The Measurement and Management of Foreign Exchange Risk in Emerging Markets

By Russell Thompson

This article addresses some of the issues related to the management and measurement of foreign exchange risk as applied to emerging markets. It is, however, just as relevant to any asset classes and asset managers with nonnormal distributions, particularly those with fat tails such as hedge funds.

The measurement and management of foreign exchange (FX) risk, particularly in emerging market portfolios, often is not fully understood. The Bank of International Settlements has proposed a "shorthand" method and a "simulation" method to measure FX risk. Essentially shorthand risk measurement refers to an outright notional expression of risk, loosely measured by some form of leverage calculation, and simulation risk refers to a simulation of the possible impact on a foreign exchange portfolio by measuring historical price action to give some degree of insight into potential profit-and-loss scenarios based on this price action. This approach commonly is referred to as Value-at-Risk (VaR) analysis.

The shorthand method of risk measurement is appropriate to only simple FX portfolios and has many flaws in its approach. Leverage is a poor measurement of risk when referring to foreign exchange because it is not related directly to any underlying variables that impact the real returns at risk in a portfolio. It completely ignores factors such as correlation among currencies and the inherent volatility underlying various currencies, which directly impacts the actual risk being run. This is a very obvious conclusion and I shall not dwell on it any further here.

> "VaR typically is the risk measure of choice for FX managers and risk departments because it expresses a portfolio's risks in a coherent and logical manner."

Value at Risk (VaR)

VaR typically is the risk measure of choice for FX managers and risk departments because it expresses a portfolio's risks in a coherent and logical manner. It is expressed in real profit-and-loss terms and can directly tell a risk manager the potential risks inherent in a portfolio based on varying degrees of statistical confidence. VaR traditionally is measured in the following three ways:

1. historical simulation
2. variance/covariance (parametric)
3. Monte Carlo simulation

Each method produces a statistical measurement of VaR that is calculated using an historical data assumption to give a level of confidence that is determined from the historical price action. Each method differs in complexity and has advantages and disadvantages.

Historical simulation assumes that the past is a good predictor of the future and that the volatility of the analyzed currencies will remain stable, within the parameters observed in the past. It uses real historical data and therefore importantly does not assume that the returns are normally distributed. It is, however, computationally intensive and completely dependent on historical price movements, and therefore it can seriously underestimate "tail risk." (Tail risk is a measurement of the probability of an event occurring at the extremes of a given distribution, the reasons for this will be explained later in this article.) Historical simulation is also dependent on the quality and depth of the input data, which can be problematic for emerging market currencies.

Variance/covariance, sometimes known as parametric VaR, is computationally easier because historical data is used to calculate the standard deviation of the changes of risk factors and the correlations between them. It is heavily disadvantaged by an assumption of the linearity of risk (the assumption that risk vs. reward is linear in nature, which is not the case with more-complex financial instruments such as options), that correlations are stable over time, and that returns are distributed normally. Products such as options introduce nonlinear risk parameters, correlations can quickly and dramatically uncouple, and as we have seen in the real world, distributions are far from normally distributed.

Monte Carlo simulation utilizes historical data by randomly creating many scenarios for different currency rates and then using VaR to estimate a worst-case scenario by pulling apart the assumption of linearity of returns and looking at thousands of potential portfolio outcomes, without using any predetermined assumptions about correlation between the currency rates. It does, however, have the same disadvantage as parametric VaR, because to predict VaR it is usually necessary to assume that the distribution of returns is either normal or log normal. The real world, however, has many types of distributions, and they are not necessarily normal or log normal. This is discussed in more detail below. Monte Carlo simulation also is very computationally intensive.

FIGURE 1: U.S. DOLLAR / INDONESIAN RUPIAH EXCHANGE RATE 1997–2005

Source: Bloomberg

Why is VaR Flawed, and Why Should You Care?

VaR suffers from three major drawbacks, some or all of which apply to each measurement of VaR.

Assumption of Normal Distributions

Both parametric and Monte Carlo VaR typically assume that returns are distributed normally. The calculation of VaR is not the only measurement that is flawed in this way; most risk and return measurements have the same problem. If you are a fund-of-funds manager who is comparing the Sharpe ratios of your managers for allocations, then you are assuming that their returns are all distributed normally. The problem here is that the Sharpe ratio uses the standard deviation, which is a flawed risk measure for nonnormal distributions (Cascon and Shadwick 2006). Unfortunately, in the real world, distributions are rarely if ever normally distributed, and therefore using a calculation based on the standard deviation with a nonnormal distribution can impair significantly the accuracy of any confidence interval assessments, especially if the distribution has fat tails. Yet risk managers still insist on utilizing standard deviation to express risk confidence levels. Numerous studies have been carried out proving that distributions of returns are nonnormal in the real world, so this assumption is plainly ludicrous. As Cascon and Shadwick (2006) pointed out (tongue in cheek): Take a given set of measurements of height from a population sample. Calculate the mean and the standard deviation. Depending on the size of your data set, there will be a level of confidence that will predict the likelihood of a human being with negative height.

Size and Reliability of the Data Set

Historical VaR calculations do not assume a normal distribution of returns, because the distribution of returns is derived from the actual return and rate observations. The problem is that the VaR in this form either overestimates or underestimates the potential risk, because it is strongly influenced by recent data observations. Two-hundred and fifty data observations are typical, but how far back should you go? In a period of relative calm, 250 data observations will strongly underestimate the potential VaR. Many foreign exchange practitioners will remember the Asian and Russian currency crises in 1997 and 1998–1999. We know and have seen that the Indonesian rupiah can move by hundreds of a percent in a short period of time. A typical VaR model will seriously underestimate the probability of such an event happening again, because in all likelihood the period of high volatility will be outside of the utilized data set.

For example, as shown in figure 1, utilizing data that did not take into account the volatility between late 1997 and 1998 would seriously underestimate the potential VaR inherent in the U.S. dollar vs. Indonesian rupiah exchange rate.

Why should you care? In emerging market currencies the distributions of both the underlying currencies, and the returns of the managers that trade them, are clearly nonnormal, because these markets are by implication less efficient, less liquid, and more prone to event risk

(Malkiel and Saha 2005). Managing a portfolio of emerging market currencies with an assumption that tail risk is a function of a normal distribution can seriously underestimate the potential risk in your portfolio. Utilizing historical returns will exacerbate this problem, because in the vast majority of cases, particularly after a period of relative calm, the historical data will underestimate the potential tail risk and further lull an investor into a false sense of security.

More bad news awaits. Emerging markets, in times of stress, exhibit strong signs of clustering. Your risk manager will probably base risk calculations on assumptions of correlation between asset classes and within asset classes. These relationships can go out the window when stress hits, and an imprecise knowledge of your real distribution of returns, combined with an inaccurate assumption of your real tail risk, can leave you especially vulnerable.

Liquidity

VaR takes no account of liquidity conditions. Even in foreign exchange, the most liquid of emerging market asset classes, liquidity can evaporate in times of stress. VaR, particularly in large portfolios, typically will not take account of the explicit practicalities of liquidating a position. Feel free to check that statement with the risk manager in Société Générale, or Amaranth, or Long Term Capital Management. Whether it is oil and gas calendar spreads or Russian bonds, it is vitally important to overlay explicit concentration limits based on regularly reviewed liquidity conditions available to you in the market.

How to Measure the Real Distribution of Returns and Compare Performance to a Normal Distribution

Explicitly Measuring Tail Risk

It is possible to measure the tail risk explicitly using the CS character (Cascon and Shadwick 2007). The first CS character is the ratio of the standard deviation to the standard dispersion. The standard dispersion of a distribution is a statistic that is defined in terms of the slope of the Omega function at the distribution's mean. It is calculated by taking the average of the sum of the absolute individual differences between the returns and the mean of the returns divided by two (Cascon and Shadwick 2007). A normal distribution will have a CS character of 2.51. The standard dispersion is a direct measure of the degree of concentration of returns around the mean—so when two distributions have the same standard deviation, the one with the smaller standard dispersion will have returns more-concentrated around the mean. For a fixed standard deviation, the smaller the standard dispersion the greater the degree of concentration about the mean and the higher the CS character. Return distributions greater than 2.51 have a greater degree of concentration around the mean, but they have fatter tails than normal distributions and therefore more tail risk. Events far from the mean are much more likely than in the normal case. The movement over time of a manager's CS character can and should be relevant to investors, particularly those invested in emerging market currencies, where tail and event risk is intrinsically higher. The clear implication here is that managers should dynamically monitor the tail risk inherent in their distribution of returns and adjust risk exposures appropriately as tail risk rises or falls over time. A CS character, however, does not tell us about the specific variation of the distribution through the entire spectrum of returns or how the likelihood of a lefthand tail event compares with a righthand tail event.

The Omega Function

A good place to start to try overcoming this latter problem is to compare your observable distribution of returns directly to a normal distribution. In this way we can ask if at certain levels of return (x) whether the distribution is superior or inferior to a normal distribution. How do we do this? We start by plotting our observed Omega function, developed by Con Keating and William Shadwick (2002). In recent years many investors have begun looking at Omega functions to gain a better understanding of their exhibited distribution of returns.

In simple terms, the Omega function of a distribution corresponds to the amounts considered as a win multiplied by their corresponding probabilities, divided by the sum of the amounts considered as a loss multiplied by their various probabilities (Keating 2004). Conceptually you can view the Omega function as the ratio of all the upside relative to return (x) to all the downside relative to return (x). Or it can be viewed as the ratio of the value of a virtual call at return (x) to the value of a virtual put at return (x). This is the best way to evaluate the Omega function at return (x) because it removes the need to evaluate the upside and downside deviations separately.

The Omega function can be plotted for a manager who trades emerging market currencies and compared with a normal distribution at any value of return (x). Figure 2 shows the log Omega function of one currency strategy, the Cambridge Strategy Asian Markets Alpha Programme (for which the author of this article is responsible for the trading and management of risk), and compares it to a normal distribution. Note the following:

- At any particular return (x) the larger the absolute Omega number the higher the quality of the bet, i.e., the virtual win when compared to the virtual loss is better for the distribution with the higher Omega function.
- The Omega function is positive and monotonically decreasing. It goes from infinity on the left to zero on the right. A function is superior if it is above and to the right.

- It takes the value 1 at the mean of the distribution.

In simple terms, the flatter the Omega function, the fatter the tails of the distribution. A rational investor will prefer to allocate to a distribution with an Omega function that dominates another Omega function above and to the right. This has powerful implications for risk management in that the Omega function utilizes all the information from a distribution and does not require estimating any moments except the mean.

From this analysis we create a firm foundation. An Omega function with a steeper slope under the value of 1 (the mean) than that of a normal distribution with the same mean dominates the normal distribution, and it can be assumed safely that any VaR analysis that utilizes an assumption of normal distributions will result in a worst-case scenario, because the put-call parity for all Omega functions less than 1 will be superior to that of the normal distribution. In other words, the return distribution of the Asian Markets Alpha Programme has a lower likelihood of producing lefthand tails than a normal distribution. In contrast a flatter Omega function for any value greater than 1 will have a higher probability of producing fatter righthand tails than a normal distribution with the same mean. In this case the Asian Markets Alpha Programme has a higher likelihood of producing more returns in the righthand tail than the corresponding normal distribution.

This, however, provides a new problem. We now are comfortable that our VaR analysis is not overstating our risk compared to a normal distribution. However we often are faced with the problem that we soon run out of data values at the extrema of our time series; i.e., when the observed returns finish, there is no statistical relevance. To get around this, we use some extreme value theory technique to fit the observed data to a piece-wise distribution (using maximum likelihood estimation techniques), then sample from that many times at various points (Pickands 1975; Balkema and de Haan 1974). This would prove statistically much more robust, but ultimately we still are extrapolating into the unknown. This is problematic when talking about extreme events, which can have a significant impact on our risk calculations.

FIGURE 2: THE OMEGA FUNCTION OF THE CAMBRIDGE STRATEGY'S ASIAN MARKETS ALPHA PROGRAMME AGAINST A NORMAL DISTRIBUTION

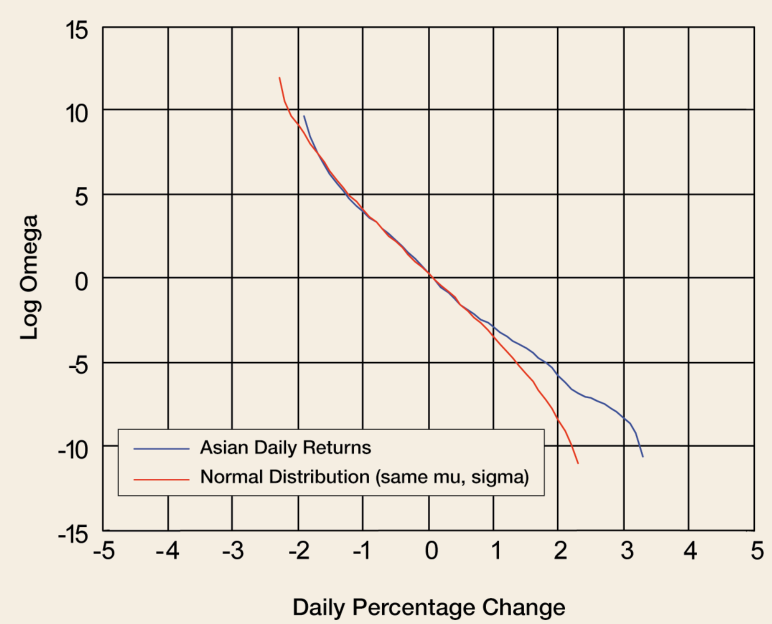

It would be possible, after calculating the Omega functions for two separate distributions with the same mean, to weight their slopes and various values above and below return (x) to provide an objective score that could be used to directly compare distributions and allocate to managers that exhibit higher scores. These scores would encompass information inherent over the entire distribution of returns, account for tail risk, and be superior to using Sharpe ratios.

Conclusion

Emerging markets exhibit specific characteristics that investors and risk managers must be cognizant of when measuring, monitoring, and evaluating risk. VaR is a powerful tool and it encompasses a great deal of important information, but it should not be the end of the risk management story. Explicit tail risk needs to be measured and monitored, and it is strongly recommended that investors consider using explicit risk parameters that measure uncorrelated VaR, where clustering phenomenon may break down long-held correlation relationships. Investors also need to carefully assess liquidity constraints, and limits should specifically be put in place to limit exposure to illiquid currencies and assets. This requires making assumptions about liquidation policy, which dictates capacity constraints.

It is sensible to proactively monitor and manage tail risk and adjust appropriately to account for this risk. A manager may have an exceptional Sharpe ratio with limited drawdowns but still have significant tail risk inherent in the distribution of returns. A manager who

Continued on page 30

Thompson

Continued from page 24

exhibits an Omega function inferior to that generated by a normal distribution with the same mean will contribute significantly more VaR than suggested by his Sharpe ratio. In this case, investors should expect more tail events than predicted by the posted VaR limits. Depending on the type of VaR methodology used, this could compound the problem. Buyer beware.

Russell Thompson is chief investment officer of the Cambridge Strategy (Asset Management) Limited, a foreign exchange asset management firm based in Mayfair, London. He earned a B.Sc. in economics from Loughborough University. Contact him at russell@thecambridgestrategy.com.

References

Balkema, A. A., and L. de Haan. 1974. Residual Life Time at Great Age. *The Annals of Probability* 2, no. 5 (October): 792–804.

Cascon, Ana, and William F. Shadwick. 2006. The CS Character and Limitations of the Sharpe Ratio, *Journal of Investment Consulting* 8, no. 1 (summer): 36–52.

———. 2007. The Standard Dispersion and Its Application to Risk Management. *Journal of Investment Consulting* 8, no. 2 (summer): 38–54.

Keating, Con. 2004. Risk and Reward. Presentation at The Institute of Risk Management Forum 2004. Available at http://www.theirm.org/riskforum/presentations2004.html.

Keating, Con, and William F. Shadwick. 2002. "A Universal Performance Measure." The Finance Centre. Available at http://www.isda.org/c_and_a/pdf/GammaPub.pdf.

Malkiel, Burton G., and Atanu Saha. 2005. Hedge Funds: Risk and Return. *Financial Analysts Journal* 61, no. 6 (November/December): 80–88.

Pickands, J. 1975. Statistical Inference Using Extreme Order Statistics. *The Annals of Statistics* 3, no. 1: 119–131.