



What fuels your algorithms?

Modern execution algorithms are complicated, but Darren Toulson, head of research at LiquidMetrix offers a few guiding principles to designing execution algorithms that can fully realise their intended outcome.

Venue fragmentation, Dark Pools, Broker Crossing Networks, HFT liquidity providers (and predators) and the many other changes to the execution landscape in recent years all present a mix of opportunity and danger that the smartest execution algorithms should be able to turn to their advantage, or at least navigate their perils.

The job of coming up with smart new ways to seek liquidity, capture spread or generally outperform other execution strategies is most often given to 'quants'. As the name suggests, quants approach this task by first analysing historical/real time market data and based upon observations/heuristics (and maybe back-testing), they will devise a suite of algos with different execution objectives.

Implicit in this modelling is usually a set of market data and statistics that the algorithms have access to at 'run time' so the execution models can make optimal decisions.

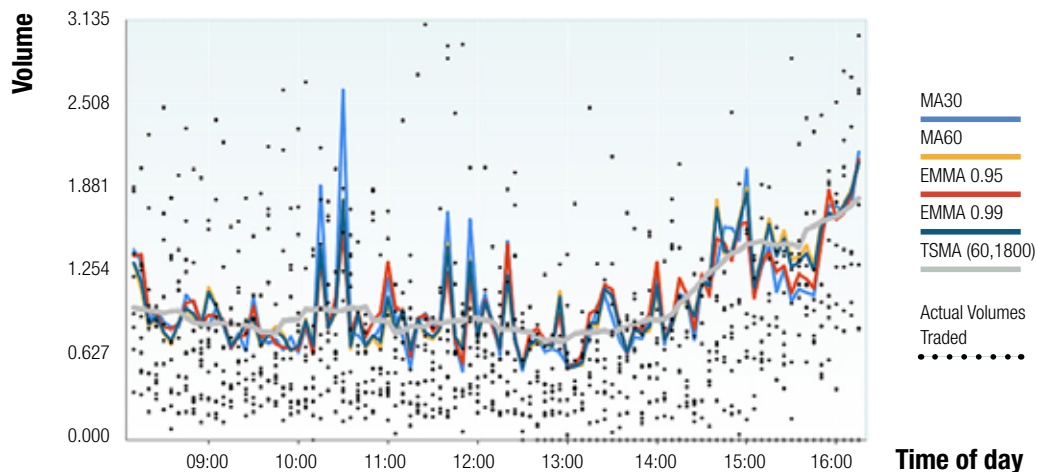
But this can be where things can get messy from a practical perspective.

Consider a simple algorithm

that is attempting to replicate or beat a market wide VWAP benchmark over the trading day, what kind of input data might such an algo require?

- Based on a start of day estimate of today's intraday 'market' trading volume profile, the execution algorithm should try to closely match this profile, executing trades at a fixed percentage of the current day's trading activity and thus closely match the day's VWAP. We may or may not wish to include start and end of day auction trading.
- As the trading day progresses, we may find that the actual market volume curve significantly diverges from our start of day estimate. We may need to tweak our target trading curve during the day, preferably based on some kind of conditional dynamic model that predicts what the rest of the day is likely to be based on the day's trading so far.
- The simplest way of executing would be to 'aggress' the lit markets each time we need to trade in order to closely track the target volume profile. However, this would mean 'paying the spread' for every trade and will ultimately result in underperforming VWAP by about half the spread. Additionally, even worse performance can result if we aggress more than one level down the order book and there is a price impact followed by short-term mean reversion in the market.
- Therefore most modern VWAP algorithms will try to trade as much as possible in mid-point matching dark pools or use passive limit orders on lit markets. This should improve overall average performance versus VWAP (by capturing some or all of the spread) but as the timing of execution of passive orders is outside our control, it makes the task of exactly

Intraday volume predictions



Estimating intraday volume profiles is not as straightforward as one might think. The chart above shows various intraday volume prediction models, each giving different estimates. The black dots show actual volumes traded, illustrating the natural intraday variability and difficulty of robust estimation.

replicating the intraday volume profile more complex and may lead to more deviation (risk) in our VWAP performance.

To summarise the above, inputs our execution algorithm would ideally have access to include:

- Start of day estimates of the intraday volume curves (including estimates of auction volumes)
- Dynamic models for updating these curves based on actual trading.
- Estimates of intraday bid/offer spreads and short-term impact/reversion models so the algorithm can model the likely cost of going aggressive.
- Estimates of execution probabilities for passive orders on lit markets so the algorithm can model the likely risk of non-execution.
- Intraday models of price volatility.

To be truly optimal these data sets should be instrument specific and venue specific as we want to know not only the probability of getting executed passively on any lit venue or dark pool but also which venues for a given instrument we are most likely to be executed on so we can set preferences on which venues to favour.

With this type of detailed run-time data, a cleverly designed VWAP algorithm should be able to make sensible decisions throughout the day on the speed at which it should trade, what venues to post passively to in order to capture spread, when to switch from passive to aggressive orders and the right 'size' to send to lit markets to minimise impact/reversion.

The practical challenges faced by algo operators are twofold:

1. Calculating or sourcing instrument/venue level statistics. The challenge here is maintaining detailed and accurate tick level databases and having processes to calculate the statistics on a daily basis.
2. Feeding these statistics into their run time execution logic (and integrating this data with other real time data that should be used to drive decisions). How easy this is will depend upon the degree of control provided by the execution software.

The prize, for firms that get this right, will be execution algorithms that fully realise the intention and cleverness of the people who design them. Otherwise, no matter how clever the algo design or impressive the real time execution architecture, Garbage In, Garbage Out... ■