Grinding costs, hidden models

Sequencing trades optimally to reduce costs is a 50-year-old problem that is still unresolved because the costs themselves are tricky to predict. Academics at New York University apply an intuitive Bayesian technique to simplify the process. Nazneen Sherif introduces this month's technical articles

Trading is not just about knowing which positions to take. Markets introduce friction in the form of transaction costs and larger trades have a tendency to push the market away – an effect called market impact. Sequencing trades optimally to reduce these costs is a computationally intensive process and most of the time firms end up doing it trade by trade instead.

In our first technical, *Multiperiod portfolio selection and Bayesian dynamic models*, Petter Kolm, a professor of mathematics at the Courant Institute of Mathematical Sciences at New York University (NYU), and Gordon Ritter, a senior portfolio manager at GSA Capital who is also an NYU adjunct professor, propose a framework that optimises trades for multiple periods in the future in the presence of unwieldy transaction costs.

"The trading cost problem is extremely important, especially for firms that trade a lot. If they get it wrong, it can form the difference between a profitable and an unprofitable strategy," says Ritter.

Typically, firms avoid transaction costs altogether or assume them to be linearly related to investment volume. "In practice, firms tend to use the naive, but accommodating, assumption of linearity of transaction costs, which turns out to be insufficient under some circumstances, like strategies, asset classes or market phases that imply significant portfolio turnover, for instance," says Hamza Bahaji, head of engineering and quantitative research at Seeyond, part of Natixis Asset Management in Paris.

An intuitive model is the answer, according to the authors. In order to deal with the transaction-cost problem and accommodate multi-period modelling, the authors use Bayesian statistics in the form of hidden Markov models to link ideal portfolio positions to the end solution – the sequence of optimal trades leading up to the future states of the portfolio.

A hidden Markov model is a variant of the Markov model in which future states of a system are predicted, based on the current state only. While Markov models use observed states, the hidden version has unobserved ones, called hidden states, which are particularly useful for modelling intuitive and uncertain situations or, in this case, the future states of the portfolio.

The authors see this as a break from traditional thinking. "For each investor there is a sequence of things they would plan to do in the future if there were no trading costs. Our model takes these as inputs and finds the optimal sequence of trades in the real world after trading costs are considered," says Ritter.

Optimising over hidden states using the Markov model allows the production of sample trading paths that are very likely to be in the neighbourhood of the true optimal path and in line with investors' preferences, according to Ritter. After generating the paths that are clustered around realistic positions, the authors apply a technique typically used to solve hidden Markov problems, called the Viterbi algorithm, to find the best path in the random sample.

Additionally, the authors of the paper show how it is perfectly valid to optimise the trading path of each asset independently of the others in a multi-asset portfolio, without accounting for interaction. This means that while simulating the path of one asset, one can hold the other trading paths fixed, and it finally leads to the optimal solution. Although this seems an oversimplification, the authors observe it still leads to the global solution, or the one that considers all possible paths.

"I spoke to some experts and this fact is not widely known. Those who were doing it this way may have been viewing it as an approximation without realising they were finding the true solution," says Ritter.

"The algorithm can be easily implemented. It allows computational time savings in the sense that the multi-asset case can be reduced to the single-asset case" Hamza Bahaji, Seeyond

The result is a model that is fast enough to preserve realistic transaction costs in a multi-period setting. "The algorithm can be easily implemented. It allows computational time savings in the sense that the multi-asset case can be reduced to the single-asset case," says Seeyond's Bahaji. This helps to achieve a dimension reduction of roughly N times for an N-asset portfolio. Usually, multi-period optimisation is ignored due to the computational power it demands.

In our second technical, *Two measures for the price of one*, Harvey Stein, head of regulation and credit risk in the strategic risk research group at Bloomberg in New York, proposes a way of calculating exposure that speeds up the process of simulating exposure in the real-world measure and repricing it in the risk-neutral measure. This method is faster by orders of magnitude and obviates the need for shortcuts such as reducing the number of paths or horizons in the simulation, which leads to inaccuracy and dependence on expensive technology. **R**