High-frequency trading: how great is the need for speed?

Just how important is speed? Risk managers and traders are weighing the value of high-frequency market data and trading technologies against their costs. Gallagher Polyn examines the debate over using high-frequency data in risk models, and profiles one leading high-frequency trading firm, BNP Cooper Neff

The growing availability of tick-by-tick market data from a variety of liquid markets has prompted one influential economist to preach the virtues of incorporating high-frequency data into risk systems to improve their performance. But risk managers at banks are unconvinced; they say there is no benefit in this approach.

Does it really add value? Is it really worth the investment? Why haven't people made a billion dollars with it yet? Last month's audience at the Irving Fisher Committee conference on central bank statistics issues, sponsored by the Bank for International Settlements in Basel, wanted answers to these questions. In response, Richard Olsen, a Swiss economist, presented some of his findings from his work on high-frequency market data over the past 20 years.

Olsen argues that, properly calculated, intra-day data can yield significantly improved estimates of volatility over longer-than-one-day time horizons. The challenge is to develop models that simultaneously characterise the short-term and the long-term behaviours of time series, Olsen says. This approach to financial markets began in 1963 with mathematician Benoit Mandelbrot's analysis of fractal behaviour in the cotton market.

Most important in the new analysis is the treatment of market time as intrinsic, or a process in which higher volatility periods are more heavily weighted than lower volatility periods. Though much work remains to be done, Olsen says the case for devoting more resources is already clear.

Olsen argues that high-frequency data modelling yields better options pricing. In addition to academic evidence, Olsen says his own business experience supports his idea. Among the

Dan Rosen, Algorithmics

several financial markets enterprises Olsen manages out of his 'incubator' company Olsen Ltd in Zurich, BoxOptions offers intra-day options trading online for select currencies and European stock indexes. Institutions do not use BoxOptions – the lot size is $200 – though Olsen says his friends in the trading community have been trading them on their own account.

What's important, Olsen says, is that on identical simultaneous options positions, BoxOptions under-prices dealers by 15%. Though partly due to automation of the service, he says another reason is the use of high-frequency data modelling in its options pricing.

Olsen argues that the promise of high-frequency data risk management has already been demonstrated in academic work from the late 1990s, showing 20% greater accuracy in Garch-based volatility estimation.

The perfect model

Others agree, to an extent. "If you really believe continuous time was right and you had the perfect model, then using any observations for small intervals would scale perfectly for very large intervals as well," says Dan Rosen, who heads marketing and has designed systems at risk management vendor Algorithmics, based in Toronto.

But what happens when, instead of just, say, two currencies, you want to calculate your risk on a portfolio with many risk factors? Is there a point to high-frequency data value-at-risk as the portfolio grows complex? Risk managers say no.

One objection is that for a global bank doing a historical VAR estimate with hundreds or thousands of risk factors, a time series with more frequent observations is not compatible with a time series with fewer observations. Just because the price is changing more visibly in one asset does not mean it is not also changing in the less frequently observed asset.

Gregg Berman, head of market risk at systems vendor RiskMetrics in New York, considers the problem of bootstrapping the high-frequency time series of multiple instruments for fixed-income VAR. "If one futures contract updated while another lagged behind by five minutes, the resulting curve would be all but useless for risk analysis," he says.

> "I know a few who will swear that by going all the way to modelling microstructures will actually lead to macro estimates. I think that work is very interesting, but I haven't seen anything practical" **Dan Rosen, Algorithmics**

Olsen responds that exposures at most banks are dominated by just a few risk factors. "Limit it to instruments where you have big exposures – you don't have to run it across your whole book," he says.

Another objection is that bootstrapping intra-day market behaviours over longer horizons assumes mean and variance do not vary over the time measured. Olsen believes even the standard stationarity assumption over one-year horizons is too strong. "It's because we have so little in terms of daily data that we believe it's true," Olsen says. "It's not true over the daily horizon either, so let's tackle this issue."

Algorithmics' Rosen is familiar with Olsen's work in high-frequency finance, and sees some promise. "I think Richard's point is fairly right for a small number of risk factors and for very specific applications in fairly liquid markets," he says.

Algorithmics has been exploring some of the statistical techniques Olsen favours. "I know a few who will swear that going all the way to modelling microstructures will actually lead to macro estimates," says Rosen. How useful the new analysis will be is uncertain. "I think that work is very interesting, but I haven't seen anything practical," he says.

Concerns

In the final analysis, risk managers say that even if the theoretical grounds for scaling market features were clear, more 'bang for the buck' is available elsewhere. "I don't believe that on the list of 10 things to do that this is the thing a lot of people would be focused on," says Noel Donohoe, head of firmwide risk at Goldman Sachs in New York. Tail-risk estimation, evaluation of illiquid positions and the aggregation of credit risk with other types are all mentioned as leading concerns at the moment.

Curiously, the fact that high-frequency market data is available may itself discourage risk managers from seeking to model it for scaling behaviours to improve VAR. "High-frequency data is available, in some senses, to parts of the portfolio that I'm least worried about," said one risk manager in New York speaking on condition of anonymity. "Because if nothing else it implies liquidity."

Olsen emphasises that research is only in its initial stages, but based on consulting work he has done for dozens of banks in the US and Europe and his own experience as an asset manager, he thinks real improvements in capital allocation can be made. "My estimate is that it's a conservative assumption that you can lower your capital requirement by 10%," Olsen says. He claims he has achieved this for his own $60 million fund, invested in a few currencies and with some interest rate exposure, which has a Sharpe ratio around two and tail-event risk equal to the fund's annual return – about 20%. ■