IF EVERY TRADER GETS A TROPHY: WHY VWAP HURTS MANAGER PERFORMANCE…

Institutional trading has become commoditized

Equity trading has evolved into a profession, which, all too often, awards trophies to most, if not all, of the participants. Benchmarks such as VWAP and other participation based metrics are "Pass Fail" metrics that fail to encourage excellence. While it is hard to prove, given the lack of proper attribution of trading costs by investment consultants or managers, it is likely that poor trading is part of the explanation for the sustained poor performance by active managers. Sadly, this is the predictable result of bundled services from brokers…

VIABLE MKTS

July 2017

Why has VWAP hurt trading performance?

What metrics and strategies should managers use?

Author: Dave Weisberger

ViableMkts

Dave@viablemkts.com

Bundling created need to justify “trading as a commodity”

For many decades, equity commissions have been allocated by a process called the “Broker Vote” which considers research, corporate access, investment banking, and other services when determining which firms to trade with. This means that trading, is not usually “bought” based on the quality or price (commission rate) of the trading itself. It is, instead, earned by providing value through those other services on a “bundled” basis, with the only caveat being that trades need to be “good enough”.
It has been in the collective interests of both buy and sell side firms to justify this “bundling”; the buy side keeps the “free” services it receives and the sell side preserves the relatively high margins that “bundling” provides. For both, bundling provides barriers to entry against smaller competitors, so it should surprise no one that VWAP (Volume Weighted Average Price) is the most used benchmark for institutional trading, as well the most used algorithm type. VWAP, you see, is a benchmark that forces traders into “herd like” behavior, which makes it very hard to distinguish a firm’s ability to source liquidity or mitigate information leakage. It is also among the easiest benchmarks to measure, which means that asset managers get the twin benefits of not having to spend a lot of money to evaluate trading and, at the same time, wont often see results that force them to change the brokers they use.

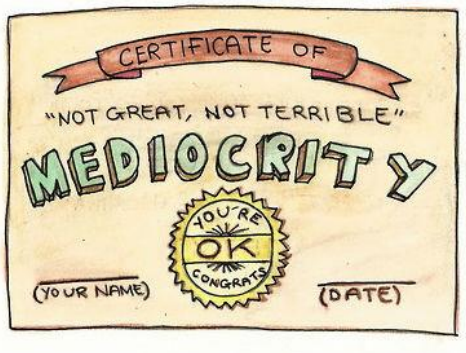

Unfortunately, while there are some valid reasons for this arrangement, the subpar returns of active management, which I have previously written about, provide an excellent reason to question brokers and trading strategies. The core problem is that the prevailing approach assumes that trading is a commodity, with little competitive differentiation. In the real world, however, trading is extremely differentiated, with some firms better able to provide / source block liquidity, and others having superior quantitative technology including smart routing and sophisticated algorithms. That said, so long as the buy side uses metrics like VWAP, they will be on the outside looking in.

The "Solution" to Mediocrity: Grade on a CURVE!

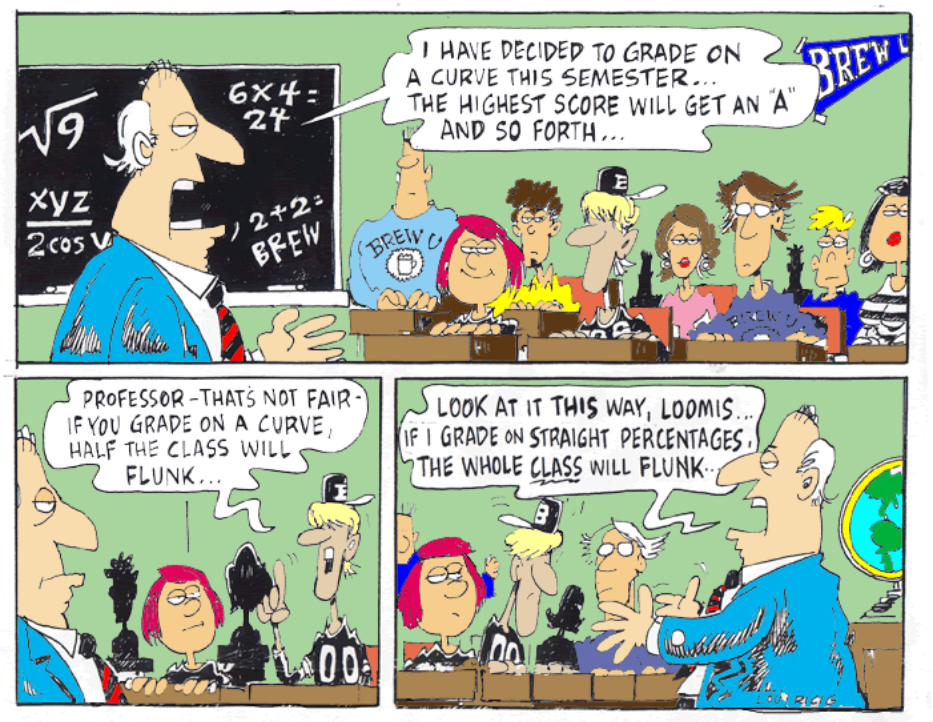

VWAP is the TCA (Trade Cost Analysis) version of grading on a curve, which guarantees that if there are enough failing performers, some can get a passing grade. Consider an example where a trader pays much more for a stock than the portfolio manager expects, BUT the rest of the market follows and trades a lot of volume at the same prices. In that case, the trader could match the VWAP, but the portfolio's returns will suffer, when the stock inevitably reverts to the mean, post trade completion. At least the trader, however, received their "passing grade" …

That said, the damage is not limited to the promotion of mediocrity. Trading, when held to the VWAP benchmark, biases intra-day trading costs higher due to asymmetry between key cost drivers of trading and the volume profile of most stocks, and impedes liquidity discovery by restricting traders to proportional trading of each order instead of finding blocks as well.

If traders are forced to ignore costs, can they do their job?

The following chart is a representation of a cross section of liquid equities throughout the trading day:

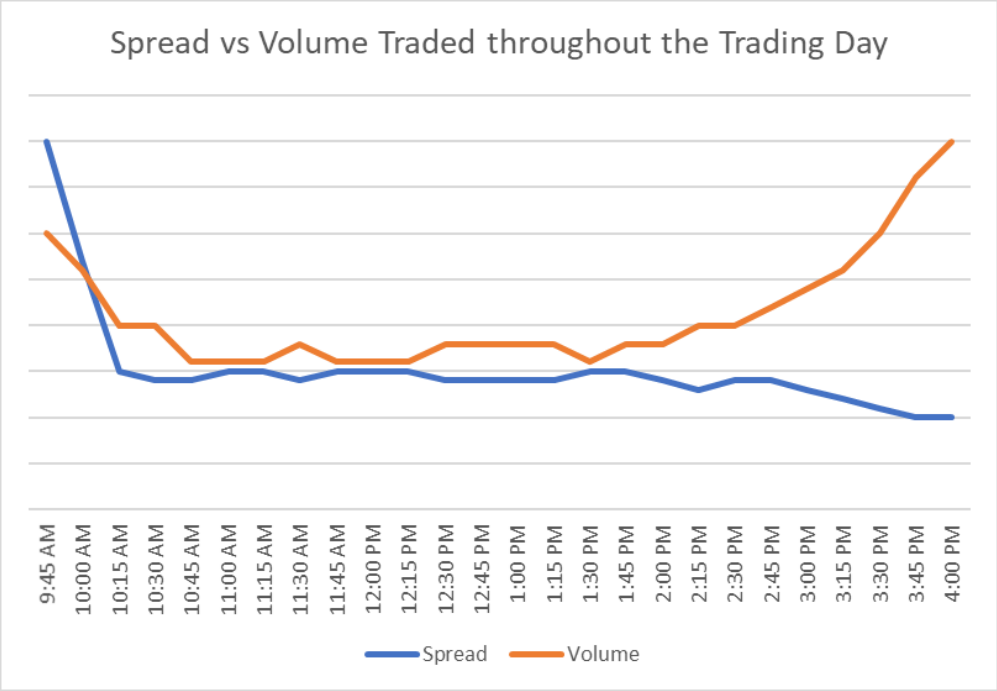

Bid offer spreads, for most stocks, fall by roughly 2/3rds from the first 15 minutes until the end of the first hour and volatility (not pictured) falls by a similar amount. Displayed sizes are also smaller early in the day, although not to as large of a degree. All three cost drivers fall again in the last hour of trading, particularly as it approaches the close. At the same time, however, the average volume traded has starts proportionally higher early in the day, falls over the first hour, stabilizes, and then rises into the close, where it reaches its highest point of the day. Thus, by forcing traders to match the VWAP, the higher volume traded early in the day compels traders to pay more to execute that substantial, early part of an order since costs are higher at that time. At the same time, by limiting the ability of traders to opportunistically capture those higher spreads or capitalize on early volatility either in the market or via block trades, rigid adherence to VWAP reduces the ability to achieve outperformance. In statistical terms, we call that truncating the "good" side of the distribution of potential outcomes.

Do Managers REALLY want DIFFERENT incentives than the traders who are working for them?

"Try this—I just bought a hundred shares."

Last, but not least, VWAP creates a mismatch of incentives, essentially pitting the interests of traders at odds with the portfolio managers. The cause of this is information leakage.

Participation based metrics such as VWAP, allow traders to impact the metric by their own trading, since their own orders and trades become part of the metric. This creates a conflict of incentives between asset managers and traders.

The reason for this, is that the optimal strategy for trying to perform better versus VWAP is to create market impact, which moves the benchmark and, thereby, makes it easier to beat. While obvious, it worth stating that increased market impact translates directly into increased trading costs, meaning that managers will do a poorer job of capturing alpha.

To illustrate how VWAP creates incentives for traders to create market impact, consider the following crazy order handling strategy, which, while extreme, helps to make the point:

> Trader receives an order to buy 250,000 shares of a stock that trades 1.5 million shares a day on average, with, on average less than 1000 shares displayed on the bid or offer at any point in time.

1. Trader enters a visible bid, on the bid side of the market, for all 250,000 shares.
   ➔Most likely, little to no shares are purchased, but the market moves higher

2. Trader then cancels that bid and enters a new bid for 100,000 shares at the new bid

➔Most likely, the same thing happens and the stock moves higher

3. Trader continues to cut the quantity bid for, until the orders start to get filled, which would indicate that the impact has run its course.

4. From that point on, trader trades "normally", meaning using dark pools and much smaller sized orders to capture the reversion in the stock price that is likely to occur.

➔The Trader "knows" that his order is not THAT large compared to ADV, while his previous behaviors "looked" to the market like there was a much larger buyer. That is why reversion is expected, when the buying pressure does not materialize.

This type of obvious malfeasance is highly unlikely (due to compliance departments doing their job), but it illustrates how a benchmark, which can be manipulated during the trading of the order itself, is vulnerable. VWAP and other participation based benchmarks rely on trading that is yet to happen, which means that IF traders can trade in a way to push prices higher, they have an incentive to do so. This encourages info leakage on purpose, so traders can capitalize on reversion. For my money, if I were a buy side portfolio manager, I would not want to provide incentives to traders to do anything other than minimize my cost of trading, and VWAP is the opposite…

So, what should managers do?

Match their trading strategies to the reason for each trade!

- Trades to reduce/manage risk should be executed as soon as practicable, since they are often crowded trades. These trades should use the arrival price to the trading desk and managers should ensure that they are delivered expeditiously.

- Trades to implement inflows / outflows should probably be traded into the close. This assumes that the new investors/exiting investors will receive the NAV at the end of the day, which can be assured by trading into the closing auction.

- Trades to capture predicted alpha should be tailored to the evaluation price used by the portfolio manager, analyst, or model, as well as the magnitude and time-frame of the predicted alpha. For long-term alpha, minimizing actual trading cost is more important than immediacy, but for higher conviction or shorter term predicted outperformance, immediacy is desirable. It is worth noting that, in practice, alpha capturing trades are typically measured against the price at which the order arrives on the trading desk, perhaps adjusting that price for moves in the market.

Lastly, managers should ensure that the venue or broker analysis they use, incorporates the context of the orders and their purpose in the analysis. The most important point is to understand which orders need to demand liquidity and which do not need to do so. For large institutional orders, for example, attempting to achieve "spread capture" is a likely source of poor performance. Who cares if you can buy at the midpoint or on the bid side of the market, when your goal is to buy a substantial amount of stock. In that case, the goal should be finding liquidity and minimizing the market impact of trading until the order is complete. That is why ***it is vital*** to measure the ***opportunity cost of not trading*** as part of any overall analysis, or, quite often the wrong conclusion is reached. (see "what you don't know CAN hurt you" for more details on this point). This, in turn, requires both better benchmarks and analysis of all the orders routed.

Bottom Line: Changes to Process & Measurement is Needed, or Performance Will Continue to Lag

In conclusion, for active managers to improve their performance, they need to stop pursuing mediocrity, and move away from using VWAP and similar benchmarks. Portfolio managers should provide their trading desks with more precise parameters and the rationale for trading, while the trading desk should use more appropriate benchmarks and strategies for each type of order. Portfolio managers should have an expected trading cost for every order, and should utilize statistical techniques to evaluate how well their firms trading process matches those expectations. (They should also use those cost expectations when determining the size of the positions they want to hold.) Heads of trading, meanwhile, should segregate their trading based on the reasons for each trade and review how well their process matches those goals. Proper TCA can be built (or bought) to analyze their own and their broker's trading processes, but such statistical measurement needs to be in the proper context of the orders and must include both explicit and opportunity costs.

My hope is that movement towards unbundling, as MiFID II is pushing in Europe, will spur managers to examine their trading processes and improve them. If this happened, it would likely herald major changes for the relationship between institutional investors, and the broker dealers that trade on their behalf. It would mean that trading services will be "purchased" based on the quality of the service, instead of other reasons. That, in turn, should encourage improvements to both trading and measurement processes, which will improve active management performance.