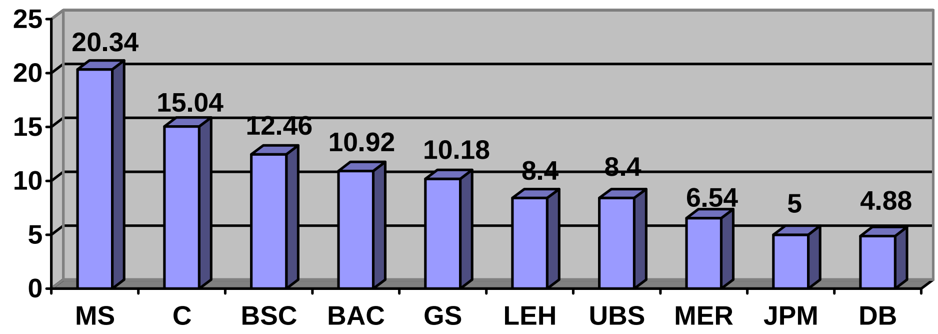

Source: Company filings, Morgan Stanley compilation. Based on 5 years of quarterly data ending 3Q2005, annualized.

The worst of the group, DeutscheBank, posts a 4.88 Sharpe ratio. An asset manager posting a 4.88 Sharpe ratio over five years would be in the top percentile if he or she existed. Morgan Stanley's 20.34 Sharpe ratio is simply not seen over five-year periods in investment portfolios.

Of course this is not a fair comparison. The numerator includes revenues (Figure 1C.1) or profits (Figure 1C.2) that are a steady stream due to franchise value and are not part of the risk-taking activity of the firm, whereas an asset manager's Sharpe ratio reflects in the numerator the rewards from risk-taking activity and in the denominator the extent of risk taken.

But that is our point. It appears from these numbers that the risk-taking activity of sell-side firms is mixed with the collection of franchise rents for intermediary activities that are vastly more lucrative than the rewards to their risk-taking activity. Thus it would be natural that sell-side risk management would be aimed at preserving the ability to collect these rents, while buy-side risk management would be aimed at allowing profits from risk-taking.

1D. Degrees of precision

Many finance "quants" have PhDs or other advanced degrees in fields like physics, mathematics, and engineering. Although I have no statistics to back this up, it appears to me that the majority of quants involved in sell side risk modeling – much of it derivatives-related – were not originally trained in finance. In the course of getting their PhDs, many quants observe "The Unreasonable Effectiveness of Mathematics in the Natural Sciences" [Wigner 1960]. If you restrict your attention to derivatives, you might believe that mathematics is unreasonably effective in explaining financial phenomena. Models that price derivatives are generally more precise than models that price primary instruments.

This is a key gap between buy side and sell side risk management. Consider, for example, [Neville 2006], a *Risk* magazine article that discusses the differences between mark-to-model and mark-to-market. Neville informs us "…that not all pricing can be explained by theoretical constructs." This seems to be breaking news for the sell side, but the buy side has always found this painfully obvious. Near the beginning of a book on fixed income risk management, buy side risk managers [Golub and Tilman 2000] note that "The theoretical values of fixed income securities are often different from their actual prices observed in the market."

In this article, Riccardo Rebonato, global head of market risk and global head of the quantitative research team at Royal Bank of Scotland says "You don't need a model to tell you what the price of a Treasury will be, and if you want to trade a Treasury you do so based on your views of macro and supply and demand issues. It is a stretch to say that is analogous to the derivatives market."

Dr. Rebonato – his PhD is in physics – is the author of the book *Interest-Rate Option Models.* He is a distinguished risk manager who is no doubt more familiar with rate models like Black-Karacinski, Hull-White, and Heath-Jarrow-Morton than just about anyone. He is certainly familiar with the Taylor rule. I therefore asked him what he meant by saying "You don't need a model to tell you what the price of a Treasury will be." I pointed out that buy-side quants spend a lot of time developing models to do exactly what he said you don't need a model to do. He agreed that this was true, and explained that what he meant was that buy side models are not like sell side models: they have much more basis risk – a polite way of saying they don't work as well – and you can't form perfect hedges with them. Sell side models have a degree of precision that minimizes basis risk and allows near-perfect hedging.

This precision decays as the instruments get more complex – witness the meltdown in pricing models for the structured credit markets in April and May 2005. The relationships between tranches of structured credit products, usually priced with a normal copula model, broke down to the point where hedging lower tranches with delta-adjusted quantities of higher tranches produced counterintuitive results, losing money for many correlation traders.

[Neville 2006] tells us that "…the consensus opinion [sic] about what happened last year in the structured credit market is that the models should be cleared of most of the blame – what occurred was the result of market dynamics rather than model risk." Saying that the market is wrong but the model is right is a perennial problem. In 1994 Askin Capital Management failed as

> The market for CMOs had deteriorated to the point where CMOs were quoted with spreads of 10 percent, which is enormous. As one observer put it, "dealers may be obliged to make a quote, but not for fair economic value." Instead of using dealer quotes, Askin simply priced his funds according to his own valuation models… he was sanctioned by the Securities and Exchange Commission for misrepresenting the value of his funds. [Jorion 2001, p. 19]

The money made or lost due to market dynamics looks remarkably similar to money made or lost due to any other reason, and most buy-side managers recognize that the job of a market participant includes anticipating market dynamics, perhaps even with a model.

A sell-side risk manager could gain the belief that he or she, with the use of a precise model, can hedge away virtually all variability – that is, the sell side lives in a world of Knightian risk, or wants to. If market dynamics can be dismissed when they intrude, then the world is a more tractable place.

A buy-side risk manager is quickly disabused of this belief. We live in a world of Knightian uncertainty. This does not mean, however, that we have no knowledge of the future – only that we operate in an area where we cannot obtain as much precision as a sell side risk manager.

There is a continuum between certainty and imprecision in hedging activities:

No arbitrage occurs when we look at objects that must be worth a certain amount on a certain date. For example, a conversion (long underlying, short call, long put) must be worth the common strike price at the common expiration date (absent some minor adjustments for the possibility of early exercise). If we know with certainty what a combination of positions will be worth on a given date, then that combination must be equal to a risk-free instrument with the same characteristics – otherwise there is a risk-free arbitrage. Pricing that relies on no-arbitrage is usually very tight – for example, since US options exchanges have matured, it is very difficult to find material conversion mispricings on them. There are well known limits to no-arbitrage (perversely usually called limits to arbitrage) involving liquidity, noise trading, short sales constraints, and other factors.

Convergence is a little less strict form of no arbitrage. The difference between on-the-run and off-the-run Treasurys is an example of a convergence trade. The US Treasury issues 10 year notes every three months. The just-issued 10-years are called the on-the-runs, and the ones that have been issued for more than a month are called off-the-runs. There is more liquidity in the on-the-runs, and other reasons why investors prefer them. However, over time the difference between the just-issued 10-year and the 9 year and 11 month issue that was just-issued last month must become minimal. Therefore there is convergence, but an investor ignores the disruptions to convergence at his or her own peril (as Long Term Capital Management learned).

Model parity occurs when a model tells us there is reason to believe that two things are the same. Delta-hedging (and more generally option replication using Greeks) is an example of model parity. If we believe Black-Scholes, then varying a stock-cash position according to the Black-Scholes delta should give us the same results as a call option. This requires us to believe the Black-Scholes model and get through market frictions.

High correlation relies on past or projected correlations between different items to form an effective hedge. A portfolio of US corporate credits may be assumed to be correlated with a credit default swap on a basket of high grade credits, together with a fixed/floating swap and cash. The actual correlation between the portfolio and the hedging instruments will vary depending on a number of things, including how spreads on the credits in the portfolio differ from those in the basket; changes in the shape of the yield curve; and differing default rates.

The basis risks of hedges that rely on no arbitrage and convergence are typically in the single basis point range. For example, a conversion (using at-the-money first month options) displays about 5bps/year tracking error to appropriate maturity Treasurys. In Table 1D.1 below we see an index fund run by MSIM that produces 20bps of tracking error – the main cause of which is different pricing between MSCI, the index provider, and the custodian bank in which the actual securities are housed. This is essentially a convergence item because over time the prices of the items in the index fund have to converge between the two pricing sources.

Table 1D.1 – basis risks

Long	Short	Basis risk
Single equity security	long call plus short put	5bps
10 year on-the-run	10 year off-the-run	5-7bps
MSCI World index fund	MSCI World index	20bps
Vanguard S&P 500 Index Fund (VFINX)	S&P 500 Index	72bps
MS SICAV Emerging Markets Debt Fund	JP Morgan EMBI Global Index	90bps
SPDRs S&P 500 ETFs	S&P 500 Index	130bps
Actively managed US credit portfolio (569bps absolute std dev)	US Treasury futures	213bps
Diversified group of US equity seed capital investments (1249bps absolute std dev)	S&P 500, Midcap, Russell 2000, and NASDAQ futures	277bps
Van Kampen Comstock Fund (US equity)	Russell 1000 Value	359bps
Vanguard Small-Cap VIPERS	MSCI US 1750 Small-Cap Index	482bps
MS Japan Small Cap Fund	MSCI Japan Small Cap Index	1058bps

Basis risks are in annualized tracking error terms, except OTR Treasurys which give the month 1 annualized yield spread. This table is provided for illustrative purposes only and should not be deemed a recommendation.

2. Some current practices in buy side risk management

2A. Separation of duties

An analogy used by many business managers is that an asset management company is like a manufacturing business. Portfolio managers take raw materials (risk and information), put it through their factory, and manufacture streams of returns with desired characteristics. There are sales, distribution, and servicing functions connected with the manufacturing effort.[4]

In any company, success is not possible unless duties are carefully matched with skills. Too much specialization may lead to overly narrow focus, while too little may lead to the inability to develop expertise in any one area. Despite the dictum from Pericles quoted at the beginning of this article, asset managers are in the business of predicting the future, but it has become clear that different skills are required to predict different aspects of the future. Several kinds of prediction can be delineated:

- Predictions of the direction and magnitude of asset price movements;
- Predictions of overall market volatility in various markets;
- Predictions of relative volatility (portfolio or security to benchmark); and
- Predictions of relationships between parameters (e.g. response of a portfolio to changes in interest rates)

In many asset management organizations, the first item is the province of portfolio managers, while the last three are done by an independent risk management group.

Why this separation? Asset managers have always known that risk and reward were inextricably linked. All of these predictions must be unified to produce a final product. Ultimately, we must implicitly or explicitly maximize a function that looks like

$$Return(\vec{w},\vec{m}) + f(Variance(\vec{w},\vec{m})) + g(\vec{H}(\vec{w},\vec{m})) \qquad (2A.1)$$

over possible weight vectors w, subject to constraints. The vector m contains current and past information on each security, the markets, and other relevant variables describing the state of the world. The functions capture our predictions of arithmetic mean return, variance, and higher moments. When the portfolio consists of simple instruments with little optionality (either explicit or implicit through dynamic option replication), then the vector function H embodying the higher moments is often ignored. The scalar functions f() and g() indicate the relative importance of variance and higher moments, respectively. f() is usually negative.

The objective function is respecified to allow parameter reduction as follows:

$$\begin{aligned} &Return(Factor_{ret}(\vec{m}), Char_{ret}(\vec{w},\vec{m})) + f(Variance(Factor_{var}(\vec{m}), Char_{var}(\vec{w},\vec{m}))) \\ &g(\vec{H}(Factor_{higher}(\vec{m}), Char_{higher}(\vec{w},\vec{m}))) \end{aligned} \qquad (2A.2)$$

[4] Portfolio managers generally detest the plebian implications of being mere factory workers. Some insist that a better analogy is that an asset management company is more like a hospital, in which the portfolio managers are skilled surgeons.

Each component is a function of two parts. The Factor vector (or matrix) takes the state of the world in m and produces behaviors (returns, variances, correlations) of a small relevant number of factors. This vector simply describes the expected behaviors of a parsimonious set of factors without reference to the contents of the portfolio. The Char(acteristics) vector gives the characteristics of the portfolio w that indicate the reliance on the factor variables. The components (return, variance, higher moments) could take the dot product of the characteristics vector with the factor behavior vector; they could use a quadratic form multiplying the characteristics vector into the factor matrix; or they could use a structure like a decision tree to put them together.

For example, we might assume we can predict return based only on effective duration, effective convexity, spread duration, and prepay duration. We would then have five components of the factor return vector and five characteristics. Or we might assume that we can predict the higher moments based only on five or six option Greeks. [Citigroup 2005], describing their YieldBook tracking error methodology, indicates that each sector of the fixed income markets has "…its own set of risks and exposures to changes in the marketplace such as the general level of interest rates, the volatility of rates, the spread the market demands for credit, liquidity, or hedging risks, and, for multi currency portfolios, exchange rate risk. The first step… is to identify those risks for the securities we are able to model."

In models like the Capital Asset Pricing Model, the Arbitrage Pricing Theory, and Fama-French, the return factors and the variance factors are the same except for idiosyncratic items – there are no unpriced systematic risk factors and there is only one riskless return factor. However over practical time periods, long-term equilibrium arguments are ineffective and many organizations pursue the return prediction function and the variance prediction function using very different methods.

(2A.2) embodies separation of duties and parameter reduction.[5] Many organizations have found that different skills and methods are useful in estimating higher moments of distributions than are useful in estimating the first moment. Traditional non-quantitative portfolio mangers may be unfamiliar with the quantitative techniques used to estimate higher moments, and even why this activity is useful. Quantitative portfolio managers will often direct more of their skills to estimating the first moment than the higher moments.[6]

We also find it is important to understand what degree of predictive power we can bring to the components of (2A.2). It is more difficult to predict absolute risk than it is to predict relative risk. Figure 2A.1 shows the VIX™ index – the option-implied volatility of the S&P 500.

[5] There may also be idiosyncratic components for each security, but we assume that these are either very simple in form or small in magnitude or both. If the idiosyncratic components are large and complex, there is no point to this formulation.
[6] Managers who trade volatility and correlation as assets are of course exceptions. In cases where managers are directly focused on estimating these statistics, a risk management group is useful if it can provide a second opinion with a different model, thereby providing some check against model risk.

Figure 2A.1: VIX and S&P 500

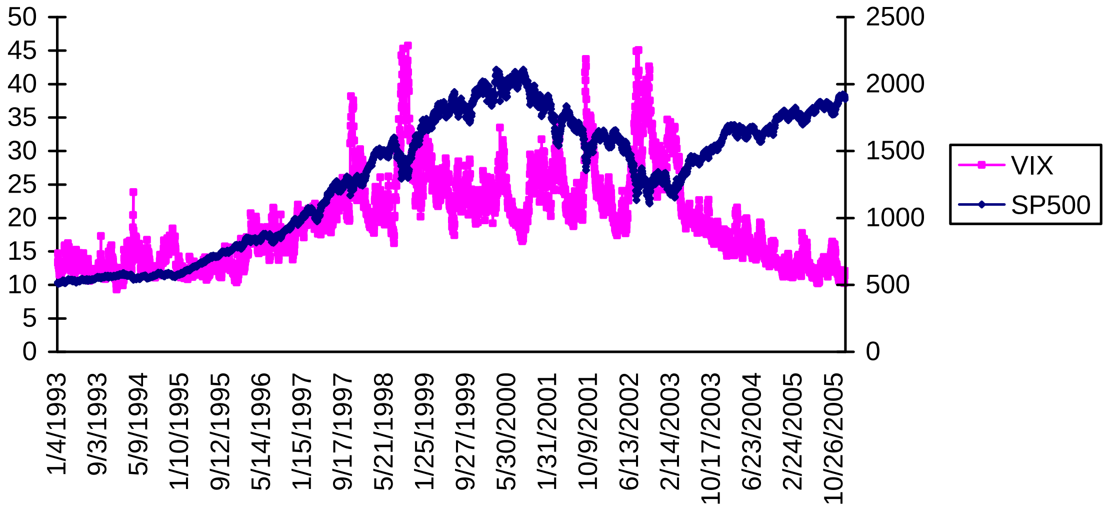

In section 2E below we will look in more detail at how well various methods predict S&P 500 volatility. We find that simply using the previous period's sample volatility to predict the next period's sample volatility does best – better than a GARCH(1,1) model; better than the VIX™, better than commercial risk models. This is more or less inevitable: if we could do much better then by trading variance swaps we could make as much money as we wanted. The fact that variance swap traders do not mint money indicates there is a natural upper limit on the ability to predict variance.

It is therefore useful to decompose variance predictions. If we are predicting absolute variance, we can write

$$Variance(Portfolio) = Variance(Market)\left[\frac{Variance(Portfolio)}{Variance(Market)}\right] \qquad (2A.3)$$

The shocks and patterns that affect market variance are largely different from the factors that affect the ratio of a diversified portfolio's variance to the market variance. Thus we may predict the two independently, or hold market variance out of our predictions, leaving it as a scenario variable

With tracking error, we can write

$$Variance(Portfolio - Benchmark) =$$
$$Variance(Benchmark)\left[\frac{Variance(Portfolio)}{Variance(Benchmark)} + (1-2\beta)\right] \qquad (2A.4)$$

The term in square brackets is more insulated from market volatility, so we can concentrate independently on calibrating the accuracy of our predictions of this relative term.

In addition to taking advantage of different skill sets, an independent buy side risk management group makes sense given the current structure of the asset management industry and the legal and regulatory environment in which it operates. Asset managers are fiduciaries, entrusted with other people's money. Fiduciaries can be held to a standard of prudence but not of perfection. Risk transparency – the clear ex ante understanding of all the risks a portfolio is being subjected to – allows asset managers to take more of the right kind of risk and to avoid taking the wrong kind of risk.

Risk transparency is the language that fiduciaries can use to communicate with their clients. The effects of their risk-taking efforts may not be apparent in the current net asset value, but by telling clients what risks are being taken, the manager can make sure that it is following the client's wishes. This doesn't mean the manager will make more correct directional calls, but it does mean the manager will be directing his or her efforts to making the kinds of directional calls the client wants.

2B. Risk Budgeting

Risk budgeting is the process by which statistics on investment risk are carefully evaluated and aligned with rewards. One interpretation of risk budgeting [Scherer 2004] is that it is simply portfolio optimization, where a manager uses a quantitative tool that "attempt[s] to trade off expected return against expected risk." Others mean by risk budgeting the allocation of a fixed and limited amount of risk to investment opportunities.

The risk budgeting process starts with a decomposition of risk. If our risk statistic (such as VaR, tracking error, absolute risk, beta) is k-homogenous (that is, $f(ax) = a^k f(x)$ for any constant a) then we have

$$f(x) = \sum_i \left(\frac{1}{k} \frac{\partial f(x)}{\partial x_i} \right) x_i = g'x \qquad \text{(2B.1)}$$

where g is the gradient as implicitly defined here. The vector x can be weights (or differential weights) of securities, or weights or loadings of factors such as those contemplated in (2A.2).

When tracking variance to a benchmark whose weight vector is b is produced by a quadratic form

$$TrkError(x)^2 = TrkVariance(x) = (x-b)'C(x-b) = g_b'(x-b)$$

it is 2-homogeneous; tracking error is 1-homogeneous[7]. The (2B.1) format breaks our statistic into a dot product of two vectors. This allows us to look at each component separately and to assess the marginal effect (g_i) of changes to the i^{th} variable. If we are estimating β to a benchmark, the problem is linear

$$\beta(x,b) = \sum \beta_i x_i = \sum \left(\frac{Cb}{b'Cb} \right) x_i$$

[7] Here our quantities are defined relative to benchmark, so for example g_b is the gradient C(x-b).

Combining with (2A.4), we can form tracking variance by multiplying an absolute and a relative part:

$$TrkVariance(x) = Variance(Benchmark)\left(1 + (h - 2\beta)'x\right)$$

where h is the absolute gradient divided by benchmark variance. This allows different estimation techniques to be used for the absolute and relative components.

Our statistics may be produced by Monte Carlo simulations, trees, or other non-closed-form methods. In this case we can approximate the gradient through differences

$$\frac{\partial f(x)}{\partial x_i} \cong \frac{f(x + \Delta x_i) - f(x - \Delta x_i)}{2\Delta x_i}$$

Here Δx_i is a small perturbation in the i^{th} variable. As nonlinearities become more important, this first-order approximation may not suffice.

A non-additive, but still informative, approach to risk statistics such as tracking error, VaR, and absolute risk is to form the incremental statistic

$$h_i(x) = f(x) - f(x - x_i)$$

where by "x-x_i" we mean the vector x with the i^{th} component zeroed out.[8] For example, if x is the vector of differential exposures to a benchmark, we might test what would happen to tracking error if we matched the benchmark's effective duration.

Table 2B.1 shows an example of a risk decomposition (of tracking error) for a global equity portfolio.

Table 2B.1 – Risk decomposition

Factor Name	Active Factor Sensitivities	Factor Volatility	Contrib to TE	Marginal Contrib to TE
MSIM World Momentum	0.007	13.563	0.019	2.920
MSIM World Size	0.079	8.783	0.184	2.326
MSIM World Value (PE)	0.023	8.249	0.005	0.223
MSCI IG Household & Personal Products	0.016	10.327	-0.046	-2.874
MSIM World Value (PCF)	-0.014	8.453	0.048	-3.358
MSCI IG Food, Beverage & Tobacco	0.005	10.658	-0.021	-3.791
MSCI IG Pharmaceuticals & Biotechnology	0.010	13.429	-0.048	-4.839
MSCI IG Health Care Equipment & Services	-0.007	12.617	0.037	-5.567
MSCI IG Utilities	0.025	13.066	-0.159	-6.428
MSCI IG Energy	-0.022	17.370	0.157	-7.150

The "active factor sensitivities" column is essentially what we have been calling x_i; the marginal contribution to TE column is the relative TE gradient component g_i. (To switch between tracking variance gradient and TE gradient, divide by overall TE.) The contribution to TE column is their product, $x_i * g_i$. If we sum all the contributions in this

[8] If there is a budget constraint we may renormalize x.

column (only ten out of a larger number are shown here) we will form tracking error f(x) as in (2B.1).

In our experience, the vast majority of professional money managers (excluding index funds) do not use optimization. Traditional portfolio managers, and even many portfolio managers who use quantitative techniques, do not produce expected return numbers for each security they follow. They may have a general idea of what they expect a particular security or a particular theme (overweight pharmaceuticals, e.g.) to do, but they rarely can produce a vector of expected returns. It is not possible for these managers to do full portfolio optimization.

In this case, risk budgeting becomes a process that allows us to bring to bear the discipline of quantitative methods in the context of fundamental or qualitative management. All agree that unrewarded risk is harmful. If we can remove a risk from the portfolio without affecting return, we move closer to the efficient frontier.

Thus the risk manager can produce a list of risks like Table 2B.1. A portfolio manager looks at this list and indicates which risks are associated with rewards. For example, the Table 2B.1 manager may have no opinion on size, but there is a large cap tilt to this portfolio that is contributing 18.4bps of tracking error. On the other hand, the manager may be content to take the 15.7bps of risk incurred by the underweight in energy.

Many risk managers find that when they first begin the analysis process they can find up to 10% unrewarded risk that can be removed from an equity portfolio with little cost to return.[9] Removing these risks without interfering with the return-generating process is often very beneficial.

After the obviously unrewarded risks have been removed, risk/reward tradeoffs must be made. The risk budgeting framework provides a way to do this without fully estimating all the parameters needed for portfolio optimization. Suppose we are evaluating portfolios on a relative-to-benchmark basis, so we have a portfolio excess (over benchmark) arithmetic mean return α; a tracking variance σ^2; and a risk tolerance λ with an objective function

$$\alpha - \lambda\sigma^2 / 2 .$$

The rate of change of this objective with respect to the i^{th} parameter is

$$\frac{\partial(\alpha - \lambda\sigma^2 / 2)}{\partial x_i} = \alpha_i - \lambda\sigma g_i$$

We may try to recover an implied λ from a reverse optimization, but this can be unreliable – among other things, we are assuming the starting portfolio is inefficient. Instead, in this less-than-full-optimization setting, the risk manager concentrates on providing good estimates of g_i to a portfolio manager and avoids estimating parameters that are more difficult to estimate.

[9] Fixed income portfolios usually do not have this kind of opportunity because the factors driving returns are smaller in number and more well known.

If g_i is negative, then increasing x_i will lead to decreased risk, so the manager knows that increasing exposure to any negative gradient item with non-negative expected α_i will increase efficiency. The manager does not have to know λ, σ, or anything about α_i except that it is expected to be non-negative. More information is necessary when α_i is expected to be negative along with g_i. Here the manager needs a rough idea of the size of α_i. If the ratio α_i/g_i is close to zero, then increasing x_i will probably improve the objective function. So, again without knowing λ, σ, or much about α_i, a manager can guess that when g_i is a very large (in magnitude) negative number, increasing x_i will improve efficiency.

This line of logic does not allow full optimization, but it allows decisions to be made about the most easily correctible inefficiencies without having to estimate parameters too closely – we need only have a reasonable sign and order of magnitude of the gradient g_i (provided by a risk manager), and some general information about expected α_i (provided by a portfolio manager). Some reasonable decisions would be:

Table 2B.2 – moves toward efficiency

Gradient g_i	Expected excess return α_i	Action
Negative	Non-negative	Increase exposure
Very large negative	Small negative	Increase exposure
Positive	Non-positive	Decrease exposure
Very large positive	Small positive	Decrease exposure

The decision grid in Table 2B.2 has some ramifications for the information in Table 2B.1 The manager is overweight exposure to the industry group "household and personal products, indicating a positive view on this industry. The marginal contribution to tracking error (gradient) is negative (-2.874). There is probably an opportunity for efficiency improvement here – the manager's overweight probably means a positive alpha is expected, so this exposure can be increased (negative gradient, non-negative α_i) without any more specific parameter estimation.

Even the least quantitative portfolio manager is keenly aware of his or her active factor tilts, although the units we express them in may be unfamiliar. But the rightmost columns of Table 2B.1, and the decision logic in Table 2B.2, allow efficiency improvements through the risk budgeting process. As the manager becomes more oriented to making specific alpha forecasts, more precise use of the process can be made, and ultimately a full portfolio optimization can be done.

A stringent form of risk budgeting insists that portfolios be managed with a fixed finite amount of risk – perhaps measured by value at risk or tracking error. In this interpretation, taking risk in one area may preclude risk being taken in another area even if both areas are expected to lead to reward. As we noted in (2A.3) and (2A.4) above, risks can be decomposed into a market risk component and a relative component. The same US equity portfolio will look riskier when the VIX™ is 30 than it does when the VIX™ is 10. What we do about that is part of our management decision making. We might conclude that the fluctuation in market volatility is temporary and take no action. Indeed, if the

VIX follows a GARCH(1,1) process, then unusually high or low volatility regimes will tend to relax back to the long-term average.

[deMarco and Petzel 2000] suggest that risk tolerance varies and depends on objectives and opportunities. They make an analogy between risk budgeting and capital budgeting. In capital budgeting, a project is taken on if it is net present value positive given the discount rate. deMarco and Petzel suggest that varying market risk is like varying the discount rate – NPV-positive projects may become unattractive with a higher discount rate, and risky investments that were attractive may no longer be. We must also factor in the potential for risk to change during the life of the "project," so a decision on making a risky investment now can change not because the investment itself changes, but because the environment changes. Taking on more risk in an unusually low risk environment in order to comply with an unchanging risk budget can backfire if risk snaps back up.

2C. Value at Risk

The Value at Risk (VaR) of a random scalar process X(t) is defined as follows:

$$\textit{If } \Pr(X(t) \geq v) = p, \ \textit{then } VaR(p,t) = v. \qquad \text{(2C.1)}$$

In other words, VaR is a quantile of a distribution. In Table 1C.1 above, we saw 99%, 1-day VaRs. Morgan Stanley, for example, expects to lose less that $81 million on 99 out of 100 trading days. Putting it another way, we expect to lose $81 million or more on 1 out of 100 trading days.

In making forward estimates of VaR, most practitioners follow the separation of duties discussed in section 2A above. Assuming X(t) is generated by a normal process with mean μ and standard deviation σ, we have

$$VaR(p,t) = \mu t - N^{-1}(p)\sigma\sqrt{t} \qquad \text{(2C.2)}$$

Here $N^{-1}(p)$ is the inverse cumulative normal function. This would require us to estimate μ. Most practitioners assume μ=0, or perhaps some non-forecast amount taking into account carry and rolldown on a fixed income portfolio. For daily calculations, μ is usually so small that it makes no material difference whether or not it is included, but over longer periods it can become substantial. None of the firms publishing the VaRs collected in Table 1C.1 used the firm's economists' projections, or any other directional projections, about interest rates or equity market direction over the coming year to change the center of the distribution.[10]

Figure 2C.1 shows the histogram (bars) and cumulative frequency (line) of S&P 500 daily log-returns from 1993 to 2005:

[10] An exception is [Rees 2000], who suggests that VaR is a superior measure (to tracking error) for fund managers because it can incorporate the first moment as well as higher moments of the distribution. Rees discusses factor methods to estimate the first moment but does not discuss the relative accuracy of first moment estimates compared to second moment estimates.

Figure 2C.1 - S&P 500 daily return histogram

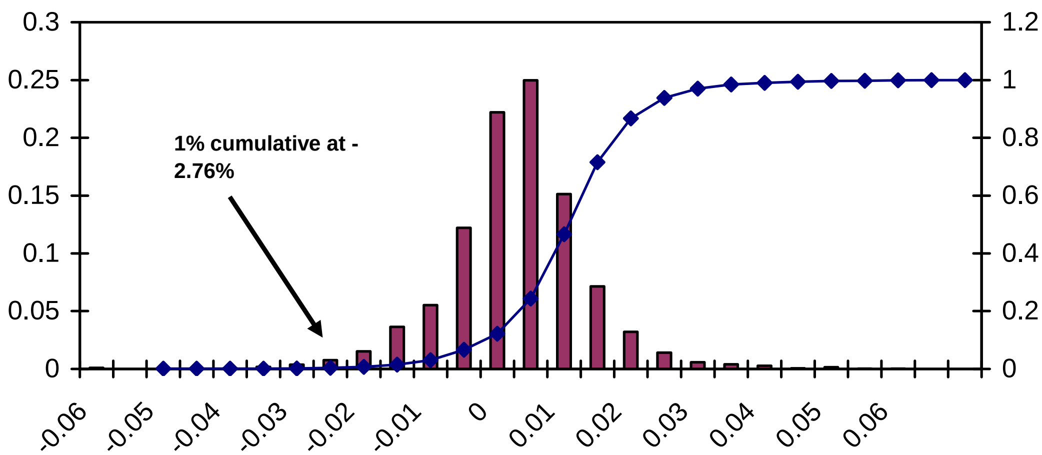

The arrow points to the place where the cumulative line gets up to 1%, namely at -2.76%, so the 99% 1-day VaR of the S&P 500 in this period is -2.76%.

In cases when there is no reason to expect significant optionality or kurtosis, normality or lognormality is assumed and VaR is simply a multiple of standard deviation. Since financial time series are usually leptokurtic [Mandlebrot 1963], this usually understates the true tail probability. For example, the standard deviation of the S&P 500 log-returns in the period shown above was σ=1.047% a day. The inverse of the cumulative normal distribution at 1% is $N^{-1}(.01) = -2.33$. Multiplying these gives -2.44%, which is less extreme than the fat-tailed actual distribution produces (-2.76%)[11]. The excess kurtosis was 3.87.

[Jorion 2001] credits Till Guldimann of JP Morgan with creating the term value at risk in the late 1980s. Since there is nothing new about quantiles of distributions, why is VaR so important and popular? VaR embodies a number of ideas and concepts other than computing a quantile of a distribution.

1. **Market values**. It appears that prior to VaR, banks and the sell side did not primarily consider variability of market values as the appropriate measurement of risk. Instead they looked at fluctuations in earnings, although not to the point of constructing a distribution and extracting statistics about this distribution. After VaR became popular, the sell side circled back and applied the other VaR techniques noted below to earnings to form a Daily Earning at Risk (DEaR) figure.

2. **Defining risk as a volatility-related statistic**. This idea goes back at least to Harry Markowitz in 1952, but was not widely used on the sell side until VaR became popular.

[11] The average of daily log-returns was 4bps, which is negligible in this calculation.

3. **Ability to express risk in common, simple terms**. The idea that currency risk, equity risk, interest rate risk, and others could all be expressed in common terms is an appealing feature of VaR. Markowitz had put things on common terms – standard deviation – in the 1950s, but some on the sell side found this statistical concept intimidating. The packaging of VaR as the answer to the question "how much can I lose in a bad period" (especially when expressed in dollar terms) made it an appealing way to communicate to corporate managers.

4. **Ability to aggregate different asset classes**. By constructing a covariance matrix, a multivariate (usually normal) distribution could be estimated and a single VaR number produced.

5. **Forward-looking estimates**. VaR was specified as an estimate of future behavior rather than a statistic drawn from past behavior.

All of these are desirable properties but it is hard for someone from the buy side, raised on Markowitz's 1952 insights and Barr Rosenberg's 1970's work on estimating covariance matrices to understand what was new about VaR in the late 1980's. As [deBever, Kozun, and Zvan 2000] noted, "VAR and the 'risk capital budgeting' metaphor pour old wine into new casks."

It appears that simple ignorance of previous work accounted for much of VaR's initial popularity. For example, Jorion uncritically quotes the director of Chrysler's pension plan as saying "We have a standard deviation for our total plan, but it only tells you what happened in the past. By contrast, VAR looks forward." More than 20 years earlier, [Rosenberg and Guy 1975] introduced a paper on estimating beta by saying "There has been a pronounced tendency to overlook the distinction between these two functions: historical evaluation and prediction. The purpose of this paper is to point out that estimates of systematic risk that are ideal for historical evaluation are far from ideal for predictive purposes, to explore criteria for optimal risk prediction, and to develop the methodology for risk prediction based upon fundamental concepts." [Rosenberg and Marathe 1976] went on to develop forward-looking methods of covariance matrix estimation. Nor was this a mere academic exercise – Rosenberg founded BARRA, which offered these concepts commercially.

JP Morgan started its Riskmetrics group to measure VaR, and eventually spun it off as a separate company. The five items we have listed above that apparently contributed to the enormous appeal of VaR on the sell side are either not substantive or not new. Nonetheless, the package of concepts was popular and a tremendous amount of effort went into VaR, so that eventually new and substantive work associated with VaR did emerge:

6. **Estimation techniques**. VaR was originally estimated with the delta-normal method, making it a simple multiple of crudely estimated standard deviation. This gave way to much more sophisticated Monte Carlo and historical sampling techniques.

7. **Focus on tail behavior**. Because of the need to explain the unusual portion of the distribution, attention to deviations from normality led to new techniques to get at rare but catastrophic events.

8. **Valuation methodologies**. Spurred by the focus on tail behavior, the huge body of work on valuing fixed income securities and derivatives was brought into the VaR framework.

Although VaR gained popularity because of its promise as a single, aggregated measure of risk (points 3 and 4 above), it soon was recognized that the weakest link in predicting an aggregate VaR was the set of assumptions made about interrelationships between different assets. Risk reports were soon disaggregated and decompositions of VaR into pieces that were more homogenous and thus more reliable were sought. Most VaR reports contain a decomposition that estimates components of VaR along the lines of (2B.1).

For most practitioners, VaR means the entire approach and body of associated work noted above, not just a quantile of a distribution. The sell side is clear in its conviction that VaR is the answer. For example, Subu Venkataraman is chief risk officer of Highbridge Capital Management, a hedge fund manager owned by JP Morgan. He is a highly respected risk manager with a sell side background. Recently he gave a talk [Venkataraman 2005] titled "Buy Side vs. Sell Side Risk: To what extent can buy-side managers integrate into their organizations the same risk management tools and processes used by the sell side?" He concludes that "There is no fundamental difference in the dimensions of risk that matter to buy versus sell side firms."

We need to examine this assumption carefully. The Basel Committee's risk framework seeks to protect the world financial system. In this context, tail risk is the right thing to focus on – if a rare but catastrophic event can so deplete a firm's capital that it is unable to honor its counterparty commitments, there could be a cascading meltdown that would destroy the world financial system and return us to living in animal skins in caves like a post-apocalyptic science fiction film. While many of us would welcome the ability to go business casual every day, the majority thinks this is an undesirable outcome with nearly infinite disutility. Thus intense focus on tail risk - VaR – was imposed by the Basel Accords on the largest multinational banks.

More recently, due to a European Union directive, the necessity to manage market, credit and soon operational risk using the Basel framework was extended to multinational non-bank financial services firms like Morgan Stanley, Goldman Sachs, Merrill Lynch, Lehman, and Bear Stearns. The concern is that the dealing desks will get in trouble and not be able to honor counterparty commitments, leading to animal skins. To prevent this, banks and sell side firms keep on hand a pool of regulatory capital that is sized according to their market risk (VaR); more recently credit risk and operational risk have been added. Keeping this pool of capital is onerous since it cannot be deployed in high-return projects – it must be kept in safe short-term instruments earning a low return. Thus firms seek to maximize their trading profits while minimizing their deadweight regulatory capital.

The sell side often express VaR in dollar terms as seen in Table 1C.1. In a typical month in 2005, Morgan Stanley grossed about $966 million, which even after expenses dwarfs its $81 million 1-day 99% VaR. So why wouldn't Morgan Stanley take more risk after it has built up some trading gains? Put another way, why do banks and securities dealers express VaR in dollar terms rather than in percentage terms? One reason is that these gains are not kept on hand; they are distributed to pay expenses and dividends, and to invest in other parts of the business. Only the pool of regulatory capital is retained.

The buy side, on the other hand, retains its capital in the pool of managed assets. Thus the amount of risk it can take is proportional to its assets and should be expressed in percentage terms rather than dollar terms. But using percentage VaR rather than dollar VaR still does not suffice for the buy side. While it is clearly a bad thing for a pool of assets managed for long term appreciation to have a particularly bad period, it is not (more or less literally) the end of the world. As we noted in our business model discussion in section 1C above, the sell side needs to stay in business in order to maintain the flow of franchise rents. The buy side needs to take risk in order to produce rewards. If the buy side focuses only on the tail risk embodied in VaR, it may not take the reward-generating risk it needs to.

This logic changes when we are managing for different objectives than long-term appreciation. When we have outcome-based portfolios where there is a full or partial guarantee of certain kinds of results, the shape of the outcome distribution can be dominated by implicit or explicit option behavior. When we have endowments or pension funds where surplus-at-risk is a key consideration, then VaR-like concepts may be appropriate. When we run long/short portfolios, the chance of bankruptcy or poor performance leading to being fired by the client becomes a dominant consideration. As we will see in Part 3 below, when we are running long/short portfolios a buy side model can contemplate absorbing barriers where we will be fired by a client if we deliver poor results. For long/short portfolios with leverage, this possibility can become quite large. But we must evaluate reward in the presence of a point at which we will be fired by a client, rather than focusing only on a constraint on tail risk.

A focus on tail risk over short horizons can become dominated by kurtosis. Asset managers need to find an appropriate strategy for an investor's horizon. Long-term wealth maximization (see e.g. [Thorp 1971]) presents a different problem than managing a daily trading book.

A single-day loss of 5% in an equity portfolio should not occur in anyone's lifetime if returns are lognormal, but in fact 5% single-day losses happen every few years. A client whose horizon is five years shouldn't be concerned about this, but if the client is focused too narrowly on short-term results, they can take money away or add money at the wrong times.

Kurtosis appears to attenuate over longer periods. Consider Table 2C.1:

Table 2C.1: S&P 500 Excess Kurtosis

Start Date	End Date	Daily	Weekly	Calendar	Calendar	Calendar

				Month	Quarter	Year
12/31/1949	12/31/1959	7.58	0.79	-0.51	-0.49	-0.50
12/31/1959	12/31/1969	10.36	1.64	0.25	3.67	-1.72
12/31/1969	12/31/1979	2.57	0.76	0.97	2.08	0.06
12/31/1979	12/31/1989	81.87	6.39	6.03	2.63	0.08
12/31/1989	12/31/1999	5.52	1.73	2.40	1.42	-1.22
12/31/1999	2/16/2006	2.49	1.93	0.08	-0.02	-0.61
12/31/1949	2/16/2006	36.83	3.38	2.46	2.17	-0.14

The sample variance of kurtosis is 24/n, so the decade annual observations are not significant. Over the 66 years shown in the last (shaded) row, kurtosis declines as the period increases.[12] This effect generally occurs in other equity markets as shown in Tables 2C.2 and 2C.3.

Table 2C.2: TOPIX Excess Kurtosis

Start Date	End Date	Daily	Weekly	Calendar Month	Calendar Quarter	Calendar Year
12/31/1949	12/31/1959	13.60	4.25	1.05	-0.02	0.95
12/31/1959	12/31/1969	2.90	2.05	0.68	-0.52	-1.53
12/31/1969	12/31/1979	13.71	4.88	1.21	-0.32	3.78
12/31/1979	12/31/1989	48.61	3.88	2.02	1.56	0.69
12/31/1989	12/31/1999	5.12	2.43	0.69	-0.06	2.93
12/31/1999	2/16/2006	2.04	0.57	-0.67	0.00	-1.52
12/31/1949	2/16/2006	11.28	3.13	1.20	0.39	1.86

Table 2C.3: FTSE 100 Excess Kurtosis

Start Date	End Date	Daily	Weekly	Calendar Month	Calendar Quarter	Calendar Year
1/3/1984	12/31/1989	24.53	12.41	6.25	3.82	0.39
12/31/1989	12/31/1999	2.37	1.25	0.15	0.18	0.75
12/31/1999	2/16/2006	3.28	6.96	0.66	0.56	-2.23
1/3/1984	2/16/2006	9.33	6.70	3.89	1.46	-0.12

It is clear from this that focus on one-day VaR when horizons are longer can be misleading. But even if we extend the horizon and define a monthly, quarterly, or annual VaR, the nature of the problem changes. At these longer periodicities, kurtosis is not as dominant.

Thus the actual VaR statistic – tail risk – especially when computed over short periods, may be inappropriate for the buy side. As we noted there are cases where it does apply – for example, in surplus risk calculations for pensions and endowments. In these circumstances, the focus on tail risk is appropriate in a way that is similar to the bank/sell side/Basel framework.

Currently there is little debate among the majority of institutional consumers of buy side management services – they want to look at tracking error to a benchmark and not VaR. At MSIM, a large number of our investment management agreements with clients, representing hundreds of billions of dollars of assets, specify that we will attempt to

[12] We note that if we include the 1920's in the US, this attenuation of leptokurtosis gets murkier.

manage the client's money to within specified target bands of tracking error to a benchmark. A minority specify VaR relative to a benchmark as a target to which we attempt to manage. This focus on tracking error is a result of the policy portfolio framework in which managers are hired to deliver results in a predetermined slot. For mandates that look at entire plan asset allocation solutions, these constraints are usually not present.

While the statistic itself may be inappropriate, the body of work surrounding VaR, especially our points 6 and 8 above, is substantive and useful. Many on the buy side have eagerly embraced these techniques, but careful buy side risk managers have applied to them in appropriate ways – adopting derivatives valuation techniques, Monte Carlo analysis, and historical sampling techniques, but using them to evaluate entire distributions or aspects of distributions that are more appropriate to the portfolio's objective.

2D. Scenario Analysis

The statement "In times of stress, the only things that go up are correlations" captures the fact that diversification often fails exactly when it is most needed. Since risk models that aggregate risks across different assets rely heavily on estimates of correlation, they have a tendency to break down when correlations undergo shocks. [Kritzman, Lowry, and Van Royen 2001] approach this problem by outlining a procedure to identify quiet and turbulent regimes using a switching regression model, and to estimate two sets of parameters for these two regimes. The regime switching approach still relies on a model to deal with times of stress, albeit a more flexible one.

More generally, scenario analysis can be used as one method of dealing with failure in risk models. In the same way we used partial risk budgeting in section 2B above to deal with partial information, scenario analysis is a more forensic than prescriptive technique that allows us to obtain some knowledge without having a full model. It looks at the consequences of certain future events without attempting to determine their likelihood.

This has two benefits: one is simply determining whether a portfolio will react as desired. In many cases, instruments and portfolios are so complex that a portfolio manager cannot determine how they will respond to economic events. A large portfolio containing a lot of optionality can be beyond the ability of the human mind to handle. A computer is needed to handle the vast amounts of data and calculations that determine how a portfolio might react to changes in inputs such as yield curves, credit spreads, volatility surfaces, currencies, equity market factors, and commodity prices. In the framework of (2A.2), we may be able to reduce portfolio return to a function of a small number of known factors. It may be difficult to estimate the parameters of the multivariate distribution generating those parameters, but we can avoid that problem by simply specifying some future values and seeing how the portfolio behaves. We can, for example, try a number of possible future prices of oil and see how the portfolio behaves at each price level. A manager trying to position a portfolio to benefit from increasing oil prices might be surprised to find that the portfolio barely reacts to higher oil prices.

A second benefit of scenario analysis is the discovery of unexpected behavior in extreme states of the world. For example, in 1995 the Derivative Policy Group (a consortium of large derivatives dealers) recommended that derivatives and portfolios containing derivatives be priced under the following extreme scenarios:

- Parallel yield curve shift of ±100bps
- Yield curve twist of ±25bps
- Equity index changes of ±10%
- Currency changes of ±6%
- Volatility changes of ±20%

Nonlinear behavior that is not apparent in the central part of a distribution that is likely to be part of any past sample can be made apparent under these scenarios. If there is a bad reaction to a scenario, a manager can make a risk/reward assessment of how much it would cost to reposition the portfolio to have a better reaction, compared to the likelihood of the scenario.

A simple scenario analysis run with Citigroup Yield Book software on the Citi Broad Investment Grade index at the end of February, 2003 shows the results of a full revaluation of the portfolio under -100, -50, 0, 50, and 100bps instantaneous yield curve shifts. These are compared with the calculated response

$$\text{Return} = -(\text{Duration})*(\text{Shift}) + ½*(\text{Convexity})*(\text{Shift})^2 + \text{Carry}$$

Table 2D.1 – Parallel yield curve shifts, instantaneous, Citi BIG (2/28/2003)

Effective Duration	3.986				
Effective Convexity	0.096				
Instantaneous Parallel Shift (bps)	-100	-50	0	50	100
Rate of Return % (1-month scenario)	4.02%	2.12%	0.32%	-1.53%	-3.54%
Rate of Return % (duration/convexity/carry)	4.31%	2.31%	0.32%	-1.67%	-3.67%

Another form of behavior testing through scenario analysis is motivated by Santayana's warning that those who do not learn from history are doomed to repeat it. Historical catastrophes – the Russian debt crisis, Long Term Capital's demise, September 11, 2001 – are analyzed and the responses of important parameters (yield curve, volatility surfaces, etc.) during the event are collected. They are then replayed through scenario analysis on the current portfolio. As we noted, this technique is forensic – the manager uses this information as part of a qualitative assessment of the portfolio rather than as part of a fully specified model. Nonetheless, we have seen clients asking that portfolios be able to withstand a suite of historical disaster scenarios with no more than some prespecified loss. An advantage of using historical scenarios is that reality has already translated the event – say the Russian debt crisis – into changes in key economic variables, removing the need for a human analyst to guess how these variables will behave in an actual similar situation.

An example of scenario analysis where human analysts did guess how to translate projected situations into projected variable changes was produced in the runup to the 2003 Iraq invasion. Citigroup's Yield Book group (apparently using Citigroup's economists' projections) defined the following scenarios for the three months forward from March 10, 2003:

Table 2D.3 – Narrative Iraq war scenarios, March 10, 2003 (source Citigroup Yield Book)

Short War No Recovery	Plus - short war; Minus - economy cannot shake off the problems, dismal corporate earnings, high oil price, layoffs. No correction in yields. No rate cut from Fed feeding enough stimulus to turn econ around. 2s settle in the 1.6% and 5s to underperform w/yields closer to 3%. 10s and 30s follow suit, backing up 30 bps and 10 bps as curve flattens. No change in swap spread and vol will decline in response to resolution of the biggest near-term uncertainty
Short War Recovery	Short war with confidence going through the roof, biz spending, Tsy sell off and vol falls like a rock, 2s jump past 2% and 5s crushed to 3.5%. Curve flattened from 5s on out. 10s shoot to 4.25%, bonds just shy 5%, Not much change in spread with widening pressure countered by upswing in Tsy issuance.
Long War Equity Crash	Skyrocketed oil price, no confidence, no biz expansion, the bottom falls out of equity. Fed cuts rates by 50bps, the relentless Tsy rally , curve steepens across all spectrum. no econ turnaround, 2s dip below 1%, 5s to 2% 10 to 3.3%, bond in mid 4%. 10-yr swap sprd blow out to 80bps, yield vol spikes reflecting declining yld and high bp vol
Long War Malaise	High oil price, biz gun shy, but equity does not fall out. Fed cautious w/ 1cut. Tsy bull flatten to intermediate as investors reach out longer, 5s and 10s yields decline by 30 bps and outperform the rest of curve. Sprds tightens as war costs fuel the prospects of ever-rsiing Tsy supply further. Yld Vol climbs due to decline in rates.
Forward No Change	Ongoing UN inspection with no war, yields stay at the same levels implied for forwards.

This was translated into:

Table 2D.4 – Iraq war scenarios quantified

	March 10	Short War No Recovery	Recovery	Long War Equity Crash	Malaise	Fwd
Fed funds	1.25	1.25	1.25	0.75	1.00	1.25
2-Yr Treasury Yield	1.33	1.63	2.04	0.89	1.08	1.48
5-Yr Treasury Yield	2.48	2.87	3.36	2.10	2.19	2.65
10-Yr Treasury Yield	3.58	3.89	4.26	3.30	3.32	3.71
30-Yr Treasury Yield	4.65	4.75	4.93	4.49	4.50	4.71
Fed funds/2s	0.08	0.38	0.79	0.14	0.08	0.23
2s/5s slope	1.15	1.24	1.32	1.21	1.11	1.17

5s/10s slope	1.10	1.02	0.90	1.20	1.21	1.11
2s/10s slope	2.25	2.26	2.22	2.41	2.24	2.23
10s/30s slope	1.07	0.86	0.67	1.19	1.18	1.00
1x10yr swaption yield vol	29.14	25.81	22.44	33.62	32.36	27.56
10-Yr Swap spread (bp)	43	42.5	43.5	80	38	43.6
5-Yr Swap spread (bp)	42	41.5	42.5	79	37	42.6

These scenarios were applied to the Citigroup Broad Investment Grade index. Valuations were made one month forward and three months forward from March 10, 2003. (0bps means no change in variables from March 10.)

Table 2D.5 – Iraq war scenarios applied

3-month Iraq War Scenario (as of Mar 2003)	SWNoRec	SWRec	LWEquity	LWMal	Forward	0bps
Rate of Return % (1-month)	-0.21%	-1.37%	1.69%	1.48%	0.28%	0.70%
Rate of Return % (3-month)	-0.61%	-4.06%	5.16%	4.49%	0.83%	2.12%

On June 10, 2003 the market thought there had been a short war – the S&P 500 had risen from 807.48 on March 10 to 984.84 on June 10. But the rising rate scenario contemplated did not occur – the UST 10 year dropped to 3.30 from 3.58. In fact we correlated the actual variables on June 10 with the scenarios in Table 2D.4 and found that the "long war malaise scenario" fit closest, with 0bps (no change) coming in second. In that sense the scenario analysis worked - the Citi BIG index was 921.3 on March 10, 2003 and 948 on June 10, up 2.90%, between the long war malaise scenario and the 0bps scenario.

Thus scenario analysis worked in the sense that the key input variables were identified in Table 2D.4 and their effects were properly computed. High grade fixed income portfolios respond more precisely to scenario analysis than do equity portfolios (with lower quality or emerging markets debt portfolios somewhere between), since idiosyncratic security behavior in equity portfolios plays a much larger part in equity than it does in high grade fixed income. Scenario analysis didn't work in our Iraq example in the sense that the translation of narrative scenarios in Table 2D.3 to key input variables was not correct.

A variant of scenario analysis is a technique I call "pessimization." This consists of using optimization techniques to minimize portfolio return over reasonably possible scenarios. A risk manager may thereby be able to determine that an unplanned confluence of events can bring disaster to a portfolio. The portfolio manager may then decide whether this scenario is so unlikely as to be unworthy of consideration, or likely enough that repositioning is in order. A form of pessimization using a fixed number of prespecified scenarios called "stress bands" is described in [Crouhy, Galai, and Mark 2001], pp. 233ff.

2E. Predictive power and commercial models

Figure 2E.1 shows sample annualized standard deviations of the S&P 500 index taken quarterly from daily log-returns:

Figure 2E.1 - S&P 500 volatilities from quarters of daily data

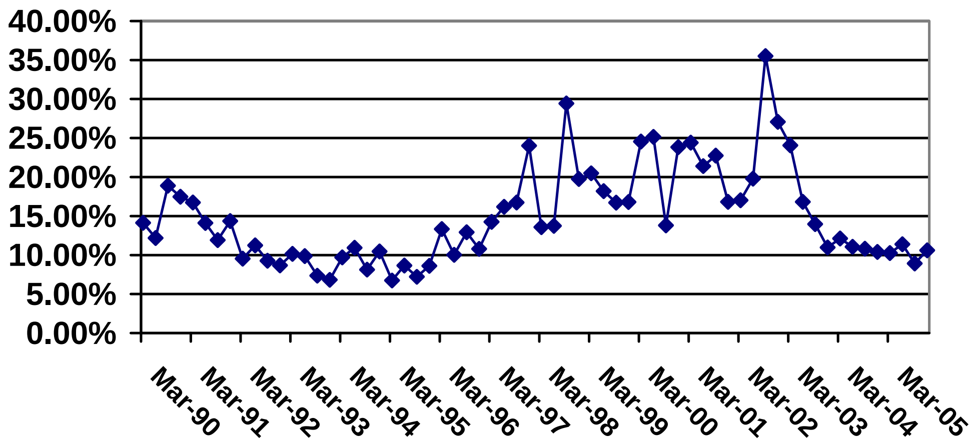

The points on this graph are computed from non-overlapping time periods. For example, the high point of 35.51% (annualized) was formed from the log-returns on the 64 business days in the third quarter of 2002. The 19.79% just before that was obtained from the second quarter of 2002 and the 27.07% following was obtained from the fourth quarter of 2002.

The distribution has positive skewness (1.09) and positive excess kurtosis (1.02), leading to a Jarque-Bera test statistic of 16.22, with a p-value of 99.97% rejecting normality. Figure 2E.2 shows the results of using four kinds of information available at the beginning of a quarter to predict the subsequent quarter's sample S&P 500 standard deviation. We took the means of the absolute deviations from between the beginning-of-quarter prediction and the sample standard deviation during the quarter.

Figure 2E.2 - S&P 500 standard deviations and predictions

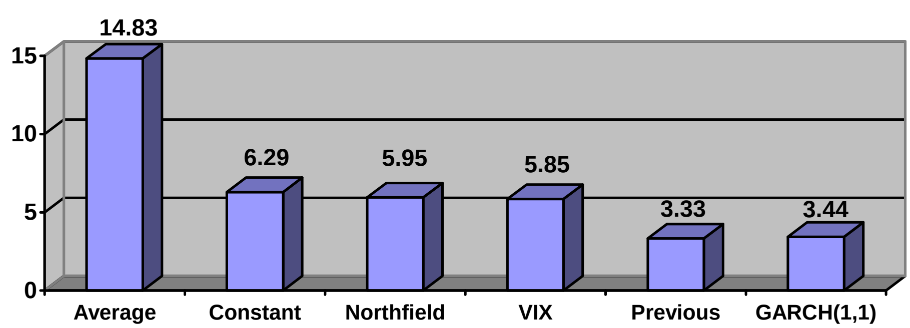

"Constant" means we simply used a constant (18.6%, the average quarterly standard deviation over the three years 1987-1989) as a predictor every time. "Northfield" refers to Northfield Information Services, who bravely supplied us with their US Fundamental model S&P 500 volatility predictions going back before 1990[13]. We also used the VIX™ index as of the end of the previous quarter, and finally just the previous quarter's sample volatility to predict the next. Clearly that did best.

It is difficult to do much better than this. A GARCH(1,1) model heavily weights the previous observation, consistent with Figure 2E.2's showing the previous was best, and does slightly worse than simply using the previous quarter.

Depending on the volatility sensitivity (vega) of the portfolio, getting to within 3.33% of the absolute level of volatility may not be too much of a problem. In the framework of (2A.3) and (2A.4), we find that while factor models like Northfield do not predict overall volatility well, they are useful to surface relative exposures. Northfield and other vendors use variations on [Rosenberg and Marathe 1976] where an equity covariance matrix C is estimated as follows:

$$C = L'FL + D$$

where C is an nxn symmetric matrix and n is large (order of 7000, for example). L is a kxn matrix of factor loadings indicating the characteristics of the securities in C – items such as industry membership, country or regional membership, capitalization, and price ratios. k is much smaller than n, usually of order 50-100. F is the kxk matrix containing estimates of the covariances between these factors. Finally, D is an nxn diagonal matrix containing specific (non-systematic) variances.

Another approach to estimating covariance matrices is principal components. This is used in both equity and fixed income. While there are many variations, the basic idea is to take a historical covariance matrix C and form its eigen decomposition

$$C = EVE'$$

where all matrices are nxn; E is a matrix of eigenvectors so that E'E=EE'=I, and V is a diagonal matrix of eigenvalues arranged in decreasing order. Since covariance matrices are positive semidefinite, the diagonal of V is nonnegative. The idea behind principal components analysis is that extracting the eigenvectors associated with the largest eigenvalues serves as a noise reduction technique. The smaller eigenvalues are called "scree," a geological term for debris. The prediction of the future covariance matrix is

$$\hat{C} = E_k V_k E_k'$$

where k is small compared to n. E_k is the nxk matrix consisting of the first k columns of E, and V_k is the upper left kxk part of V. Unlike the [Rosenberg and Marathe 1976] approach, this estimation technique uses no information other than past returns or yields. For example, the first three components of the Treasury yield curve have been shown to cover well over 90% of the variability in Treasury yields and roughly correspond to combinations of level, slope and shape [Litterman and Scheinkman 1991].

[13] Northfield's predictions are systematically too high – they average about 4.4% higher than the sample average. Removing this bias drops the mean of absolute differences to 4.34%. We used the long-term Northfield model here – they and other commercial services have short-term models that are designed to be nimbler in catching short-term changes.

We used the APT, Inc. global model, which is based on principal components, to predict EAFE volatility (taken from daily log-returns each three months) over the period 1997-2005.

Figure 2E.3 - EAFE volatilities from quarters of daily data

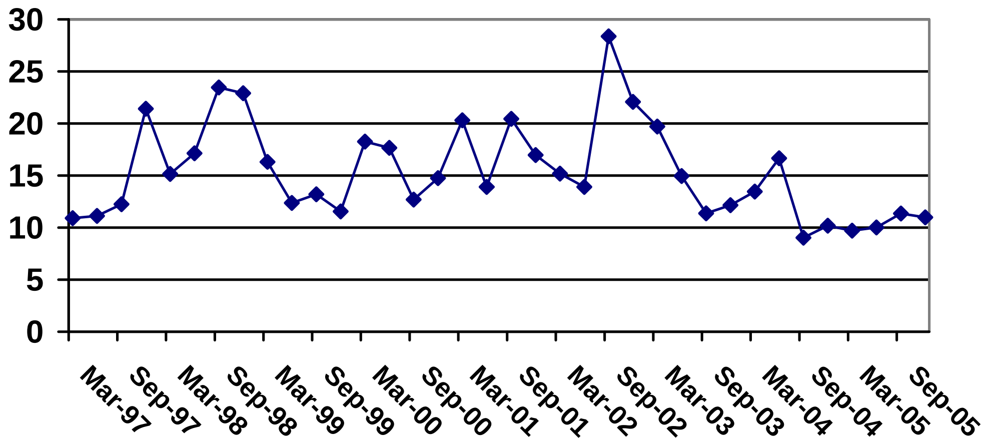

Figure 2E.4 - EAFE standard deviations and predictions

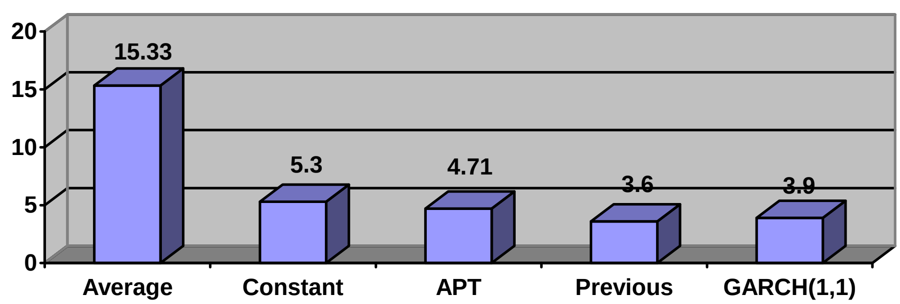

We used 10.12% as the constant based on the three years ending December, 1996.

A similar predictive pattern can be found using the Citigroup Broad Investment Grade index.

Figure 2E.4 - Citi BIG volatilities from months of daily data

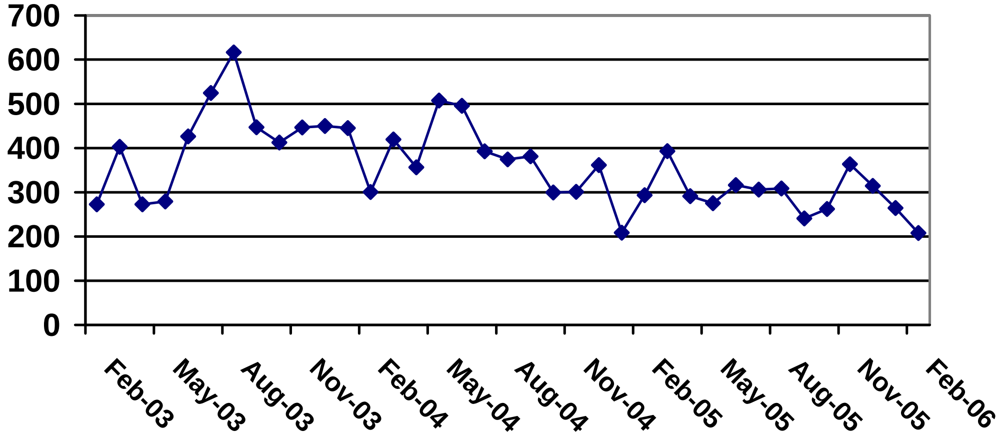

The units are basis points of annualized standard deviation. The distribution of observations during this short period is not significantly non-normal.

Figure 2E.5 - Citi BIG standard deviations and predictions

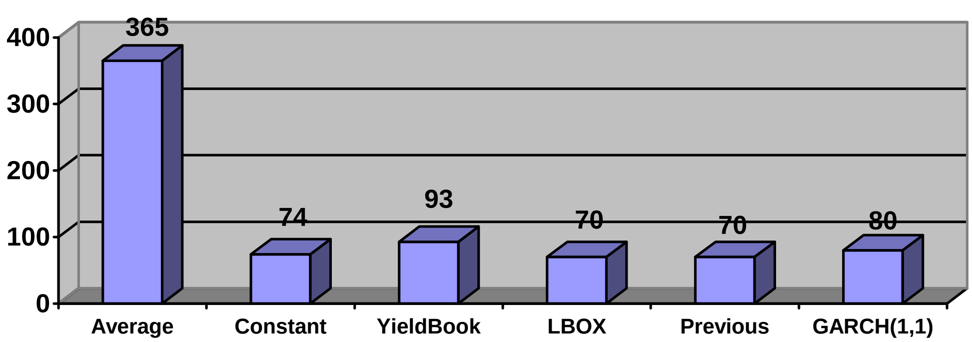

Using the previous one-month Citi BIG volatility to predict the next month ties using the LBOX (the Lehman index of swaption-implied volatilities) as a predictor. As with equity, the previous sample value had a mean absolute value of prediction error that was about 20% of the average. For many fixed income portfolios, vega is such that a 70bps error out of 365bps is not significant; for others with significant optionality this error can be material. We also used Citigroup's YieldBook tracking error module to produce monthly forecasts, and found those forecasts to be the worst predictors of the ones shown here. But as with equity systems, systems like YieldBook are useful to discover relative exposures.

YieldBook uses Monte Carlo simulation on over 800 risk factors, constructing 10,000 states and pricing a portfolio (or portfolio and benchmark) at each state [Citigroup 2005]. The resulting distribution is then displayed and its tracking error calculated.

2F. Model-free controls

Even the best risk models break down at times. For example, in the years before and after the TMT bubble bursting in 2000, no equity risk model did a very good job of predicting either relative or absolute risk statistics.

Figure 2F.1 - S&P monthly |Realized(i)-Realized(i-1)|

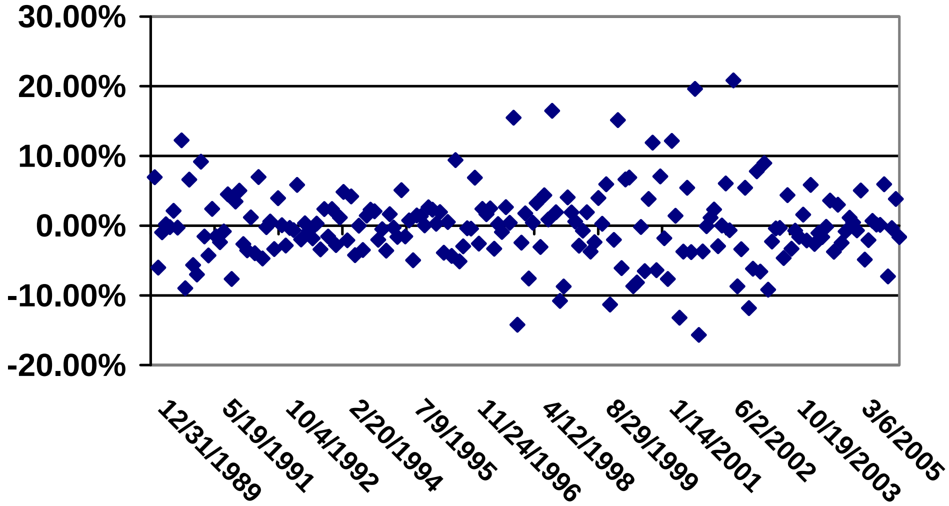

Breakdowns of risk models are associated with certain indicators. The current general level of volatility does not appear to be one of them. However, implied correlation is (inversely) associated with risk prediction breakdowns. Implied correlation is derived from the difference between the variance of a portfolio and the variances of the items in the portfolio. Intuitively, the more highly correlated items in the portfolio are, the bigger the portfolio variance, while less correlated portfolios produce lower overall variance because of diversification. This allows us to look at the difference between overall portfolio variance and individual variances and extract an average correlation without actually observing the correlations, whose number is quadratic in portfolio size. We do this by forming the ratio

$$\rho_{avg} = \frac{\sigma^2 - w'v}{(w's)^2 - w'v}$$

where w is the vector of weights of securities in the portfolio; v is the vector of individual variances and s is the vector of individual standard deviations; and σ^2 is the overall variance. We can observe the overall variance σ^2 either from option-implied volatilities or from historical calculations. We can observe the individual security variances σ_i^2 in the same way. This involves 2n+1 items (weights, individual variances, and overall variance) and produces an implied average of n(n-1)/2 correlations.

Figure 2F.2 - S&P 500 implied correlations

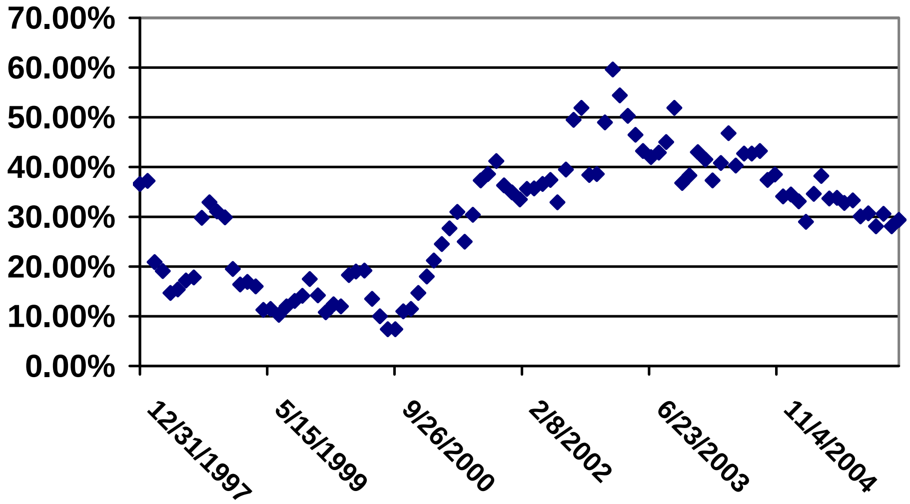

Source: Credit Suisse Derivatives Research

The high deviations in Figure 2F.1 are associated with low implied correlations in Figure 2F.2. Cross-sectional volatility is also associated with low predictive power and low implied correlations. In these regimes that stock-picking or choosing individual credits works best. In a very high implied correlation market such as early 2003 (in the runup to the Iraq war), individual name selection is less rewarding because everything is responding to the same macro factor.

If return opportunities to individual name selection are greatest when risk models work worst, buy side risk managers must find ways to allow return generation while their models are flying blind. In these periods, more primitive methods can be used to control risk. For example, if we are trying to control tracking error to a benchmark, we can set a maximum deviation δ on a security by security basis, so that if the percentage invested in security i in the benchmark is b_i, then the percentage invested in the security in the portfolio can be no less than $\max(0,b_i-\delta)$ and no more than $\min(1,b_i+\delta)$. As δ gets smaller, the portfolio will perforce behave more like the benchmark until finally, when $\delta=0$, we have an index fund.

The disadvantage of this method is that it handcuffs managers when their opportunities may be greatest. A manager liking Royal Dutch and hating Exxon cannot overweight the former and underweight the latter too much with this method. A risk model might note that both have similar characteristics and, under the theory that the Royal Dutch/Exxon spread is uncorrelated with other items in the portfolio, allow more leeway. In normal times, this works, but when risk models break down the assumption can fail. This method is also not generally useful for fixed income portfolios, where holding the exact issues in a benchmark can be difficult or impossible.

Thus we might move to not entirely model-free, but model-freer, control methods. One method involves making characteristics ranges or grids. For example, a fixed income portfolio could have a grid with credit ratings on the vertical and maturities or partial durations on the horizontal:

	0-1	1-3	3-5	5-7	7-10	10-20	20+	Sum
AAA	2.37	-1.41	1.25	0.65	1.33	-2.11	-0.41	1.68
AA+	1.94	-0.94	0.34	0.34	-1.48	-1.17	-0.37	-1.33
AA-	0.93	-1.98	-2.32	-0.81	-1.11	-0.97	-1.13	-7.40
A+	-0.48	-1.52	-2.36	2.48	-0.97	0.81	2.05	0.02
A-	-1.27	1.41	1.43	1.52	-2.39	2.36	1.91	4.98
BBB	1.27	0.60	1.15	0.15	-0.10	-0.42	-0.60	2.06
Sum	4.76	-3.83	-0.51	4.34	-4.72	-1.49	1.45	0.00

This shows, for example, that the portfolio is overweight 1.52% in A- bonds in the 5-7 year maturity range. By limiting the permitted deviations, we can force the portfolio's characteristics closer to the benchmark's. We still rely on the assumptions that credit ratings and maturity are the drivers of returns. We could easily form a portfolio that matched the benchmark exactly in these characteristics, causing the grid to contain all zeroes, but that didn't overlap the benchmark in a single issue. On the other hand, a risk model may say that we only need consider three principal interest rate components and two spread duration components rather than the 42 numbers in this grid. Thus the grid approach can be more restrictive.

Measures of concentration, both relative and absolute, are also used to force a portfolio to diversify without making model assumptions. The Gini index is based on the Lorenz curve, which in this notation is obtained by sorting portfolio weights or characteristics such as industry or country weights so that $x_1 \leq x_2 \leq \ldots \leq x_n$. Let Q_i be the partial sum of the first i x's (in sort order):

$$Q_i = \sum_{j=1}^{i} x_j$$

The x_i satisfy the budget constraint $Q_n=1$. The Lorenz curve is the monotone increasing graph of Q_i on the vertical axis against i on the horizontal axis. If the Lorenz curve is a straight line, then all the x_i are equal to 1/n and the portfolio is minimally concentrated. If the Lorenz curve is an "L" shape (or backward L) then all the Q_i equal zero except Q_n, which equals one, and the portfolio is maximally concentrated.

The Gini index is formed by looking at the difference between the line y=x on the Lorenz curve graph, and the Lorenz curve itself. In the minimally concentrated case, this area will be zero since the Lorenz curve is the 45 degree line y=x. In the maximally concentrated case, this area will be a triangle with base n and height one, or area n/2.

The Gini index is the ratio of the area described in the previous paragraph to the area of the full triangle (n/2). In the minimally concentrated case it is zero, while in the maximally concentrated case it is one. In general it is

$$Gini \equiv \sum_{i=1}^{n} 2\left(i/n - Q_i\right)/n = \frac{n+1}{n} - \frac{2}{n}\sum_{i=1}^{n} Q_i$$

It is also possible to form a relative Gini index, where the weights sum to zero instead of one and represent differentials between a portfolio and a benchmark – this can also be used for a long/short portfolio. We can modify the formula above so the relative Gini index is zero if the portfolio matches the benchmark and one if the portfolio and benchmark are invested in two completely different securities.

If we look at the levels of control mechanisms during periods when our models predict well, we can use them to manage risk when our models break down. For example, if a portfolio's Gini index has been 8% plus or minus 3% for several years when its tracking error was in a reasonable range, then we might get alarmed if its Gini index goes to 30% even if our risk models are not working. Similarly, if a portfolio has kept its industry weights to within 7% of its benchmark, in a low-confidence period a 15% deviation might be cause for alarm.

3. Long/short portfolios[14]

With the increased interest in hedge funds and other forms of long/short portfolios, existing risk management tools have become strained. The usual assumption that a portfolio's returns are lognormally distributed is clearly violated with a long/short portfolio that has the possibility of bankruptcy or worse if the value of the shorts exceeds the value of the longs. Even if both the long side and the short side are lognormal, the difference is not – the difference of two lognormal distributions is not a tractable distribution.

We assume there is a long portfolio, L, and a short portfolio, S, that follow correlated geometric Brownian motion according to

$$\frac{dL}{L} = \alpha_L dt + \sigma_L dZ_L \quad \frac{dS}{S} = \alpha_S dt + \sigma_S dZ_S \quad dZ_L dZ_S = \rho dt \ . \qquad (3.1)$$

We will assume that both the long and short portfolios are themselves long-only portfolios – they consist of non-negative holdings of assets that have a minimum value of zero. The actual portfolio we will hold at time T is *L(T)-S(T)*. We will assume an initial normalizing budget constraint *L(0)-S(0)=1*. We will use the letter λ, usually with an implicit argument of time T, to denote leverage: *λ=(L(T)+S(T))/(L(T)-S(T))*. Thus initial leverage is the sum of the longs and the shorts.

[14] This section is excerpted from [Hewett and Winston 2006].

In Table 3.1, we show sample parameters we will use to generate examples. These parameters could be typical of a managed futures portfolio:

Table 3.1 – Sample parameter set

Parameter	Description	Value
α_L	Long side drift	3%
σ_L	Long side volatility	20%
α_S	Short side drift	-2%
σ_S	Short side volatility	15%
ρ	Long/short correlation	0
L(0)	Initial long	4.5
S(0)	Initial short	3.5

Consider the ratio $\frac{L(T)}{S(T)}$. Bankruptcy (or worse) may be characterized by the condition $\frac{L(T)}{S(T)} \leq 1$. A ratio of lognormals is itself lognormal. Applying Itô's lemma to the ratio f=L/S, we obtain

$$\frac{df}{f} = Adt + \Sigma dZ$$

where

$$A = \alpha_L - \alpha_S + \sigma_S^2 - \rho\sigma_L\sigma_S$$
$$\Sigma^2 = \sigma_L^2 + \sigma_S^2 - 2\rho\sigma_L\sigma_S$$
$$\Sigma dZ = \sigma_L dZ_L + \sigma_S dZ_S$$

Thus the long/short ratio's distribution is given exactly by

$$P(L(T)/S(T) \leq r) = P\left(Z \leq \frac{\ln\left(\frac{r}{f(0)}\right) - (A - \Sigma^2/2)T}{\Sigma\sqrt{T}} \right) = N(D_1) \qquad (3.2)$$

where N(.) is the standard cumulative normal distribution function and

$$D_1 = \frac{\ln\left(\frac{rS(0)}{L(0)}\right) - (A - \Sigma^2/2)T}{\Sigma\sqrt{T}}$$

Figure 3.1 illustrates the response of 3.2 to volatilities at a particular set of parameters.

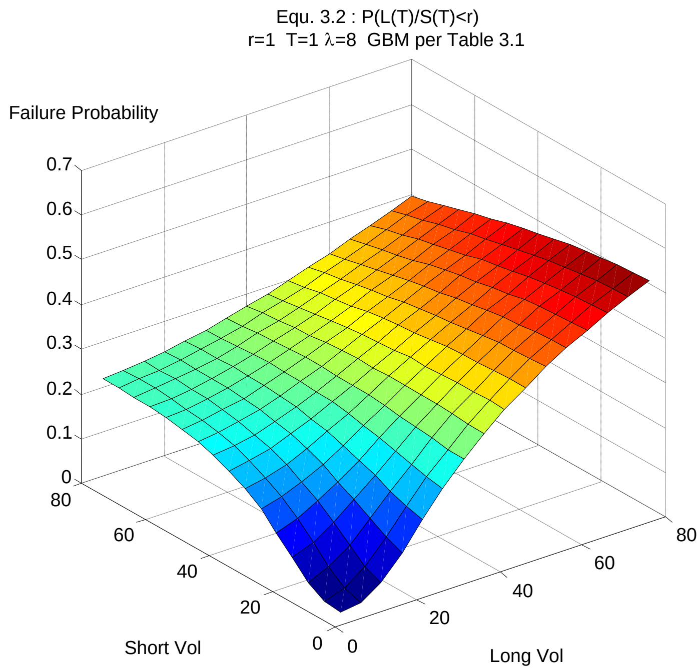

Figure 3.1

The probability of failure expressed in (3.2) does not reflect economic reality since it is a terminal value that does not take into account the absorbing barrier of bankruptcy[15]. (3.2) includes paths in which the long/short ratio goes below one. We calculate the stopping time for drawdown to the threshold, r, namely:

$$\tau_r = Inf\{t : t > 0, L(t)/S(t) \le r\}$$

Using a result from [Musiela and Rutkowski 1998], we obtain the probability the long/short ratio hits the absorbing barrier r sometime in the interval [0,t]:

$$P(\tau_r < t) = N(D_1) + \left(\frac{L(0)}{rS(0)}\right)^{1-\frac{2A}{\Sigma^2}} N(D_2) \qquad (3.3)$$

where D_1 is as defined above and $D_2 = D_1 + 2\frac{(A - \Sigma^2/2)\sqrt{t}}{\Sigma}$. Thus the probability of hitting the absorbing barrier is the failure probability (3.2) – that is, $N(D_1)$ – plus another term that is always positive.

Figure 3.2 shows the additional probability that is added to (3.2) by the presence of the absorbing barrier – that is, the difference between (3.3) and (3.2). The common parameters are the same as those in Figure 1. The surface is complex, but generally shows increasing incremental probability in short volatility. In long volatility, incremental probability increases to a point, and the tails off at very high levels.

---

[15] [Kritzman, Lowry, and Van Royen 2001] made a similar observation in the context of their regime-switching method.

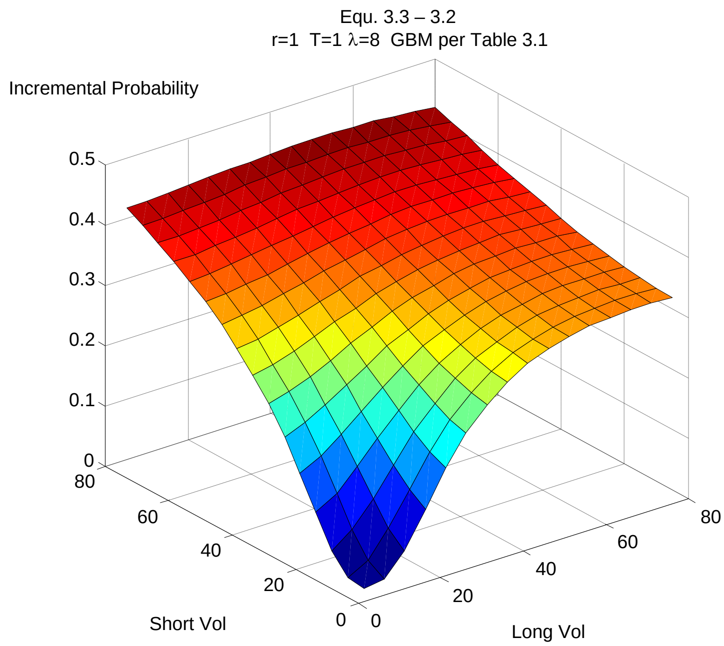

Figure 3.2

Analyzing the ratio L(T)/S(T) gives an exact solution for bankruptcy (r=1), but is not realistic for barriers that don't go all the way to full bankruptcy. A fund that starts with $100 million of capital and gets to L(T)=$.02, S(T)=$.01 may look acceptable from a ratio standpoint, but is effectively bankrupt. Most hedge funds would lose all their clients long before bankruptcy – in some cases, at very small drawdowns. It is therefore important to look at the difference L(T)-S(T). As we noted above, this is not a tractable distribution. This problem is similar to pricing a barrier option, and [Kirk 1995] has produced a useful approximation that we use. With this approximation, we can make a transformation to the formulas for the ratio to obtain formulas for the difference. Monte Carlo simulation confirms that these transformations work very well.

Our approximation gives

$$P(L(T) - S(T) \le k) = P(L/(S+k) \le 1) \approx N(\hat{D}_1) \tag{3.4}$$

where

$$\hat{A} = \alpha_L - SF\alpha_S + SF^2\sigma_S{}^2 - \rho\sigma_L\sigma_S SF$$

$$\hat{\Sigma}^2 = \sigma_L{}^2 + SF^2\sigma_S{}^2 - 2\sigma_L\sigma_S\rho SF$$

$$\hat{D}_1 = \frac{\ln\left(\frac{S(0)+k}{L(0)}\right) - (\hat{A} - \hat{\Sigma}^2/2)T}{\hat{\Sigma}\sqrt{T}}$$

where SF=S(0)/(S(0)+k). It is likely that a portfolio experiencing a 50% drawdown in a year would go out of business, so Figure 3.3 shows the probability surface with parameters similar to those in Figure 3.1, but with k=.5 instead of k=0 (r=1).

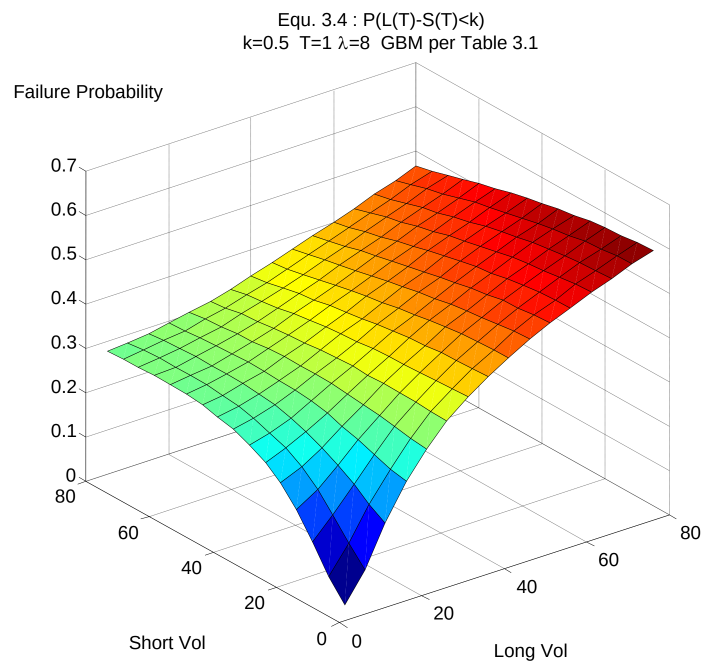

Figure 3.3

Importing the logic leading to (3.3), we define

$$\tau_k = Inf\{t : t > 0, L(t) \le S(t) + k\}$$

and obtain

$$P(\tau_k \le t) \approx N\left(\hat{D}_1\right) + \left(\frac{L(0)}{S(0)+k}\right)^{1-\frac{2\hat{A}}{\hat{\Sigma}^2}} N\left(\hat{D}_2\right) \qquad (3.5)$$

where $\hat{D}_2 = \hat{D}_1 + 2\frac{\left(\hat{A} - \hat{\Sigma}^2/2\right)\sqrt{T}}{\hat{\Sigma}}$.

As we saw with the ratio analysis, the difference between (3.5) and (3.4) is a positive quantity. Figure 3.4 graphs the added probability of failure due to this term; that is, due to the presence of an absorbing barrier. Comparing Figure 3.4 to Figure 3.2, one notes higher incremental probabilities.

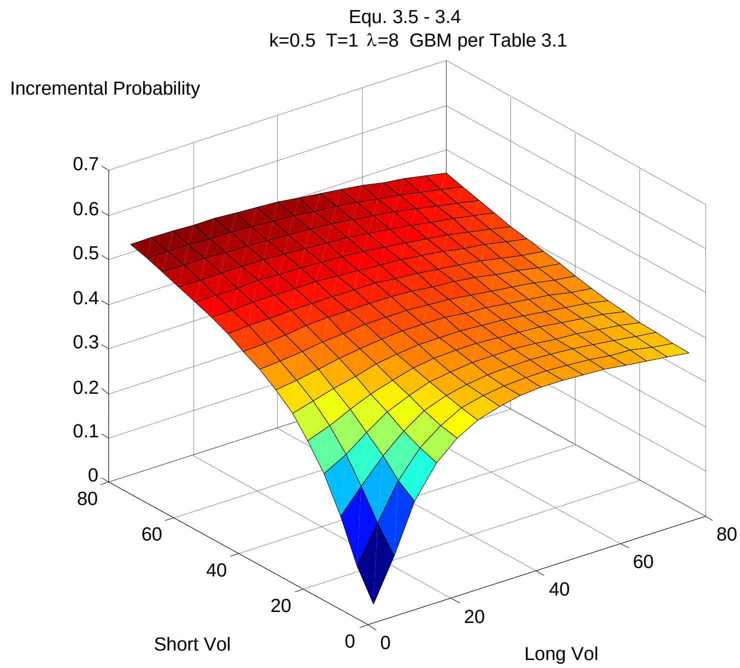

Figure 3.4

Figures 1-4 have been based on the parameters of Table 3.1 with leverage of 8, but Figure 3.5 shows that the added probability of failure can be significant with lower leverage of 2, especially at higher volatilities.

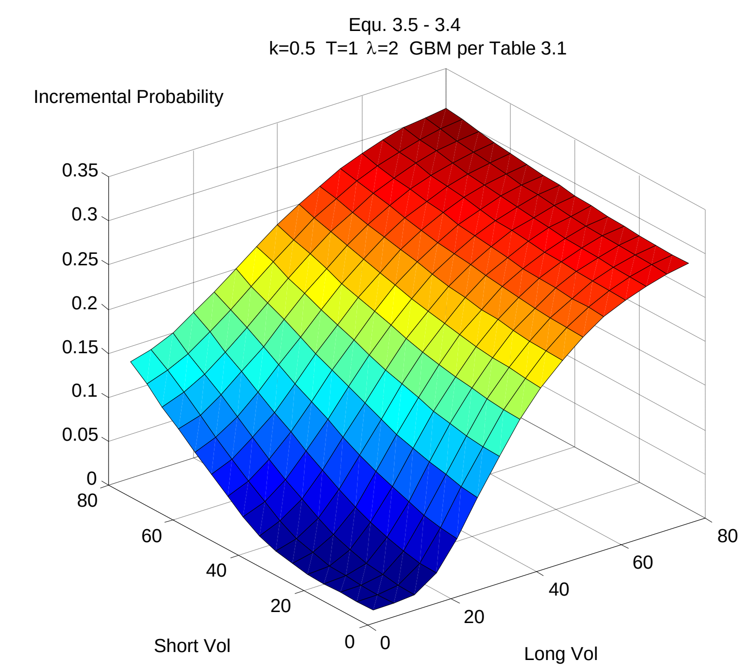

Figure 3.5

As we expect, negative correlation is harmful in long-short portfolio:

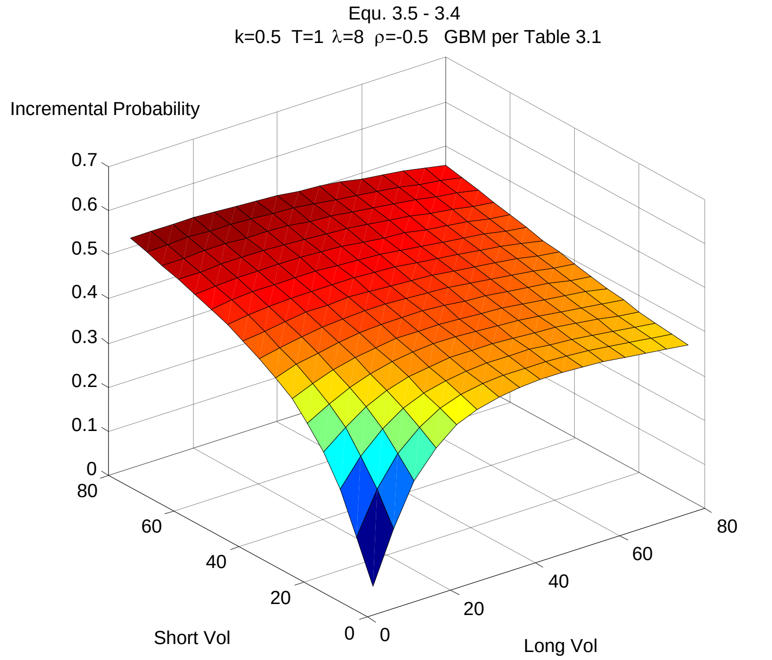

Figure 3.6

(3.5) is a useful formula, but managing long/short portfolios using only (3.5) may lead us to excessive risk aversion, as it concentrates only on failure. We now characterize the terminal distribution of results given that we avoid failure. That is, we compute the joint probability

$$P(K,k) = P(L(T) - S(T) \le K, \forall t \in [0,T] : L(t) - S(t) > k)$$

where we assume that $k < L(0) - S(0) < K$. Lower case k characterizes failure, and upper case K is the parameter giving us the terminal distribution of results assuming we avoid failure.

Using results from [Lee 2004] derived for the purpose of pricing barrier options contingent on multiple assets, we obtain

$$P(K,k) \approx \Phi_2\left(\hat{D}_1(K), -\hat{D}_1(k), -\hat{R}\right) - \left(\frac{L(0)}{S(0)+k}\right)^{1-\frac{2\hat{A}(k)}{\hat{\Sigma}(k)^2}} \Phi_2\left(\hat{D}_1(K) + \frac{2\hat{R}\ln\left(\frac{L(0)}{S(0)+k}\right)}{\hat{\Sigma}(k)\sqrt{T}}, \hat{D}_2(k), -\hat{R}\right)$$

(3.6)

where $\Phi_2(a,b,R) = P(Z_1 \le a, Z_2 \le b)$ is the bivariate cumulative distribution function for a standard normal 2-vector (Z_1,Z_2) with correlation R. Here the correlation is given by

$$\hat{R} \equiv \frac{1}{\hat{\Sigma}(K)\hat{\Sigma}(k)}\left[\sigma_L^2 - \rho\sigma_L\sigma_S S(0)\left(\frac{1}{S(0)+K} + \frac{1}{S(0)+k}\right) + \frac{\sigma_S^2 S(0)^2}{(S(0)+K)(S(0)+k)}\right]$$

In Figures 3.7 to 3.9, we show P(K,k) as a function of K both using formula 3.6 and using a Monte Carlo simulation. These are the lower lines; it is hard to see the difference. The upper lines in these figures show the sum of formula 3.6 and the stopping probability in formula 3.5, which approach one as K approaches infinity.

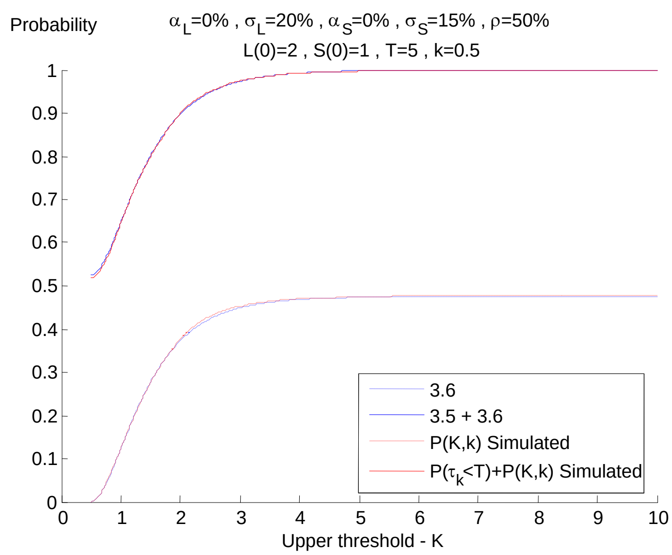

Figure 3.7

In Figure 3.7, a manager with no skill and leverage=3 has a 52% chance of not surviving five years without dropping to half the original capital. For the vast majority of the other 48%, the best result is a tripling of original capital (K=3), with doubling or less occurring about 39% of the time, and breakeven or less about 15%.

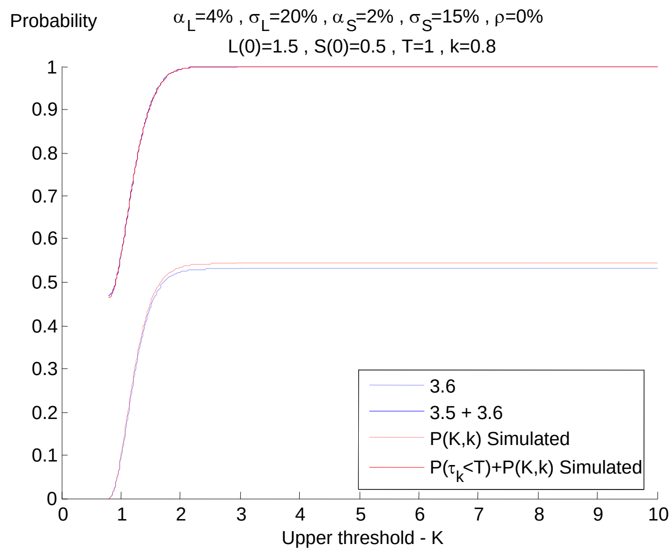

Figure 3.8

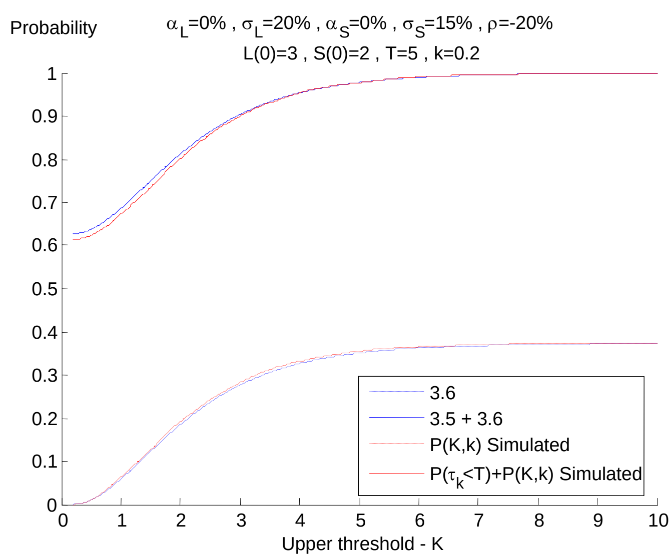

Figure 3.9

The probability of "success" i.e. of avoiding k and reaching K is given by $S(k, K, T) = 1 - P(\tau_k < T) - P(K, k)$. An example of this surface is given in Figure 3.10.

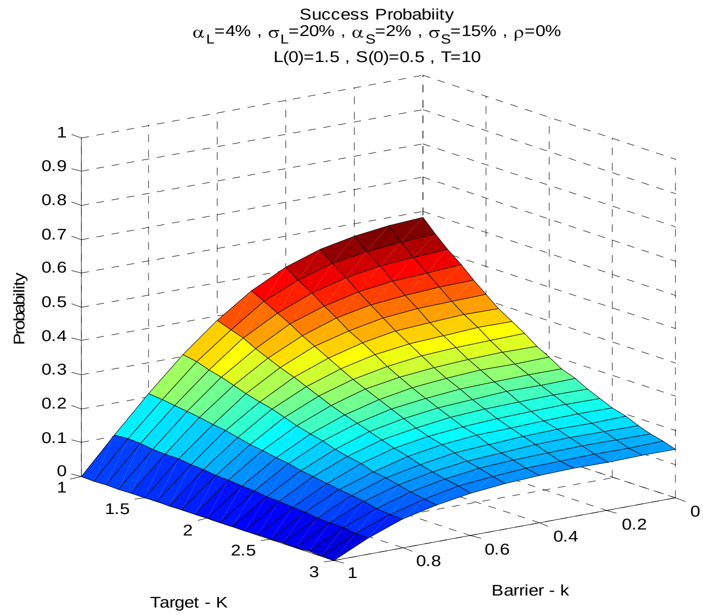

Figure 3.10

Figure 3.10 contemplates a long/short manager using low leverage and having positive skill, but over a long time period (ten years) the chance of success is dauntingly low. This is why long/short portfolios need continual risk management. In practice no long/short portfolio would go unmanaged for ten years, or even one year. Figures 3.7-3.10 all have failure probabilities that are unacceptably high. To lower the failure probability 3.5 one could change the portfolios generally by increasing long/short correlation, lowering leverage, lowering long volatility, and usually lowering short volatility. Some managers rebalance to constant leverage, which makes the portfolio once again a tractable geometric Brownian motion, but may not be the optimal risk control strategy. Another option is evaluating the portfolios at more frequent intervals and dynamically managing the risk, for example by delivering as we get closer to the barrier.[16]

Thus we see that long-short portfolios present different risk management challenges than long-only portfolios.

[16] We have extended our analysis to include jump diffusions, which only exacerbates the failure probabilities.

References

deBever, Leo, Korzun, Wayne, and Zvan, Barbara (2000) "Risk Budgeting in a Pension Fund," in Rahl, Leslie (ed.), Risk Budgeting: A New Approach to Investing, Risk Books, pp. 293-297.

Citigroup Global Fixed Income Analytics (2005) "A Common Framework for Estimating Tracking Error of Fixed Income Portfolios."

Crouhy, M., Galai, D., and Mark, R. (2001) Risk Management. McGraw-Hill.

Falkenstein, Eric (2003) "Alpha Deception," www.riskcenter.com, October 8, 2003.

Golub, Bennett and Tilman, Leo (2000) "Risk Management: Approaches for Fixed Income Markets." John Wiley & Sons.

Hewett, T. and Winston, K (2006) "Long-short portfolio behavior with barriers: Part 1 – Mechanism" Morgan Stanley working paper.

Jorion, Philippe (2001) *Value at Risk,* Second Edition, McGraw-Hill.

Keynes, J.M. (1937) "The General Theory of Employment," *Quarterly Journal of Economics*, Vol. 51, p.209-23.

Kirk, E. (1995) "Correlation in Energy Markets," *Managing Energy Price Risk*, London: Risk Publications and Enron pp. 71-78.

Knight, F.H. (1921) *Risk, Uncertainty and Profit*. 1933 reprint, London: London School of Economics.

Kritzman, M, Lowry, K., and Van Royen, A. (2001) "Risk, Regimes, and Overconfidence" in Levich and Figlewski, eds., *Risk Management: The State of the Art*, Kluwer Academic Publishers, pp. 129-144.

Lee, H. (2004) "A Joint Distribution of Two-Dimensional Brownian Motion with an Application to an Outside Barrier Option", *Journal of the Korean Statistical Society* 33, pp. 245-254.

Musiela, M. and Rutkowski, M. (1998); *Martingale Methods in Financial Modelling*;Springer Verlag.

Neville, Laurence (2006) "Model risk: Questioning the numbers," *Risk magazine*, January 2006.

Litterman, Robert (1996) "Hot Spots™ and Hedges," *The Journal of Portfolio Management*, Special Issue, pp. 52-75.

Litterman, R., and Scheinkman, J.A. (1991) "Common Factors Affecting Bond Returns," *Journal of Fixed Income*, 1.

Mandlebrot, Benoit (1963) "The Variation of Certain Speculative Prices," *Journal of Business* 36, pp. 394-419.

Rees, Stephen (2000) "VAR for Fund Managers," *Risk magazine*, January 2000 pp. 67-70.

Rosenberg, Barr and Guy, James (1975) "The Prediction of Systematic Risk," Berkeley Research Program in Finance Working Paper 33.

Rosenberg, Barr and Marathe, Vinay (1976) "Common Factors in Security Returns: Microeconomic Determinants and Macroeconomic Correlates," Berkeley Research Program in Finance Working Paper 44.

Thorp, Edward (1971) "Portfolio Choice and the Kelly Criterion," *Business and Economics Statistics Section Proceedings of the American Statistical Association,* pp.215-224.

Venkataraman, Subu (2005) "Buy-side vs. Sell-side Risk," presentation at GARP 2005 convention, February 1, 2005.

Wigner, Eugene (1960) "The Unreasonable Effectiveness of Mathematics in the Natural Sciences," *Communications in Pure and Applied Mathematics*, vol. 13, No. I (February 1960).