How to model VAR over more than one day

The simplest option – scaling by the square root of time – is the best, argues David Rowe

When the Bank for International Settlements (BIS) and national banking regulators sanctioned the use of internal market risk models for calculating regulatory capital, they had to establish basic ground rules. The main aim was to establish a common standard for the risk measure being calculated. A secondary objective was to mandate a framework that would produce a level of required capital that was acceptable to senior regulatory officials.

This resulted in a multiplicative factor to be applied to the model-based risk estimates in arriving at the required capital. This factor is set equal to three if model performance is acceptable based on back-testing, but it may be set higher for specific institutions at the discretion of national regulators. When asked to justify the value of this multiplier, I have heard officials describe it as a "regulatory comfort factor". To keep this factor from appearing too large, however, the parameters for the value-at-risk estimate were set higher than was typical for pre-existing internal risk estimates. Specifically, the confidence level was set at 99% and the simulation period was set at 10 days. This left open an important question, namely: "How should VAR be modelled over a period of more than one day?"

It is widely recognised that standard overnight VAR estimates incorporate a questionable assumption, namely the application of one-day market changes to a static portfolio. In reality, traders adjust their positions throughout the trading day. Given that avoiding losses is very much in a trader's interest, it is likely that assuming a static portfolio for 24 hours is a conservative procedure. It is to be expected that intra-day adjustments will consistently narrow the range of losses from any given daily change in market conditions relative to a risk estimate based on the static portfolio assumption. Nevertheless, it is generally felt, correctly I think, that the static portfolio assumption is not unreasonable. In particular, if significant changes in prices can occur from close to open, based on overnight events, then the traditional procedure should be conservative enough to capture this phenomenon.

If we extend the simulation horizon beyond one day, however, the static portfolio assumption becomes untenable. For one thing, some transactions will mature during the simulation period. It is quite possible to perform a multi-step Monte Carlo simulation and still hold to the static portfolio assumption. The maturing transactions are simply allowed to run off so that they do not contribute to the change in portfolio value on subsequent days.

David Rowe is president of SunGard's Infinity business unit
e-mail: david_rowe@infinity.com

The result of such an exercise, however, is almost meaningless. This is because traders do not set hedges based on maturity matching over a multi-day horizon. They always know that the portfolio hedge will be adjusted dynamically as new trades come in, as market conditions change and as deals mature. As a result, a multi-day, multi-step simulation based on a static portfolio assumption gives utterly meaningless distributions of value change that have no connection to likely losses over the same horizon.

Some would argue that the solution is to develop internal model rules for how the hedges will be adjusted as the multi-day period unfolds and the hedge ratios deteriorate. Such an exercise is, at best, an interesting intellectual effort with little or no practical value. The result is primarily a reflection of the rehedging rules embedded in the model rather than the true risk of the trading operation. In my experience, there is virtually no chance that such endogenously modelled rehedging rules will bear any meaningful relation to what actually happens in the trading room.

Furthermore, if such rules did represent actual rehedging patterns, it would be criminally negligent to pay traders big compensation packages to implement what could be accomplished much more cheaply by simply following the rules laid down in the model.

Some also argue for multi-step Monte Carlo as a means of capturing hidden risk in oversized but balanced option positions with nearly equal strikes and expiry dates. The VAR of such positions will be small as long as the time to expiry is sufficiently long, but can become extremely large if the expiry date approaches with the market price very close to the common strike price. An example would be a large written put and an equal sized written call struck at the same price, and with expiry dates two or three days apart. The value of these deals would move symmetrically until expiry approaches. If the market price is close to the strike at that point, however, the sensitivities begin to differ and the risk increases. Obviously the extreme is when the first option expires and the second is so large that an effective hedge cannot be found due to market illiquidity. The question is whether multi-step Monte Carlo simulation is the appropriate method for detecting such risks. I believe the answer is no.

All possible paths

For such a procedure to capture this type of risk reliably would literally require simulating "every possible future path". Needless to say, that is a practical impossibility. In addition, without careful attention to liquidity issues, endogenous rehedging will further mask the risk. A more reliable approach is a search procedure to tabulate and report on the total underlying amount of options with strike prices and expiry dates within some specified ranges. This would give consistent early warning of such concentrations based on the structure of the book, and would allow risk and trading management to review the situation before it became critical with limited time to react.

How then should we proceed to arrive at a meaningful VAR estimate for periods longer than one day? I believe this is a case where the simple answer is also the best answer. The usual procedure is to multiply the one-day VAR estimates by the square root of 10 (or roughly 3.16) to arrive at 10-day VAR estimates. This is rooted in the assumption that changes in market variables are (at least approximately) statistically independent from day to day. In effect, this procedure says that if traders rehedged in such a way as to maintain the current day's VAR for 10 consecutive days, the potential loss over the full 10-day period would be 3.16 times the VAR for each individual day.

Averaging such estimates over a month or a quarter is the most realistic way to estimate average market risk calibrated to a 10-day time horizon. Such an estimate embodies all the weaknesses of one-day VAR itself, but those weaknesses are significantly magnified by attempting a multi-day multi-step simulation with endogenous rehedging rules. ■