Optimal Use of Algorithms: An Approach for the Buy Side

STEPHEN ENGDAHL AND HARISH DEVARAJAN

STEPHEN ENGDAHL
is the Director of Product Management at Charles River Development.
StephenEngdahl@crd.com

HARISH DEVARAJAN
is the Managing Director at Deep Value, Inc.
hd.ii@deepvalue.net

A BRIEF LOOK BACK

Trading algorithms originated as tools by the sell side for the sell side to automate trading activities of their own desks. Development of an algorithmic trading engine not only made use of available technology but also required detailed understanding of market microstructure and complex statistical modeling and mathematical theory.

Two major factors spurred brokers to open their proprietary algorithmic engines for use by their investment management and hedge fund customers. First, electronic trade communication via FIX became more prevalent between the investment manager and broker. Second, and simultaneously, brokerage commissions plummeted. Offering algorithmic trading enabled brokers to manage their own costs by decreasing sales trader headcount while providing a "premium product" to clients to bolster commission levels.

For an investment firm to utilize a broker algorithm, it must communicate a set of parameters that dictate how the algorithm should operate for each particular order. Broker suppliers thus needed a presence on the buy side trader's desktop to capture those inputs. Brokers quickly discovered that partnering with the order management system (OMS) already resident on the buy side trader's desktop was a better means to this end than building proprietary applications or acquiring frontend applications.

Despite the explosion in deployment of broker-supplied algorithmic capabilities and despite the gradual growth in use of those algorithms by the buy side, the type of usage remains pedestrian. For example, the most common algorithm, the volume-weighted average price (VWAP), appears to serve a utilitarian need, but in many cases it neither advances the cause of advanced trade execution nor warrants the attention the industry has devoted to this segment.

Unfortunately, the reality is that with the exception of advanced quantitative trading houses, buy side trading desks are generally not equipped to build and deploy their own algorithms. They typically lack both the academic bandwidth to develop a truly innovative algorithm as well as the appetite to invest in the technology required to rapidly execute an algorithm in the market. In contrast, brokerage firms often colocate their algorithmic servers directly at the exchange to minimize latency in trade execution. They are also continuously advancing the theory behind their algorithms. We therefore expect broker provision of algorithms to continue—but how will the buy side know whether they are getting their money's worth?

A Better Mousetrap: Integration of TCA to Improve Algorithmic Selection

With a variety of algorithmic trading options from which to choose, the investment managers are actively looking for processes, tools, and workflows to manage the complexity of choosing a particular algorithm for a

particular purpose. The increasing interest in transaction cost analysis (TCA) offerings, which ultimately measure algorithmic performance, has been a predictable development. However, in the current state of the art, TCA products have been integrated only in a rudimentary fashion into the process of algorithmic vehicle selection. The next frontier will involve tightly coupling TCA with algorithmic selection to make optimal algorithmic trading decisions.

The algorithmic trading landscape that confronts buy side desks is one where the offerings are relatively undifferentiated ("How is this VWAP different from that one?"), not very transparent ("Cool name, but what is really inside that black box?"), and hard to evaluate ("With eight different VWAP algorithms available on my desktop, which is best for this order?"). Even if investment managers intend to be data driven, having each desk perform the experiment with each algorithmic vehicle is inefficient. Desks may not have the resources, skill sets, time, and/or enough order flow to get good data, even if they had the inclination to do so.

By integrating TCA systematically with trade decision making, the buy side can apply data-driven processes to distinguish among vendor offerings based on their historical and ongoing performance. OMS providers can supply traders with means to apply data-driven smart routing, in real time, between and among the various algorithms that are available to them on an order-by-order basis. This will give the buy side better control of how and when algorithms are invoked in their trading process.

Once there is commitment to trade, the remaining things to optimize are fill rates, volatility of performance, and commissions and fees, in addition to minimizing market impact. The existing disconnect between TCA products and the manual process of algorithm selection would be all right if algorithms didn't vary too much along these four dimensions. However, the reality of algorithmic trading is considerably messier.

Even similar algorithms can produce dramatically different results over time, based upon a particular order's stock-level characteristics (market capitalization, liquidity), the amount of aggressiveness used in implementation, the execution venue (OTC versus listed versus ETF versus block destinations), market conditions, and commissions and fees (pricing plus differences in habits of posting versus taking liquidity).

Given the large range of variation these factors introduce, the right way of choosing between eight VWAP vehicles must be data driven, automated, real time, and particular to each individual order. Platform vendors, OMS vendors, and sell side houses are actively rolling out or acquiring TCA offerings. While they tend to compete on feature richness, with multiple tabs, drill-downs, Technicolor pie charts, and 3D graphs, such features are entertaining but often not very actionable. The ideal system would say "Vendor A does well on midcap listed names where more than 5% of the ADV is demanded, whereas Vendor B does better, commissions and fees included, on liquid OTC names," and hence sends large orders in listed midcap stocks to A and the liquid OTC to B automatically.

Backing it Up: Pay for Performance

Sell side cash equity desks nearly universally get paid per share executed. This omnipresent aspect of the agency model insulates to a substantial degree such desks from the direct cost of providing poor performance. One could argue that the costs of sub-par performance are borne primarily and first by the buy side desks, and are passed on to the sell side vendor only via reduced trading and/or use of their services.

At the industry-wide level, pay-per-share models select (in the evolutionary sense) sell side firms whose internal incentive structures, organizations, and expertise are honed to get the buy side to trade *more* instead of *better*—as increases in buy side P&L don't get passed on to the sell side, and only extra commissions from increased volume do. One potential alternative would be to follow a pay-for-performance model, where at least a portion of the commission on the trade was derived from the actual cost savings achieved as a result of using the algorithm.

Such a model requires consent and cooperation from the buy side: a buy side desk would need to present, in an independently verifiable way, the cost basis, measured by the slippage from submission time prices that it suffers in the trading it does. A random sample of historical trades would be bucketed into different categories based on trade difficulty, and the cost basis would be a set of numbers, one per bucket, determined based on realized performance. To the extent that usage of a particular algorithm is able to reduce the overall costs of trading versus the predefined cost basis, the sell side vendor would charge a modest percentage of cost savings achieved.

From the sell side perspective, the model may at first glance seem to be just another turn of the vise, yet in actual use it has yielded, possibly counterintuitively, substantial value. In cases where this model has been deployed, buy side trading desks have demonstrated a greater

willingness to share some elements of the alpha driving their order flow. Buy side and broker incentives are aligned, and the buy side has been able to improve performance and share the value created with the sell side.

While not all broker and buy side trading agreements will move to a pay-for-performance model, there is telling evidence that there would be less of a need for best execution regulation if natural market forces were already driving all participants in that direction. Best execution would become more of a business driver than a legal risk.

Implicit costs in trading (i.e., the market impact of trading—a cost not invoiced, and hence termed "implicit") typically destroy far more P&L than explicit costs (i.e., commissions and fees). As explicit costs fall with margin compression, there is substantially more value in lowering implicit costs and sharing in the unlocked value—rather than in obtaining value merely from ever-declining explicit costs. This is easily understood: the greatest remaining untapped source of value in the cash equities value chain is in what has always been the hardest part of trading well—lowering impact and being able to do so systematically.

CONCLUSIONS

The market for algorithms is far from mature. Much room for improvement and innovation remains in the science of algorithmic behavior. However, to continue to increase the degree to which investment managers are attracted to, use, and benefit from algorithms, the providers of algorithms and their OMS vendor partners must be able to demonstrate the value provided by particular algorithms, and assist the buy side in choosing the right ones at the right times.

To order reprints of this article, please contact Dewey Palmieri at dpalmieri@iijournals.com or 212-224-3675.