Case Study: Effects of Execution Strategy on Statistical Arbitrage Performance

I. INTRODUCTION

Many investment managers decouple portfolio management—the process for portfolio selection that defines an investment strategy—from the trading process that implements that strategy. Sometimes this is a result of institutional structure, but in some cases it reflects lack of awareness of how much is actually at stake in execution. This separation can result in unnecessary costs and suboptimal strategy performance, which is especially true for high-turnover strategies: the higher the turnover, the more significant the effects of execution quality on investment performance. This case study applies various tools for analyzing a statistical arbitrage trading program—a relatively high-turnover middle-frequency trading strategy—and demonstrates what one can expect from examining the execution combined with the investment strategy.

Statistical arbitrage trading strategies are built on the premise of buying or selling mis-priced securities and profiting from the reversion of the security price to the efficient price. The most common approach is to build a factor model or co-integration model that prices each security based on the price of other securities [1]. When the security price deviates significantly from this benchmark, the trader bets on reversion of this security price to the efficient price predicted by the model.

The portfolio manager of a statistical arbitrage program needs to overcome two hurdles. First is to find a model that generates useful signals, i.e. identifies securities that will revert to some efficient price in a statistically significant way. Second, the reversion needs to be large enough to overcome the transaction costs involved in initiating and liquidating the positions. Often, these costs cause significant degradation in the performance of the strategy and even make an apparently profitable strategy unprofitable. The higher the turnover of the strategy, the more sensitive it becomes to market frictions and the more value there is to the strategy's performance in optimizing execution quality.

In this short paper we describe an analysis we have conducted for a statistical arbitrage program in order to improve the strategy performance via better trading results, to help the portfolio manager capture more alpha.

TABLE I Basic information about the trading basket.

METRIC	VALUE
SHARES TRADED	26.6 M
VALUE TRADED	$846 M
NUMBER OF TRADES	16,727
NUMBER OF SYMBOLS	839

This case study applies various tools for analyzing a statistical arbitrage trading program—a relatively high-turnover middle-frequency trading strategy—and demonstrates what one can expect from examining the execution combined with the investment strategy.

II. TRADE PROCESS

A) GENERAL INFORMATION AND TRADE PROCESS

Our data includes the trading history from April through June, 2011. The data consists of date, ticker, side, quantity, and average execution price for each trade executed during that period. The program optimized the portfolio, i.e., generated the basket to be traded, at 2pm each day. The basket was traded on the close using Market On Close (MOC) orders.

In Table I and in Figure 1, we provide some basic statistics about the trading during the sample period. A total of 839 names were traded over 16,727 executions. Figure 1 illustrates that the trading was cash-neutral, while the size of the trading basket changed, both in terms of value and number of names. The overall value traded is correlated with the number of names traded.

B) OBSERVATIONS

The delay between the 2pm program optimization and the 4pm trade time may provide an opportunity for the fund to extract additional alpha from its model. Possible benefits of trading over the entire interval include capture of alpha that decays during the waiting period and reduction of transaction costs by spreading the trading out over time. In addition, trading over the interval may increase the capacity of the strategy, as the market impact of concentrating the trades in a MOC execution may become problematic more quickly because of increased market impact.

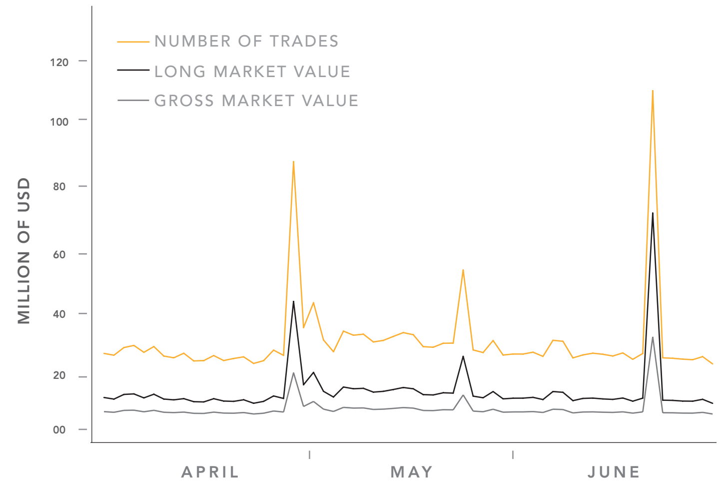

FIGURE 1 Time series for the gross and long market value and the number of symbols traded each day.

III. PORTFOLIO ANALYSIS

A) INTRADAY ALPHA

In Figure 2, we depict the value-weighted intraday return of all the trades as a function of time. The X-axis represents the time of day, and the Y-axis shows cumulative returns in basis points with respect to the price at 2pm on the trade date. The cumulative returns to the close price and next day's open price are indicated in the rightmost two points on the time series (time not depicted to scale). All the returns have been normalized to buys. The figure also depicts a two-sigma confidence bound for the average returns with dashed lines. As the figure shows, the average shortfall from optimization (at 2pm) to market close is about 0.27 basis points. In addition, the figure shows that during the first hour after the decision time the return on the basket is negative, i.e. the price becomes more favorable by about 1 basis point before it starts to revert. This momentum and reversion represents an opportunity: trading over that first hour might deliver some price improvement. A ten-fold cross-validation analysis (not provided here) demonstrates that this pattern is stable across the sampling period.

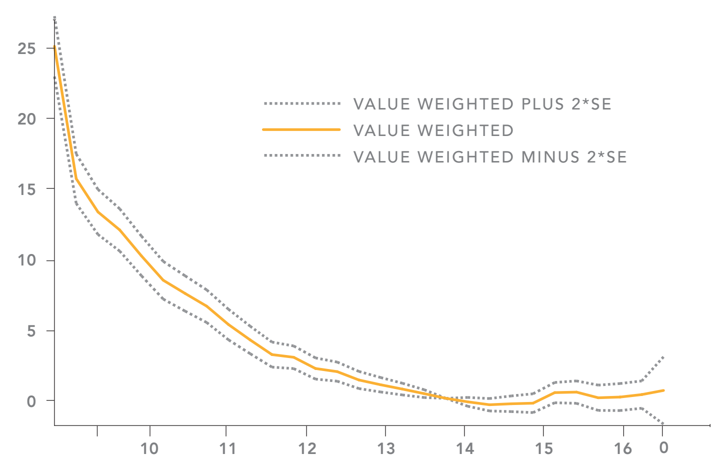

FIGURE 2 Value-weighted returns in basis points (buy-normalized) between the optimization time, 2pm, and various points in time during the trading day.

Shortfall Time series

In Figures 3 and 4, we depict the daily time series of the value-weighted average shortfall between reference price, which is the 2pm arrival price, and the close price. On any given day the average close price is relatively close to the arrival price, with some deviation that occasionally reaches 5 bps. In Figure 4, we depict for the same time series the best and worst shortfalls each day within the basket. The best and worst shortfalls each day are on the order of 200 basis points. To summarize, every day the execution price of some individual stocks deviates significantly from the arrival price, but these deviations on average mostly cancel each other out.

B) ORDER SIZES AND IMPACT

In Figure 5, we depict the cumulative distribution function of the order sizes relative to both the volume transacted between 2pm and 4pm and also relative to the size of the closing print. As can be seen from the chart, about 20% of the trades are 10% of the close print or larger. In addition, about 10% of the trades are larger than 1% of the volume traded during the 2-4pm period. In Figure 6, we depict a similar ratio but now we replace the contemporaneous volume with the average volume over the last 20 days. In summary, the trades are not large, and hence the expected impact is very low. As expected, in Figure 2 we do not observe any price reversion between the close price and the next day's open price. Having said that, tripling the size of the program while still concentrating execution at the close might put a strain on the current MOC execution strategy. A larger number of trades would be a significant fraction of the close print, and incur noticeable market impact.

IV. SIMULATION STUDY

Figure 2 suggests that trading in a continuous fashion between 2pm and the close has potential benefits. In order to understand these benefits, we have conducted a trading simulation study with different parameters. In particular, we have simulated trading each day's lists starting at 2pm with 15, 60, and 120 minutes interval VWAP strategies. In Table II, we provide the results. In addition to the simulated shortfall of each strategy (labeled alpha loss), we also provide the estimated net pass-through costs or rebates from the execution venues—assuming that on average "providing" generates a rebate of 25 mils ($0.0025) per share and that taking costs, on average, 15 mils per share. (These rates are conservative estimates based on our experience with such flow, so the realized profit might be larger.) Note that the average MOC pass-through fee is estimated at 8 mils per share. We assume a cost-plus model, where a broker charges a fee on top of the pass-through fees generated by the execution strategy, but that this fee does not vary depending on the execution strategy.

The best strategy, 60 minute VWAP, results in an estimated improvement of .37 bps in shortfall by completing the trade during the period when negative alpha is still in effect before the price starts to revert. (The VWAP shortfall is similar among all 3 candidate VWAP

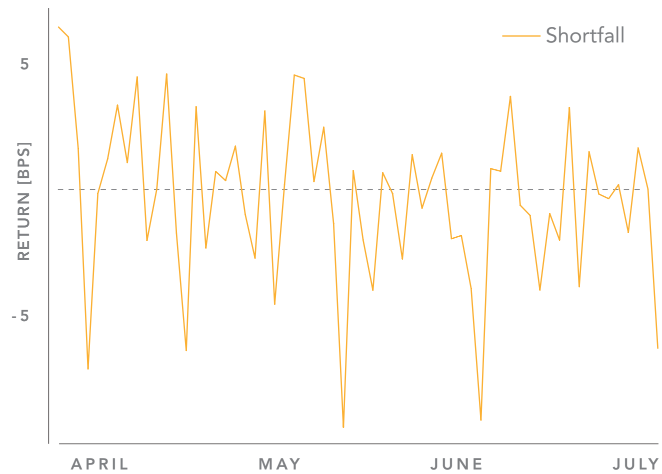

FIGURE 3 Value-weighted returns (buy-normalized) between the optimization time, 2pm, and the close as a function of time.

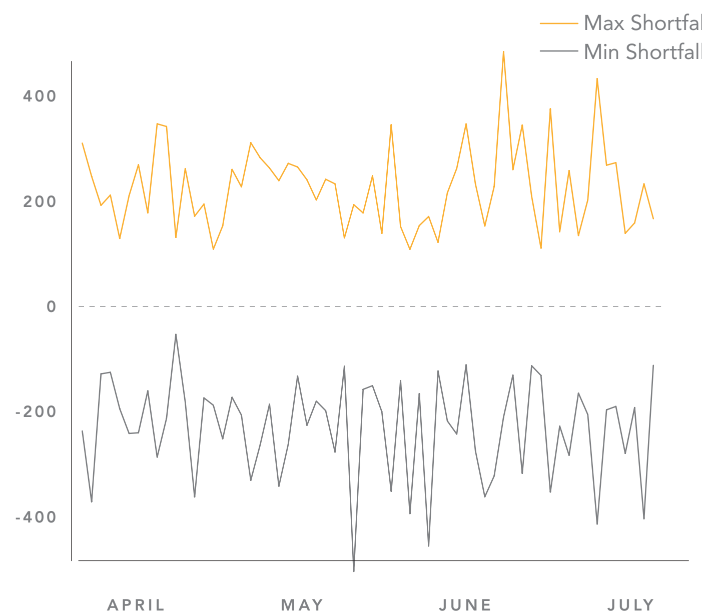

FIGURE 4 Best and worst shortfall in basis points (buy-normalized) between the optimization time, 2pm, and the close as a function of time.

TABLE II Simulation results for trading the list with various-length WVAP strategies.

STRATEGY	α LOSS (gain) [bps]	VWAP sf [bps]	LIQUIDITY COST (rebate) [mils/sh]	NET COST (gain) [mils/sh]
MOC	0.27	NA	8	16
15M VWAP	0.1	2.3	-1	2.2
60M VWAP	-0.1	2.4	-13.4	-16
120 VWAP	0.14	2.3	-18.6	-14.8
60M VWAP MOC (NET)	-0.37	NA	-21	-32

strategies, consistent with the orders being small and the overall rate of trading being modest even over a 15 minute trade interval). The explicit execution cost of the 60 minute VWAP strategy is a net rebate of about 13.4[1] mils/share versus a cost of 8 mils per share for the MOC strategy, a swing of 21 mils/share or about .64 bps. In total, we estimate potential saving of 32 mils per share, or about one basis point.

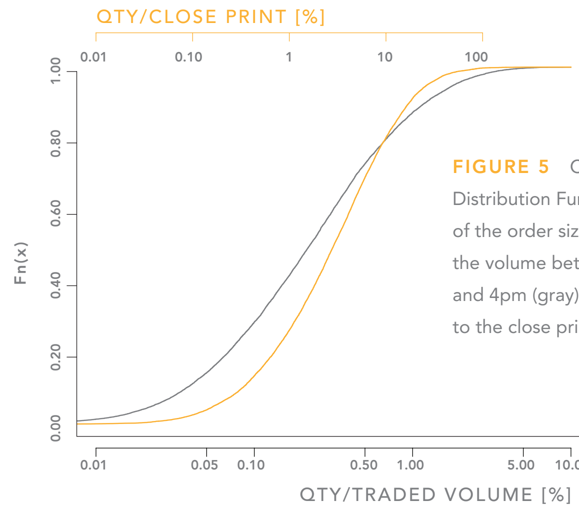

FIGURE 5 Cumulative Distribution Function (CDF) of the order size relative to the volume between 2pm and 4pm (gray) and relative to the close print (yellow).

V. CONCLUSIONS

Changing the execution strategy of the portfolio from market-on-close execution to continuous trading between 2pm and the close should provide several independent benefits:

- CAPTURE A SYSTEMATIC PRICE TROUGH BETWEEN 2PM AND 3PM.
- CAPTURE NET REBATES FOR PROVIDING LIQUIDITY TO THE MARKET.
- INCREASE THE CAPACITY OF THE PROGRAM.

We have used high-frequency trade simulations to quantify these benefits, which we estimate to be on the order of one basis point on a leveraged basis per turn around. This result illustrates the power of high-frequency trading simulation as a tool for traders, and the potential benefits to be realized in a systematic approach to considering execution quality as an important aspect of the investment strategy.

[1] In the table, VWAP shortfall is relative to the interval VWAP.

REFERENCES [1] M. Avellaneda, J.-H Lee, "Statistical arbitrage in the US equities market," Quantitative Finance, 2010, vol. 10, issue 7, pages 761- 782

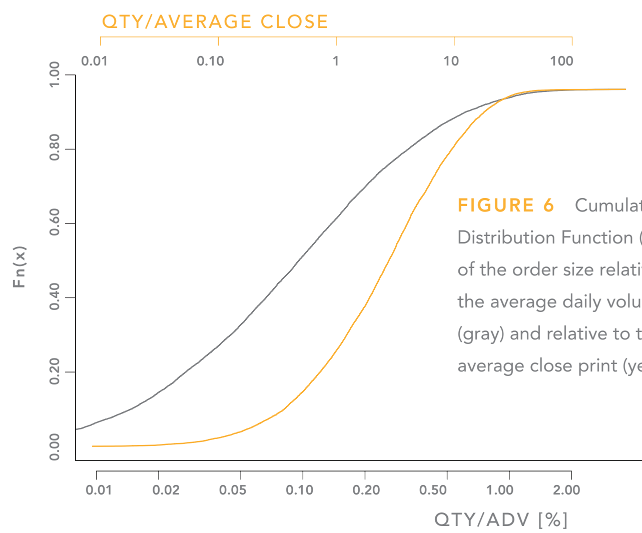

FIGURE 6 Cumulative Distribution Function (CDF) of the order size relative to the average daily volume (gray) and relative to the average close print (yellow).

For questions or comments, please email Dr. Eran Fishler, Director of Research (technotes@pragmatrading.com).
Pragma provides comprehensive broker-independent trade cost and trade process analysis services.