What will the next generation of algorithms offer?

Correcting the structural shortcomings of today's portfolio execution algorithms

Peter Fraenkel and Lee Maclin***

I. Introduction

Today's execution algorithms were designed to address the needs of sell-side institutions, and as such they can be inappropriate or even damaging to buy-side portfolio managers and investors. The agency task, reflected in traditional implementation shortfall algorithms, is to minimise the mean and variance of execution shortfall of the basket being traded. By contrast, the goal of proprietary traders and the buy-side should be to minimise the risk and maximise the alpha of an entire portfolio. Unfortunately, the typical use of algorithms can force buy-side traders to behave as if they had sell-side priorities. To optimise the complete portfolio strategy, an algorithm needs two key features:

1. True portfolio level risk measurement, including positions not changing during a given trade – not just the marginal contribution of unexecuted shares.
2. Trajectory optimisation with accurately estimated time profile of alpha.

Below in section II, we illustrate how restricting risk measurement only to unexecuted shares creates trading pressures that degrade portfolio performance. In section III, we discuss the role of alpha in optimal trading and its elevated importance once portfolio risk is accurately calculated. In section IV, we touch on the mathematics behind today's algorithms and the main incorrect assumption that leads to their suboptimal results.

II. The optimisation problems of the portfolio manager and the trader

Despite the traditional institutional separation of portfolio managers and traders, modern finance theory has shown that both groups

* **Peter Fraenkel,** director of Quantitative Services,

** **Lee Maclin,** director of Research,

Pragma Financial Systems

Innovation in algorithmic design

"The risk perceived by the trader differs from the actual risk of the portfolio. Mitigation of this perceived risk degrades overall strategy performance."

perform conceptually similar tasks. In a celebrated article in the 1952 *Journal of Finance*, Harry Markowitz[1] framed the managers' problem as an optimisation: maximise the expected mean of returns, while minimising the risk from their variance. Nearly 50 years later, Robert Almgren and Neil Chriss[2] showed that trading can also be viewed as an optimisation: minimise the expected impact of executions, while also minimising the variance from the unexecuted positions. Today, skilled practitioners in both groups accomplish their respective optimisations with the help of sophisticated tools – constrained quadratic programming tools for the managers and execution algorithms for the traders.

The question we address here is whether the result of these two separate optimisations is in fact optimal for the total performance of the portfolio. In general, the answer is no, and the following simple example shows one of the reasons: the risk perceived by the trader differs from the actual risk of the portfolio, and mitigation of this perceived risk degrades overall strategy performance. (In section IV, below, we formalise and generalise this intuitive result in a mathematical notation.)

Suppose we have a pre-existing portfolio consisting of a \$10 million position in IBM. We wish to sell this, replacing it by an equivalent \$10 million position in MO (Altria). Note that, while this example is somewhat contrived, its risk characteristics are close to what happens during a typical rebalance of an ongoing investment strategy.

First, look at this problem from the trader's point of view. He knows nothing of the current position in IBM, and has simply been told to execute a buy/sell basket containing negative \$10 million IBM and positive \$10 million MO.

Assuming that the two stocks are uncorrelated with identical volatilities of 20%, the total volatility risk of this basket at the beginning of the trade is:

$$V_{trader,\ initial} = \sqrt{(2 \times 10^2 \times 0.2^2)} = \$2.8 \text{ million}$$

1 Markowitz H., 'Portfolio Selection', *The Journal of Finance*, Vol 7, No. 1 (Mar. 1952), pp 77-91.

2 Almgren R. & Chriss N., 2001, 'Optimal Execution of Portfolio Transactions', *Journal of Risk*, 3, 5-39.

Figure 1:
A trader's perception of risk during an execution compared to actual total portfolio risk during an exchange of one asset for another

A trading trajectory that is optimal given the trader's perceived risk will not be optimal for the portfolio as a whole.

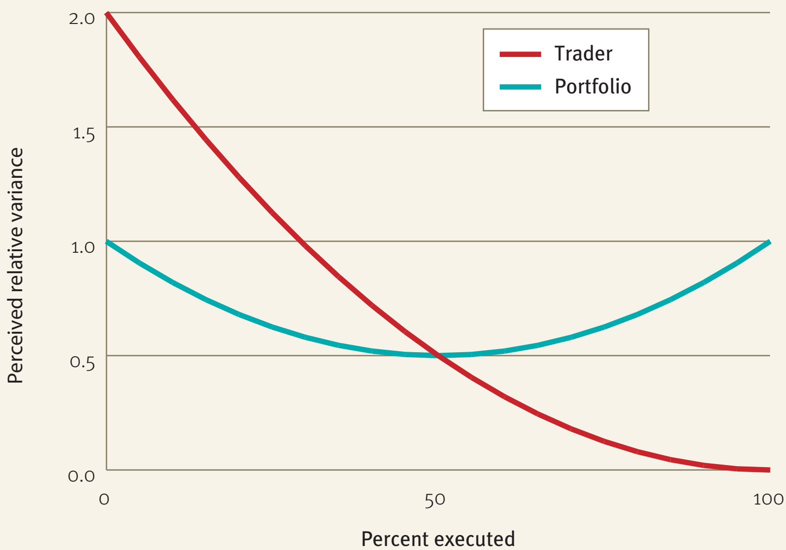

Halfway through the trade, he has \$5 million of each, so:

$$V_{trader,\ half\text{-}way} = \sqrt{(2 \times 5^2 \times 0.2^2)} = \$1.4 \text{ million}$$

And of course when he's done:

$$V_{trader,final} = 0$$

The expected variance of execution shortfall will be proportional to V^2 and to the amount of time it takes to perform the trade. By contrast, the expected market impact will increase for short execution times. In a standard implementation shortfall algorithm, the higher is V_{trader}, the shorter the execution time and the greater the market impact that results from the optimisation.

Now consider the risk from the true portfolio, containing +10 million of IBM at the beginning and +10 million of MO at the

Innovation in algorithmic design

"Basket algorithms take into account the correlative interactions of stocks within the trading order but ignore the effect of other stocks in the portfolio, including those whose positions are not changing."

end. The total volatility risk V_{port} of the portfolio is the same at the beginning and end of the trade:

$$V_{port,initial} = V_{final} = \sqrt{(10^2 \times 0.2^2)} = \$2 \text{ million}$$

Half way through the trade, when we own $5 million each of IBM and MO, the volatility risk reflects the diversification of this mini-portfolio:

$$V_{port,half\text{-}way} = \sqrt{(2 \times 5^2 \times 0.2^2)} = \$1.4 \text{ million}$$

Figure 1 shows that, while the trader perceives a risk that starts out large but decreases to zero by the end of the trade, the true risk of the portfolio starts out lower than the trader's perception, dips slightly and then returns to its initial value. This discrepancy motivates several important observations:

1. The trader will perceive a higher urgency to trade than the overall portfolio requires, and as a consequence will be willing to incur higher market impact.
2. This unnecessarily high market impact, repeated over multiple rebalances, will degrade the overall performance of the portfolio strategy while not actually making it less risky. The result is a reduced Sharpe ratio.
3. Conventional transaction cost analysis (TCA), evaluating trading in the absence of portfolio context, will not detect the damage to the portfolio's performance and could even penalise 'truly' optimal trading.
4. The trader was likely acting rationally, in accordance with his own compensation structure, which also ignores portfolio context.

Consider now what a portfolio-aware execution optimisation algorithm would have done in the situation described above. The range of true variance is ¼ the swing in the trader's perceived variance, and it reaches a minimum halfway through execution, rather than at the end. Accordingly, the optimal trajectory will be much slower

Figure 2:
Typical exponential realisation of alpha in stock prices, with corresponding exponential decay in the alpha drift rate of those stocks

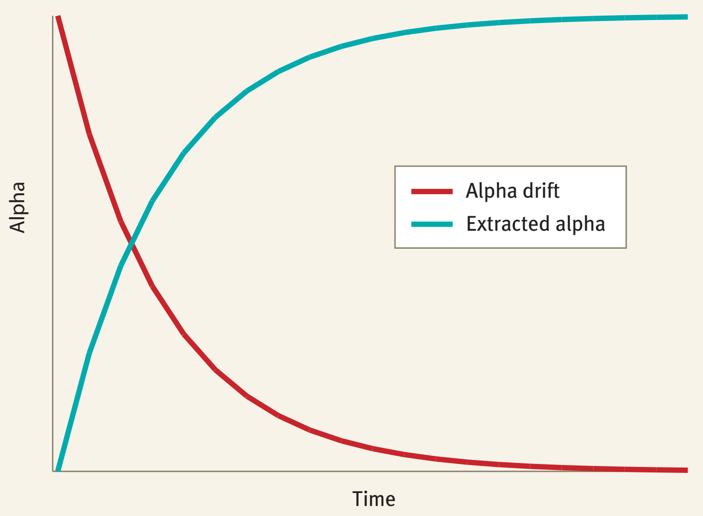

than conventional implementation shortfall solutions, and it will tend to linger in the middle where the variance is lowest. This is, of course, just one example and there are in fact situations where inclusion of portfolio dynamics will have precisely the opposite effect on the optimal rate of trading. There is no rule of thumb that substitutes for a complete algorithmic solution.

Note that the need to consider the portfolio context goes beyond the abilities of so-called basket algorithms. These take into account the correlative interactions of stocks within the trading order but ignore the effect of other stocks in the portfolio, including those whose positions are not changing. To correctly assess risk, minimise impact and measure performance, it is necessary to include the full set of interactions.

III. The time dependence of alpha and variance

If there were no reason to trade other than variance reduction, execution in the above example would simply stop

Innovation in algorithmic design

> "Trading and portfolio construction are both essentially optimisation activities that have traditionally been performed separately. We now know that this institutional separation of duties creates false incentives for both parties."

halfway through, at the point of maximum diversification. Obviously, there are other reasons, the most significant among them being the need to act on alpha signals. That is, we're trading because we have information that MO will outperform IBM, so:

1. There is an opportunity cost to letting returns accrue before we've traded the assets.
2. The rate of return accrual is expected to diminish as the returns are realised through the trading of market participants.

Point 2 is especially important, because it implies that the magnitude of returns will be higher during the execution period than over the life of the portfolio. Alpha extraction may thus dominate other considerations during optimal execution, especially in cases such as the example above, where true risk is lower than apparent trader's volatility. Figure 2 shows the typical behavior of alpha over the time period following the arrival of a trading signal. Initially, the rate of change (or drift) in the stock price is rapid, but as the price incorporates realisation of alpha, the rate decreases. This alpha decay exhibits exponential time dependence.

For example, an asset in a portfolio that returns 10% per year will on average accrue 0.04% in a trading day. However, if trading reflects fresh information, a far greater than average fraction of that accrual will occur during the execution period. Depending on the type of information, we might see an expected rate of return in excess of 50bps per day during execution – over ten-times higher than the mean rate of the portfolio as a whole. Such amplified drift rates during the period of execution can match or dominate competing algorithmic factors like market impact. In general, the time profile of alpha can be quite complicated, even changing sign for some counter-trending strategies.

Alpha is not the only trading variable whose time-dependence affects optimal trading. In addition, the variance of equity returns changes significantly over the course of the day. For example, Almgren *et al*[3] measure volatility varying intraday by more than a factor of three. Since volatility enters both the calculation of market risk and the estimation of market impact, this time-dependence can have a significant effect on the shape of the optimal trading trajectory.

IV. Optimal utility in portfolio selection and trade execution

We can now pinpoint the mathematical assumption that is implicitly made when optimisations for portfolio selection and execution are performed separately. Following the notation of Robert Engle and Robert Ferstenberg[4], we write the total utility function of the combined optimisation problem as:

$$U_{total} = \sum_{t=1}^{T} \left(-\Delta x'_t \, \tau(\Delta x_t) + x'_{t-1}\mu_t - \lambda x'_{t-1}\Omega_t x_{t-1}\right)$$

Here, t represents discrete time intervals ending at T, x is a vector of position holdings over time, τ is a function for temporary market impact due to position changes Δx, $_t$ is the coefficient of drift (expressing alpha), λ is the risk aversion, and Ω_t is the covariance matrix of asset returns. This multi-period utility function entangles the goals of portfolio manager and trader, and its maximisation results in a time-varying vector of asset holdings that achieve the optimal risk and impact-adjusted returns.

With the assumption that and Ω are constant over the trade and for the holding period that follows, Engle and Ferstenberg derive a relationship among , Ω and the terminal portfolio holdings x_T

$$\mu = 2\lambda\Omega x_T$$

and substitute it into U_{total} (after expanding the Ω term and completing a square) to split out independent optimisations for the portfolio manager and trader:

$$U_{port} = x'_T\mu - \lambda x'_T\Omega x_T$$

$$U_{trade} = \sum_{t=1}^{T} \left[-\Delta x'_t \, \tau(\Delta x_t) - \lambda(x_T - x_{t-1})'\Omega(x_T - x_{t-1})\right]$$

Since the first utility depends only on terminal holdings x_T, while the second depends only on the unexecuted trade list, the optimisations seem to be independent of each other, and, what's more, it appears that drift disappears completely from the trading problem. Unfortunately,

3 R. Almgren, C. Thum, E. Hauptmann and H. Li, 'Direct Estimation of Equity Market Impact', ***Risk***, July 2005.

4 Engle, Robert F. and Ferstenberg, Robert, 'Execution Risk', April 2006, ***NBER Working Paper No. W12165***, National Bureau of Economic Research.

Innovation in algorithmic design

these simplifications depend critically on the assumptions of constant drift rate and constant variance, which we know to be strongly violated in most trading situations.[5]

V. Conclusion

Trading and portfolio construction are both essentially optimisation activities that have traditionally been performed separately. We now know that this institutional separation of duties creates false incentives for both parties – with corresponding false utilities for their respective optimisations – resulting in solutions that in combination are suboptimal for the performance of the portfolio as a whole. Thus, great traders with great trading software may nonetheless fail to serve the interests of the investor. The correct, global optimisation takes the entire lifecycle of holdings into account, including the true risk of the total dynamic portfolio, the market impact due to its continuous evolution, the time-profile of alpha and the intraday volatility pattern. As investors come to realise the costs of performing separate, incomplete optimisations, vendors will respond with unified algorithms, and trading institutions will evolve the organisational changes required.

5 'Risk Aversion in Optimized Execution', 2007, Pragma Financial Systems whitepaper.