



V I S I O N A R Y B R I E F S

FROM LEADING INVESTMENT MANAGEMENT FIRMS FOR PLAN SPONSORS WORLDWIDE

HOW TO IDENTIFY, MEASURE, MONITOR AND
MANAGE RISK
IN GLOBAL MARKETS

I N V E S T O R S P R E S S , I N C .

Making VaR Useful for Plan Sponsors

*by Eric Stubbs, V.P. Head of Strategic
Investment Advisory Group*

*Francis Gupta, Associate
Strategic Investment Advisory Group*

Over the last decade, the sell side has developed a robust and sophisticated structure for market risk reporting and management based on Value at Risk (VaR), a statistical measure of the expected normal markets loss that may occur for a given set of positions (portfolio) at a defined frequency. The 95% daily VaR, for example, defines the expected normal loss on one day out of twenty. Typically, the sell side couples VaR reporting as a descriptive measure of risk with a number of prescriptive tools, including loss limits and maximum drawdown, and performance standards, such as capital at risk, daily P&L and highest risk positions.

In contrast, the buy side has not so far been successful in either defining a standard approach to risk reporting and management or in adapting VaR to meaningfully convey actionable information about portfolios. One important element of this failure has been our inability to connect risk measures with investment guidelines and policy on one end of the investment process and with tactical investment decisions at the other end. Put differently, we have failed in investment management to provide satisfying answers to two fundamental questions:

1. Is the risk profile of the plan sponsor's investment allocation supportive of overall plan goals?
2. Is the plan sponsor getting paid for the risks the manager is taking on its behalf?

In this article we propose a first step to a framework for market risk management that creates a context for VaR within the investment process. The first section contrasts the underlying motivations on the buy side versus the sell side, where VaR originated, to develop a rationale for buy side risk measurement and management. This comparison gives insights into the relevant characteristics that a buy side risk methodology must address. The second section examines the relationship between risk measurement and asset allocation.

**A COMPARISON OF
THE SELL SIDE AND
BUY SIDE CAN
PROVIDE INSIGHTS
TO A RISK
MEASUREMENT
METHODOLOGY**

Because the buy side and sell side are two sides of the same coin, it is tempting to think that risk management would play similar roles in both. In fact, for a variety of institutional and functional reasons, this is not the case. Indeed, consideration of the contrasts is not only important, but essential if we are to make VaR and the sell-side approach to risk reporting and management useful to the buy side.

The differences in the institutional environments on the buy side change not only the *kinds* of risk decisions that the plan sponsor will make, but potentially, even the *reasons* for the decisions. The buy side cannot simply graft on sell-side tools and techniques, such as VaR, and expect to have a useful risk management capability.

As a starting point to the discussion, **Chart 1** compares the buy side and sell side in terms of characteristics that define market risk management issues:

The contrasts in **Chart 1** reveal important differences in the underpinnings of risk management on the buy side versus the sell side. To tailor a risk management to the buy side, the risk management framework must be modified to take these differences into account.

Chart 1

The Buy and Sell Side of Risk Management: A Comparison		
Characteristic	Buy Side	Sell Side
	Plan Sponsor	Investment Manager
• Investment decisions	• Delegated	• Direct
• Traditional risk measures	• Asset-level allocation • Tracking error • Sponsor peer comparisons	• VaR • Stress tests
• Purpose of risk controls	• Maximize allocative efficiency • Monitor relative risk	• Loss avoidance
• Control methods	• Benchmarking • Diversification • Rebalancing	• Capital rationing • Inventory control • Regulatory requirements
• Access to underlying information	• Diversification • Sell rules	• VaR limits • Stop loss • Capital allocation
• Pricing	• Delayed and indirect, at best	• Rapid and detailed
• Typical horizon	• Smoothed, fair value • MTM	• MTM
• Definition of positions	• One year plus	• Monthly, quarterly
	• Outright • Perceives low turnover	• Outright • Lower turnover
		• High turnover • Long/ short net strategy • Highly leveraged

- **Time horizon:** The nature of sell side trading, with its high turnover, incentive to manage inventories by laying off risks, and relative absence of restrictions on holding periods outside of the treasury functions, makes short-term VaR an ideal risk measure. Generally, the sell-side VaR is set with an eye to three measures of portfolio time horizon: typical holding period, typical unwind period for the product (a measure of the market's liquidity), and profit and loss measurement period. All three measure between a day and a week in most sell side contexts.

VaR is much less of an all-purpose tool for the buy side because turnover is typically slower and investment horizons are usually measured in months or years, rather than in days. Because a VaR number will grow with the square root of the time horizon, the VaR statistic will usually become so wide for horizons longer than a month that it ceases to convey any useful information about portfolio risks.

To convey useful information, the time horizon for VaR should be relatively small. Because investment committees typically think in terms of longer time frames, VaR by itself is not an adequate risk measure for the plan sponsor.

- **Purpose of risk measures:** Banks on the sell side use VaR risk measurement for three primary reasons: to ration risk capital across trading businesses in an effort to maximize profits; for inventory control in market making; and to satisfy regulatory capital adequacy requirements. The parallels within the investment management process are, at best, very weak.

In addition, there are important differences in the underpinnings of how these two communities structure their investment problems. On the sell side, capital allocations (or rationing) is decided according to the proprietary interests of the bank and its views on where to find the best financial opportunities. Proprietary sell side organizations are pretty nimble in reallocating capital as market conditions, possibly signaled by VaR, dictate. The consequence is that the tools developed for one may not address issues in the other.

Investment managers, acting largely as agents, do not have the same discretion to rapidly reallocate across asset classes or even subclasses because discretion to reallocate is usually vested in the plan sponsor's investment committee. Moreover, portfolio allocation is considered a strategic decision and is revisited only every few years as a function of plan goals and secular market changes. This is a consequence of the Capital Asset Pricing Model (CAPM) framework that underpins nearly all investment thinking today and contrasts markedly with the tactical trading tradition of the sell side. Finally, even the investment manager's freedom to reallocate across sub-asset classes can be restricted by the potential cost of deviating significantly from his or her benchmarks.

By itself, a VaR measure will not address the plan sponsor's requirements for longer-term risk control, which requires that VaR be linked with more traditional portfolio allocation concepts.

- **Optimal responses to risk:** Because the sell side measures risk to allocate capital on a daily or weekly basis and to satisfy regulatory compliance requirements, its responses to the discovery of excessive risk relative to performance are simple and clear: reduce risk and reallocate capital to trading lines where the return is better.

In comparison, on the buy side, we still do not have satisfactory answers to:

- How much risk is too much?
- Having identified too much risk, what, if anything, should we do about it?

In part, this is a result of the separation within investment management between the principal, the plan sponsor, and the agent, the investment manager, and a justified reluctance on the part of most plan sponsors to micromanage their agents' investment processes and strategies. The separation has a second effect. Unlike the sell side principal who can receive and analyze trade information almost in real time, the plan sponsor usually does not receive the same level of detail on account holdings, and when they do get information, it's already stale. Understanding this difference, however, does not answer the questions. The reality is that on the buy side we do not have a well-reasoned approach to the allocation of responsibility for short-term risk management between the principal and agent.

The consequence is that far from having too few measures of risk, the buy side arguably has too many. These include prudent person rules, guidelines, strategic allocations, tactical bounds, compliance requirements, derivative usage reports, benchmarks and statistical measures such as information ratios and VaR. At some point, transparency is probably diminished rather than enhanced by this multiplicity of interlaced measures.

However, short-term VaR measures can be particularly useful to the plan sponsor for manager and benchmark VaR comparisons. Significant differences in either direction may be a signal that the manager is roaming. But VaR is not an evaluative tool when used this way. Until it is coupled with performance data, it will not reveal whether the plan sponsor is getting paid for risks over and above the benchmark.

Buy side risk management would be enhanced through clear definitions of roles and responsibilities between plan sponsors and managers, coupled with reliance on fewer and more eloquent risk measures.

- **Impact of market risk:** On the sell side, the translation from market risk and performance to bottom-line corporate performance is transparent. Apart from the complications of leverage and income tax provisions, the connection between return and risk, and corporate value is straightforward.

Not so on the buy side. Particular funding rules, calculations of contributions and positive and negative expenses complicate the connection to plan and firm corporate value. The marginal impact of a dollar earned on corporate value by a plan with a 99% funded ratio may be very different from that of a plan with a 150% funded ratio.

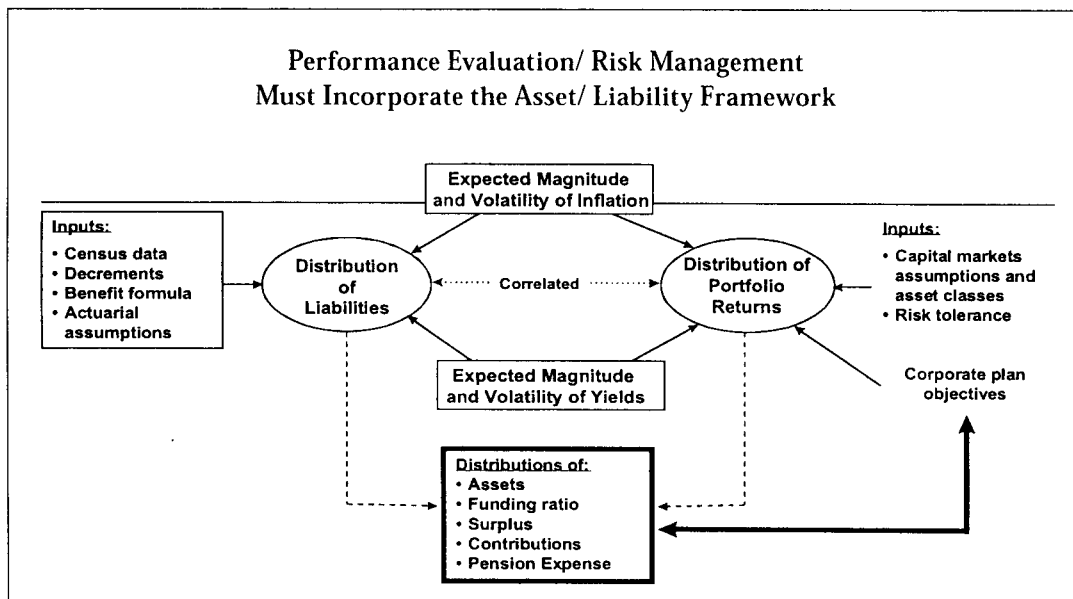
While VaR may usefully measure market risk, it must be interpreted within the institutional regulations of pension plan finance for the plan sponsor to adequately assess the actual impact of the risk on performance. Identical dollar risks may have very different implications for different plans, or to the same plan at different times.

- **Incorporate liabilities:** When the sell side estimates market risk, the calculation usually encompasses both the outright (long) position and its funding or cost of capital (short position). The funding is a liability that corresponds to the asset position. Profit and loss are measured net of the trader's ability to carry the position profitably. Similarly, it is the risk of the surplus – assets minus liabilities – that is most important.

The buy side has not thought in terms of funding costs because funds have an outright net asset position. However, pension funds have a liability in the form of the ABO or PBO that can serve a role that parallels that of funding costs on the sell side. The market cost of this liability is a hurdle that the plan sponsor must jump, much like a trader must overcome their funding hurdle.

Like sell side funding, the value of buy side liabilities changes with the market. Consequently, when computing market risk, it is the risk of the net position or “surplus risk” that should be paramount to the plan sponsor. Unlike the sell side, there are significant uncertainties associated with the calculation of the ABO and PBO liabilities because factors such as future salary growth, inflation, employment growth, retirement patterns and workforce demographics that affect these long-term liabilities are difficult to predict. As a first approximation, the plan sponsor can model the liabilities side as a short bond position; however, these uncertainties suggest that many liability streams are not actually very “bond-like” and a more precise ABO/PBO calculation would serve better. **Chart 2** displays the relationship between assets and plan liabilities, which determine the distribution measures such as assets, contributions and pension expenses, etc. These, in turn, are closely related to corporate plan objectives.

Chart 2



Plan sponsors should focus on their net positions – assets minus plan liabilities – when calculating market risk.

Three key conclusions for the design of a buy side risk management framework follow from the comparisons above:

- The buy side needs to link a short-term risk measure, such as VaR, with a longer-term asset allocation approach to risk management.
- Movements in asset-side VaR may be useful in signaling significant changes in manager behavior; however, to make that information actionable, plan sponsors and managers need a clear agreement on their respective roles in risk management.
- To be fully relevant, a plan sponsor’s risk measures must incorporate the liabilities side of their portfolio along with the assets.

EVEN FOR A WELL-DIVERSIFIED PORTFOLIO, ACTUAL MARKET RISK PROFILES CAN VARY DRAMATICALLY FROM THEIR STRATEGIC TARGETS

The first line of defense in creating a buy side risk profile appropriate for a plan sponsor is to ensure that the overall asset allocation is consistent with the plan sponsor’s investment goals. The simple process of identifying and quantifying investment goals is an important risk management step, and for most investment committees, no mean feat in itself.

However, the plan sponsor should not be lulled into false security by setting a strategic allocation consistent with their goals, however well thought out they may be. We have found in our analysis that the actual volatility of a given asset allocation can have quite a wide range. Thus the actual portfolio-level risk can vary dramatically from the risk that your investment committee targeted as their “comfort range”, without any changes in your asset allocation.

In the example shown in **Chart 3**, we took 25 years of historical data for a 5/35/40/10/10 —cash/bonds/S&P 500/Russell 2000/EAFE portfolio and calculated the VaR using monthly data over rolling two-year periods. The portfolio was rebalanced to its strategic allocation every month. The straight line in **Chart 3** represents the expected long-term VaR based on strategic asset assumptions—in this case, calculated using historical data over the 25-year period.

Chart 3

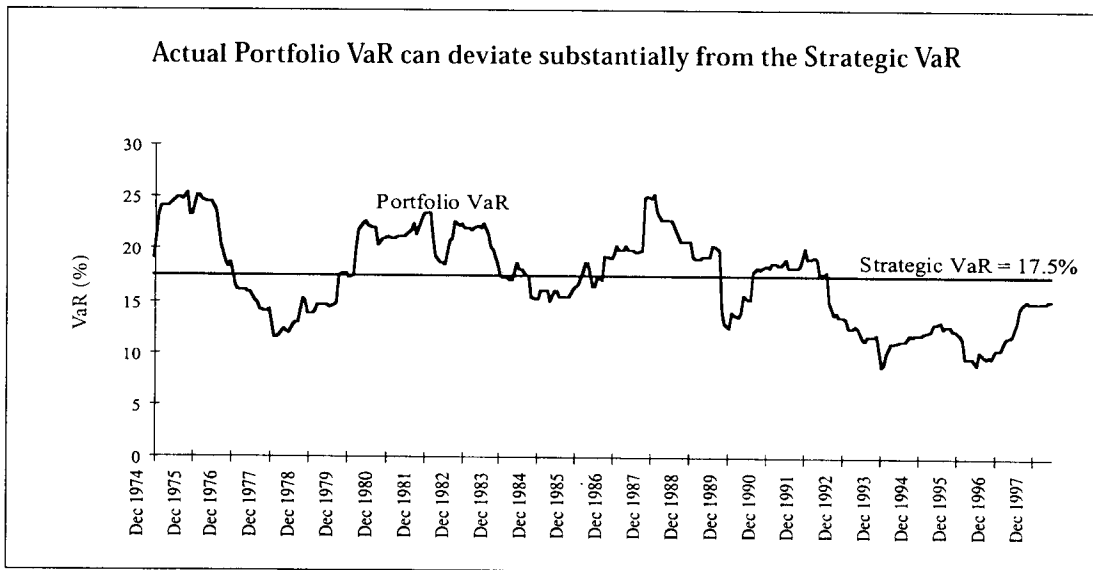



Chart 4 reveals that actual VaR can be as small as 60% of the VaR of the strategic mix and can be as large as one and a half times. In 11% of the months, the actual VaR is greater than 24.5% (1.4 times the strategic VaR of 17.5%). This means that in more than one month out of every ten, the actual VaR exceeds the strategic VaR by more than 40%.

Chart 4



If the plan sponsor's maximum tolerance for actual VaR is 1.4 times that of the strategic VaR, any month in which the VaR hit 24.5% would flash a warning signal and the plan sponsor would seek a response to remedy the situation. One method of putting reins on VaR is to make tactical shifts in the asset allocation. However, to do this right, it's important to understand the contribution of each asset class to total VaR. **Chart 5** portrays the percent contribution of each asset class to actual VaR.

It is evident from the chart that the relative contribution of each asset class to total VaR can vary significantly. Bonds, for instance, in some months make up much as 40% of the total VaR, whereas in other months they comprise only about 5% of it. This in itself suggests that there is a scope for constructing optimal tactical strategies to combat VaR.

Chart 5

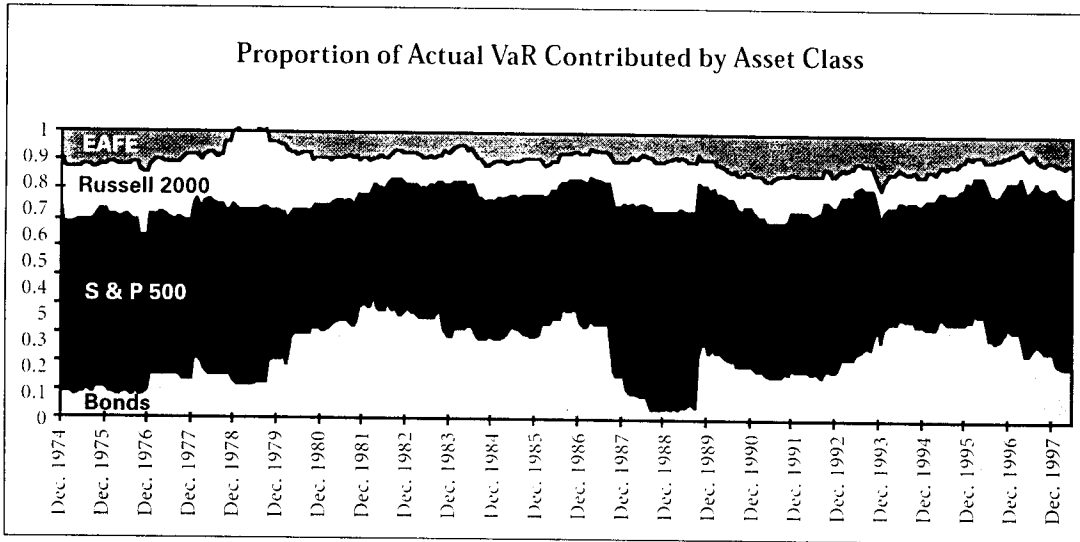



Chart 6 (see page 48) presents the percent contribution of U.S. fixed income and large cap equity to actual VaR, but as deviations from their respective percent contribution to strategic VaR. A high positive value on this chart means that the asset class is contributing a disproportionately greater amount to the portfolio's overall risk during that period. A chart of this type is useful in identifying the asset classes that are deviating the most from their strategic levels, for any given month.

The range of the deviations of each asset class depends on three factors: the share of the asset class in the portfolio, the intrinsic variability of the asset class and the correlation between riskiness in this asset class and the others.

For instance, ignoring the weights of the asset classes in the portfolio would lead one to conclude that bonds have the largest variability. The absolute ranges are 40% for bonds, and 32%, 21% and 20% for S&P 500, Russell 2000 and EAFE, respectively. However, when corrected for the relative shares of each asset class in the portfolio (thereby, reflecting only the intrinsic variability of the asset class itself), the relative ranges are 115% for bonds, 79% for large cap equity, 206% for small cap and 195% for EAFE. Given that each asset class contributes in different proportions to the portfolio's overall riskiness, it is very difficult to devise a straightforward strategy for tactical reallocation to control actual portfolio risks.

Chart 6



Note: Cash, Russell 2000 and EAFE were excluded from this chart to keep it comprehensible. But a similar chart may be generated for those omitted classes.

It is apparent from these results that your asset allocation tells you less about your overall risk profile than you probably think. Corollaries are:

- It is a good practice to monitor portfolio VaR at least quarterly even for a constant mix (no tactical deviations) that is invested passively. Under active management, it may make sense to monitor VaR even more frequently.
- Even though the VaR of your strategic mix may be well within your tolerance, don't let it provide you with a false sense of security. Actual risks can vary dramatically over time and have a good chance of exceeding your appetite for risk.

It may be better for your Investment Committee to first identify their risk comfort level as a function of plan goals and corporate finance considerations, then derive an asset allocation consistent with that risk level. The asset allocation should be reanalyzed periodically (perhaps annually) as market conditions change, adjusting it to keep the portfolio risk level in range.



SMART PILLS



- Remember that your asset allocation is not a good measure of the risks of your investments because actual risks can very easily deviate substantially from the level you thought you targeted with your strategic allocation.
- Begin the risk management process by ensuring that there is investment committee agreement on what fund and corporate financial goals should be and that these goals are reflected in the asset allocation.
- Suggest that your investment committee articulate risk preferences by setting a long-term target rate of return for the fund, identifying the asset allocation and risk level necessary to achieve it with some probability, and also how long in months or years they can tolerate falling behind the target.
- Check the VaR of your allocation at least quarterly and adjust asset allocations annually to keep overall risks within the investment committee's "comfort zone".
- Consider market risks from the perspective of net asset liability management and not as an asset-only issue. For these purposes, absolute VaR is more meaningful than relative VaR.
- Use relative VaR, calculated on your investment management accounts on a short horizon of a week or a month, as a useful signal of profile or style deviations by your managers.
- Link manager-level VaR estimation with performance attribution to get a better read on whether you are getting paid for the decision risks associated with your managers' investment processes.
- Your Investment Committee may find it useful to identify a trigger point for VaR, perhaps 25% higher than the strategic level, at which allocations are re-examined.
- Once you identify a significant increase in VaR, you need to examine the marginal contribution to risk and return of each asset class to determine the best way to roll back risk.