

Only For Super Nerdy Eyes



At the time, I actually didn't know - but I do now.

"Dave, it's always fun talking to you, because as smart as you are, you are sooo naive. Don't you realize that the buy-side will never pay more for trading with computers than with actual traders, it doesn't matter if they work better?"

This conversation, 15 years ago, between yours truly and the head of a major buy-side trading desk, epitomizes the feelings among the "old guard" on Wall Street. The punch line in the conversation came later, however...

In the fall of 1999, I had just designed the first computerized trading algorithm at Salomon Smith Barney. That algo, aimed at meeting the Volume Weighted Average Price (VWAP) benchmark, had been tested on our portfolio trading desk, in competition with two dedicated human traders. The results were startling; the computer outperformed the humans by roughly 4 CENTS per share. Back then, stocks still traded in increments called sixteenths or 6.25 cents per share, and the average spread was higher. The improvement was roughly equivalent to 30-35% of the bid offer spread.

I was, of course, really excited about our findings and asked a friend, who was a buy-side trader, if Salomon would be able to charge an extra cent per share for delivering 4 cents of outperformance...

After telling me how naive I was, he went on. "You know, as well as I do, that execution quality is subjective. While we use VWAP as a "tick the box" compliance requirement, it is not that clear that it translates into real trading costs. As you know, the decisions on who to trade with and what commissions we pay, are mostly about relationships, research and banking."

At the time, I actually didn't know - but I do now.

And, despite all the changes in market structure and the, relentless, drive of technology to reduce spreads and trading costs, it's *still* true. What's hard for people to understand, is that it actually does make sense in many cases. It is completely reasonable for an asset manager to knowingly trade with a broker dealer that will deliver inferior execution quality, as long as the following holds true:

1. The broker dealer must deliver value to the manager that exceeds the excess cost incurred by trading with that broker dealer

and

2. The asset manager should be able to quantify the costs of executing, with each broker dealer, so that they can make optimal trading decisions.

How can it make sense to trade with a broker that delivers poor execution quality?

The answer:

Institutional brokers can add value to asset managers in several important areas, including access to research and corporate executives, data, trading ideas and, of course, access to the IPO calendar.

Earlier this year, I had a conversation with a senior person

at a regulator where I explained the following:

"Consider an example where you have two broker dealers that offer trading algorithms at the same commission rate, but broker "A" delivers performance that is 0.001% worse than broker "B" (1/10th of one basis point)."

(He asked me if that is a typical variation, I told him that typical performance variations range from that, to a few times that, but that it would serve for the example.)

"Next, consider that Broker "A" has an investment bank and that this asset manager, who trades about \$10 billion, in notional value per year, is interested in being allocated shares of company IPOs."

(Since the day of this conversation was the day that Shake Shack went public, I used it as my example.)

"So, if broker "A" were able to allocate a manager 25,000 shares of Shake Shack in the IPO, they would have added over \$500,000 in value to the investors, in that manager's funds that day. This means that, despite broker "A" costing about \$100,000, over the course of the year in inferior performance, the manager should have traded with that broker if such trading "earned" them access to the IPO.

His eyes widened in recognition. The next part of the conversation was why I included the second caveat. "Obviously, the real situations faced by managers are not that clear cut", I continued. "More often, managers allocate a meaningful percentage of their commission dollars based on internal votes from portfolio managers and traders. These votes could pertain to research, corporate access, and IPO allocations as well as the quality of the trading relationship." After digesting this, he asked me if I believed that there was a consistent method that the buy-side used to measure trading value. My answer was that most firms had some type of voting process, but that there was a lot

of variability in how different items are valued. In addition, other than a handful of extremely quantitative firms, few were able to quantify the performance differences between their brokers and the individual trading strategies they utilized.

The reason, I explained, is that most buy side firms need to upgrade the transaction cost analysis (TCA) that they employ. TCA, properly constructed, can provide actionable information to evaluate trading strategies and the capabilities of individual broker-dealers. In order to do so, however, asset managers should consider three key areas where improvement may be needed.

First, pre-trade benchmarks that measure “implementation shortfall” that have been constructed using industry averages, should be upgraded to be more specific to individual firms. Parameterizing such analysis to consider the nuances of asset managers own trading will make it more accurate.

Second, TCA that evaluates trading compared to participation based benchmarks, such as VWAP, while providing some value, is not sufficient. The goal should be to compare execution quality to the pre-trade expectation of cost.

Third, and most important, to create actionable information, TCA platforms should analyze individual executions and orders, in the appropriate context, based on the type of order and the subsequent market movement.

The best pre-trade benchmark is the cost that the portfolio manager assumes will be incurred when implementing their decision to trade. In a well-run investment management firm, trading decisions are based on a combination of four factors: alpha (the expected price performance of the assets), beta (the desired correlation to the funds index benchmark), systematic risk (the variance created by excessive correlation to individual factors), and expected trading cost.

To understand why trading costs should be part of the initial trading decision is simple. If the predicted trading cost to enter and exit a position becomes greater than the predicted “alpha” (outperformance of the benchmark) then the trade would make little sense. Quantitative fund managers refer to this phenomenon as understanding the “capacity” of the fund to invest. They are careful to have a market impact model built into their investment process, to avoid such concentration risk. To bring this back to the topic of TCA, the most relevant pre-trade benchmark should be to match whatever model each manager uses to predict their individual trading costs. (If a portfolio manager does not have such a model, we would be happy to provide insight on how to create one).

While it's important to analyze parent orders given to individual brokers, it's hard to derive actionable information from them. The trading cost incurred on those orders are derived from multiple factors and each of them should be measured separately. In addition, it is hard to attribute costs at the parent order level to specific system behaviors, since both the time period for trading and the method of slicing up orders are often constrained. Many algorithms such as “VWAP” and “Participation” are essentially forced to trade based on either historical or that day's volume pattern, so brokers have little flexibility. In addition, many brokers offer parameters to their clients that limit the venues, types of counterparties, or “aggressiveness” of the algorithm. The impact of these constraints is to make it hard to either attribute performance to individual brokers or their strategies or compare different strategies.

So, what can be done to provide actionable data? The answer, is to focus on a “bottoms up” analysis of the key trading methods that all trading systems rely upon: smart order routing, and order placement.

Routing can be analyzed in two critical ways; routing efficiency and venue specific liquidity analysis. Routing efficiency of a smart order router (SOR) can be measured by comparing the available liquidity, in the market as a whole, to the volume that the router was actually able to

execute. This type of analysis would show how well SORs were able to find hidden liquidity, and would also quantify potential information leakage caused. This analysis requires knowledge of the orders sent to the SOR and not only the actually executed orders sent by the router. This would give a picture of the impact of so called “swing and miss” orders routed to various exchanges, dark pools and external liquidity providers. The resulting analysis will show both the efficiency by which available liquidity was accessed as well as capturing the impact of adverse price moves that might have been caused by sub-optimal routing strategies. (as described [previously](#))

Order placement decisions can also be evaluated by understanding the quality of the executions that result and by the market movement in response to placing such orders. As a general principle, it is better to have higher execution percentages and lower subsequent adverse price moves. It's also important to measure the market impact of particular venues and display strategies, particularly when evaluating posting orders in dark versus lit venues. (The commonly held belief that orders placed in dark pools create less impact, is why such venues are often used.)

As a final note on TCA, participation based benchmarks, such as VWAP, do provide perspective on how orders are handled, It is, however, more important to analyze the market impact of such orders, by measuring the VWAP, in the period subsequent to the completion of an order. The analysis of what quantitative traders call “reversion,” is relevant, as it shows the market impact created by the order. Large reversion metrics can be indicative of strategies that are either too aggressive or that leak too much information.

In order to bring my story full circle, it is clear that many asset managers are still treating execution quality subjectively. There is a widespread focus on qualitative measures in the community, which has resulted in most large brokers receiving routing and dark pool questionnaires. There seems to be a commonly held idea that buy-side traders should choose for their brokers

where they should route orders, and what counterparties their brokers should trade with when they route. This idea is an oversimplification at best, and counterproductive at worst. It would be far better for buy-side traders to quantitatively measure the performance of their brokers routing methodologies instead of constraining them. If properly conducted, careful measurement of trading strategies, combined with a quantitative approach, can provide substantial insights and improvement in the trading process.

David Weisberger is the Managing Director and Head of Market Structure Analysis of RegOne Solutions, by Markit
[Email the Author](#)

The views expressed are my own and do not necessarily represent the views of Markit or any of Markit's divisions or employees.