

Transaction cost analysis in fragmented markets

Can TCA remain relevant in an environment where the performance of the trades across multiple venues is often not comparable?

*Robert Kay**

Trading is the implementation of investment ideas. The costs of implementation for institutional asset managers can be significant and fall into three categories. First is the bid-offer spread between the price at which the security can be bought and sold. Second is the impact on the underlying price brought about by trading; pushing the price up on buying and down on selling. Finally there is the cost that arises from movements in price that occur between the time of the investment decision and the final completion of the trade. The contributions of each of these factors to the total cost of implementation are not independent of each other making analysis complex.

Transaction cost analysis (TCA) supports asset managers by measuring each of these different costs in respect of their own transactions and also anonymously

comparing their results with those of other managers completing similar orders.

Something new

Into this already subtle process now comes the growing fragmentation of liquidity in many European markets. Following the early success of Chi-X, recent months have seen the launch of Turquoise, Nasdaq OMX, NYFIX Millennium and BATS Trading Europe, as well as announcement of development and launch of many new MTFs, including some to be created by existing stock exchanges in direct competition with each other as well as their 'upstart' rivals. Table 1 highlights the liquidity fragmentation already visible in a number of major listed equities. With more MTFs due to launch in the near future, and the potential for existing exchanges to launch MTFs of their own, the

*

Robert Kay
managing director

GSCS Information
Services

Trade analysis

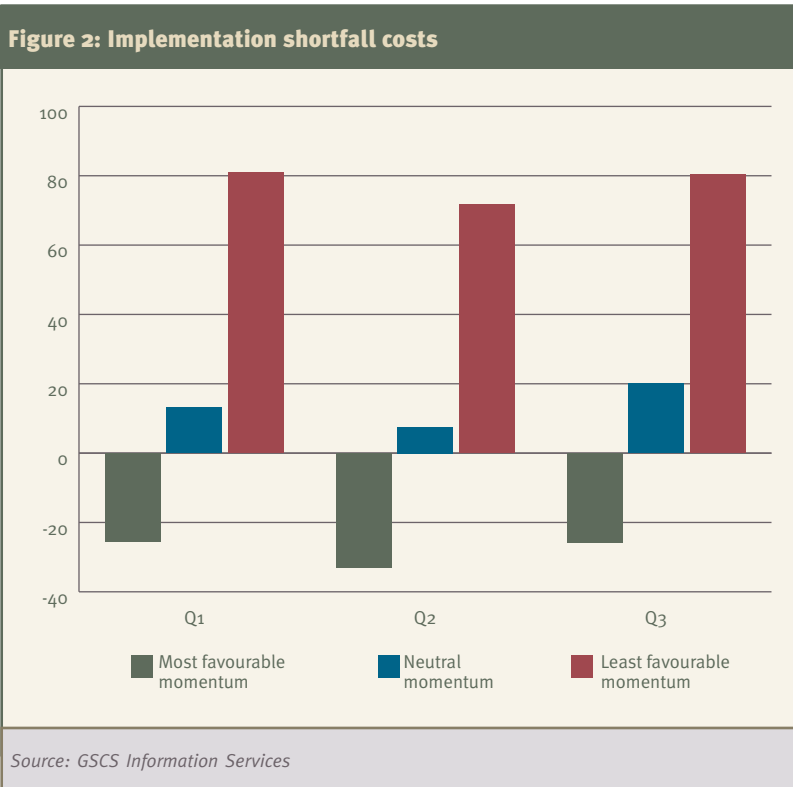
Figure 1: Fragmentation of liquidity for leading European stocks

Security	Trading August 25-29	Period September 22-26
	% Fragmentation	% Fragmentation
BP	25.67	24.43
Royal Dutch Shell	19.58	28.76
HSBC	16.44	13.61
Total	17.85	13.91
BNP Paribas	19.12	11.58
AXA	18.74	11.17
Allianz	11.46	11.19
Deutsche	19.54	12.39
Siemens	13.66	10.59
Credit Suisse	7.60	4.18
UBS	6.41	4.23
Nestle	6.67	8.01
ING Groep	18.79	14.30
Fortis	15.55	11.41
Unilever	20.10	19.65

Source: GSCS Information Services

fragmentation is expected to get greater in the near future. Fragmentation of liquidity has two effects on TCA, one real and important and one of much less consequence. The real effect is that fragmented liquidity does indeed alter the way in which trades actually get executed. This is not simply a question of use of one venue versus another. In many cases trades may be completed using more than one

execution venue. The second less consequential impact is how the performance, whether of broker counterparties, buy-side firms or indeed traders themselves, should be measured. The real risk in terms of TCA is that the real changes lead individuals, the companies that employ them and consultants, to alter the latter process of performance measurement in ways that are not simply expensive but actually wrongheaded.



Any review of how TCA might evolve or recommendations concerning how it should evolve need to understand the basic components that are applied within TCA, how these have responded in the past, and where the drivers of evolution have come from. In this connection it is important to recognise that there are two significantly different approaches towards TCA, both of which have firm adherents and strong

opponents. Interestingly, these two approaches parallel different approaches to measurement of investment performance. They may usefully be alluded to as ‘relative’ and ‘absolute’.

Implementation does not always cost

Using implementation shortfall as a benchmark against which to assess execution performance is the TCA equivalent akin to the ‘absolute

Trade analysis

return' approach of alternative asset managers. This methodology compares the price at which the trade was completed against the price at a specific point in time in the execution process (most usually when the trade is 'released' by the portfolio manager to the buy-side trader, or when the trader in turn 'releases' the order to the counterparty broker or market). Buy-side traders and the brokers that work on their behalf seek to minimise this implementation by optimal balancing of the cost of delay with market impact cost.

The problem with this approach is that success in beating the benchmark is heavily influenced by the nature of the portfolio for which the trades are being executed. If the portfolio is 'momentum' driven then by definition it is going to be buying securities whose value is on an upward trend. Given the likely high cost of delay, this makes it virtually impossible to 'beat' the benchmark. Similarly, if the portfolio is value driven, it will be common to find it executing trades that provide liquidity to the market (selling stocks into a rising market) and in these cases 'beating' the benchmark requires little skill.

Figure 2 illustrates this problem by showing the breakdown of cost of delay results based on the underlying momentum of the trade

in question. The results come from the GSCS universe of data covering millions of transactions involving more than 100 managers and 200 hundred brokers. As the bar graph shows, when executing trades with the most favourable momentum, the average trade is executed at a price that is more than 20bps better than the price prevailing when the trade was released. Equally, where the momentum is most unfavourable (i.e. the manager is 'chasing' the market) the likely cost will be 80bps or more. Clearly the scale of this factor dominates any modest improvement in performance that can be achieved by a buy-side trader or the sell-side counterparty through superior judgment, trading skill or lower commission cost.

The 'handicap' system

In an effort to deal with this problem, brokers and vendors have tried to establish what is in effect a 'handicapping' system. There are two main ways in which this 'handicap' is established. One is to look in a fairly generic way at the type of trade (size, liquidity and market situation) and see what similar types of transactions have cost in the past. The second is to use the pre-trade cost estimates available from brokers.

Irrespective of the handicap used, TCA performance using

implementation shortfall methodology is usually measured after adjusting for the handicap. Therefore the calculation of the handicap arguably becomes the most important factor in assessing performance. In fact it turns a theoretical ‘absolute’ performance back into a ‘relative’ performance assessment.

Fragmentation – handicap or excuse

The further problem that arises from liquidity fragmentation is that it will make it even harder to measure two key elements. First, in assessing the price at a point in time, it will be possible to include or exclude some or all execution venues. This will make a small difference to the actual result but will lead to extensive discussion and debate about whether the ‘correct’ figure has been used to calculate the benchmark

The second, much more important effect will be on the calculation of the appropriate ‘handicap’ that should be applied. Existing generic handicap systems suffer from the heterogeneous nature of the transactions that are included in the different averages. The question for the future will be whether the execution venue on which trades were completed should also be a factor in the creation of the relevant ‘handicap’

groups. It is already clear that some venues attract a particular type of client and business and hence trade. Is the performance of these trades genuinely comparable with those completed on other venues?

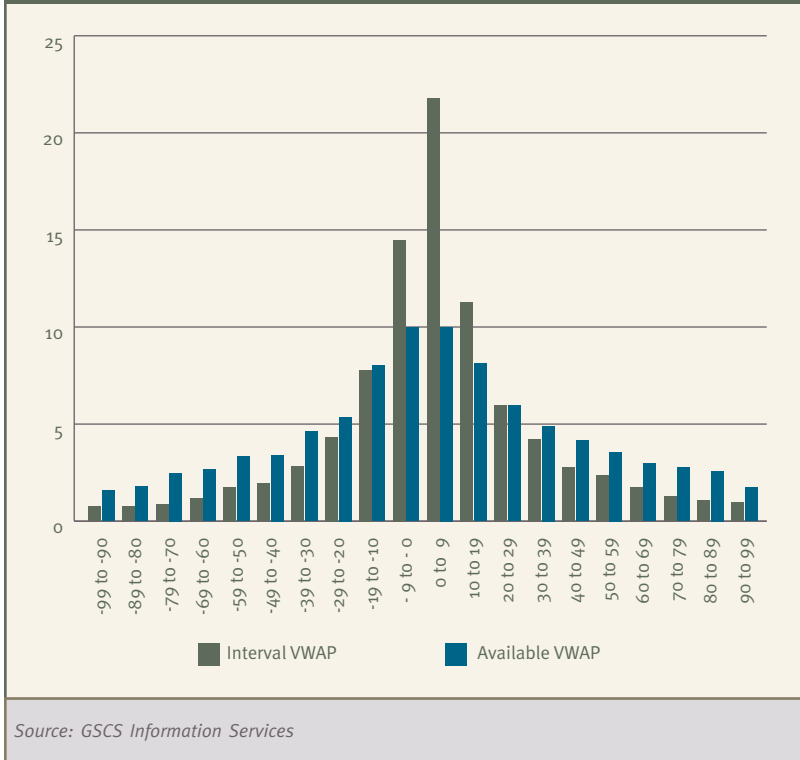
Similarly pre-trade cost estimates, are likely to vary even more in the future as different dealers take more or less account of venues to which they are connected and weight more prominently results from the venues they use most. This will make it more difficult for clients to determine the most appropriate pre-trade cost estimates. Generating consistency in an overall approach given this additional complexity will be virtually impossible. Without consistency however, TCA becomes less and less useful as a means for determining whether traders or counterparties have performed well. The number of ‘excuses’ for apparently ‘poor’ performance will grow in line with the growth of complexity in the process. Actionable analysis will become even harder to achieve.

Is everything relative?

The equivalent ‘relative’ approach in terms of TCA is the use of various volume weighted average price (VWAP) benchmarks. Usually these include ‘interval VWAP’ and ‘available VWAP’.

Trade analysis

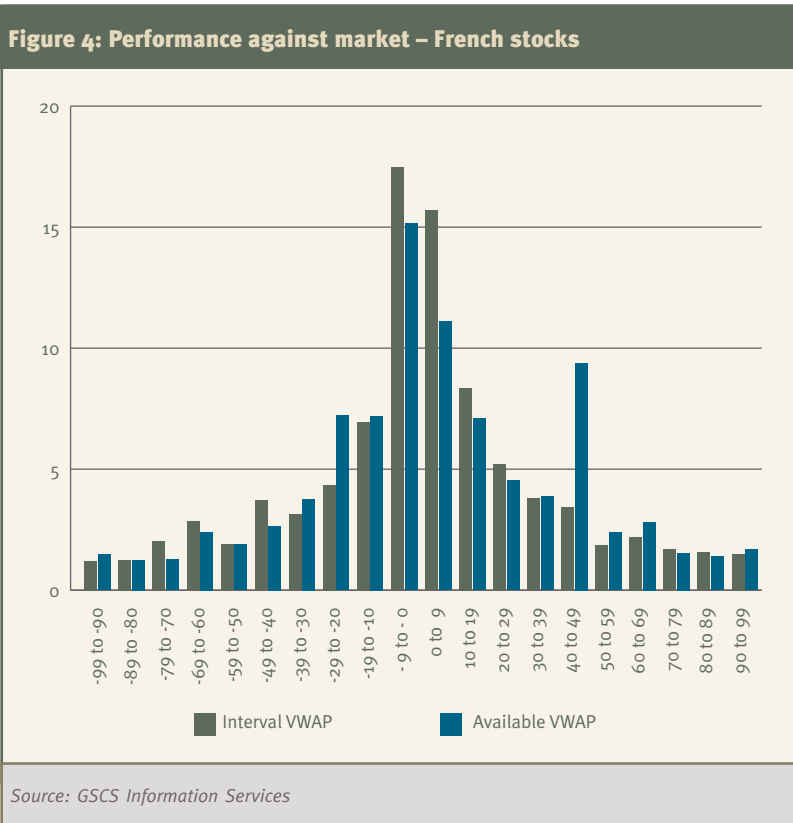
Figure 3: Performance against market – UK stocks



In some cases, where only limited time-stamp data is available, the ‘daily VWAP’ may be used as a benchmark. This methodology is similar to traditional asset managers assessing their investment performance against the return achieved by a single market, regional or global index.

The problem with this approach parallels that of the relative approach in fund management; it

is possible to perform well against the benchmark while watching the price move dramatically against the interests of the portfolio. So delay in execution of a momentum driven order would yield a poor execution price but be seen as providing adequate or in some cases good performance against the relevant VWAP measures. To many traders, and portfolio managers, these problems have

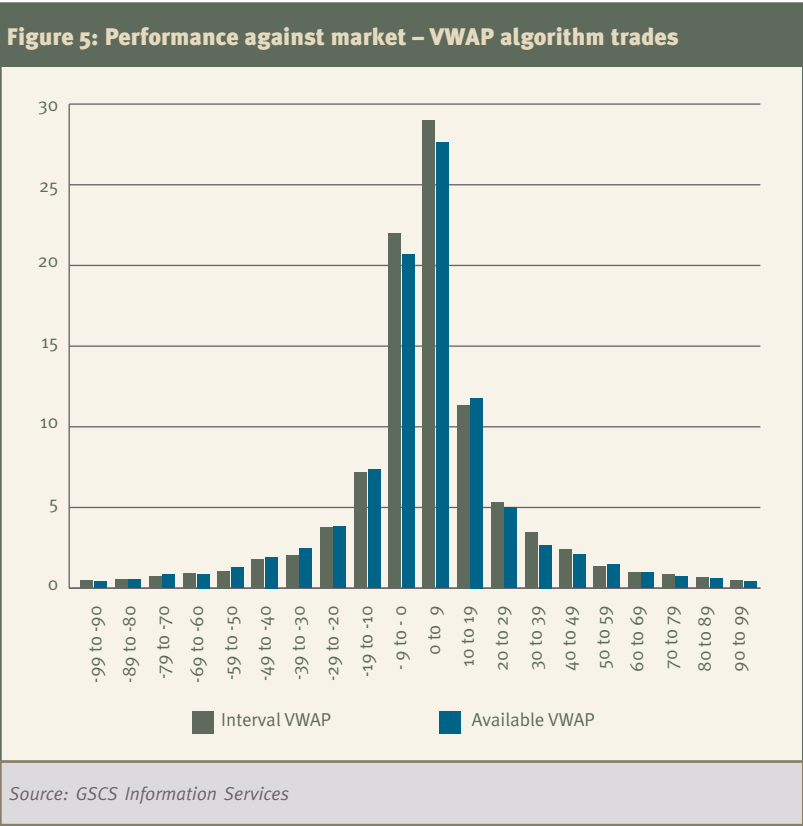


rendered ‘relative’ measures unattractive as part of an execution process.

Unloved but still widely used
 Nonetheless, VWAP algorithms remain easily the most popular with traders and other ‘in-line’ execution strategies continue to be pursued by many buy-side traders. While not followed as an execution strategy, many

traders, portfolio managers, compliance officers and not a few clients, like to confirm that their trades were done ‘as well as the market’. This means that VWAP measures, often in conjunction with implementation shortfall, remain a factor in the TCA process. Figures 3 and 4 highlight performance within the GSCS universe of trades with different markets measured against interval

Trade analysis



and available VWAP. These show trades ‘normally distributed’ around the mean performance, which is similar in each case. However, the standard deviation of the performance against the available VWAP measure is slightly higher.

Figure 5 shows the results of VWAP algorithms against the same benchmarks. This data is taken from a study conducted by

GSCS on behalf of TradingScreen Inc., a major vendor of multi-asset class execution management systems to hedge funds and other asset managers around the world. This also shows mean performance comparable with the measures of ‘human trading’ in the GSCS universe, but with less standard deviation. Such comparisons will no doubt lead to a desire for further refinement

Figure 6: Primary exchange vs alternative exchanges – VWAP comparison

Security	Exchange better for purchases	Average bp Improvement	Alternate better for purchases	Average bp improvement	No difference
BP	69.57%	5.81	30.43%	15.83	4.35%
HSBC	38.10%	4.43	42.86%	11.98	19.04%
Royal Dutch Shell	41.67%	5.01	49.82%	9.17	8.51%

Source: GSCS Information Services

in measures to allow more ‘excuses’ for apparently indifferent or relatively poor human performance. Fragmentation will once again be a useful tool in further complicating an already obfuscated picture.

Figure 6 shows some results of VWAP calculations for different securities on a number of execution venues, most notably Chi-X and Turquoise and the stock exchanges on which the primary listing exists. Given differences between venues, fragmentation means that clients may want to limit the VWAP calculation to those exchanges in which they are able to participate, or alternatively to some other subset of the total level of activity.

Once again, however, fragmentation will mean that performance results can be influenced more by the selection of the benchmark than the actual contribution to the results achieved by the traders

and counterparties. The more legitimate alternative measures that are available, the harder it is to ascertain what component of the supposed outperformance or underperformance is down to the measurement criteria and what is down to the skills and expertise of the individuals and institutions that are being evaluated.

Given that much TCA is made available by broker-dealers in support of their desire to demonstrate superior execution, this additional complexity runs the risk of making the analysis, whatever its true motives, appear to be simply a self-serving sales and marketing tool.

Where now?

Irrespective of the underlying methodology used by traders and brokers using TCA and the precise performance benchmarks that are used, there are two clear conclusions to be drawn in terms of the impact of fragmentation.

Trade analysis

First is that fragmentation could add to the complexity of the traditional processes that are used by all parties to ‘adjust’ performance to take account of transaction characteristics of the executions being evaluated. Second is the fact that this complexity is likely to mean more time is spent debating the merits of the process than trying to determine actions that could be taken to actually improve performance. These two conclusions, taken together, risk further undermining the viability and credibility of TCA as a business.

Fortunately an elegant solution is available. The question that needs to be considered is why has fragmentation occurred and what is its purpose? The European regulators who created the framework (under MiFID) for competition intended that national stock exchanges should no longer enjoy a monopoly in trading so that costs would be reduced and performance improved. Similarly, those institutions that have launched and/or supported the new venues have done so in the expectation of the venues winning business by offering a superior product, possibly at lower prices than incumbent exchanges. Finally, those institutions that are using the alternative venues are

presumably doing so because they are better, cheaper or both. In all cases however, there is nothing that requires participation in new venues unless they are superior. Participation is voluntary and therefore the only reason to participate in the new MTFs is because they deliver superior execution and/or lower costs.

The right answer

Given this rationale for participating in new venues, the right way to benchmark performance must be to compare the execution result achieved using a benchmark or benchmarks calculated based on the prices of trades done on the indigenous primary exchange. If new MTFs offer lower cost execution, then the performance of those that use them will show up favourably against these ‘traditional’ benchmarks.

Such an approach has three obvious advantages over any other. First, it is cheaper since it does not involve any material change to existing processes and procedures. Second, it will show to all parties the extent to which they themselves have improved their trading by using alternative venues. To the extent that many fund managers or brokers are satisfied their execution performance is enhanced,

whatever methodology they choose to measure it, they will determine the long term success or failure of the new MTFs. Finally, it reflects the reality of the fact that ‘price discovery’ still takes place on the primary exchange.

Making improvements not making excuses

To the extent that performance appears better or worse using alternative venues compared with the option of using incumbent exchanges, then further more detailed analysis will definitely be required. This approach combines a surveillance role across all trades (“are we performing to a generally acceptable standard?”) with a more detailed multi-

venue analysis of transactions that appear exceptionally good or bad. The analysis will help ascertain whether the performance differential is the result of the characteristics of the transactions themselves, the skills of the traders and technology they use, or the underlying market microstructure, including costs. Such analysis has the potential to identify and quantify the benefits accruing to clients from fragmentation in ways that allow replication across more transactions and hence a progressive improvement in performance. This combination of surveillance and continuous enhancement of execution quality is exactly what TCA should be seeking to offer its clients. ■

GSCS Information Services

applying science to the art of transaction management