May 2011

INVESTMENT TECHNOLOGY GROUP

Algorithmic Trading Usage Patterns and their Costs

Ian Domowitz
Managing Director, ITG, Inc.

Henry Yegerman
Managing Director, ITG, Inc.

212.444.6300
info@itg.com
www.itg.com

Executive Summary

Using algorithmic trading data across seven strategy types over 2009 and 2010, we examine usage patterns and performance for a sample of buy-side firms served by a multiplicity of brokers. Strategy usage is categorized by demand for liquidity, volatility, and concentration of orders traded. The data suggest employment of dominant strategies for the majority of firms, and shifts in strategy use are marginal across time and market conditions. In terms of performance, dominant strategies constitute a sensible approach at two ends of a spectrum: for easy orders and for situations which are extremely demanding in terms of liquidity and volatility. Performance matters, but does not distinguish individual strategy types in either regime. In all other circumstances, strategy shifts are possible and potentially profitable, given performance differences.

I. Introduction

In 2005, an internet search engine produced about 25,000 hits on the term *algorithmic trading*, rising to 325,000 by 2007. The growth rate in interest continues today, with over 2.4 million references available. The utilization of algorithmic trading has grown by 36 percent since 2007, now accounting for 30 percent of institutional share trading in the U.S. [1]

Despite all the references, hard evidence on algorithmic trading usage and performance remains elusive. Even amongst eight papers in a special section on algorithms in *Journal of Trading*, performance is quantified in only one piece, and then only as an example.[2] It should come as no surprise that the majority of traders perceive no differences in algorithms, with only 21 percent viewing performance as a differentiating factor. [3]

There are several ways to think about usage. Utilization rates for algorithms as a group are a common output of survey evidence, as are relative usage percentages across providers. Breakdowns of algorithm commission spend also are available, targeted specifically towards elements of service provided by brokers. In contrast, we focus on usage by strategy type, across seven categories. Our interest is in use by volatility regime, by relative demand for liquidity, and in the concentration of strategy use by buy-side firms.

The question behind this part of our investigation is very simple: *once a decision has been made to trade via an algorithm, what choices are being made?* The 'why' is not always obvious, and phrased this way, the results may appear more interesting to vendors than to consumers of the service.

The question addressed for the algorithm user also is straightforward: *do these choices make sense?* We avoid explicit broker comparisons in this paper, which means that the answer does not depend on degree of customization, broker-specific service levels, and the like. Our interest is in relative performance, measured in terms of the distributions of transaction costs across strategy types and within individual strategy categories.

Algorithm usage by demand for liquidity and volatility regimes is summarized. The behavioral patterns suggest a portfolio of strategies, but shifts in strategy use are marginal, on a desk-by-desk basis. The vast majority of users have dominant strategies, and the percentage of orders executed by a single strategy often exceeds 70 percent. The choice of dominant strategies is almost invariant to the demand for liquidity. Although there is some switching from one dominant strategy to another as volatility conditions change, it is limited.

The focus then shifts to whether the adoption of invariant dominant strategies makes sense, using transaction costs as a metric of algorithm performance. Dominant strategies constitute a reasonable approach for very small orders done in normal environments, and, perhaps surprisingly, for situations that are extremely demanding in terms of liquidity and volatility. Performance does not distinguish strategy choices in either regime, but in all other cases, strategy shifts should be motivated by algorithm performance. We go beyond average costs in making this assessment, by looking at certainty of outcome as described by the performance distributions for each strategy type. Although equivalence between most strategies is affirmed, the probability of an extremely bad outcome rises sharply for a subset of strategies in difficult markets.

---

1 Tabb Group, *US Equity Trading* 2010/2011: *Outflows, Outrage, and Balance,* December 2010, exhibit 1.

2 *Journal of Trading*, Volume 4, Number 3, 2009.

3 Tabb Group, *US Equity Trading* 2010/2011: *Outflows, Outrage, and Balance*, December 2010, exhibit 18.

Dominant strategies exist within an environment characterized by deviations from expectations, or surprises. The examination of performance is extended to analyze cost in the presence of large surprises in volatility and volume. Strategy usage does not change when surprises occur, reinforcing the stability of dominant algorithms. This makes sense in the context of volume differing from its historical norms. A case for the maintenance of dominant strategies across volatility regimes can be made only in markets characterized by volatility lower than expectations. The distributions of performance outcomes generally suggest the advisability of abandoning dominant strategies in market conditions which do not conform to expectations.

II. Algorithmic Trading Strategies and Data

We examine trading strategies used by 20 large institutions, for which algorithms are provided by approximately 30 brokers. Two sample periods are used. The first, covering 799,772 orders and 22,921,786 trades, ranges from the fourth quarter of 2009 through the fourth quarter of 2010. The second time period represents a sustained level of high volatility relative to the first, and generally uncertain conditions, over the first two quarters of 2009.[4] The gap between periods is deliberate, as we observe a drift to a distinctly different volatility regime over the third quarter of 2009.

Algorithms are aggregated across brokers into seven categories. We will sometimes refer to *scheduled strategies,* which include VWAP, TWAP, and Participation algorithms. These require no special explanation, as does trading at the *Close*, a common technique now implemented through algorithms. *Dark* denotes a liquidity-seeking strategy amongst dark pools, while *Liquidity Seeking* is more general, representing opportunistic strategies geared to liquidity provision. Unlike Dark, Liquidity Seeking strategies often attempt to source liquidity in both dark and lit venues. There are generally a range of options to determine how aggressively liquidity should be sourced. *IS* is shorthand for implementation shortfall, creating execution patterns based on cost and risk minimization, and employing urgency levels to determine aggressiveness of the strategy while attempting to match the decision price.

We analyze performance along the dimension of transaction costs. The benchmark for all orders and all algorithms is arrival price, defined as the mid-quote at the time of arrival of the order to the algorithm. Cost is defined as the difference between execution and arrival prices, expressed in basis points (bps).[5] Since VWAP, TWAP, and Participation constitute part of the sample, no adjustments are made for relative trade difficulty. We will, however, present results with respect to the distributions of cost outcomes in terms of relative demand for liquidity and volatility.

The work is restricted to an examination of unpriced market orders. Orders with limits merit additional study, but are not included here.

4 553,956 orders and 12,665,475 trades.

5 Underperformance against the benchmark is represented as a negatively signed number (trading cost), while over performance against the arrival price is represented as a positive number.

III. Strategy Usage at the Level of the Buy-Side Desk

We begin with algorithm usage sorted by demand for liquidity, which is proxied here by the percentage of median daily volume (MDV) for the order. The vast majority of orders in the sample are under five percent of MDV. We sort these orders into three categories relative to MDV: less than 0.5 percent, 0.5 to 1 percent, and 1 to 5 percent, in Figure 1 below.

Figure 1: Algorithmic Trading Strategy Usage by Liquidity Demand

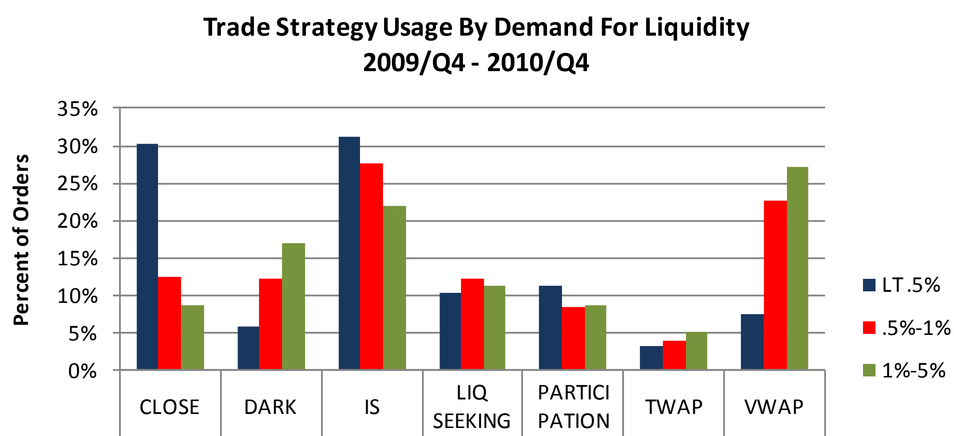

	CLOSE	DARK	IS	LIQ SEEKING	PARTICI PATION	TWAP	VWAP
LT .5%	30%	6%	31%	10%	11%	3%	8%
.5%-1%	13%	12%	28%	12%	9%	4%	23%
1%-5%	9%	17%	22%	11%	9%	5%	27%

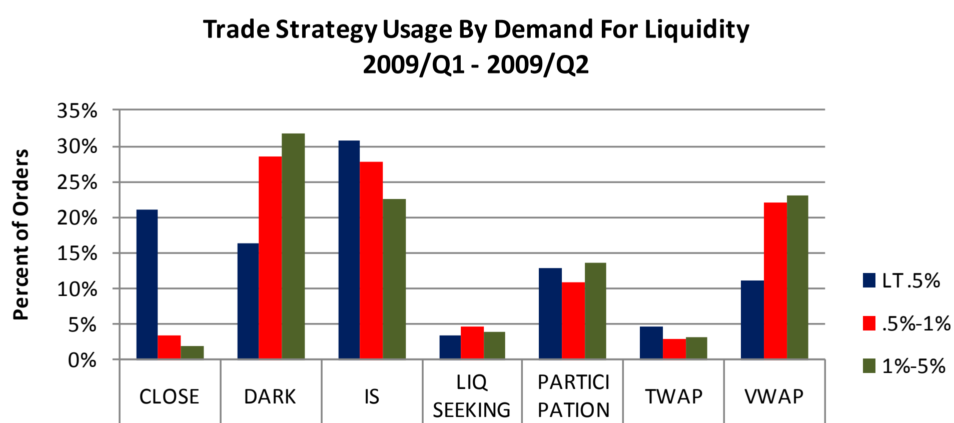

	CLOSE	DARK	IS	LIQ SEEKING	PARTICI PATION	TWAP	VWAP
LT .5%	21%	16%	31%	3%	13%	5%	11%
.5%-1%	3%	28%	28%	5%	11%	3%	22%
1%-5%	2%	32%	23%	4%	14%	3%	23%

The general format of Figure 1 is repeated in various contexts below, in terms of time periods and their interpretation. The top panel contains results from the last quarter of 2009 through the last quarter of 2010. The bottom panel represents two quarters of activity, at the beginning of 2009. During the latter period, what we would normally term "high volatility" is much higher than during the time period represented in the first panel. Average volatility for the first two quarters of 2009 is 480 basis points, while the later period exhibits volatility on the order of 203 basis points. We will sometimes refer to the first two quarters of 2009 as "high volatility," which should be interpreted in a relative sense.

The use of Liquidity Seeking, Participation, and TWAP strategies rises due to a shift from high to normal volatility, but in each regime, usage is almost invariant to changes in the demand for liquidity. In comparison, VWAP use is relatively immune to volatility shifts, and jumps sharply as one moves from very small to medium size orders. The employment of Scheduled strategies as a whole follows the VWAP pattern.

IS strategies exhibit a different pattern. For each level of liquidity demand, the percentage of orders sent to IS is almost identical across volatility regimes. There is a decline in usage, however, as the demand for liquidity increases. The decreases are large in percentage terms. Regardless of volatility regime, as an order goes from less than half a percent of MDV to upwards of five percent, use declines by almost 30 percent.

Such declines in use are more pronounced for the Close. During the first two quarters of 2009, an increase in order size from below half a percent of MDV to upwards of five percent results in a tenth of the use observed for very small orders. That decline is about 70 percent in more normal markets, but the pattern remains the same.

Dark searches are used more in the high volatility regime, by a substantial margin. This result is inconsistent with the commonly accepted notion that traders move away from dark pools when volatility rises, but may be due to the specific time period studied. Usage of dark algorithms increases with order size, reflecting the original rationale for anonymous crossing systems.

These patterns suggest employment of a portfolio of strategies, which would be no surprise. The surprise lies in the fact that any shifts in strategy are relatively marginal, desk by desk. We illustrate this phenomenon in Figure 2.

Figure 2: Dominant Strategies by User and Type

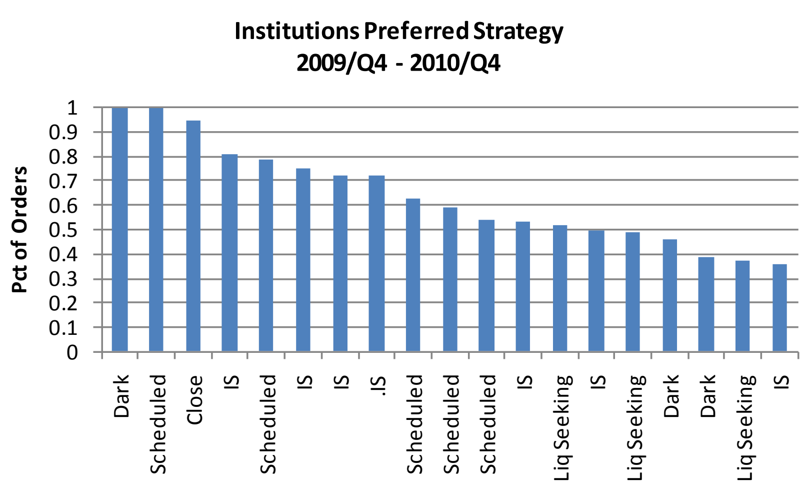

Figure 2: Dominant Strategies by User and Type Part 2

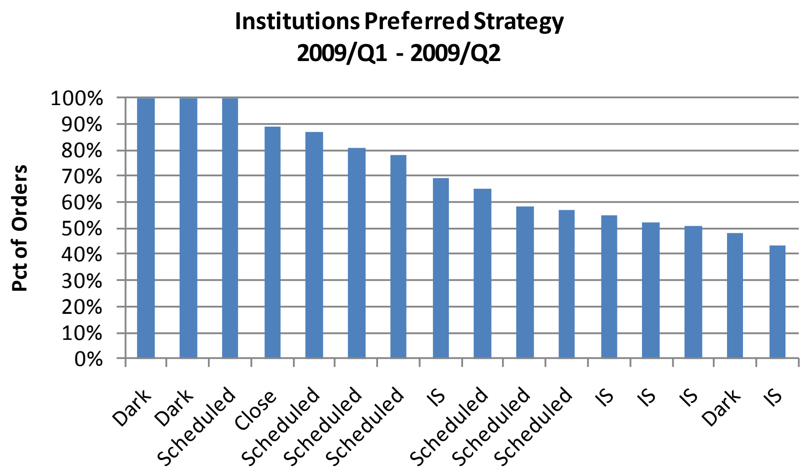

The percentage of orders is on the vertical axis. Each bar on the graph represents a buy-side desk, and is labeled with the strategy most used by that desk. We call any strategy accounting for at least 50 percent of order activity a *dominant strategy*. For this purpose, we group VWAP, TWAP, and Participation as *Scheduled* strategies.

The vast majority of users have dominant strategies; the percentage of orders executed by a single strategy exceeds 70 percent in many cases. In the high volatility environment, 18 of 20 users in the sample possess dominant strategies, with three users exhibiting use of a single strategy, and seven employing one strategy at least 80 percent of the time. Dominant strategies fell, to about 75 percent of users in 2010, but remain significant.

The employment of dominant strategies is almost invariant to the demand for liquidity, as suggested by the very high concentration of orders in Figure 2. There is some switching from one dominant strategy to another across volatility conditions, but it is limited. Six of twenty desks in the sample made a switch between the first two quarters of 2009, and 2010. There is no pattern in the nature of these changes, although half of the changes are from Scheduled strategies to something else. The number of these changes may be overstated in qualitative terms, however. Of the six desks that switched, two of them have dominant strategies representing close to 50 percent of their order activity in the early period, and the new dominant strategy was second in line in 2009, with almost half the activity. In other words, for two of six changes, two major strategies are in use, with a small change that altered the ranking in terms of dominance. Taking those desks out of consideration implies a move only 20 percent of the time, as the markets change dramatically.

IV. Algorithmic Strategy Performance

We now turn to the question, do these choices make sense? Our answers are based on transaction costs as a means of evaluating trading performance. The benchmark for all orders and all algorithms is arrival price, defined as the mid-quote at the time of arrival of the order to the algorithm. Cost is defined as the difference between the execution and arrival prices, expressed in basis points (bps). All analysis is at the level of the underlying order, as opposed to individual executions. We begin with average costs, by strategy type and relative demand for liquidity, illustrated in Figure 3.

Figure 3: Average Transaction Costs by Strategy Type

Trading Strategy Costs By Demand For Liquidity: 2009/Q4-2010/Q4

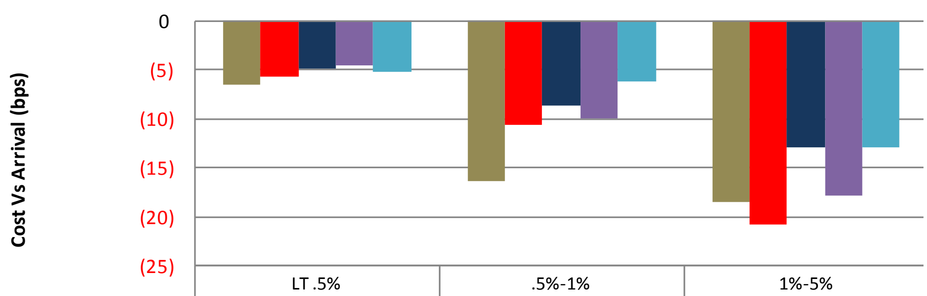

	LT .5%	.5%-1%	1%-5%
CLOSE	(6)	(16)	(18)
DARK	(6)	(11)	(21)
IS	(5)	(9)	(13)
LIQUIDITY SEEKING	(5)	(10)	(18)
SCHEDULED	(5)	(6)	(13)

Trading Strategy Costs By Demand For Liquidity: 2009/Q1-2009/Q2

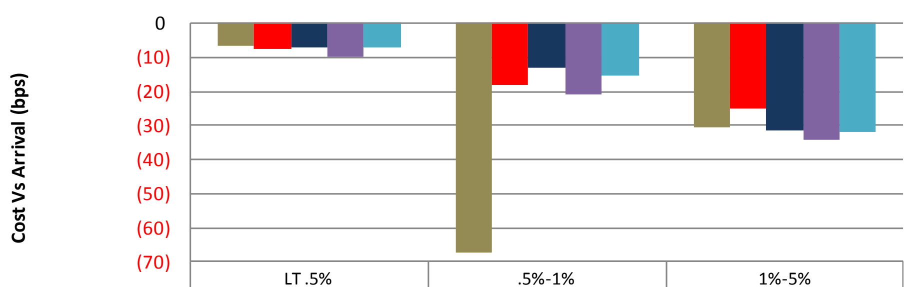

	LT .5%	.5%-1%	1%-5%
CLOSE	(6)	(67)	(30)
DARK	(8)	(18)	(25)
IS	(7)	(13)	(31)
LIQUIDITY SEEKING	(10)	(21)	(34)
SCHEDULED	(7)	(15)	(32)

Survey evidence with respect to the commoditization of algorithms and lack of attention to relative performance is corroborated for very small orders in a normal volatility environment.[6]

6 Tabb Group, *US Equity Trading 2010/2011*: *Outflows, Outrage, and Balance*, December 2010, exhibit 18; 52 percent of respondents say that the algorithms are commoditized, and only 21 percent believe that performance is a differentiator.

Costs are quite low, at 5 to 6 basis points, and do not vary significantly across strategies. Costs for very small orders rise by 43 percent, on average, in the high volatility period. Differences across strategies remain clustered within a basis point, however, with the exception of liquidity seeking algorithms.

For very small orders, invariant dominant strategies make sense, and the choice of strategy is not well informed by performance characteristics. Commoditization, in the context of performance, is not limited to cross-vendor comparisons; it appears to apply across strategies as well.

This conclusion begins to break down with a small increase in demand for liquidity. In normal volatility environments, and for orders up to one percent of average volume, Dark, IS, and Liquidity Seeking strategies remain clustered within a couple of basis points. Scheduled strategies retain their low cost characteristics, while trading at the Close becomes quite expensive; the spread between the two is ten bps. Differences become obvious, once one considers orders up to five percent of volume. The exception is large orders during the first two quarters of 2009. In that case, performance clustering again is evident, although costs are substantially larger than in 2010.

The data imply that dominant strategies constitute a sensible approach for very easy orders, and for situations which are very demanding in terms of liquidity and volatility. In the latter case, however, no strategy type does well, with average costs on the order of six times those experienced in low liquidity demand, normal volatility markets. Performance matters, but does not distinguish individual strategy types in either regime. In all other circumstances, strategy shifts are possible, and potentially profitable given performance differences.

Average costs do not tell the entire story. Certainty of outcome, expressed by the distribution of costs, contributes more information. We approach this in a couple of ways, beginning with the empirical distribution of cost outcomes, show in Figure 4.

Figure 4: Distribution of Algorithmic Strategy Performance

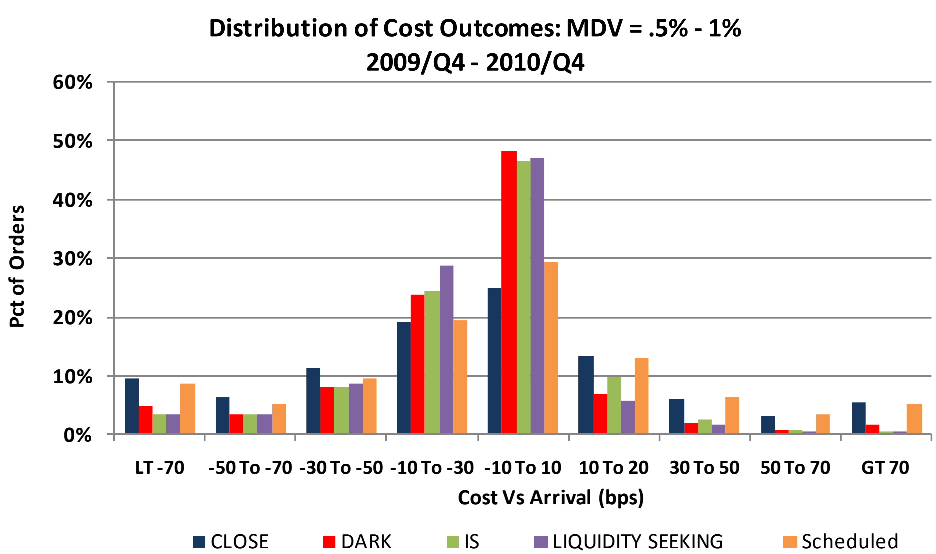

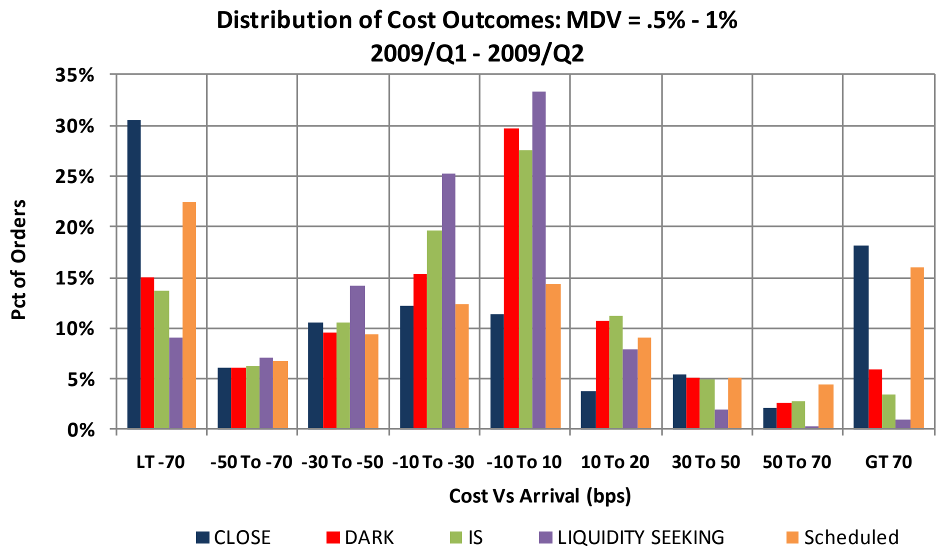

Figure 4 illustrates outcomes for the middle of liquidity demand groups.[7] Aggregated across strategies, the figure represents the overall distribution of algorithmic trading costs. In normal environments, the shape of the distribution is what one might expect: outcomes are clustered in the range of 30 bps under-performance to 10 bps over-performance. There are three interesting points, however, relating to commoditization and performance.

With the exception of the Close, the distributions of outcomes in the -70 to -30 range are basically the same across strategies. In other words, the probabilities of obtaining a bad outcome in that range are the same. This is yet another way in which relative performance is not a driver, and dominant strategy use may make sense.

The second point relates to the Close specifically. The probability of obtaining a poor outcome is much higher than for the other strategies. Although this strategy does not stand out as poor on average, the likelihood of a bad outcome for any given order is relatively high. This result updates earlier research on the dangers of trading at the close. [8]

Finally, the equivalence across strategies in terms of certainty of outcome fails for very poor outcomes, defined here as costs in excess of 70 bps. Scheduled strategies do well on average, but have a spike with respect to very bad results. IS and Liquidity Seeking strategies perform best in this respect, while the probability of doing very poorly at the Close rises to roughly ten percent.

Two items also stand out from examination of the cost distributions in the high volatility period, aside from the obvious fact that probabilities of poor outcomes increase across the board. The equivalence between strategies breaks down somewhat, but is roughly maintained in the range of -30 to -70 bps. The exception now is Liquidity Seeking algorithms, for which we observe a relatively higher increase in the probability of a bad outcome. Second, the probability of obtaining a very bad outcome rises sharply for both the Close and Scheduled strategies. The odds of getting an order done at more than 70 bps of cost rise to over 30 percent for the Close, and are roughly 23 percent for strategies like VWAP.

---

7 Space considerations preclude showing all distributions. For very small orders, these distributions are more 'peaked,' as might be expected, while for large orders, the distributions become increasing diffuse, exhibiting greater 'fat tailed' behavior, also as one might conjecture.

8 See , for example, David Cushing and Ananth Madhavan, "The Stock Returns and Institutional *Trading at the Close,"*

Regardless of what the averages say, the increase in risk from using such strategies in a difficult market is high. Eight of seventeen dominant strategies in the first two quarters of 2009 belong to one of these two categories. This represents a large degree of risk taking in a high volatility environment.

A different view of the performance of Close and Scheduled strategies is provided in Figure 5.

Figure 5: Probability of Negative Cost Outcomes for Close and Scheduled Strategies

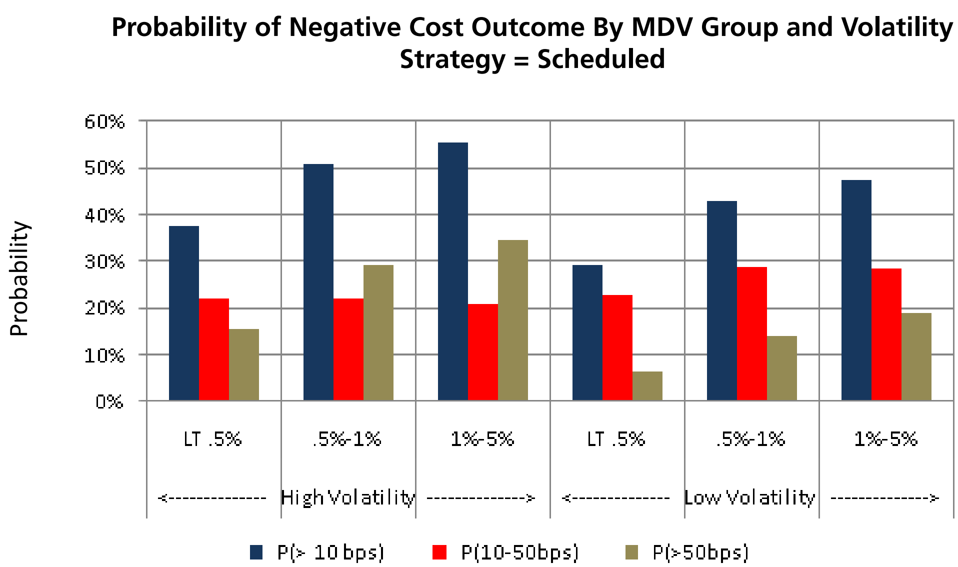

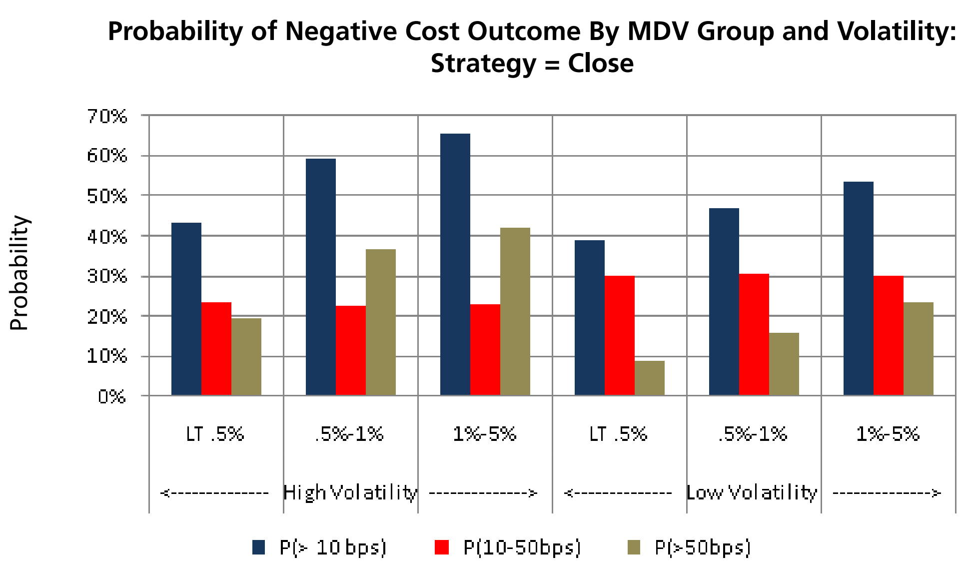

The figure combines high and low volatility periods, and replaces the empirical distributions by probabilities across liquidity demand groups.

For the Close, the probability of obtaining performance in the range of the averages (a loss of 10 to 50 bps) is the same across liquidity groups, rising only as one goes from high to low volatility conditions. The same cannot be said for performance loss exceeding 50 bps. In that case, the probabilities rise with demand for liquidity. The probability of a bad outcome in normal environments rises 144 percent, or 13 percentage points, as demand for liquidity increases. In a high volatility world, that increase is 22 percentage points, or roughly double, as one moves from very small to larger orders. The odds of obtaining a bad outcome for this strategy in a high volatility regime are now over 40 percent.

Similar behavior is observed for Scheduled strategies. The probability of obtaining outcomes in the range of the averages is flat across liquidity demand in high volatility, but does rise somewhat as demand increases in low volatility. The latter changes are relatively small however, and are confined to the move from very small orders to the group representing one percent of median daily volume. Behavior across volatility regimes, in terms of poor outcomes, is sharply different, however. For example, poor performance for the one percent MDV category in high volatility has a probability approaching 35 percent, dropping to under 15 percent in a more normal regime.

Both of these strategy groups appear to have unacceptable odds of a bad result in many conditions, and during high volatility regimes, in particular. Moving away from a dominant strategy within these two groups therefore does not necessarily pay dividends in terms of certainty of outcome as order size increases. On the other hand, shifts in strategy based on volatility, assuming a feasible alternative, seems to be a good idea.

Space considerations preclude including such charts for all strategies, but we illustrate one of the feasible alternatives in Figure 6 below, relating to IS strategy types.

Figure 6: Probability of Negative Cost Outcomes for IS Strategies

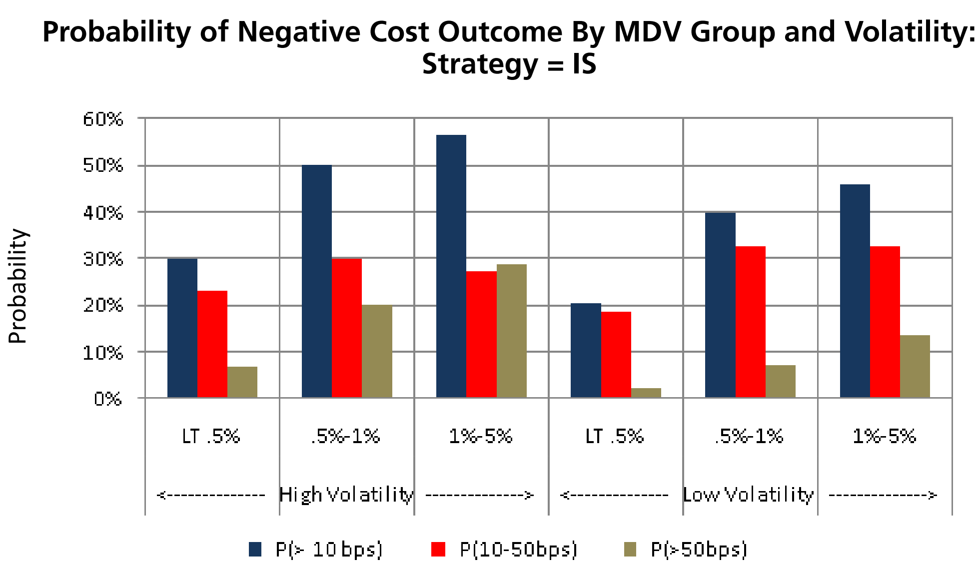

The similarities and differences between this chart and those for the Close and IS are of some interest.

The probability of observing performance in the range of the averages rises somewhat as one moves from very small orders to those in excess of 0.5 percent of MDV, but is otherwise stable across liquidity demand. This behavior mirrors that of Scheduled strategies, and is almost as stable as for the Close. The odds of a bad outcome also rise sharply with both volatility and demand for liquidity. On the other hand, those odds also are greatly reduced relative to the Close and Scheduled. For example, Scheduled strategies exhibit a probability of a poor outcome which is 85 percent greater than for IS, in the middle range of liquidity demand in low volatility. In the higher volatility regime, that figure drops, but is still a significant 40 percent.

V. Dominant Strategies and Surprises in Market Conditions

Dominant strategies exist in a world characterized by deviations from expectations, otherwise known as surprises. Surprises in market conditions do not constitute a new theme in performance analysis; a notable example is the work of Abrokwah and Sofianos.[9] Their study concentrates on surprises in transaction costs, relative to pre-trade estimates, as a function of surprises in price, volume, volatility, and spreads. Our goal is different, in that we continue to pursue the question, do observed strategy choices make sense? The underlying theme is the same: expectations matter, but deviations from expectations are a driver of performance.[10] We describe the methodology in terms of volatility surprises; the same techniques are used for volume.

Expectations are simply defined as historical averages, stock by stock.[11] Using the empirical distributions of deviations from the averages, on daily and intraday frequencies, a percentile rank is used to determine the extent of deviations. A day for which a stock experiences volatility in the top 20th percentile is said to have "high volatility" relative to expectations. If the volatility surprise is in the top 5 percent of occurrences, we term this "very high volatility." Approximately 20 percent of orders executed experience high volatility conditions relative to expectations, while 5 percent fall into the very high volatility category.[12]

Strategy usage is invariant, relative to volatility and volume surprises; we show this for volatility in Table 1.

---

9 Kwaku Abrokwah and George Sofianos, "Shortfall Surprises," Journal of Trading 2, Summer 2007.

10 Data considerations force us to change the sample relative to the previous sections. For the empirical work here, we confine analysis to August 2010 through January 2011, with respect to 480,000 unpriced market orders leading to over 8 million executions.

11 We deliberately avoid using forecasting models to produce expectations, preferring the simplicity of historical trends to uncertainty arising from the use of any specific forecasting paradigm.

12 The distribution of surprises is a bit fat-tailed but symmetrical.

Table 1: Usage of Trading Strategies by Volatility Surprise Regime

	Volatility Relative to Expectations					
Algo Type	Very Low	Low	Normal	High	Very High	Total
Close	4%	4%	3%	3%	3%	3%
Dark	4%	6%	5%	5%	5%	5%
IS	20%	23%	22%	23%	27%	22%
Liquidity Seeking	6%	7%	7%	7%	8%	7%
Participation	13%	11%	10%	8%	10%	10%
TWAP	3%	6%	6%	6%	6%	6%
VWAP	50%	44%	47%	49%	42%	46%
Total	100%	100%	100%	100%	100%	100%

The figures in the table represent the percentage of orders done via a particular strategy type. Despite some small movement in the case of VWAP, there is no significant variation in usage across volatility surprise regimes, further reinforcing the dominance of strategy usage.

As trading moves from expected market conditions to higher-than-normal volatility, the costs of all strategies increase, sometimes sharply, illustrated in Table 2.

Table 2: Transaction Costs and Volatility Surprises

Cost (bps)	Volatility Value Category				
Strategy Type	Very Low	Low	Normal	High	Very High
CLOSE	(4)	(3)	(4)	(9)	2
DARK	(8)	(7)	(10)	(17)	(33)
IS	(5)	(8)	(8)	(12)	(34)
LIQUIDITY SEEKING	(5)	(7)	(11)	(13)	(21)
PARTICIPATION	(5)	(3)	(8)	(23)	(13)
TWAP	(6)	(3)	(4)	(15)	(36)
VWAP	(1)	(7)	(10)	(25)	(49)

A case for the maintenance of dominant strategies can be made only in environments characterized by volatility lower than expectations. This conclusion is consistent with our earlier findings that dominant strategies make most sense for 'easy orders' executed through an algorithm. The spread of costs between strategies operating within high volatility surprise days is 16 bps, rising to 36 bps for very large surprises. [13]

13 We exclude the gain in 'close' strategies from this calculation. The anomalous figure for very high deviations from expectations is due in part to sample size at the close, and may also be an artifact of the way in which volatility impacts the arrival price relative to the closing price in this case.

In contrast, strategy performance is almost invariant across regimes characterized by volume surprises. Extremes in volume relative to expectations do not provide a meaningful rationale for a change in strategy type based on cost performance.[14]

These results on the degradation of strategy performance complement those of earlier work, in which volume is less of an issue than volatility, and strategy switches are suggested based on the rate of performance decline, not just on absolute comparisons.[15] We elaborate on this theme in Figure 7, with respect to VWAP and IS strategies, in particular.

Figure 7: VWAP and IS Strategy Performance with Surprises

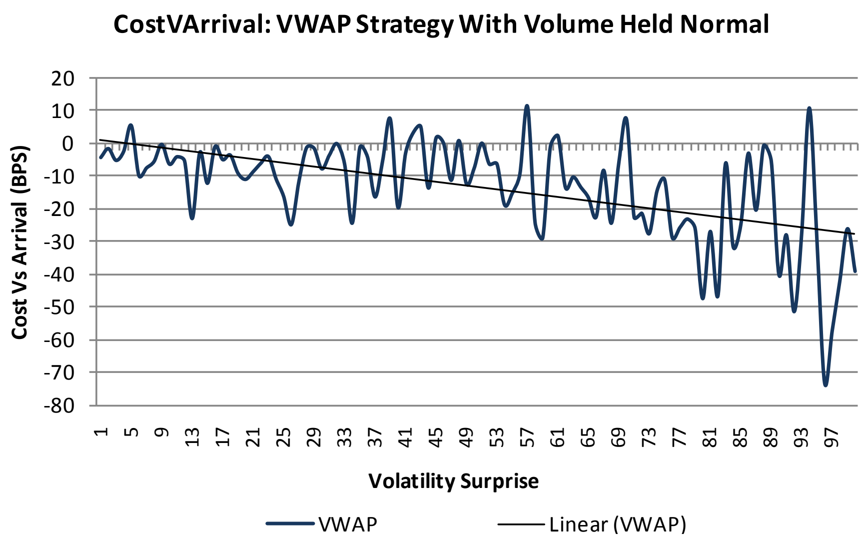

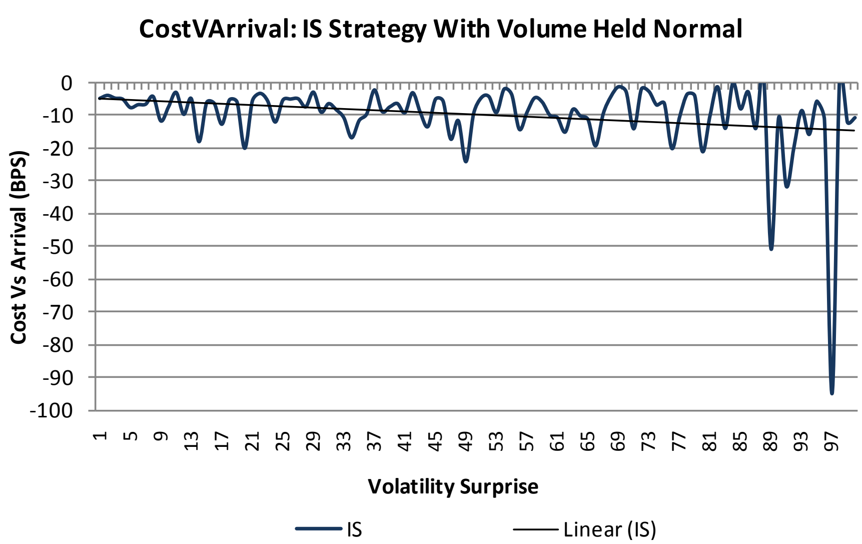

14 Although volume and volatility are typically characterized as highly correlated, the correlation between volume and volatility surprises in our data is only about 20 percent. High volume, high volatility days drive performance along the volatility dimension, as opposed to the volume axis.

15 Yossi Brandes, Ian Domowitz, Brett Jiu, and Henry Yegerman, "Algorithms, Trading Cost, and Order Size," in *Algorithmic Trading: A Buy-Side Handbook*, The Trade Publications, 2nd edition, 2007.

The straight lines in these graphs are statistical trend lines in cost as the volatility surprise moves from favorable to extremely high volatility conditions, relative to expectations. In similar charts for volume surprises, these lines are virtually horizontal for all strategy types; a zero slope of this line indicates no performance change across surprise regimes. Larger slopes, in absolute value, are measures of a decrease in performance as regimes become less favorable.

The slope of the VWAP strategy is roughly three times that of IS. VWAP may appear to be a candidate for a dominant strategy in favorable environments, but its performance drops relative to IS quickly as conditions deteriorate. VWAP and IS costs are similar only in the case of no volatility surprises. As one approaches a high volatility environment relative to expectations, IS becomes a preferred strategy in terms of absolute performance. The way in which a strategy reacts to abnormal conditions matters, and trading cost is one way of measuring its effectiveness. In the case of this particular comparison, IS incorporates volatility information in some form, while the scheduling inherent in VWAP effectively does not.

The story does not end with an investigation of average costs. Figure 8 illustrates a different sort of comparison of the two strategies, based on the distribution of observed outcomes.

Figure 8: VWAP and IS Cost Distributions with Surprises

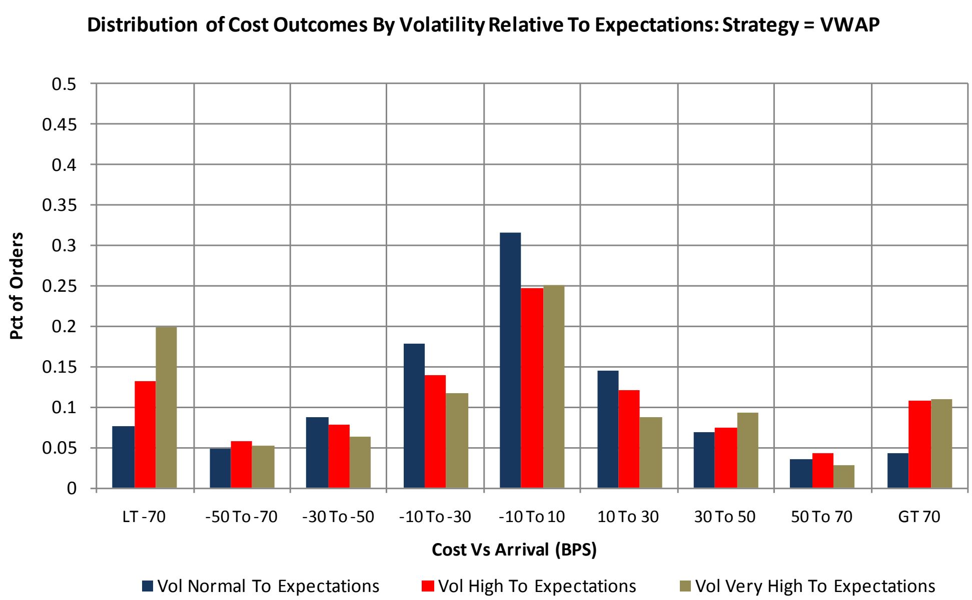

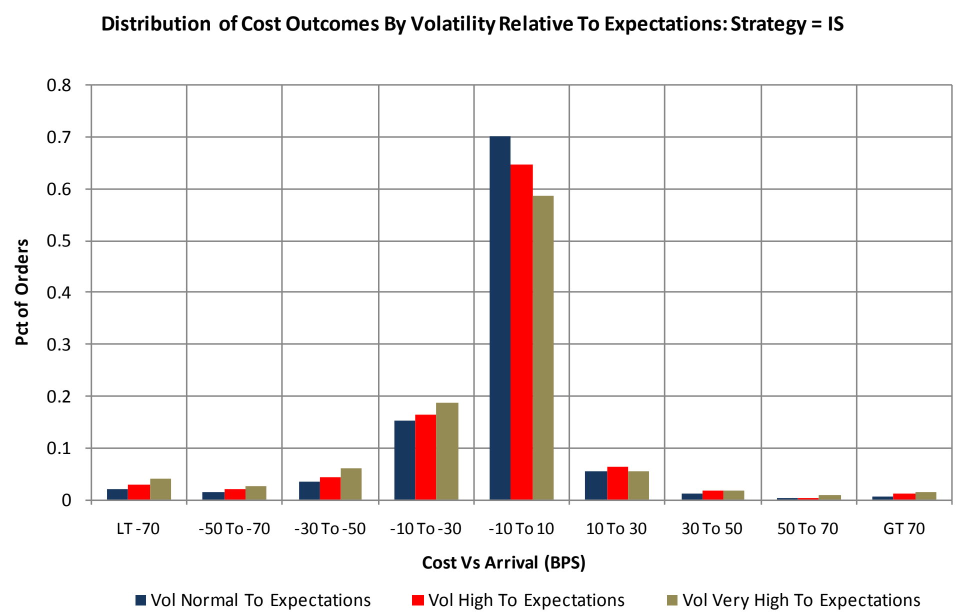

The distribution of IS performance is more concentrated than that of VWAP, implying a lower probability of a poor outcome regardless of surprise in market conditions.

The interesting comparison, however, lies in the shape of the distributions as volatility worsens relative to historical norms. In the case of the IS strategy, moving from normal to very high volatility occasions some increase in the probability of negative outcomes, but those increases are very small. The situation is quite different for VWAP. As one moves from normal to very high volatility conditions, the probabilities of negative outcomes appear to shift all the way to truly poor performance. We would characterize such outcomes as outliers, in common parlance, but there is a difference here. The probability of obtaining a VWAP cost greater than 70 bps rises from around 7 percent in normal markets to about 20 percent when volatility truly increases relative to expectations. The entire tail of the VWAP distribution is shifting to very bad results.

IS and VWAP together constitute polar extremes with respect to the strategy types considered here. Various permutations of these results are evident in the others. For example, Participation strategies show the slow degradation of IS, while the probability of obtaining a bad outcome increases sharply with the size of volatility surprise. In most cases, however, the data suggest the advisability of abandoning dominant strategies in market conditions which do not conform to expectations.

VI. Conclusions

We began with the intention of providing empirical evidence on algorithm usage and the costs associated with current practices. Two questions were posed: what choices are being made, and do these choices make sense?

Breakdowns of strategy implementation, across demand for liquidity, volatility, and time period, appear to suggest the use of a portfolio of strategies. Shifts in strategy are relatively marginal, however, at the level of any one buy-side desk. The vast majority of users employs dominant strategies, and tends not to change them as market conditions change. Dominant strategies are clustered in scheduled strategy types (e.g., VWAP) and cost minimization strategies, characterized as implementation shortfall algorithms.

There is a variety of reasons for the choice of few strategies by a desk. The most obvious, as expressed by several traders, is that it is only possible to know one or two "black boxes" extremely well. Trade characteristics such as fill rates also play a role. Most providers offer a complete suite across the spectrum of algorithms surveyed here, so the choice of dominant strategy has little to do with the vendor.

The choice of dominant strategies also has little to do with performance and market conditions. Dominant strategies constitute a sensible approach for very easy orders, and for situations which are extremely demanding in terms of liquidity and volatility. In the former case, there is no difference in performance as measured by transaction costs. In the latter, no strategy type does well, although performance is closely clustered across strategies. Performance matters, but does not distinguish individual strategy types in either regime. With some caveats relating to the Close, examination of the empirical distributions of performance suggests that dominant strategies make sense in market circumstances that are not extreme or otherwise unusual. We pursue this line of thinking by analyzing performance in situations which deviate from expectations and the historical record. In that case, strategy shifts are desirable, whether it be to minimize average costs or to decrease the probability of a poor outcome for any given order.

We now see the sense in which a majority of survey respondents believe algorithms to be commoditized, while a small fraction of them believe that performance is a differentiator. The commoditization of algorithms appears to extend beyond comparisons of a single algorithm across multiple providers; it applies across strategies themselves under some conditions. Ignoring performance may simply come from the use of dominant strategies in the first place. Turning a blind eye to performance as a differentiator makes sense only in special market environments. Indirectly, we recognize that the usage of an algorithm, not the algorithm itself, is the important point. This recognition suggests potentially large incremental value in education with respect to strategies and their use, in a world characterized by self-directed trading.

References

Abrokwah, Kwaku, and George Sofianos, "Shortfall Surprises," *Journal of Trading* 2, Summer 2007.

Brandes, Yossi, Ian Domowitz, Brett Jiu, and Henry Yegerman, "Algorithms, Trading Cost, and Order Size," in *Algorithmic Trading: A Buy-Side Handbook*, The Trade Publications, 2nd edition, 2007.

Cushing, David, and Ananth Madhavan, "The Stock Returns and Institutional Trading at the Close," *Journal of Financial Markets* 3, 2000

Special Section on Algorithmic Trading, Journal of Trading, Volume 4, Number 3, 2009.
Tabb Group, *US Equity Trading 2010/2011*: Outflows, Outrage, and Balance, December 2010.