Algorithms, trading costs and order size

How does the impact of increases in order size on trading cost vary by strategy?

Yossi Brandes, Ian Domowitz**, Brett Jiu*** and Henry Yegerman*****

Two years ago, a trip to an internet search engine yielded approximately 25,000 hits on the term *algorithmic trading*. Today, that number exceeds 325,000 references. Surveys suggest that 77% of buy-side firms now employ the technology, with decision-support tools ranked top of the 'wish list' of most 'needed improvements'.[1] That survey ranking is an interesting statistic: despite intense activity in the area, there remains little information regarding the performance of algorithmic trading engines since we began researching the topic some time ago.[2]

Traders have become more sophisticated in their choices. The use of price-based algorithms has increased by about 50% since 2005, edging out VWAP-style strategies in terms of shares traded, while the adoption of liquidity-seeking algorithms has grown three-fold in the last two years. Not only has the number of black box strategies gone up, the number of ways in which a strategy may be tuned for an individual order has increased as well. Algorithmic trading begins to look like computer-assisted human trading, as perhaps was originally intended.

Such considerations make it difficult to speak about differences in the level of trading costs incurred across strategies. In this chapter, we limit ourselves to an easier, but still interesting, question: how does the impact of increases in order size on trading cost vary by strategy? We also address the role of intraday

* **Yossi Brandes**, senior vice president,

** **Ian Domowitz**, managing director,

*** **Brett Jiu**, senior research analyst,

**** **Henry Yegerman**, managing director,

Investment Technology Group, Inc

1 Tabb Group, *Institutional Equity Trading in America 2006*, October 2006.

2 See Domowitz and Yegerman, 'The Cost of Algorithmic Trading: A First Look at Comparative Performance,' *Algorithmic Trading: Precision, Control, Execution*, March 2005, and Domowitz and Yegerman, 'Measuring and Interpreting the Performance of Broker Algorithms,' *Algorithmic Trading: A Buy-side Handbook*, August 2005. References to relevant performance and trading cost literature are contained therein.

volatility within the same statistical model, measuring the impact of strategy on cost as a function of both size and volatility.

We find that this understanding of the role of intraday volatility is essential in attempting to answer our primary question. Strategies differ markedly in the responsiveness of trading costs to changes in volatility, which in turn implies that one cannot simply rank the cost performance of strategies on the basis of order size. This is an intuitive result, and one way to view a contribution of this chapter is empirical confirmation and quantification of this intuition across a set of algorithmic trading strategies. For example, VWAP and volume participation strategies are shown to perform similarly in low volatility environments, in terms of transaction cost as size increases. The performance of volume participation is, however, the more sensitive with respect to increases in volatility. Although the VWAP algorithm exhibits more marked increases in cost as order size goes up, the difference in volatility response translates into VWAP appearing to be a dominant strategy, compared to volume participation, in a higher volatility environment.

On the other hand, the relative rates at which strategies generate an increase in transaction costs, as order size rises, are reasonably stable across different volatility environments. While it may be difficult to rank strategies by the level of transaction cost, the efficiency of algorithms, with respect to how quickly performance degrades with size, can be assessed. We begin this process by describing the strategies and the model.

Algorithmic trading strategies, models and data

We study algorithms offered by Investment Technology Group (ITG), aggregated into five general strategies. *Dark* denotes a liquidity-seeking strategy, looking for liquidity among visible and hidden venues, typically splitting orders among various alternative trading systems, including ITG's own POSIT Match^{SM}. *Active* provides liquidity to the markets, using slices of the order to earn the spread or take liquidity at opportune times, completing orders subject to price and volume constraints. *Volume participation* (VP) participates at a specified percentage of printed volume, while *VWAP* is the volume-weighted average price strategy, executing orders over specified time horizons, and spreading trades according to historical volume distributions. Finally, *Implementation Shortfall* (IS) schedules orders into trades to minimise execution costs and degree of risk, employing client-specific urgency levels to determine aggressiveness of the strategy. Although these particular algorithms, as represented in our data, are proprietary to ITG, generic strategies with similar goals, even bearing similar names, are distributed by a variety of vendors throughout the industry.

The sample used for analysis contains 810,920 orders over the period October 2006 through March 2007, with a total dollar value of $266 billion. Data on orders not traded via these algorithms are not represented here.[3]

We analyse relative performance only along the dimension of transaction costs. Our benchmark for all orders and all algorithms is an arrival price, defined here as the mid-quote at the time of arrival of the order to the algorithm. Costs are analysed in terms of the difference between execution and arrival prices, expressed in basis points. Since VWAP and VP constitute part of the sample, no adjustments are made for relative trade difficulty.[4] In figures presented here, we sometimes refer to this measure of transaction cost as the *relative conditional impact* (bps).

The statistical framework used for analysis is a nonlinear model, relating transaction costs to a set of stock-specific variables and to parameters reflecting the impact of each strategy on cost. There is a large body of literature on potential candidate variables with respect to influencing trading costs. Given the inclusion of strategies, we find that name-specific average daily volatility and market capitalisation, combined with half-hourly time-of-day variables, suffice for our modeling goals.[5] The strategy-specific impacts are nonlinear functions of intraday volatility and the size of the order, defined in terms of percentage of average daily volume.[6]

The model can be best illustrated by focusing on its goal: the analysis of how quickly trading costs increase as order size increases, across strategies. Figure 1 exhibits the relationship between cost, size and volatility for one particular strategy, VP.[7]

3 In statistical terms, this implies that the analysis is subject to so-called sample selection bias, and must be interpreted as conditional upon the user already having chosen to use an algorithm in the first place.

4 Other choices are possible, of course. See, for example, Domowitz and Yegerman, 'The Cost of Algorithmic Trading: A First Look at Comparative Performance,' ***Algorithmic Trading: Precision, Control, Execution***, March 2005.

5 The cost literature is substantively covered in Domowitz, Glen, and Madhavan, 'Global Equity Trading Costs', Journal of International Finance, May 2001.

6 The mathematical formulation of the statistical model is given in a short appendix to this chapter.

7 In this figure and all charts that follow, cost on the vertical axis ranges from low to high, but we have omitted specific numbers. The reason for this is that all costs for the charts are 'fitted' costs from the model, and are scaled arbitrarily based on values used for various variables in the model which are not illustrated in the graphs. The source of all data in this chart and all graphs to follow in the chapter is Investment Technology Group.

Quantifying value and measuring performance

Figure 1:
The relationship between cost, volatility and order size

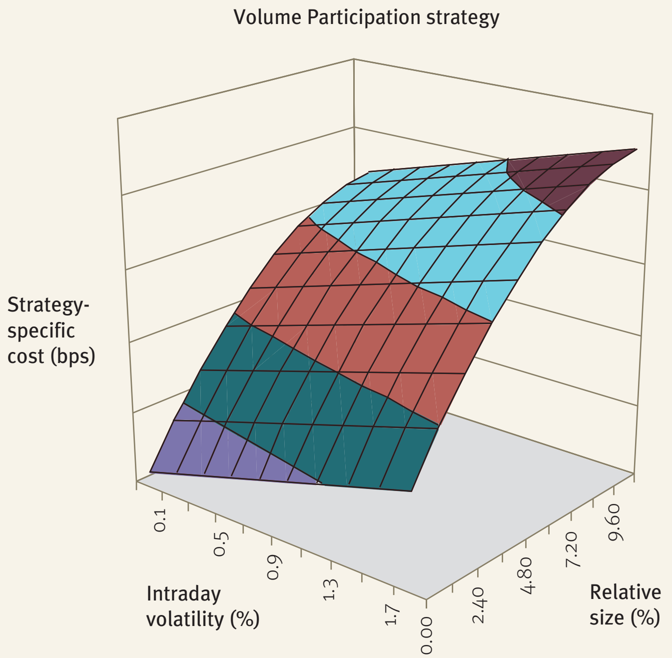

Costs rise as volatility and/or size increase – which is not surprising, of course. The tilt and curvature of the estimated surface suggests, however, that the relationship between increases in order size and cost may differ considerably depending on the nature of volatility during the execution of the order. The key is to quantify the extent of these differences for any given strategy and also across algorithms, to which we now turn.

Comparative impacts of order size on trading costs

Figure 1 suggests that an appreciation of the role of intraday volatility is central to our question concerning the effect of size on strategies' impact on cost. We illustrate the increase in cost as a function of volatility in Figure 2.

The first point is obvious from the graph: the Implementation Shortfall strategy (IS) appears completely insensitive to volatility[8]. The IS strategy explicitly embodies

8 This is true not only in terms of the values depicted, but also in statistical terms: volatility coefficients for the IS strategy are routinely statistically insignificantly different from zero in any of our analyses.

Figure 2:
Intraday volatility and cost

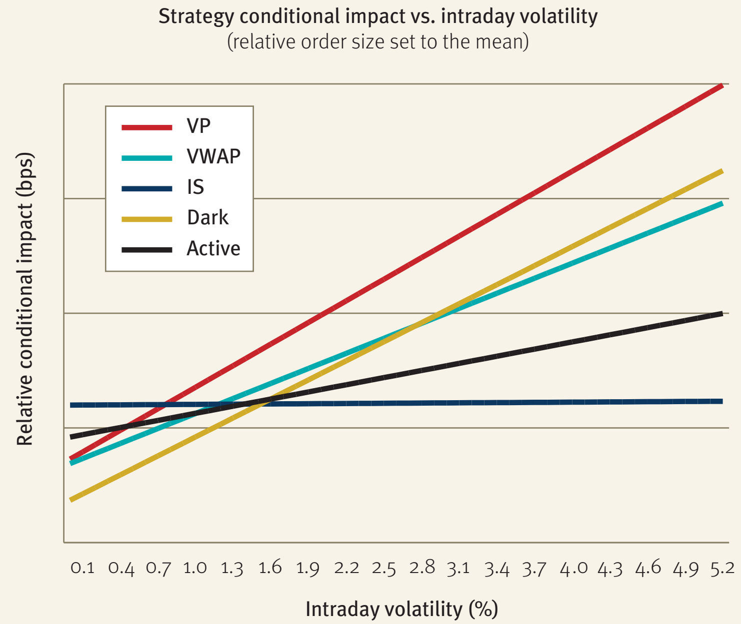

a trade off between expected cost and risk, which may be tuned by the user. In fact, trading over a fixed horizon is commonplace in IS applications, which makes trade aggressiveness sensitive to order size, but insensitive to intraday fluctuations in prices.

The remaining algorithms exhibit significant positive impacts of volatility on cost, clearly differing by strategy. The volume participation (VP) strategy appears most sensitive to intraday price fluctuations, due to its participation with printed market volume, which itself fluctuates with price movements. The dynamic version of pegging-and-discretion inherent in Active smoothes performance relative to short-run price movements, resulting in low sensitivity of costs to increases in volatility, relative to other strategies studied here.

Consistent with our appraisal of Figure 1, differing sensitivities to volatility imply that one cannot simply rank the cost performance of strategies on the basis of order size. This is evident in Figure 3, in which the impact of order

Quantifying value and measuring performance

Figure 3:
Relative order size and transaction cost

Panel A: Low volatility cases

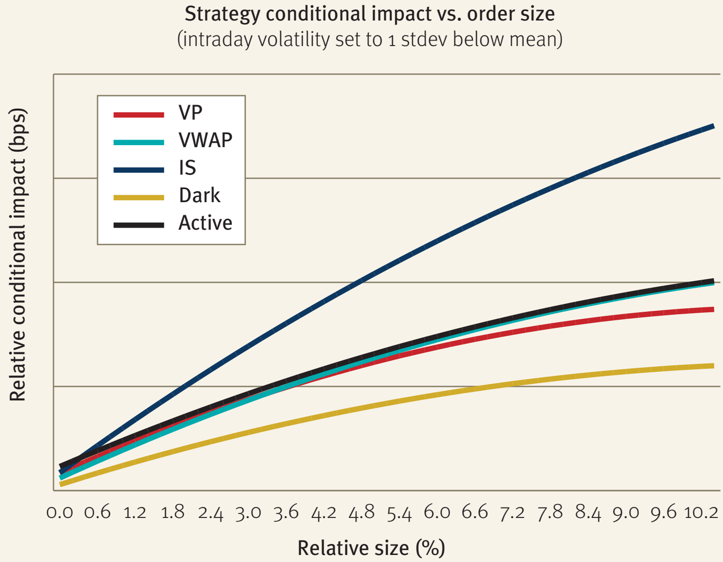

Panel B: High volatility cases

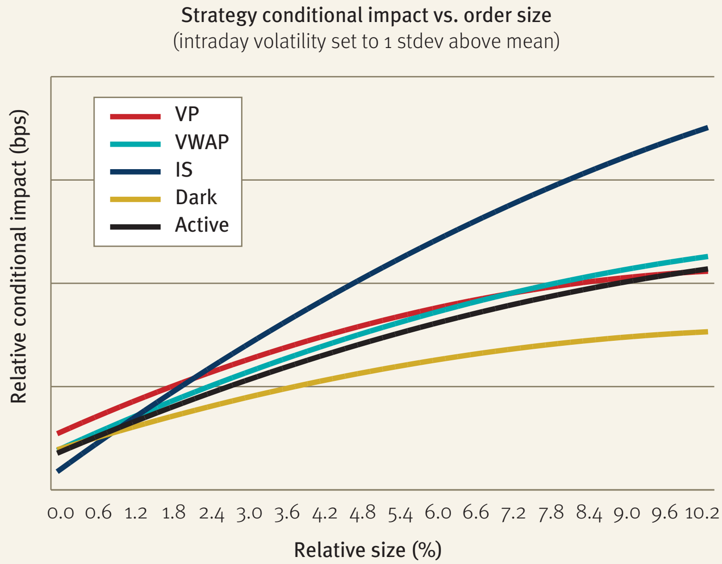

size by strategy is illustrated for volatility below its average and above its average, respectively, in panels A and B. In order to focus on the main points, we limit our discussion of comparisons to those with the VWAP strategy.

VP and VWAP strategies sometimes are thought to be similar, but are shown here to exhibit quite different behaviour, depending on volatility. In Panel A (low volatility environment), the two are almost indistinguishable in terms of the rate of change in cost as size goes up and in terms of the level of cost, but only for smaller order sizes. As size grows, however, the VWAP strategy exhibits more sensitivity to that increase, and eventually, for larger order size, exhibits higher cost than the VP alternative. In Panel B (high volatility environment), VWAP remains more sensitive to increases in size, but, from Figure 2, VP exhibits faster increases in cost as volatility rises. The result is that VWAP appears to be a dominant strategy in the high volatility case, up to large order sizes in algorithmic trading terms. The point at which the two strategies now look similar, with respect to cost, is no longer at small sizes but rather at large ones.

The same basic story applies in a comparison of VWAP to Active strategies. In a low volatility environment, the VWAP and Active curves virtually lie on top of each other for larger orders. VWAP is more sensitive to increases in order size, however, and to increases in volatility. The result is to change relative performance in higher volatility environments, as shown in Panel B, where the two strategies now are similar only for small size orders.

Finally, consider the case of the Implementation Shortfall strategy. From Figure 2, this strategy is insensitive to intraday volatility. Other strategies, including VWAP, change behaviour with volatility, however, and affect comparisons with the IS alternative. While IS exhibits more sensitivity to changes in order size than VWAP for all volatility environments, it may perform better in terms of cost than the VWAP alternative only in high volatility cases.

Conclusion

We have suggested previously that different algorithms occupy places in a trade structure continuum, ranging from opportunistic strategies to schedule-driven trades.[9] Opportunistic algorithms do not have pre-defined execution schedules, utilising real-time information to search for good execution, and effectively creating schedules as they go along. The Active strategy illustrated here is

9 Domowitz and Yegerman, 'Measuring and Interpreting the Performance of Broker Algorithms', *Algorithmic Trading: A Buy-side Handbook*, August 2005.

a case in point, while VWAP is the classic example of a schedule-driven trade. Between these extremes lie evaluative strategies, combining both approaches and blurring distinctions. Implementation shortfall algorithms may be viewed as such.

It is tempting to divide the world up in this fashion as a path towards algorithm selection, but specific algorithms are chosen for a variety of reasons in practice.[10] Intuition dictates that no algorithm is 'best' under all conditions, and one contribution of this chapter is to suggest that no class of algorithms, nor any algorithm in any particular class, may dominate the choice in terms of performance, at least in simple terms.

What do we mean by 'simple'? In the context of the model used here, even changes in intraday volatility alone reverse performance rankings of algorithms. As a result, the analysis has concentrated on relative changes in transaction cost as order sizes increase, across strategies. In this sense, we find some stability in how algorithms compare to one another.

For any single strategy, the results imply that we might reasonably forecast performance changes as one moves to increase order size by 25%, for example. On the other hand, one cannot fairly choose to shift strategies in the same scenario without further information. These findings do not preclude the goal of optimal algorithm selection. They do suggest, however, that the issue requires thinking beyond a simple 'envelope of possibilities', to a more complicated form of optimisation.

10 Discussions abound on this point; see, for example, Tracy Black and Owain Self, 'Choosing the Right Algorithm for Your Trading Strategy,' in *Algorithmic Trading: A Buy-Side Handbook*, 2005.

ITG® INVESTMENT TECHNOLOGY GROUP

The authors would like to thank Milan Borkovec for his comments at all stages of this project, and James Wong, both for his advice and assistance with the data. The information contained herein has been taken from trade and statistical services and other sources we deem reliable, but we do not represent that such information is accurate or complete, and it should not be relied upon as such. Any opinions expressed herein reflect our judgment at this date and are subject to change. This report is for informational purposes and is neither an offer to sell nor a solicitation of an offer to buy any security or other financial instrument in any jurisdiction where such offer or solicitation would be illegal. This report does not provide any form of advice (investment, tax, or legal). No part of this report may be reproduced or retransmitted in any manner without permission.

Appendix: The statistical model

The statistical model underlying the results of this chapter may be written as:

$$C_{it} = \beta' X_{it} + \gamma_1 S^1_{it} + \ldots + \gamma_5 S^2_{it} + \epsilon_{it}$$

Where C_{it} denotes the transaction cost for an order in name i at time t, and X_{it} is a vector of variables, including the market capitalisation of stock i, the average daily price volatility of stock i, and a set of time-of-day variables indexing half-hour intervals throughout the trading day. S^j_{it} is set equal to 1 if strategy j is used for the stock, and zero otherwise; $j = 1,\ldots,5$ denote the strategies identified as VWAP, VP, IS, Dark, and Active in the main body of the chapter. The regression error term is given by ϵ_{it}. The parameters γ_j are themselves time varying and stock specific, following equations of the form:

$$\gamma_j = \omega_{oj} + \omega_{1j}\, size_{ijt} + \omega_{2j}\, size^2_{ijt} + \omega_{3j}\, vol_{it}$$

Where $size_{ijt}$ denotes the relative size of the order and vol_{it} is intraday volatility, the former indexed by the name of the stock, the time, and the strategy followed in the trading of that name at that time, while volatility is simply time and stock specific.