Forthcoming: *Derivatives Quarterly* Vol. 5, No. 2

Value at Risk for Asset Managers [1]

by

Christopher L. Culp, [2] Ron Mensink, [3] and Andrea M.P. Neves [4]

This Draft: January 8, 1999

INTRODUCTION

Value at risk ("VaR") is a probability-based metric for quantifying the market risk of assets and portfolios. [5] VaR is often used as an approximation of the "maximum reasonable loss" over a chosen time horizon. Its primary appeal—widespread amongst commercial bankers, derivatives dealers, and corporate treasury risk managers—is its ease of interpretation as a *summary* measure of risk, as well as its *consistent* treatment of risk across different financial instruments and asset classes.

Value at risk is not nearly as well-accepted in the institutional investment community as it is elsewhere. [6] The main reason is that asset managers are typically in the business of *taking* risks, either to fund uncertain liability streams or to generate positive excess risk-adjusted returns. Not surprisingly, asset managers—mutual funds,

---

[1] We are grateful to Brian Heimsoth and Geoff Ihle for work with us on this subject. Nevertheless, the usual disclaimer applies and we alone are responsible for remaining errors and omissions. In particular, the views herein do not necessarily represent those of the State of Wisconsin Investment Board or any CP Risk Management client.

[2] Director of Risk Management Services, CP Risk Management LLC, and Adjunct Associate Professor of Finance, Graduate School of Business, The University of Chicago.

[3] Director of Quantitative Research and Risk Analytics, State of Wisconsin Investment Board.

[4] Senior Risk Management Consultant, CP Risk Management LLC.

[5] "Market risk" is the risk that the value of an asset or portfolio declines because of adverse movements in market prices such as interest rates, exchange rates, and security prices. Market risk is distinct from other types of financial risk, such as default risk or liquidity risk. In this paper, our attention will be limited to market risk, primarily because that risk is the main financial risk that VaR was developed to measure.

[6] *See* Christopher L. Culp, Merton H. Miller, and Andrea M.P. Neves, "Value at Risk: Uses and Abuses," *Journal of Applied Corporate Finance* 10(4) (Winter 1998).

private banks, hedge funds, pension plans, endowments, and foundations—often view "risk management" in general and value at risk in particular as inherently at odds with their primary business mandate. [7] Nevertheless, VaR *can be* a useful tool by which asset managers can better ascertain whether the risks they *are* taking are those risks they *want or need to be* and *think they are* taking. Investors, as well, are becoming increasingly aware of the benefits of VaR as a *monitoring* tool, thereby further prodding their fiduciary asset managers toward the regular calculation and disclosure of this measure of market risk. [8]

In this article, some of the applications of VaR to asset management are explored, with particular attention to the importance of VaR for multi-currency asset managers. [9] We first explain what VaR is and why it is so appealing conceptually. In the subsequent section, the mechanics of calculating VaR are explained, including the importance of some of the assumptions underlying the most common VaR measurement methodology. Bearing in mind the measurement difficulties with VaR, we then summarize four concrete applications of VaR to asset management. These applications involve the use of VaR to (i) monitor managers, portfolios, and hedging programs; (ii) eliminate *ex ante* transactional approval requirements; (iii) define a formal system of risk targets and thresholds; and (iv) implement a "risk budget."

---

[7] *See* Christopher L. Culp and Ron Mensink, "An Integrated Approach to Asset Allocation, Performance Evaluation, and Risk Management for Institutional Investors," Manuscript (September 1998).

[8] Throughout this article, we shall refer to "managers" as those who invest funds on behalf of outsiders. The outsiders placing capital with the asset manager are called "investors." In a pension plan, for example, the managers are those internal and external portfolio managers (and their senior supervisors) who invest capital on behalf of retirees. The retiree beneficiaries, in turn, are the investors. At a hedge fund, mutual fund, or private bank, the managers again represent the portfolio manager(s) or general partner, whereas the investors are the outside depositors/purchasers/limited partners.

[9] For a more specific discussion of the use of VaR by pension plans, *see* Christopher L. Culp, Kamaryn T. Tanner, and Ron Mensink, "Risks, Returns and Retirement," *Risk* 10(10) (October 1987).

WHAT IS VAR? [10]

Value at risk is a statistic that summarizes the exposure of an asset or portfolio to market risk. VaR allows managers to quantify and express risk as follows: "We do not expect losses to exceed 10% of the fund's net asset value in more than 1 out of the next 20 quarters." [11]

To arrive at a VaR measure for a given portfolio, a manager must generate a probability distribution of possible returns or changes in the value of that portfolio (or its component assets) over a specific time horizon. The distribution of possible portfolio returns or future values is called the "VaR distribution." [12] The VaR statistic for the portfolio is the return or absolute dollar loss corresponding to some pre-defined probability level—usually 5% or less—as defined by the left-hand tail of the VaR distribution. Alternatively, VaR is the adverse return or dollar loss that is expected to occur no more than 5% of the time over the specified time horizon.

VaR is often considered a useful summary measure of market risk for several reasons. One feature of VaR is its *consistency* as a measure of financial risk. By expressing risk using a possible "dollar loss" or "adverse return" metric, VaR facilitates direct comparisons of risk across different portfolios (*e.g.,* equity vs. fixed income) and distinct financial products (*e.g.,* interest rate swaps vs. common stock).

In addition to consistency, VaR enables managers or investors to examine potential losses over particular time horizons with which they are concerned. Any measure of VaR requires the specification of such a "risk horizon." A judicious choice of that risk horizon can aid asset managers in numerous risk management and disclosure matters. An asset manager's choice of an appropriate risk horizon may depend upon the timing of events including the following: outside manager evaluations; board or trustee meetings; performance disclosures to investors, limited partners, or third-party tracking

---

[10] Portions of this section draw heavily from Culp, Miller, and Neves, *op. cit.*.

[11] For a general description of VaR, *see* Philippe Jorion, *Value at Risk* (Chicago: Irwin Professional Publishing, 1997).

[12] The VaR distribution may be expressed in returns or in dollars. Return distributions often are empirically more tractable, and these can always be converted to a corresponding potential dollar loss given the current portfolio value.

services (*e.g.*, MAR or AIMR); regulatory examinations; tax assessments; key client meetings; and the like. [13]

Another advantage of VaR traces to its roots in probability theory. With whatever degree of confidence a portfolio manager wants to specify, VaR allows a specific potential loss over the risk horizon to be associated with that level of confidence. A 95% confidence level with a one-month risk horizon, for example, tells the portfolio manager, strictly speaking, that returns can be expected to dip below, say, X% in the next month with 95% confidence. Some often go on to assume that the 5% confidence level means they stand to experience returns below X% in no more than five months out of 100, an inference that is true only if strong assumptions are made about the stability of the underlying probability distributions. [14] Either way, VaR measures are *forward-looking* approximations of market risk unlike traditional backward-looking measures of actual historical performance, such as the *ex post* Sharpe ratio that is calculated for performance evaluation purposes using *realized* manager returns. [15] VaR thus compliments rather than displaces such conventional methods.

As a related advantage of VaR, this measure of risk often appeals to asset managers because it is largely tactical neutral. In other words, VaR is calculated by examining the market risk of the individual instruments in a portfolio, *not* using actual historical *performance*. [16] Whereas typical performance measurements reflect manager performance, VaR reveals the market risk borne by an investor based *solely* on the asset mix and current security holdings. Nevertheless, VaR is not totally neutral to *all* active strategies. If a manager loads up on a higher volatility sector (*e.g.*, technology), for example, that strategy-induced volatility is picked up by VaR. Managers' tactical shifts

---

[13] *See* James V. Jordan and Robert J. Mackay, "Assessing Value at Risk for Equity Portfolios: Implementing Alternative Techniques," in *Derivatives Hamdbook*, R.J. Schwartz and C.W. Smith, Jr., eds. (New York, N.Y.: John Wiley & Sons, 1997).

[14] This interpretation assumes that asset price changes are independently and identically distributed—*i.e.*, that price changes are drawn from essentially the same distribution every period.

[15] *See* William F. Sharpe, "The Sharpe Ratio," *Journal of Portfolio Management* (Fall 1994).

[16] This is the source of some confusion in the institutional investment community. When the actual returns of a manager or portfolio serve as the basis for calculating a summary statistic, the resulting metric can be used for performance evaluation but *not* for VaR.

and market timing decisions, by contrast, are *not* reflected in VaR measures—nor, as we shall see later, should they be. [17]

THE MECHANICS OF VAR ESTIMATION

Creating a VaR distribution for a particular portfolio and a given risk horizon can be viewed as a two-step process. [18] In the first step, the price or return distributions for each individual security in the portfolio are generated. These distributions represent possible value changes in all the component assets over the risk horizon. Next, the individual distributions somehow must be aggregated into a single portfolio distribution using appropriate measures of covariation, such as correlation. The resulting portfolio distribution then serves as the basis for the VaR summary measure.

Methods for generating both individual asset and portfolio VaR distributions range from the simplistic to the incredibly complex. The more realistic approaches to VaR measurement generate the VaR distribution by re-valuing all the assets in a portfolio for the most realistic market risk scenarios possible. These "full re-valuation" methods, however, can be very costly computationally. Accordingly, simpler approaches often make a number of statistical assumptions intended to reduce data requirements and computing costs. One such assumption is to assume that small changes in portfolio values around its current value are representative of larger potential value changes. This is known as the "partial re-valuation" approach to VaR measurement, and is analogous to using a bond's duration as an approximation of risk while ignoring its convexity. [19] Other simplifications made with the intention of relieving computing costs include a variety of distributional assumptions, most of which will be discussed below.

To date, most asset managers that implemented VaR using the more simplified computational approaches. Accordingly, we explore in this section the mechanics of simplified VaR calculations using an international equity portfolio as an example. We

---

[17] *See* Culp and Mensink, *op. cit.*

[18] In practice, VaR is not often implemented in a clean two-step manner, but discussing it in this way simplifies our discussion—without any loss of generality.

[19] The "full re-valuation" and "partial re-valuation" methods are compared in JP Morgan, *RiskMetrics Technical Document*, 4th ed. (December 1996).

begin with a simple historical calculation of VaR that relies exclusively on historical time series data. We then discuss the most common method of VaR estimation—called "variance-based" VaR measurement—popularized by JP Morgan by its 1994 introduction of "RiskMetrics™". We conclude the section with a short discussion of the shortfallings of the variance-based approach and a survey of other alternatives.

Historical VaR

One way to calculate VaR is simply to assume that the future will behave precisely like the past. VaR then can be calculated using a sample time series of past security returns. An example will help illustrate.

Figure 1 shows a frequency distribution (*i.e.,* histogram) of monthly returns on the FTSE 100 stock index from January 1988 to January 1995, as well as the summary statistics for that time series. The horizontal axis shows returns, and the vertical axis shows the percentage of the total sample of historical returns associated with each return interval on the horizontal axis. By assuming the future behaves like the past, we can call this frequency of past returns the "probability" that these same return levels will be realized in the future. The right-most bar, for example, indicates that approximately 1% of historical monthly returns were greater than 14%. [20] Assuming the next month in the future behaves like prior months, we can interpret this to mean that there is a 1% probability of a return in excess of 14% next month.

As the summary statistics and Figure 1 show, this return distribution is positively skewed—*viz.,* more probability lies in the right-hand tail than the left-hand tail. In other words, the probability of a return X% above the mean is *higher* than the probability of a return X% below the mean. We also can see that the excess kurtosis is positive, indicating that the distribution shown in Figure 1 has fatter tails than the normal distribution. [21]

[20] The x-axis labels on the figure correspond to the upper tick. For example, the 15% label lies between two tick marks, and the bar above that number implies that approximately 1% of the historical data was less than or equal to 15% and greater than 14%.

[21] Excess kurtosis is the amount by which the kurtosis of a distribution exceeds the kurtosis of the normal distribution, which is 3. A positive excess kurtosis statistic indicates that the distribution has a more peaked center and fatter tails than the normal distribution.

Suppose a British asset manager has an equity portfolio with a current market value of £1 million invested entirely in FTSE-100 stocks. Suppose further that the manager holds the equities in exactly the same proportion as their portfolio weights in the FTSE-100 index—*i.e.*, the fund is an index fund. [22] If the manager calculates the VaR for this portfolio at the 95% confidence level over a one-month risk horizon, the resulting VaR measure will reveal the value below which the portfolio is expected to exceed in five of the next 100 months. Using only the historical data in Figure 1, this VaR can be calculated as follows:

$$\text{VaR} = £1{,}000{,}000 \cdot R^{.05}$$

where $R^{.05}$ denotes the 5th percentile historical monthly return for the FTSE-100. This 5th percentile return, indicated with a bold vertical line on Figure 1, was –6.87% in the 1988-1995 sample period. So, the sterling-denominated one-month VaR for this portfolio is £68,700. [23] We thus expect the net asset value ("NAV") of this portfolio to decline by more than £68,700 in only five of the next 100 months.[24]

For large portfolios with numerous assets and exposures, the historical approach quickly becomes intractable. The data requirements alone can render this simplistic approach virtually worthless for large portfolios. Suppose, for example, we want to calculate the VaR of a portfolio comprised of FTSE-100 stocks but in which the portfolio weights for each stock are now different from the weights used to calculate the FTSE-100 index. In that case, each stock would need to be examined directly—*i.e.*, obtain historical time series for each of the 100 stocks and calculate the VaR using the new portfolio weights and the historical correlation matrix.

Variance-Based VaR

JP Morgan and Reuters greatly simplified the data problems of the pure historical approach by proposing a more simplified method of VaR measurement. This

[22] For the purpose of this example, we thus treat the FTSE-100 "portfolio" as a *single asset*.

[23] By convention, we drop the negative sign when reporting VaR.

[24] VaR, however, does not give any indication of how far beyond this amount the portfolio could decline; it could be £68,701 or £1,000,000.

RiskMetrics™ VaR calculation methodology circumvents the data problems associated with using actual historical data in two ways. First, the RiskMetrics approach relies on "primitive securities" rather than actual security-level portfolio holdings. Primitive securities are securities intended to *represent* actual security holdings for the purpose of VaR calculations. The cash flows on a given portfolio of actual securities are "mapped" into corresponding cash flows on primitive securities in order to perform the VaR calculation. Second, RiskMetrics makes available the relevant data on those primitive securities so that virtually no historical data collection is required. As an example, RiskMetrics does *not* furnish data on coupon-bearing U.S. Treasury securities, but it *does* on zero-coupon Treasuries. So, a portfolio of coupon-bearing Treasuries (*i.e.,* actual holdings) can be mapped into a portfolio of corresponding zero-coupon Treasury securities (*i.e.,* primitive securities) for which data *is* available. The primitive securities available through RiskMetrics include government bonds, swap rates (which double as corporate bond rates), exchange rates, equity index levels, commodity prices, and money market rates, all of which are available for various countries and currency denominations.

Aside from addressing the data problems associated with historical VaR, the RiskMetrics approach makes a number of statistical assumptions that also simplify the *computational* side of the VaR calculation. Most importantly, the approach assumes that all primitive security returns are distributed normally. An attractive property of the normal distribution is that it is symmetric. Mean and variance thus are "sufficient statistics" to fully characterize a normal distribution—*viz.,* the variance of an asset whose return is normally distributed is all that is needed to summarize the risk of that asset.

So simplistic is Morgan's calculation method that many managers now use it with primitive securities of their own—*i.e.,* without using the RiskMetrics data set. The approach thus has come to be known more generally just as the "variance-based approach." Because this is by far the most predominant—and affordable—method of calculating VaR, the following subsections explore the mechanics of this methodology in more detail.

Single Asset, Single-Period VaR

To calculate VaR using the variance-based approach, we rely on the fact that the probability in the left-hand tail of a normal distribution is a known function of the standard deviation of the distribution. Five percent of the normal distribution, for example, lies 1.65 standard deviations below the mean. Again, the mechanics of this approach are best illustrated by example.

Consider again the FTSE-100 index portfolio from the standpoint of the sterling-based asset manager. From the summary statistics provided in Figure 1, we know that the average monthly return is 0.76% and the monthly variance is 0.21%. With a current portfolio NAV of £1,000,000, the one-month VaR at the 95% confidence level is calculated as follows:

$$\text{VaR}_e = £1{,}000{,}000\cdot(0.0076 - 1.65\cdot 0.045) = £1{,}000{,}000\cdot(-0.068) = £68{,}012$$

On Figure 2, the historical distribution of returns from Figure 1 is super-imposed with a normal distribution whose mean and variance are equal to the historical sample mean and variance. Figure 2 then depicts this variance-based VaR estimate (expressed as a percentage return) with a vertical line. Note that the variance-based VaR expressed as a return is the same (to the second decimal) as the 5th percentile return of the actual frequency distribution. Not surprisingly, the variance-based £68,012 VaR is only trivially different from the £68,700 one-month VaR we calculated earlier using only the sample frequency distribution.

For a measure of the standard deviation used in the variance-based VaR calculation, we used the "unconditional" variance of the historical time series—*i.e.*, the standard deviation of all 85 monthly returns shown in Figures 1 and 2. Alternatively, we could have used a different method to estimate the volatility input to the variance-based VaR estimate.[25] If we chose to re-calculate this number on a regular basis, for example, we might have used a moving average of return variance as our estimate for volatility. [26] Because moving-average volatility is calculated using equal weights for all observations

[25] For a review of these methods, *see* Culp, Miller, and Neves, *op. cit.*, and Jorion, *op. cit.*.

[26] To get a moving average estimate of variance, the average is taken over a rolling window of historical volatility data. Given a 20-month rolling window, for example, the variance used for one-month VaR calculations would be the average monthly variance over the most recent 20 months.

in the historical time series, the calculations are very simple. The result, however, is a smoothing effect that causes sharp changes in volatility to appear as plateaus over longer periods of time, failing to capture dramatic changes in volatility.

In order to remedy this problem, the RiskMetrics data sets provided by JP Morgan and Reuters include volatilities for all the supported primitive securities computed at both the daily and monthly frequencies using an "exponentially weighted moving average."[27] Unlike the unconditional variance or the simple moving-average volatility estimate, an exponentially weighted moving average allows the most recent observations to be more influential in the calculation than observations further in the past. This has the advantage of capturing shocks in the market better than the simple moving average and thus is often regarded as producing a better volatility for variance-based VaR.

Multi-Asset, Single-Period VaR

The real savings in data and computing costs delivered by the variance-based VaR approach comes into play when the VaR is desired for a large portfolio of multiple assets and currency denominations. We can illustrate the simplicity of this approach—even in just the two-asset case—by assuming in our earlier example that the FTSE-100 portfolio is now run by a *dollar-based* asset manager. The total portfolio thus consists now of *two* positions—a sterling-denominated equity exposure, and a spot foreign exchange position to convert sterling into dollars.

Suppose the sterling price of U.S. dollars is £1.629/US$. The equity portfolio then is worth US$613,874 at the prevailing exchange rate. The combined portfolio of the dollar-denominated investor thus includes a US$613,874 position in the FTSE-100 (the equity primitive security) and an equivalent spot exchange rate position (the FX primitive). The one-month VaR of the equity is now [28]

---

[27] The data sets also include correlations.

[28] We could have come to the same result (save for rounding error) by taking our original equity VaR calculation and simply converting that VaR into U.S. dollars at the prevailing spot rate.

$$VaR_e = US\$613{,}874 \cdot (0.0076 - 1.65 \cdot 0.045) = US\$40{,}915$$

Because this considers the VaR of the equity portfolio *in isolation*, this risk measure is called the "undiversified VaR" of the equity position.

We now can calculate the undiversified VaR of the spot exchange rate position:

$$VaR_{fx} = US\$613{,}874 \cdot (:_{fx} - 1.65\Phi_{fx})$$

where $:_{fx}$ and Φ_{fx} are the mean and standard deviation of the monthly percentage change in the sterling/dollar rate. Again using the unconditional standard deviation (from January 1988 to January 1995) as our volatility estimate, we calculate Φ_{fx} as 0.0368 and $:_{fx}$ as –0.001. Substituting into the above equation, we ascertain that the undiversified VaR of the currency exposure is

$$VaR_{fx} = US\$613{,}874 \cdot (-0.001 - 1.65 \cdot 0.0368) = US\$37{,}888$$

We now can calculate the total portfolio VaR using a powerful property of the *bivariate normal* probability distribution. [29] Specifically,

$$VaR_p^2 = VaR_e^2 + VaR_{fx}^2 + 2\Delta VaR_e VaR_{fx}$$

where Δ is the correlation between FTSE-100 returns and monthly percentage changes in the sterling/dollar rate. We estimate Δ to be –0.2136 from our 1988-1995 sample—*i.e.*, the FTSE-100 varies *inversely* with changes in the sterling/dollar spot rate. Substituting that correlation coefficient and the other values into the above equation, we thus calculate the one-month portfolio VaR as follows:

$$VaR_p = [(\$40{,}915)^2 + (\$37{,}888)^2 + 2(-0.2136)(\$40{,}915)(\$37{,}888)]^{1/2} = \$49{,}470$$

This number represents the *diversified* VaR, or the one-month VaR (at the 95% confidence level) that reflects *both* the equity *and* currency exposures, as well as the correlation between the two. The fund manager thus should expect the portfolio to lose 8% or more of its current NAV in five of the next 100 months from its combined FTSE and sterling exposure. As one would expect, diversified VaR benefits from the lack of

[29] The property on which we implicitly rely is that the variance of a portfolio of two assets whose returns are distributed bivariate normal is a linear function of the variance of each asset return plus twice the correlation of the two returns times the two standard deviations. This result can be extended to portfolios comprised of more than two assets if returns are distributed multivariate normal.

perfect correlation—and, indeed, the negative correlation in this case—and is considerably less than the sum of the two undiversified VaRs.

Multi-Period VaR

The examples above have used monthly data to generate VaR measures for a one-month risk horizon. Because RiskMetrics data is indeed available at the monthly frequency, [30] no further adjustments need to be made if we care about a monthly risk horizon. To compute the VaR for longer risk horizons, however, the one-period VaR must be adjusted. This calculation is performed with the aid of *another* simplifying assumption in the variance-based approach—namely, return distributions are assumed to be *independent* and *stable* over time (collectively, "distributional stationarity"). This means that the VaR distribution we use to calculate one-month VaR is presumed *identical* to the distribution from which successive monthly returns are drawn. The multi-period VaR is then just the one-period VaR multiplied by the square root of the number of periods in the risk horizon.

Suppose we now wish to compute the diversified FTSE portfolio VaR at the 95% confidence level for a *one-quarter* risk horizon. To accomplish this, we need only multiply the one-month diversified VaR by the square root of three—*i.e.,* $49,470 x 1.7321 = $85,685. So, we expect losses in our international equity portfolio not to exceed $85,685 in 95 of the next 100 quarters.

Alternatives to the Variance-Based VaR Approach

As we saw in Figures 1 and 2, the undiversified variance-based VaR is almost exactly the same as the undiversified VaR computed using actual historical data. At first glance, this may seem surprising. Figure 2 and the sample statistics for the actual historical show quite clearly, after all, that the frequency distribution does *not* resemble the normal distribution.[31] Nevertheless, even when the underlying data is not normally

[30] Volatility and correlation data also is available for the daily frequency.

[31] Note in Figure 2 the fat tails of the actual distribution relative to the normal. This is consistent with the positive excess kurtosis statistic we examined earlier.

distributed, managers can sometimes get lucky—*i.e.*, the normal approximation is *sometimes* realistic *enough* for VaR calculations. But not always, and perhaps not even often.

Figure 3 shows the frequency distribution of monthly returns on the Nikkei 225 Japanese stock index from January 1988 through January 1995. As in Figure 2, a normal distribution is super-imposed on the frequency distribution whose mean and variance come from the underlying historical return data. Figure 3 and the sample statistics show this distribution is negatively skewed and has even fatter tails relative to the normal distribution than the FTSE-100. Not surprisingly, the variance-based undiversified critical VaR return — -11.68% when expressed as a return—is higher than the historical 5th percentile return of -12.62%. An undiversified VaR calculated using the assumption of normality thus *understates* the true risk of the Nikkei 225 portfolio.

As a general rule, variance-based VaR understates the risk of an asset or portfolio whenever the underlying distribution is negatively skewed and has tails fatter than the normal distribution (*i.e.*, "leptokurtic"). Quite a few asset classes fall into this category, including real estate, commodities, private placements, and some bonds. Many asset managers, not surprisingly, are unsatisfied with the assumption that variance is a complete measure of risk for all asset classes. In addition, long-term asset managers with risk horizons of a quarter or more often are equally displeased with the assumption of distributional stationarity in the variance-based VaR approach. Some investors thus opt to calculate VaR using a measurement methodology that does *not* rely on variance and distributional stationarity. [32] Two such methods include Monte Carlo simulation and historical simulation. [33]

Non-variance-based VaR can be computationally quite expensive. Advanced VaR calculation systems that do not force the user to rely on the assumption of normally distributed asset returns often cost over $1 million—sometimes $1 million *per annum* in leasing fees. Beyond that purchase/leasing price, more advanced calculation methods also

---

[32] Some investors also avoid VaR for this reason, as well, choosing instead to focus on "downside risk" measures such as below-target risk or downside semi-variance. *See* Culp, Tanner, and Mensink, *op. cit.*

[33] For examples of these approaches, *see* Jordan and Mackay, *op. cit.*

typically necessitate significantly more *data* (*e.g.,* historical security returns). In some cases, the cost of obtaining and maintaining this data can be as much if not more than the cost of the system itself.

Aside from the systems and data costs of advanced VaR measurement, the more complex the system the more costly it is to implement and maintain *from a labor cost standpoint.* A very advanced system often necessitates the creation of a full-time job to manage the inputs and outputs of the system. And the more complex the system, the more savvy its user must be.

Asset managers thus often face a relatively unpleasant tradeoff between cost and precision/realism in their VaR estimates. Not surprisingly, many managers just eschew VaR altogether, especially when the *applications* of VaR are far from obvious for many asset managers. If the job of an asset manager is to *take risk* in order to fund an uncertain liability stream and/or earn a competitive return on invested capital, after all, why should the manager spend over a million dollars on a VaR measurement system? In the next section, we attempt to answer that question by proposing some applications of VaR that neither require a huge expenditure on a VaR system nor needlessly attenuate the investment autonomy of portfolio managers.

APPLICATIONS OF VAR TO ASSET MANAGEMENT

In order for VaR to make sense to most asset managers, the investment policy must first be accepted as sacrosanct. VaR should *compliment* rather than *compete with* the primary investment management goals of the asset manager. It is a tool for helping managers determine whether the risks to which they *are* exposed are those risks to which they *think they are* and *want to be* ex posed. Value at risk will *never* tell an asset manager "how much risk to take." It will only tell a manager how much risk *is being taken*.

Taking the investment policy *as a given,* asset managers can apply VaR in at least four ways to the operation of their funds.

Monitoring

One of the primary benefits of VaR for asset managers is that it facilitates the consistent and regular monitoring of market risk. Institutional investors can calculate and monitor VaR on a variety of different levels. When calculated and monitored at the *portfolio* level, the risks taken by individual asset managers—whether internal traders and portfolio managers or external account managers—can be evaluated on an ongoing basis. Market risk also can be tracked and monitored at the aggregate fund level, as well as by asset class, by issuer/counterparty, and the like. We discuss three specific monitoring applications of VaR in the subsections below.

Internal Manager/Portfolio Monitoring

Suppose that Pension Plan Dearborn is a defined-benefit plan with a current NAV of US$1 billion invested through internal *and* external fund managers in all traditional asset classes. Pension Plan Dearborn calculates the quarterly VaR once a week for all of its portfolio managers. Suppose that External Account Manager Rush manages a $100 million international bond portfolio for Dearborn. If the VaR for Account Manager Rush is monitored each week, major departures of Rush's VaR from various established comparison risk measures should trigger an enquiry into Rush's recent investment activities—transactions that Pension Plan Dearborn's senior managers might otherwise have no reason to scrutinize.

Consider the specific case in which External Account Manager Rush has a quarterly diversified VaR that has averaged $10 million (at the 95% confidence level) over the last two years. That means that Pension Plan Dearborn expects to lose no more than 10% of its total investment with Manager Rush in 95 of the next 100 quarters. If the diversified VaR for Manager Rush is re-calculated every week and suddenly jumps to $50 million from the historical average of $10 million, the senior managers of Plan Dearborn might be inclined to wonder—and to ask—why. Because VaR is strategy neutral, only two answers are possible. First, the volatility in Manager Rush's international bond holdings has gone up (*e.g.,* following the Asian currency crisis), in which case the risk of the position simply reflects the risks of the asset allocation decision. Alternatively, Manager Rush has acquired new securities that expose Dearborn to significant additional market risks that were *not* contemplated by the asset allocation decision.

Aside from monitoring a manager's risk relative to its own historical risk profile, VaR also facilitates comparisons of the risks of one portfolio or manager with other portfolios or managers. To continue the above example, the increase in Manager Rush's VaR might *not* appear problematic if the VaR of other international bond managers engaged by Plan Dearborn rise at the same time. The same is true if the VaR of the benchmark used to evaluate the performance of Manager Rush rises proportionately. In these cases, Plan Dearborn may conclude that its market risk has increased because of its exposure to international fixed income as an asset class rather than because of its particular exposure to the investment strategies of Manager Rush.

For VaR to provide a useful monitoring benefit, precision in the measurement of VaR is *not* absolutely essential. In fact, the primary benefit of VaR monitoring comes from examining *relative* VaR, or the VaR of a manager or portfolio compared to the VaR of a benchmark portfolio, peer group portfolios, other internal or external managers, or the same manager *over time*. Even if the actual levels of VaR—$10 million and $50 million above—are imprecisely measured, the same measurement bias may affect other portfolios in the same way. [34] The theory is that these measurement errors "cancel out" when *relative* VaR is the focus instead of the absolute level of the VaR measure in question. Consequently, asset managers can derive a surprising amount of marginal benefit from monitoring even a variance-based VaR.

External Monitoring

The monitoring benefit of VaR is not restricted to internal and external portfolio management. Especially for asset managers whose portfolio holdings are not transparently available to investors at all times (*e.g.*, most hedge funds), a VaR reported to investors can help assuage any investors' concerns about market risk without necessitating disclosure of portfolio holdings. Similarly, regular reports of VaR to boards of directors or trustees can go a long way toward re-assuring these bodies that market risk is within the specified risk tolerance of the investment pool.

---

[34] This also may *not* be the case. Each situation should be evaluated on its own to identify sources of measurement error in the VaR statistic—and, in particular, whether measurement error is consistent across assets and portfolios.

Hedge Effectiveness Monitoring

Asset managers may use diversified VaR to monitor the extent to which their hedging strategies are accomplishing the desired objectives. To take a simple example, consider a $500 million mutual fund invested in domestic and international equities that hedges its exchange rate risk using currency forwards and futures. Suppose the one-quarter diversified VaR for the fund is 15% of NAV *without* including the currency hedges in the VaR calculation. The mutual fund manager can evaluate the effectiveness of the hedge (and analyze the extent to which returns and risks are affected by currency risk) by re-calculating diversified VaR *with* the hedges. If the diversified VaR of the *hedged* portfolio is 14.5% of NAV, the fund manager might question whether the hedging is worthwhile *or* whether the hedges have been properly implemented.

The use of VaR to evaluate the effectiveness of a hedging program depends strongly on the type of hedging program in place *and* on how the investment policy defines the hedging objective. To see why, consider two different examples, both of which concern a pension plan invested in international bonds and equities with a mandate to control its currency risk. [35] In the first case, suppose the pension plan specifies in its investment policy that no more than 1% of the current value *of any particular portfolio* can be exposed to exchange rate risk and that portfolio managers are left to hedge those risks on their own. The effectiveness of each manager's hedge can be evaluated by examining the diversified VaRs of each portfolio *separately* with and without the inclusion of the hedging contracts.

Now suppose the same pension plan specifies in its investment policy that no more than 1% of the current NAV *of the whole plan* can be exposed to exchange rate risk. Suppose further those managers do *not* do their own hedging, but that the plan engages a FX overlay manager to hedge its *consolidated* currency risk. The effectiveness of the overlay plan can be evaluated by comparing the plan's *aggregate* diversified VaR with and without the overlay program included.

---

[35] This situation is explored in more detail in Geoffrey Ihle, "Forward Hedges that Increase Value at Risk," *Derivatives Quarterly* 4(4) (Summer 1998).

Some interesting problems can arise when hedge effectiveness is monitored at a level different from the one specified in the investment policy as the hedging objective. Specifically, suppose the plan's international equity is invested entirely in Canadian dollar-denominated equities and that its international bonds are denominated entirely in Austrian shillings. [36] These two currencies happen to be *negatively correlated*. Consequently, if the pension plan specifies a hedging policy *by portfolio*, the independent hedging decisions of the bond and equity managers may achieve the plan's desired VaR reduction on an individual portfolio basis but may also increase the plan's *aggregate* diversified VaR. By partially hedging two portfolios whose currency risks are negatively correlated, the "natural hedge" is removed and the plan's consolidated diversified VaR rises. [37]

An overlay manager could, of course, construct a partial hedge that incorporates the negative correlation between the ATS and CAD holdings. For a plan interested in hedging its *aggregate* exchange rate risk, the overlay manager thus can achieve the desired VaR reduction. At the same time, if an overlay manager is used and the diversified VaRs of the individual portfolios are then evaluated, at least one of the portfolios will look *riskier* when the hedges are taken into account.

The usefulness of VaR in evaluating hedging programs thus depends strongly on the particular hedging objectives specified by the asset manager in the investment policy. Some asset managers may be more concerned about currency risk at the manager level than at the aggregate fund level. The fund may, for example, want to control the independent currency bets made by outside managers and thus may opt to require portfolio-specific hedging for *management* purposes. In that case, the fund would specify portfolio-level hedging requirements but then must monitor hedge effectiveness *at the portfolio level*. In the end, the plan *will* likely end up over-hedged (*i.e.*, speculating on currencies) in aggregate when some of the currencies held are negatively correlated. Such plans thus should *also* carefully monitor their aggregate diversified VaR to ensure this

[36] This example follows from Ihle, *op. cit.*

[37] This is illustrated numerically in Ihle, *op. cit.*

residual FX exposure does not become an unexpected (and possibly significant) risk factor.

By contrast, an asset manager whose appetite for currency risk is defined at the aggregate level should both hedge at the aggregate level *and* monitor the management of that risk using aggregate fund-level diversified VaR.

"What-If" Modeling of Candidate Trades

VaR also can be beneficial to asset managers that wish to eliminate transactional scrutiny by senior managers or directors and trustees. In this way, VaR can actually help give portfolio managers more autonomy than they might otherwise have without a formalized, VaR-based risk management process.

After the "great derivatives disasters" of the early 1990s, many directors and trustees of institutional investors became concerned with the risks posed by derivatives transactions. As a result, such transactions were prohibited in numerous investment policies and were subject to board-level approval in many others. Transactional monitoring using VaR can be an effective way of addressing this issue.

Suppose, for example, that the limited partners of a hedge fund are concerned about the possibility that the fund managers will engage in leveraged derivatives to augment their returns. [38] Because most hedge funds do not report their portfolio holdings on a regular basis, the limited partners might be inclined to ask the general partner to prohibit such investments. As a better alternative, the general partner might simply agree not to engage in any transactions that would increase the fund's VaR by more than X% of the fund's capital.

In order to minimize unnecessary scrutiny of particular trades, an asset manager need not require that VaR be calculated and reported *ex ante*. Especially if the cost of a

[38] Not all derivatives are leveraged. Our use of the term here refers to *formula* leverage rather than margin. A swap in which the hedge fund receives a fixed rate of 8% and pays LIBOR against a notional amount of $1 million would not be formula-leveraged. A swap in which the fund receives 8% fixed and pays LIBOR *squared* would be.

VaR system is an issue, few affordable VaR systems allow for this type of real-time computation. Nevertheless, the requirement could be instituted and enforced *ex post*, with the corresponding requirement that any trades in violation of the "maximum marginal VaR" requirement would be liquidated or hedged within, say, a week of the deal.

As a cautionary note, asset managers must be attentive to the means by which the VaR of a particular transaction is calculated in order for this application of VaR to make sense. One VaR-like statistic, proposed by Mark Garman, is called "DelVar" and examines the impact of a particular trade on the VaR of a portfolio. [39] Although quite useful in its own right, DelVaR would *not* be appropriate for the application of VaR discussed here. The particular measurement method proposed by Garman is useful for evaluating the impact of a particular security on the VaR of a portfolio *only for small presumed changes in underlying prices*. DelVar badly underestimates, however, the marginal risk of certain trades for *wide* price swings. In order to apply VaR in the manner discussed here, the portfolio VaR would have to be *fully re-estimated* with and without the candidate trade rather than simply approximated using a measure like DelVar.

Risk Targets and Thresholds

A third application of VaR to asset management involves measuring and monitoring market risk using a formal system of pre-defined risk targets or thresholds**.** In essence, risk thresholds take ad hoc risk monitoring one step further and systematize the process by which VaR levels are evaluated and discussed for portfolios or managers—or, in some cases, for the whole investment fund.

A system of risk thresholds is tantamount to setting up a tripwire around an investment "field," where the field is characterized by a fund's investment policy and risk tolerance. This tripwire is defined in terms of the maximum tolerable VaR allocated to a manager or portfolio and then is monitored by regularly (*e.g.*, weekly) comparing actual VaRs to these pre-defined targets. Investment managers are permitted to leave the field when they wish, but the tripwire signals senior managers that they have done so. When a

---

[39] Mark Garman, "Improving on VAR," *Risk* 9(5) (1996).

tripwire is hit (*i.e.*, a VaR threshold is breached), an "exception report" is generated, and discussions and explanations are required.

Risk targets can be specified in terms of absolute or relative VaR. A private bank might conclude, for example, that a particular client's capital should never be placed at risk above a certain amount *regardless* of the risks taken by other clients or managers. In that case, the traders on that client's account could be subject to an absolute VaR threshold. A mutual fund, by contrast, might prefer to specify its risk targets relative to the VaR of its benchmark portfolio or peer group.

The hallmark of a well-functioning risk target system is not that targets are never breached or that all exceptions are rectified through liquidating or hedging current holdings. Rather, the primary benefit of a risk target system is the formalization of a *process* by which exceptions are discussed, addressed, and analyzed. Risk thresholds thus are a useful means by which asset managers can systematically monitor and control their market risks without attenuating the autonomy of their portfolio managers. Because the primary purpose of risk limits is to systematize discussions about actual market risk exposures relative to defined risk tolerances, huge investments in VaR calculation systems, moreover, typically are not required. Even an imprecise measure of VaR will usually accomplish the desired result of formalizing the risk monitoring process.

Risk Limits and Risk Budgets

A more extreme version of risk targets and risk thresholds is a system of rigid risk limits. This application of VaR is also known as a "risk budget." In a risk budget, the fund's total VaR is calculated and then allocated to asset classes *and* specific portfolios in terms of absolute and benchmark-relative VaR, as well as shortfall at risk ("SaR"). [40] Managers are then *required* to remain within their allocated risk budget along these risk dimensions. So, whereas risk targets resemble a tripwire around a field that managers must account *ex post* for crossing, a true risk budget instead acts as an electric fence around the field that managers simply cannot cross *ex ante.*

---

[40] SaR is directly analogous to VaR with the *net* asset/liability position subjected to the risk calculation and summary measure. *See* Culp, Tanner, and Mensink, *op. cit.*

A total risk budget defined across all portfolios can create numerous problems for an asset manager. First, risk budgeting relies at some level on the absolute VaR of a fund and its portfolios. To the extent that the measurement methodology is flawed, the risk budget will be wrong. If the VaR measurement methodology is more biased for some asset classes or security types than others, moreover, some managers could be penalized or rewarded simply because of flaws in the measurement methodology. In the extreme, relatively riskier funds could be given a risk budget that is *too high*, whereas relatively safer funds could be allocated *too little* VaR.

Second, risk budgeting defined across both asset classes and portfolios can contradict and call into question the fund's asset allocation decision. This can be especially problematic when a fund manager's board must approve changes in the asset allocation *unless* hitting a risk limit in the risk budget triggers the change. Suppose a pension plan allocates capital into asset classes *annually* using traditional mean-variance asset allocation and portfolio optimization techniques. Then suppose the plan defines a VaR budget for asset classes and portfolios, where VaR is measured using a variance-based approach. If the risk budget is enforced more frequently than annually, the risk budget will call into question the asset allocation *simply because volatility changes in the markets on a regular basis.* Variance-induced changes in VaR thus prompt a shift in the asset allocation through the risk budget. Even though the *practical* consequence is a change in the asset allocation itself, the *actual* trigger is the risk budget; the board may never be consulted. [41]

To avoid this problem, risk budgeting should be limited to re-balancing funds between portfolios *within the same asset class*. Even then, asset managers contemplating a risk budget will need to allocate a considerable sum of money for the VaR calculation system to ensure that the calculation method is not biased against particular managers or financial instruments.

CONCLUSION

[41] Ironically, a risk budget may be supported by a board, which thinks the budget reduces the need for board-level micro-management. Yet, to the extent that the risk budget simply grants fund managers license to circumvent the board-approved asset allocation, the opposite intention will have been achieved.

Many asset managers have avoided or criticized VaR based on the notion that systematic measurements and disclosures of risk serve only to attenuate the autonomous nature of the investment management process. On the contrary, measuring VaR and using it as the basis for internal monitoring and risk targets, external risk disclosures, and transactional risk evaluations can actually give the asset manager *more* autonomy than if investors or senior managers are unsure of what the fund's market risk exposures actually are. A sound VaR-based risk management system should take the investment policy *as given* and should seek only to help managers and investors ensure that the risks to which the fund *is* exposed are those risks to which it *thinks it is* and *wants or needs to be* exposed.

Most cost-effective systems for measuring VaR rely on simplistic and often unrealistic assumptions. Nevertheless, the benefit of most VaR applications for asset managers traces more to how the VaR estimate is used than to the calculation methodology. Especially for asset managers with exposures in multiple currencies, even simplified VaR can be an invaluable tool for distilling market risk into one summary statistic. It is no panacea, but an asset manager that measures its VaR may better be able to manage its primary investment business than one that is largely unaware of its consolidated market risk exposures.

Figure 1: FTSE-100 Monthly Returns (1/88 - 1/95)

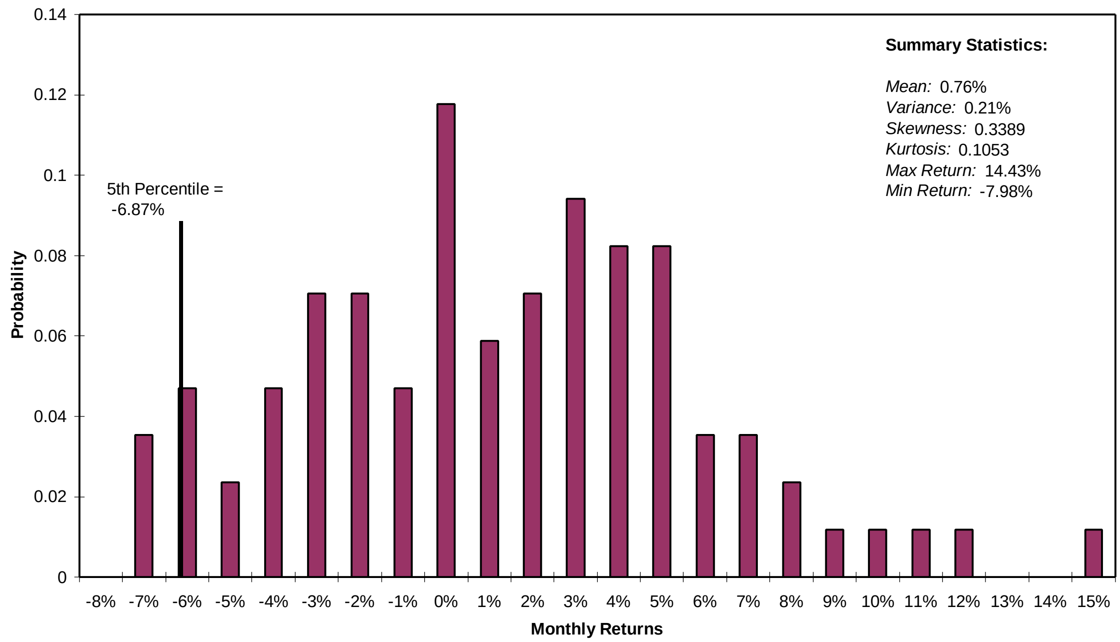

Figure 2: Monthly VaR Distribution for FTSE 100 (1/88 - 1/95)

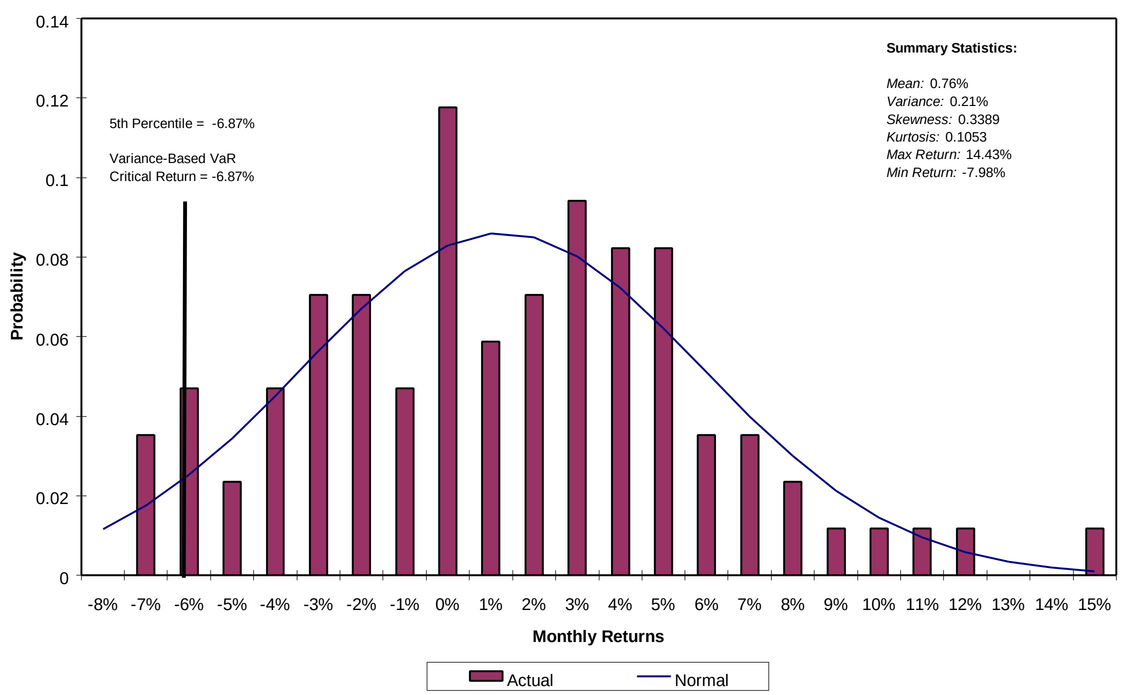

Figure 3: Monthly VaR Distribution for Nikkei 225 (1/88 - 1/95)