TechHackers'

Quant*News*

July 23, 1999

Welcome to the fourth issue of Quant*News*.

Correlation and Dependence

With the development of more advanced derivative products, which may depend on the prices of several assets, and with a renewed emphasis on risk, the study of correlations has gained new importance. While there are many aspects to correlation, here we will look at just one: the relationship between correlation and dependence.

If the correlation between to assets is zero, we often assume that they are independent. For example: if two assets X and Y have zero correlation and individually have value at risk VaR_X and VaR_Y, we would assume that the VaR for the entire portfolio is $(\mathrm{VaR}_X^2 + \mathrm{VaR}_Y^2)^{1/2}$. However, this assumption is only valid if the two assets are actually independent, not just uncorrelated.

We often treat "uncorrelated" as meaning "independent"—and in some cases, they are the same. For example, random variables with a joint normal distribution *are* independent if they are uncorrelated. And of course, normal random variables are our "favorites". But for other distributions, this relationship is not true—here is a simple example:

Consider two random variables which each (individually) take the values $-1, 0, 1$ with probability $1/3$. The joint probability distribution is given by

$$p(0,0) = 1/3,$$

$$p(-1,1) = p(-1,-1) = p(1,-1) = p(1,1) = 1/6.$$

Clearly these random variables are not independent, although they are uncorrelated. For example: consider an option which pays the difference of the two random variables if it is positive. The expected payoff for independent random variables would be

$$\frac{2}{9} \cdot 1 + \frac{1}{9} \cdot 2 = \frac{4}{9},$$

while the answer in our case is

$$\frac{1}{6} \cdot 2 = \frac{1}{3}.$$

Thus independence implies no correlation, but the reverse is not true for non-normal random variables.

The Impact of Liquidity on Market Risk

Background

Many institutions now use Value-at-Risk (VaR) to quantify the market risk inherent in a portfolio of investments. The VaR number is supposed to indicate the maximum possible reduction in portfolio market value that may occur within some time horizon, given a certain confidence interval. Time horizons are generally short, often only one day, and confidence intervals are typically 95% or greater. Thus, a VaR of $10 million on some portfolio, given a horizon of 5 days and a confidence level of 99%, would indicate that there is only a 1% chance that the portfolio would lose more than $10 million in value over the next 5 days.

There are several widely used methods for calculating VaR. Each of these methods is subject to various customizations and optimizations. Many commercial VaR applications are ultimately based on the RiskMetrics™framework, originated by J.P. Morgan. The RiskMetrics dataset, which is now published daily by Reuters, contains recent volatilities on and correlations between over 400 values that are derived from highly liquid markets.

Most importantly for fixed income investors, the dataset contains information on all heavily traded yield curves, including U.S. Treasuries and U.S. Dollar LIBOR. However, the dataset does not contain information on yield spreads between less liquid investment types (e.g. off-the- run Treasuries, US Government Agencies, MBS) and the more liquid instruments.

The Problem

Most commercial VaR calculation methods focus on the risk inherent in fluctuating yield curves, commodity and equity prices, and currency exchange rates. Recently, many risk managers have begun to focus on risks tied to impaired liquidity in financial markets. Unfortunately, most VaR models tend to understate liquidity risk, and many ignore it altogether.

First of all, a distinction needs to be made between endogenous and exogenous liquidity risk. The former is associated to the size of an individual dealer position in relation to the depth of a—typically liquid—market, while the latter is a structural—or temporary—condition of an entire market.

There are three liquidity effects that should arguably effect a VaR calculation:

- **Endogenous Liquidity Risk**: the effect that the liquidation of a position might have on the market as a whole. That is, the risk that the very act of liquidating a position (e.g. in a "fire sale") will send a "signal" to the market that will change the behavior of other participants.
- **Time-to-liquidation**: It takes longer to liquidate illiquid investments, and the VaR should reflect the volatility during this liquidation period. While most VaR models allow the user to vary the horizon for an entire portfolio, it is often not possible to vary the horizon for individual positions.
- **Spread Volatility**: The value of most investments depends not only on fluctuations in market interest rates (i.e. yield curves), but also on the spreads to those curves. The effect of the widening and

tightening of spreads, which is chiefly driven by changes in market liquidity, is completely ignored by most models. Given the difficulty in hedging spread exposure, many portfolios are significantly exposed to spread changes.

Of course, these effects are not entirely independent of one another.

Solutions

There are several approaches that can be used to account for each of these liquidity effects. These vary in their ease of application as well as in the degree to which they take liquidity into account. We will briefly outline some of them, describing their merits, and limitations.

Endogenous Liquidity Risk. One way to model this risk, following the work of Colin Lawrence, is through a cost function, which accounts for transaction costs, exposure costs and hedging costs. The value of the cost function, when minimized with respect to the time to liquidation, provides the liquidity adjusted VaR. Standard VaR takes into account only exposure cost to the liquidation horizon.

Alternatively, one can model the optimal liquidation process in terms of two stochastic quantities: the quantity discount $0 < c(S) < 1$, a non-increasing function of the position size S, and the execution lag $D(S)$, a non-decreasing function of S. Each has an associated volatility.

Both the above approaches suffer from requiring an extensive analysis of historical data.

Time-to-liquidation. A simple approximation consists in magnifying each position in the portfolio by a factor that accounts for the time it takes to liquidate that instrument from the portfolio and then compute the VaR on the new portfolio. This is like computing the VaR for the original portfolio having several correlated time horizons over a single time horizon. The determination of the multiplicative factor is very straightforward under the usual assumptions of log-normality of the returns. This approach has the drawback of not including finer details on the correlations among the different positions and their relative size. However, it already captures the most significant effects.

Another approach consists of evaluating the VaR based on a time horizon determined on the basis of a weighted average determined from the liquidity of the positions. The latter is obtained from the statistical analysis of the time it takes to liquidate each position, based on security types. This methodology has the advantage over the previous one of making it very simple to obtain the value of the VaR from a given effective time horizon to the another, as it can be needed when new data arrive and modify the estimate of the time to liquidation.

All the above methods suffer from the drawback of being *static* methods—in the spirit of the VaR. When more subtle effects need to be captured, such as the variation of the market risk due to dynamical variations in the portfolio, a scenario generation approach can be used.

Spread Volatility. Many portfolios are considerably more exposed to volatility in yield spreads than they are to volatility in underlying rates. This was made painfully obvious in the crisis of August 1998, where off-the-run, corporate, mortgage, and other spreads widened sharply, even as interest rates on on-the-run treasuries fell.

Within a VaR framework one can account for liquidity spread in treating each position as depending on the liquidity spread in addition to the other components that enter in determining the price of each instrument. Then, a delta-gamma approach can be used—where now the gamma has a liquidity spread component too—and then use the standard delta-gamma methodology.

As an alternative, one can model the liquidity spread as another stochastic factor, running thus a Monte Carlo method to generate a number of scenarios. The results of the simulation will give the distribution of the returns (as an added bonus, one is also freed by this approach from the log-normal hypothesis). Thus, the VaR to the desired confidence level is obtained.

To put it in a more formal context, the liquidity correction to the price movement of an individual security can be expressed as an additive term to the standard VaR displacement. For example, the VaR displacement at 99% is $P \cdot (1 - \exp(-2.33 \cdot \sigma))$ and the liquidity correction $0.5 \cdot P \cdot (\mathrm{sp} + a \cdot \sigma_1)$. Here P is the current price of the security, σ is the standard deviation of returns, sp is the historical average of the "percentage" bid-ask spread and σ_1 its historical standard deviation. If we further make the assumption that there is a very high correlation between price and spread fluctuations and a Gaussian distribution for the spreads, then $a = 2.33$.

This procedure can be extended at the portfolio level by calculating a portfolio bid-ask spread distribution as a weighted average of the distributions of the component securities and so obtaining a single additive liquidity correction for the entire portfolio.

The basic economic justification for the correction is that, in illiquid markets, the marking to market of the security should be done at the bid price, rather than the reference midpoint price used in the standard VaR calculation. In liquid markets, under normal conditions, the realized intra-dealer bid-ask spreads are much smaller than the quoted values (which justifies the usual VaR procedure) but this may not hold in turbulent or structurally illiquid markets, where the bid price—if any—becomes the only liquidation price.

This treatment should be seen as complementary to the analysis of endogenous liquidity risk.

Conclusions

The VaR framework can be expanded to handle the liquidity risk effects that have been traditionally ignored by many market risk analysis techniques. The several techniques differ in ease of development and use. Practitioners need to assess their needs and choose a technique that will be maintainable and usable in their business context. For example, if frequent intra-day VaR calculations are necessary, performance may be an important consideration. The availability and quality of the data from which volatility and correlation assumptions will be derived should also be considered.

Monte Carlo Simulations: Macro vs. Micro Risk Management

Part 1 of a series by Maurizio Mondello

When considering the application of Monte

Carlo (MC) to market risk management, one often concentrates on the process of macro-risk management, such as exemplified by (firm-wide) VaR calculations. Here we want to emphasize the similarities and differences between the use of MC for VaR calculations and the its use in the process of deal / portfolio hedging and marking-to-market (micro-risk management).

One important point is that the two levels need to be integrated. For example, the sensitivities obtained at the deal / portfolio level provide inputs for Delta-Gamma VaR, while recalculation of the present values (marking-to-market of deal / portfolio) is required as part of a stress-test / scenario analysis. Simulations can therefore be used at different stages (levels of aggregation) of the same risk assessment process, making issues of efficiency (which will be discussed in a following note) all the more compelling.

While the mechanics of a basic MC simulation does not greatly depend on its specific context, in order to develop effective optimization strategies it is important to understand the differences in inputs and outputs associated with different applications of the simulation method. The pricing of individual deals by Monte Carlo simulation is based on the use of the risk-neutral distribution. Under the corresponding probability measure, expected returns on all assets are equal to the riskless interest rate, which is also used for present value discounting. On the other hand, Monte Carlo simulation for VaR requires the use of the "real world" (historical) distribution of returns. Typically, here we are interested in the future value (no discounting) of a static portfolio (no dynamic hedging).

With regard to the output, in deal pricing we want to determine the average ("fair") value with respect to the risk-neutral distribution, i.e. we are performing a Monte Carlo integration under the risk-neutral measure over the life of the deal. Compare this with the VaR calculation, where we are interested in determining a quantile in the distribution of possible values for our "book" over a preset time horizon. Monte Carlo is used here to map the multivariate distribution of risk factors into the one-dimensional distribution of (future) "book" values. The dimensionality of the simulation problem (as exemplified by the multivariate distribution underlying the VaR calculation) is perhaps the single most important element to consider, when assessing the potential of different general (non product-specific) optimization techniques. The dimensionality of the problem is determined by two elements: the number of risk factors (underlyings) involved (e.g. in a basket option or VaR calculation) and the number of time steps required for each factor (e.g. in an Asian option or general path-dependent option). When both factors are present, the dimension of the (simulation) problem is given by the product of the two numbers.

In a subsequent note, we will discuss how some of these methodological considerations affect the practical choices involved in the development of a general-purpose simulation engine for market risk applications.

In Conclusion...

We hope you've enjoyed issue 4 of Quant*News*. Comments, submissions, and other requests should be sent to your editor at sjanowsky@thi.com.