

# TCA – what’s it for?

Darren Toulson, head of research, LiquidMetrix

We’re often asked: beyond a regulatory duty, what’s the purpose of TCA?

Done correctly, TCA can tell you many things about your current execution performance, including why your performance is good or bad and what you can do to improve it. Done poorly, TCA is something you run once a quarter, file away and forget about.

So how can you best use TCA, what might your TCA be telling you and what kind of questions should you be asking of it?

## **TCA performance measures: Implementation Shortfall and VWAP – Are they useful?**

The summary sections of most TCA reports are usually presented as a set of high-level performance metrics showing how good or bad your order execution was. Results are typically presented with your flow broken down by broker, region, order size and algo; the purpose being to present your overall performance as well as highlighting specific areas where you may be under or over performing.

## **Which performance metrics are generally used?**

There are two key families of TCA metrics that most people focus on. VWAP measures compare the price you achieved versus the market price over the period of your order. Implementation Shortfall (IS) examines the difference between the market price when you decided to trade (or started

trading) versus the average price you actually achieved when completely filling your order.

Within these two broad definitions there are many different ways to calculate the measures. For instance, the benchmarks used may be based on VWAP for the whole Trading Day, Interval VWAP from first fill to last fill, VWAP including or excluding certain venues or trade types, VWAP with a limit price applied or a ‘PVWAP’ simulating what price you would achieve if you had participated at a fixed percentage of market volume until your entire order was filled.

The level of complexity and customisation that has been required over the past years for the top level IS and VWAP figures hints at the fact, that for TCA to be really useful to a wide range of different participants and trading styles, it’s very hard to tell the whole story using one or two simple measures.

There’s nothing wrong with IS and VWAP. IS tells you how close your execution price was to the price in the market at the time you decided to trade; VWAP measures how well your broker/trader has managed to time the executions such that they capture market spread and trade at a constant participation rate. These are useful things to know.

But what it does suggest is that you should see any ‘top level’ TCA measures as exactly that: numbers to be explained by other statistics or breakdowns so you know why performance is good or bad.

## Moving from simple TCA performance reporting to 'Why?'

Let's consider one of the more popular, top level TCA measures: Interval VWAP. Say your TCA report tells you that for a certain batch of orders, you underperformed Interval VWAP by 2.4 BPS with a standard deviation of 8.6 BPS.

The obvious question is why? To answer this question let's consider what you have to get right in a trading strategy to match or outperform VWAP.

First, you need to trade consistently at the market rate throughout the order duration. This will reduce the standard deviation. But if you want to beat average VWAP over all of your orders you need to capture the bid/offer spread as often as possible on individual fills, otherwise on average you may expect to miss VWAP by half the bid/offer spread. To capture spread you will need to rest on various lit and dark venues rather than aggressively trading in the market. But some venues you rest on may not always provide as much spread capture as desired due to HFT participants picking you off at bad times or because of simple adverse selection.

Care must be taken when you're unable to fill passively as when aggressing the lit market, it's easy to go too deep into the order book and pay prices away from short term TWAM (Time Weighted Average Mid). Also, when accessing markets in different locations you need to be careful you don't get front run by HFT traders cancelling or trading orders in front of you.

In other words, VWAP performance is a combination of timing, spread capture, adverse selection, short-term impact/reversion, venue selection and HFT gaming. In the past, most TCA analysis glossed over this kind of low-level detail and presented just the top-level numbers. The reasons for this may have been a combination of lack of clean high frequency client and market data, relatively less powerful computing options and perhaps the fact that in the past, markets and algos offered fewer choices. But in today's markets, there really is no reason not to dig deeper and look for the 'why' behind the top level numbers.



Darren Toulson

## The TCA 'Pyramid'

We view TCA as not just a set of top-level numbers, but rather a pyramid of different metrics with each level helping to explain the performance of the level above. (See Figure 1.)

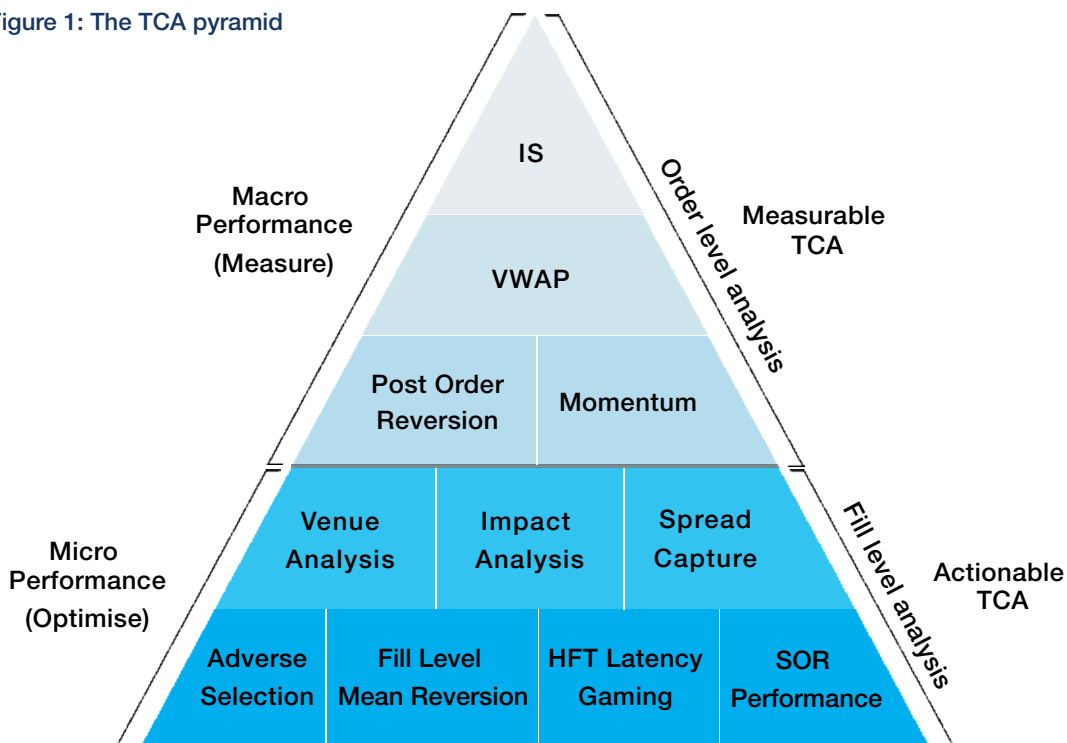
At the top of the pyramid are IS and VWAP. As we go down the pyramid we increasingly look at more and more granular aspects of trading relating to fills, venues and even the effectiveness of order submission strategies to counteract HFT gaming at the millisecond level.

Most people looking at TCA reports, especially on the buy-side, are unlikely to have direct responsibility or control over many of the micro-strategies being used to execute their orders.

However, it's important to realise that overall execution performance will ultimately depend on how well others are implementing those choices on your behalf. Brokers / Venues / Algos that perform well at a micro level should, on average over a period of time, mean better top-level performance. Mistakes made at the low level through bad venue selection, bad algo selection or bad trading connections / SORs will 'bleed' into overall performance figures. Individual mistakes at the fill level are usually trivial, but repeated over many executions can lead to death by a thousand cuts.

Another advantage of drilling deeper is

Figure 1: The TCA pyramid



to separate out luck from deserved good performance. The fact that Broker A's VWAP algorithm is better than Broker B's for a couple of hundred orders in one quarter may well be fairly random and the ranking may reverse next quarter. But if you can see that over thousands of fills, Broker A had better spread capture or that Broker B was accessing a 'toxic' Dark Pool (illustrated by sharp price impacts and bad TWAM) then the fact that Broker A is better is more likely due to skill than luck.

Drilling into this level of detail also allows buysides to get more comfortable that their broker is acting strictly in their interest. By seeing breakdowns of which venues a broker routes most orders to, can illustrate any venue biases the broker might have when compared to the overall market. For instance, if a broker sends a relatively large proportion of flow to its own broker crossing network (BCN), this is fine as long as it's clear that the performance of the flow sent to that venue

is at least as good as flow routed elsewhere. If a broker posts a lot of flow passively to an MTF with a rebate scheme, then again, this is fine as long as you can see that the strategy is effective at both the fill and order levels.

### Some practical examples

We will show two practical examples of the TCA pyramid. The first is an example of using micro metrics to explain macro performance and the second of using micro to identify good / bad routing choices. (See Figure 2.)

The view above shows a breakdown of a set of day-long VWAP orders executed via a broker. The basic top-level TCA metrics are that the orders missed VWAP by 2.62BPS and had an Implementation Shortfall versus market opening price of 35.95BPS. So performance is OK but not great.

The rest of the view breaks down these top-level numbers into factors that explain the performance:

Figure 2:



- The price chart to the left shows market prices (in the same direction as each of the orders) before, during and after the order. The key point is the large step in prices between the open and arrival time of orders: over 20BPS. So, most of the 35.95BPS of IS cost is caused by orders arriving 'late' to the trading desk. There was also a small cost (2.12BPS) for the delay between orders arriving at the trading desk and the first fill.
- Spread Capture for fills was approximately 50%, so the 2.62 BPS 'miss' for VWAP wasn't due to not capturing enough spread but more likely a result of bad timing. We can also see from the Fill Rate graph that the orders are somewhat rushed towards the end of the order (the fill rate curve is not perfectly uniform and rises at the end). Generally this will not help VWAP.
- There is clear evidence of mean reversion occurring after the last fill – about 4.5 BPS within fifteen minutes after the last fill and a similar amount again by day close.
- Price momentum the day before the orders are received is actually slightly negative, so the relatively poor performance is most likely not due to having to chase a rising/falling stock.

This only scratches the surface but does show how an overall VWAP or IS number is more useful when decomposed down to reasons why the number is good or bad.

Further down the pyramid, at fill level, the figure below gives an idea of the type of analysis that can be done at this level of detail.

The two histograms below (Figs. 3 & 4) show spread capture for resting orders in two dark pools. Pool 'A' is a near perfect EBBO mid-point matching pool and almost all fills happen at exactly 50% spread capture. Pool B, on the other hand, shows strong adverse selection; many of our fills actually happen at EBBO touch rather than mid and overall we only get about 35% spread capture.

Figure 5, below shows high frequency price movements just before and after fills happening on 3 different pools. Pool A is excellent with very little price movement on the main market before or after our trades so we can be confident of trading at a neutral time. In contrast, our fills on Pools B and C appear to happen mainly when prices are moving so we can see adverse impact followed by mean reversion.

This probably means the other side of our trade

Figure 3: Dark pools – spread capture histogram




Figure 4: Dark pools – spread capture histogram





Figure 5: High frequency price movements just before and after fills



is picking this time to trade to their advantage and our disadvantage: we are being gamed.

The point of this type of low-level analysis is to explain why top-level numbers are good or bad. A bad spread capture on a dark pool will lower average spread capture and mean we may underperform VWAP.

Mean reversion following passive fills means we will be trading away from TWAM and underperform VWAP. Price impact following aggressive fills means we are probably trading too fast and pushing prices and so will underperform IS. And so on.

### Conclusion

TCA shouldn't be something to do once a quarter and file away. Done right it can tell you not just how well you or your broker are performing but also why and how to improve.

With recent press and headlines about toxic dark pools, high frequency traders and flash crashes, there is understandable concern that these things are affecting the overall performance of buy-side funds. TCA can be the means to prove or disprove to what extent any of this might be happening, what impact it has and even suggest where to look for improvements. ■