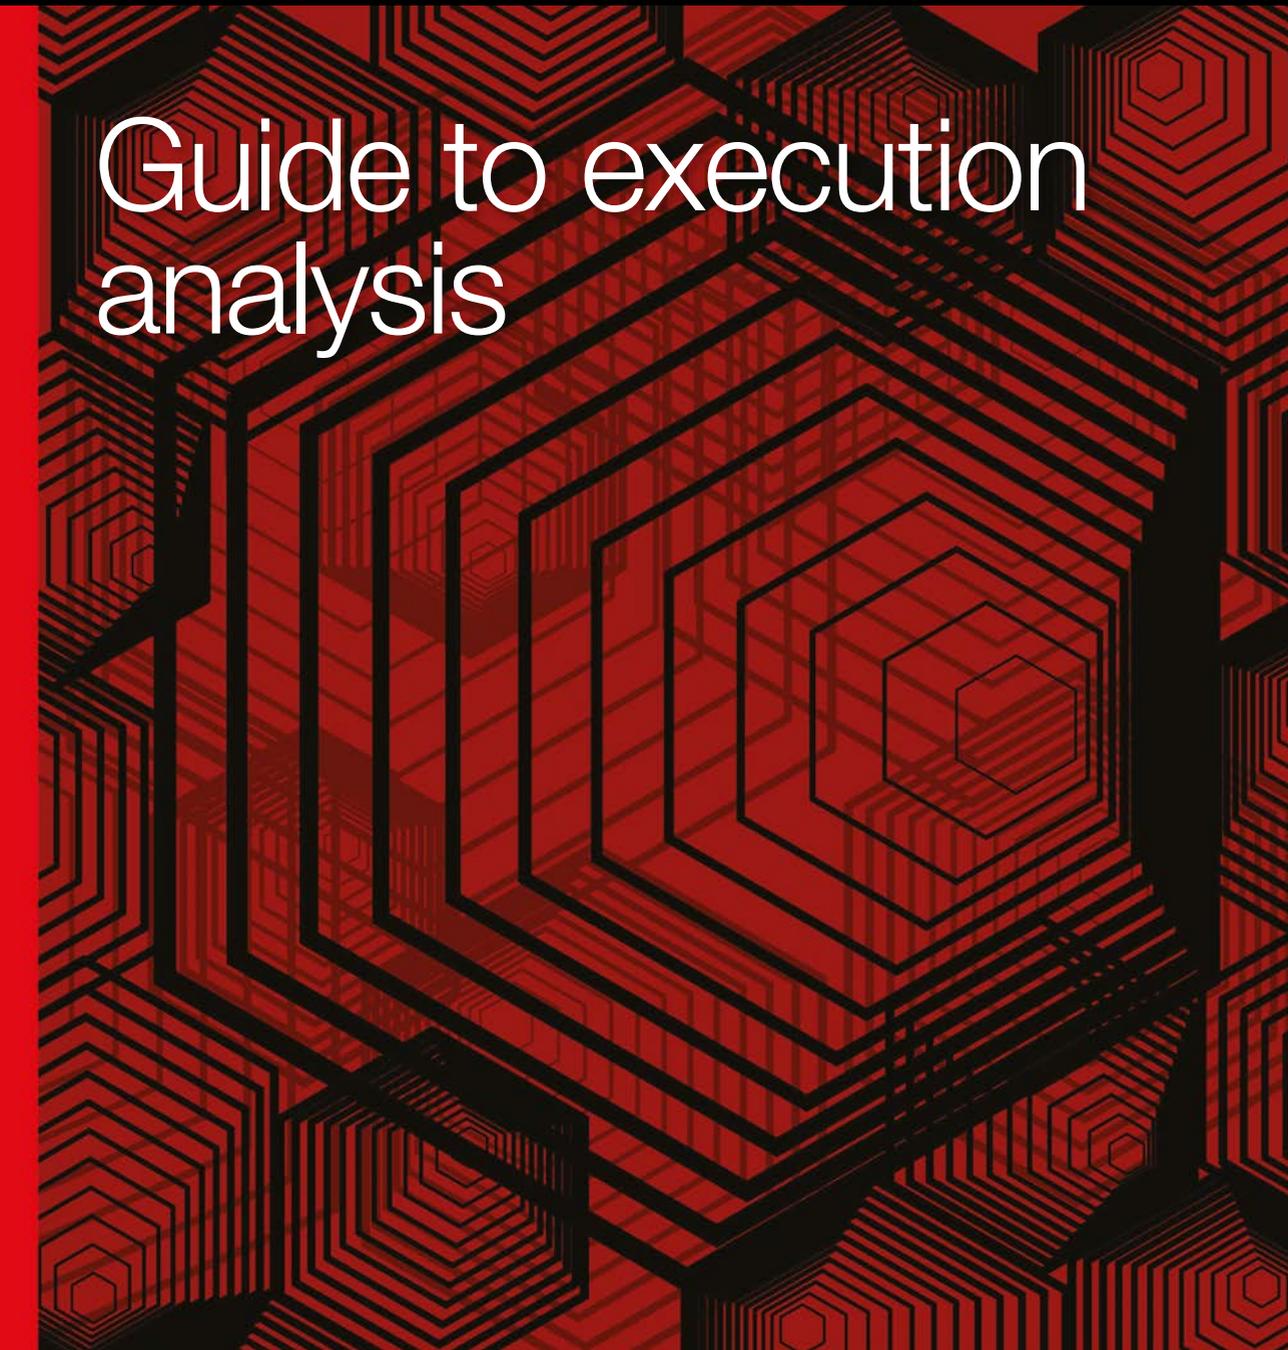


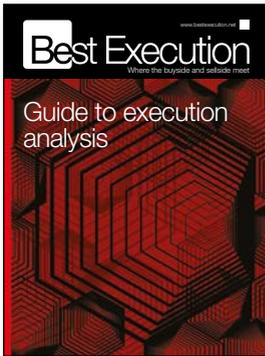
www.bestexecution.net

Best Execution

Where the buy-side and sell-side meet

Guide to execution analysis

The background of the lower half of the cover is a complex, abstract geometric pattern. It consists of numerous overlapping, concentric hexagonal shapes and lines in shades of red and black, creating a sense of depth and movement. The pattern is dense and fills the entire lower section of the cover.



Guide to Execution Analysis 2020

Managing editor: Lynn Strongin Dodds

Publisher: Ian Rycott

Commercial director: Scott Galvin

Photography:

Chris Mikami, www.mikami.co.uk

Design & production:

Siobhan Brownlow, RSB design

Best Execution

Markets Media Europe Ltd
Registered office: Suite D,
The Business Centre, Faringdon Avenue,
Romford, Essex RM3 8EN, UK
Tel: +44 (0)20 8017 1734
Email: info@bestexecution.net

Best Execution (ISSN: 2042-108X) is published quarterly.

Printed by Buxton Press, UK.

Distributed by Spatial Global Ltd,
Castle Donington, UK.

Subscription price for 4 issues:

UK £135, Europe €200, ROW \$260.
Tel: +44 (0)20 8017 1734 for full details.

© Best Execution 2020

All information and forecasts contained in this publication have been checked to the best of the author's and publisher's ability, but they do not accept any liability or responsibility for any errors, omissions or loss arising from decisions based on them. All rights reserved. No part of this publication may be reproduced without prior permission from the publisher.

www.bestexecution.net

Introduction

Best execution is no longer a tick box exercise but post MiFID II has become an integral and dynamic component of execution. A recent study by Greenwich Associates – State of Transaction Cost Analysis 2019 – showed that almost all of institutional trading desks in Europe are employing TCA for equities (95%), however fixed income remains tricky according to a separate report by Liquidnet because accurate and reliable data needed for the decision-making process is still harder to unearth.

In this publication, Kevin O'Connor, Head of Workflow Technology and Analytics and Michael Sparkes, European Analytics Business Development at Virtu explain how TCA is being applied and further developed in a multi asset class context. Dr Mainak Sarkar, Head of European Execution Advisory, and James Baugh, Head of European Equities Market Structure at Citi, on the other hand, assess the challenges and importance of choosing the right benchmarks.

Meanwhile, Henry Yegerman, Global Head of Sales, and Chris Sparrow, Head of Research at ISS LiquidMetrix advises market participants to look at the big picture and the inner workings of how an order is executed.

Banca IMI's Umberto Menconi and Carlo Contino also underline the need to work in partnership to achieve execution quality. He points out that the building blocks needed to attain best execution depend not only upon qualitative and quantitative analysis of trading behaviour, but on the resources of sellside counterparties.

Lynn Strongin Dodds,
Managing editor

In these uncertain times and as the tide of regulation rises ever higher and wider, Best Execution offers an in-depth analysis into the major trends that are shaping the financial services industry as well as providing a more detailed insight into the technology driving new products and services being developed to meet these challenges. It's where the buy-side and sell-side meet.



“Launched initially as an equity-focused compliance-driven process, the breadth and depth of analysis has expanded to cover almost every major asset class, incorporating not just implicit and explicit costs but also factors such as liquidity profiling, algo analytics and venue analysis.”

Kevin O'Connor, Virtu

Multi-asset TCA: faster, broader, deeper

By Kevin O'Connor, Head of Workflow Technology and Analytics & Michael Sparkes, European Analytics Business Development at Virtu

The scope and application of transaction cost analysis (TCA) and its various close relations has evolved dramatically over recent years, driven and enabled by changes in technology, regulation, market structure and client demand. Launched initially as an equity-focused compliance-driven process, the breadth and depth of analysis has expanded to cover almost every major asset class, incorporating not just implicit and explicit costs but also factors such as liquidity profiling, algo analytics and venue analysis.

The expanded use of TCA across asset classes has introduced a variety of challenges due to different market structures, data availability, cost, and the relevance of differing metrics – particularly for less liquid instruments. In addition to asset class expansion, the intended audience has also evolved. Its use is not just limited to traders or compliance departments but typically constitutes a series of elements and actionable outputs that can be relevant throughout the investment process targeting portfolio managers, risk managers and CIOs.

A major driver for many of these developments has been a raft of new regulatory requirements, not the least of which were introduced by MiFID II in Europe. The regulation requiring a demonstrable process for monitoring best execution across asset classes, using data from actual results as an input to future decisions, has galvanised the buy-side community to reassess their approach. While it is

not mandatory under MiFID II to use an external TCA service, it is required that a systematic method of capturing and reviewing trade data be in place. Approximately 95% of firms reported using TCA across asset classes at the start of 2019, up from 75% in 2017, according to Greenwich Associates*.

In addition to these regulatory best execution requirements, the global trend towards commission unbundling has led to a greater focus on the quality of order execution to ensure trading and routing decisions are driven by objective and quantifiable results, separated from any need to pay for research or other bundled services which used to be funded by commissions.

Many firms had only just begun to address the new European regulations when we reviewed the state of preparedness of the buy-side in our 2015 best execution article, *MiFID II and Best Execution Across Asset Classes*. Our recent update of this survey captures how the industry has responded and indicates the likely path forward. While the median or typical firm has raised their game considerably, in terms of TCA data utilisation, leading firms have, if anything, pulled further ahead of the pack.

Top firms' in-house data science teams are working directly with external TCA vendors as they consume normalised benchmarked data and load it into their proprietary databases for further analysis. The TCA vendor provides guidance not



“Over time, the continued development and refinement of new technologies should deliver even more precise data, better transparency, oversight and control of trading events in multiple venues and their impact on all asset classes.”

Michael Sparkes, Virtu

only as these teams seek to adjust or redirect trading allocations according to which brokers and counter-parties are consistently outperforming but also as they undertake more forensic testing and experimentation with different strategies to minimise costs and preserve alpha in line with their investment objectives.

Equity analysis – more granular and holistic

In 2015 it was typical for firms to monitor allocation-level costs in equities based on data from an order management system (OMS), while only a minority of firms also analysed fill-level data. Fast forward to today and we see more firms monitoring algos and venues in detail alongside the increased level of control and choices available. It is now typical for analysis at both levels of granularity, with firms adding fill-level analysis – which often

incorporates data drawn from an additional source such as an execution management system (EMS).

Additionally, changes in European market structure have accelerated the evolution of TCA. Double volume caps on dark trading, the growth of systematic internalisers, the increasing role of electronic liquidity providers and the introduction of periodic auctions are all examples of changes which buy-side firms need to monitor and consider when deciding on optimal trading strategies.

The proliferation of EMSs has also led to the increased use of real-time analytics to monitor trades in-flight and make real-time adjustments in response to pre-trade TCA alerts and changing market conditions. EMS functionality continues to expand, providing greater execution strategy options, increased order control and more opportunity to analyse the post-event outcomes.

Table 1: Qualitative current state of MiFID II compliance for the buy-side in equities

MiFID II requirement	Median firm typically have	Leading-edge firms also typically have
Measure and minimise costs	<ul style="list-style-type: none"> • TCA/best execution system in place to calculate costs both at allocation level and fill level • Data to allow monitoring of algos and use of venues • Regular (typically monthly) review of data in best execution committee or similar process • Peer-based comparisons on quarterly basis 	<ul style="list-style-type: none"> • Documented process to apply conclusions from historic data going forward in trading process. • In-house trading database used to test and analyse strategies • Dynamic A/B testing of alternative strategies • Auto-routing based on analytical results • Use of algo wheel data as part of algo measurement and selection process
Assess execution factors	<ul style="list-style-type: none"> • Factors (price, cost, speed, etc) are aggregate metrics available in reports 	<ul style="list-style-type: none"> • Factors are incorporated into cost reporting and benchmarking, and used as an input in future decision making
Increased monitoring	<ul style="list-style-type: none"> • Daily monitoring of trade outliers • Monthly/quarterly monitoring of data including brokers, algos and venues • Trade surveillance process performed at least monthly • Documentation of outlier reviews • Visualisation tools to support monitoring/ review process • Incorporation of investment process factors such as instructions 	<ul style="list-style-type: none"> • Real-time monitoring of trade outliers • Trade surveillance process performed daily • Full audit trail and summarised log of outlier reviews including external information (news events, screen shots, etc)
Top 5 venue reporting	<ul style="list-style-type: none"> • RTS 28 reports showing top-5 venues used 	<ul style="list-style-type: none"> • Detailed analysis of venues used by each broker and by each algo to identify inefficiencies or biases by trade type or market conditions

Source: Virtu Financial

Many EMS providers are now applying the same optionality used in equities to other asset classes.

The trend for more frequent post-trade analysis remains strong. While many firms are beginning to leverage pre-trade and real-time TCA to support intra-day optimisations, most review their post-trade results daily for significant outliers in equity trading. This process usually starts with the trading desk and is then checked by the compliance function. There is still an important place for monthly and quarterly processes that look at larger volumes of trades to determine any recurring trends or biases.

The definition of outliers has evolved, firms frequently use multiple filters – for instance a basis point threshold coupled with a minimum value – to highlight true deviations in performance. Similarly, multi-level analysis may incorporate information relating to portfolio manager instructions to assess trading desk performance and trader instructions to accurately evaluate broker performance.

Another hot topic for leading firms is alpha profiling. Greenwich Associates reports that more than a third of those surveyed now conduct analysis of this type, linking the execution strategy and its outcomes back to the investment decision and portfolio construction process.

Previously used by leading-edge firms exclusively, algo wheel technology is becoming main stream among buy-side clients. Virtu's Algo Wheel, originally developed as a best execution order routing solution for equities, has now expanded into FX and futures asset classes. The use of both automated routing and algo wheels are increasingly being leveraged to remove much of the process noise and (often unintended) biases introduced by humans. The data, once normalised, can provide an objective and fair comparison for counterparties and strategies.

FX TCA – transparency

Foreign exchange market structure has made tremendous advances in the last five years, supporting the evolution in the analysis of the FX market. Many buy-side firms are bringing FX trading back in-house, and in some cases this has

led to increased algo usage and request for quote (RFQ) platform adoption. The use of WM Fix related trades has also come under intense scrutiny after allegations of manipulation by market participants of these time-specific price benchmarks. Lastly, improvements in technology now make capturing accurate timestamps achievable, providing the trader with more accurate data to perform better analysis. The combination of better data and new tools has significantly enhanced the quality and scope of analysis available in the fragmented global FX market.

The implementation of FX TCA by buy-side desks conducting their own trades, rather than outsourcing to the banks, has been widespread. Using data from an EMS and/or FX trading platform allows for meaningful data examination and guides the calibration of strategies accordingly. This includes the detailed review of algo behaviour and performance, as well as strategy options such as netting. Benchmarking for FX trading has also progressed to reflect the increased data reliability. The array of metrics now includes simple mid-price and bid/ask metrics as well as more advanced calculations such as size-adjusted spread and cost impact models.

The reliance on pre-trade tools, such as Virtu's FX ACE model, has become more widespread in FX with currency pairs exhibiting patterns related to available liquidity at different times of the day. This information can be used in decision-support prior to trading and in post-trade review to help maximise liquidity, minimise spreads and to assist in comparing different strategies.

Additionally, leading firms are now leveraging peer-based comparisons, allowing them to experiment beyond standard benchmarking. The use of peer data must be handled with care to ensure true apples-to-apples comparison, with standardisation and curation to ensure the data is clean and relevant. Meaningful peer-based comparisons also require sufficient breadth and depth of data to have statistical significance. But such data can be invaluable in helping identify areas which need further scrutiny in what is a very fragmented market.

Table 2: Qualitative current state of MiFID II compliance for the buy-side in FX

MiFID II requirement	Median firm typically have	Leading-edge firms also typically have
Measure and minimise costs	<ul style="list-style-type: none"> • Daily measurement of costs for spot, forward and swap transactions • Regular (typically monthly) review of data in best execution committee or similar process • Quarterly peer group comparisons 	<ul style="list-style-type: none"> • Documented process to apply conclusions from historic data going forward in trading process. • Regular detailed FX algo data analysis • Dynamic A/B testing of alternative strategies • Auto-routing in place based on analytical results
Assess execution factors	<ul style="list-style-type: none"> • FX flows tagged by investment objective with relevant benchmark applied – e.g. 4:00pm fixing trades • Analysis of market impact by venue, aggregator or liquidity provider • Factors (price, cost, speed, etc) are aggregate metrics available in reports • Comparison of different FX strategies – netting, RFS, RFQ, time of day, algos, etc • Separate analysis for active executions and custodial trades 	<ul style="list-style-type: none"> • Use of pre-trade models and decision support tools • Analysis of execution method – e.g. full amount, sweep, all-or-none • Separate analysis of market impact from rejected orders and filled orders • Analysis of 'last look' trades versus no-last-look • Analysis of global code of conduct signatories versus non-signatories • Analysis of FX trades linked to other trades (e.g. fixed income or equity basket trades)
Increased monitoring	<ul style="list-style-type: none"> • Daily monitoring of FX outliers • Monthly/quarterly monitoring of FX data including counter-parties and algos • Incorporation of FX process factors such as instructions • Losing quote analysis • FX trade surveillance process run at least monthly • Documentation of FX outlier reviews • Visualisation tools to support FX monitoring/ review process 	<ul style="list-style-type: none"> • Real time monitoring of FX outliers • FX trade surveillance process run daily • Full audit trail and summarised log of FX outlier reviews including external information (news, screen shots, etc)
Top 5 venue reporting	<ul style="list-style-type: none"> • RTS 28 reports showing top-5 counter-parties used 	<ul style="list-style-type: none"> • Detailed analysis of FX venues to identify inefficiencies or biases by trade type or market conditions

Source: Virtu Financial

Analysis of FX trading has been used by leading edge firms to look at the total cost of trades in other instruments. In other words, if switching from one asset class to another involves a currency exchange (for instance, from bonds to equities, or from European to US equities), it is now possible to factor in the costs related to implementing the currency transaction alongside the cost associated with the underlying trades themselves. Often the cost of delaying a currency purchase can outweigh the gains made by the related trades – assuming the currency had been traded instantly. It is highly likely that this kind of

multi-asset analysis will become mainstream in the years to come.

Fixed income TCA – time to play catch-up

Compared to equities or even to FX, fixed income is a late entrant to the world of analytics. Driven by many of the same factors as other asset classes, FI TCA is finally starting to catch up – mainly due to the regulatory requirements from MiFID II and the availability of reference data. The increased use of electronic platforms in fixed income assists firms in capturing data and assessing the outcomes against a variety of benchmarks and metrics. In some

Table 3: Qualitative current state of MiFID II compliance for the buy-side in fixed income (FI)

MiFID II requirement	Median firm typically have	Leading-edge firms also typically have
Measure and minimise costs	<ul style="list-style-type: none"> • Use of FI TCA reports provided by electronic platforms and/or third-party providers • Quarterly monitoring of FI data including counter-party performance • Automated comparison of market quotes to execution rates • Monthly/quarterly monitoring of outlier trades 	<ul style="list-style-type: none"> • Daily monitoring of outlier trades • Documented process to apply conclusions from historic data going forward in trading process • Use of pre-trade decision support tools
Assess execution factors	<ul style="list-style-type: none"> • None 	<ul style="list-style-type: none"> • Fairness of price compared to similar products • Incorporation of FI process factors such as instructions
Increased monitoring	<ul style="list-style-type: none"> • Monthly monitoring of outlier trades • Documentation of outlier reviews • Trade surveillance process run at least monthly • Visualisation tools to support monitoring/review process 	<ul style="list-style-type: none"> • Daily monitoring of outlier trades • Trade surveillance process run daily • Losing quote analysis • Full audit-trail and summarised log of outlier reviews including external information (news, screen shots, etc)
Top 5 venue reporting	<ul style="list-style-type: none"> • RTS 28 reports showing top 5 venues 	<ul style="list-style-type: none"> • Detailed analysis of FI venues to identify inefficiencies or biases by trade type or market conditions

Source: Virtu Financial

cases, the analysis is primarily for regulatory and compliance purposes, but as in other asset classes, leading firms are putting a lot of effort and using sophisticated data to enhance their performance.

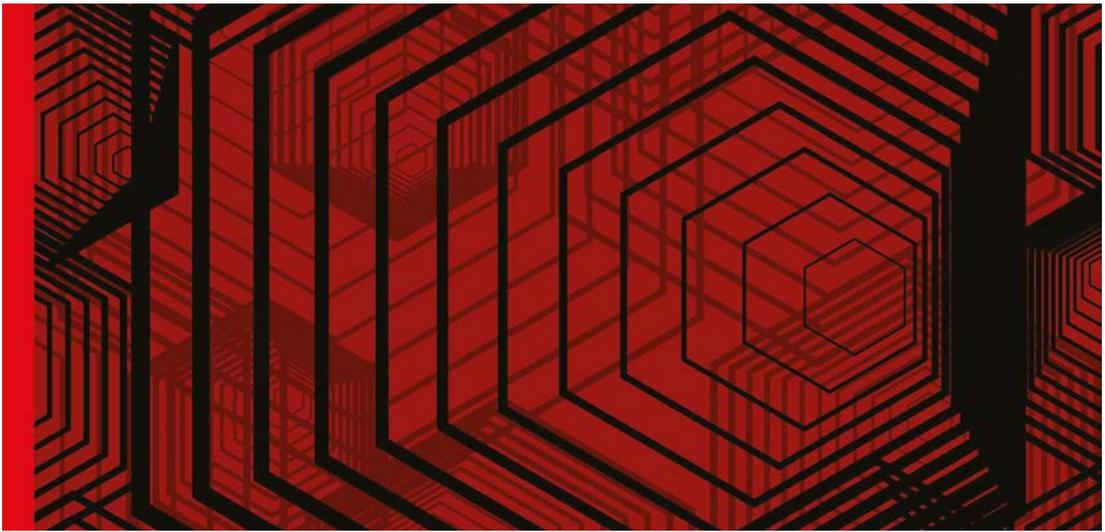
However, not all executions have moved onto such systems with a considerable volume of trading still being conducted by voice, especially in less liquid instruments. The fragmented nature of the market, as it relates to trading venues and the number and complexity of instruments, has been a challenge which has slowed the transition.

The fixed income market incorporates a variety of asset types – ranging from the more liquid government and corporate bonds to the less frequently traded instruments such as Munis, mortgage-backed, CDS, IRS and other related categories. Each class of instrument has its own unique set of trading characteristics. These require the appropriate application of relevant market data and suitable metrics designed to provide actionable analysis.

As with all trading analysis, a key ingredient is good market data against which the trade is

compared. While reliable sources are available for liquid instruments, it is considerably harder to find meaningful reference data for the more esoteric instruments. In many cases it requires the aggregation and cleaning of data from multiple sources in order to produce a realistic and fair comparison. As with all market data, the cost to obtain data has increased, in some cases very steeply. There is renewed talk from ESMA of a mandatory consolidated tape being introduced for analysing bond trading, although this is unlikely to occur quickly. In Europe there has been an ongoing discussion of a consolidated tape for equities since MiFID in 2007 and the regulators are still in the process of finding a way to implement it.

The multiplicity of venues and trading methods means that fixed income data must be sourced wherever it is available – indicative and firm quotes, historical prices, evaluated prices, dealer-to-client platforms, multi-dealer platforms and so on. One hope has been the Approved Publication Arrangement (APA) platforms, although so far these have proved less than ideal in terms of consistency



and accuracy. To incorporate and standardise a range of data sources into a coherent and consistent repository requires significant scale and is not something any individual buy-side firm can easily do. Hence the need for an external analytics provider who can step in as the aggregator, cleanser and curator of vast amounts of data from these various inputs.

Interestingly this aggregation process is also an important ingredient as part of the workflow that an EMS provides. Just as with equities and FX, it is likely that greater data aggregation and decision support will become a vital ingredient in the EMS with the subsequent execution data feeding back into the post trade analysis. We expect to see leading firms implement deeper integration of pre and post-trade analytics for fixed income as trading technologies develop further.

Derivatives TCA – a mixed bag

Derivatives TCA is hard to summarise in a simple table. The most liquid instruments, such as listed index futures, have characteristics similar to global equities. OTC instruments are much harder to measure, and the availability of data and level of sophisticated analysis is akin to what is found in illiquid fixed income. As median firms and leading-edge firms expand their analysis and monitoring

into derivatives, they will require the same tools available for other asset classes – including cost models and peer group analysis. Although data is available, the analysis and monitoring performed for derivatives by most firms lag other asset classes with the same liquidity characteristics.

Conclusion – how far we have come

The world of analytics continues to move forward rapidly, driven by regulatory priorities and spurred forward by leading firms and their quest to add alpha. Over time, the continued development and refinement of new technologies should deliver even more precise data, better transparency, oversight and control of trading events in multiple venues and their impact on all asset classes.

Given the increased complexity of market structure and the ever-changing patterns of trading, the key to improving execution remains a moving target – one that will continue to evolve, in some ways radically, for the foreseeable future. Firms will need to remain nimble and apply lessons learned from post-trade analysis to ensure they adapt and thrive in the changing landscape. ■

*Sources: "MIFID II Fuels Investor Demand for TCA" and "The State of Transaction Cost Analysis – 2019", Greenwich Associates; available at <https://www.greenwich.com/node/109726> and <https://www.greenwich.com/market-structure-technology/state-of-transaction-cost-analysis-2019>

The evolution of transaction cost analysis

By Dr Mainak Sarkar, Head of European Execution Advisory, and James Baugh, Head of European Equities Market Structure, Citi.

Introduction

The post trade workflow for most institutions on the buy side typically consists of the following steps: record all trade data and enrich it with the relevant market data in order to compute a set of diverse benchmarks, the most popular being Implementation Shortfall (IS), volume weighted average price (VWAP) and participation-weighted-price (PWP). The performance of orders measured against these benchmarks is then analysed based on stock and order level characteristics (such as the bid-offer spread, size of the order as percentage of daily volume, daily volatility of the stock and so on). This is the attribution phase. Finally the performance data is usually analysed by the broker to measure any systematic differences in execution, all of this is as mandated under current regulations.

Limitations

We start by noting the challenges associated with each step of the post trade process:

- **Data quality** issues and latency in data capture during the measurement phase. We observe considerable diversity in data infrastructure and investment amongst buy-side firms that connect to Citi as a broker. Most institutions do not maintain their own market data as it involves significant cost of setup and maintenance and usually resort to buying the same from third party providers. This can lead to an added layer of complexity and challenge. We do encounter frequent complaints by clients with regards to the quality of the data received from vendors, which can lead to inconsistent estimation of performance.
- Selecting the right **benchmarks** is critical and this can vary by the type of firm and trading objective. For instance a quant fund with short term alpha compared to passive investors have both different trading objectives and horizons and therefore utilise different performance benchmarks. A narrow focus on a particular benchmark can potentially lead to the algorithm trying to achieve that goal possibly at the exclusion of other important goals, thereby hampering best execution as broadly defined under MiFID II. For instance, slippage against arrival may be reduced by trading intensively (small orders only), however, this can lead to high levels of impact and price reversion, which is not desirable. Conversely, too many objectives can often cloud the findings regarding performance, therefore making any changes to trading difficult. As a broker we have observed significant variation in how clients approach this and various clients are grappling with this on an ongoing basis as there is no single size fits all solution to this problem.
- **Attribution** can often be challenging- for instance, if the stock moved adversely during trading, decomposing this move into market impact caused by the trading of the algorithm with the rest being natural alpha in the stock is usually challenging. Various approaches using index and sector level beta adjustment exist, however they are not entirely satisfactory in their outcome. Similarly, small sample sizes can lead to the analysis frequently being restricted (or dominated by notional weighting) to the



large capitalisation / more liquid names with significant order flow, however, the performance across brokers in the less liquid category may be more relevant but difficult to measure.

- **Evaluation** of brokers may be similarly biased due to the skewed asymmetric distribution of order flow, particularly when they are manually traded, where the allocation decision may be driven by perception or muscle memory. The algorithmic wheel setup seeks to ameliorate these problems by randomising the flow of orders across brokers. Citi's experience in this regard has been varied as we are well represented on multiple wheels across a host of clients. Wheels are frequently time consuming to setup and calibrate. The process of collecting the right data across brokers and adjust it for order differences to achieve a fair comparison can be quite challenging.
- **Statistics / machine learning** – working with large datasets (larger clients can trade thousands of orders in a day) usually requires key decisions to be taken regarding how to deal with outliers (identify, remove, winsorise, etc), how to price the unfilled portion of any orders (opportunity cost), whether to penalise under participation by brokers under adverse conditions (scaling wrongly), etc. mostly in an automated fashion – all of which require sophisticated analysis on a large enough sample of orders by a specialised team to draw meaningful conclusions. Citi has observed a significant uptick in the hiring of such professionals in recent times on the buy side. There is a trend towards using more advanced statistical analysis and machine learning techniques to analyse performance. This is an industry trend which has been strengthened by regulatory requirements.

Citi recognised the complexity of the above mentioned process of post trade analytics and created an Execution Advisory team, whose role is to engage and consult with clients to help improve their benchmark performances. There has been significant engagement across a range of clients who have benefitted from this more consultative approach.



“A significant development in the post MiFID world has been the increasing automation of the trader workflow and has seen the growth in the use of algorithmic wheels. The wheel engine randomly identifies and allocates a set of ‘easy to trade’ orders across a list of brokers.”

Dr Mainak Sarkar, Citi

Challenges for the broker

A significant development in the post MiFID world has been the increasing automation of the trader workflow and has seen the growth in the use of *algorithmic wheels*. The wheel engine randomly identifies and allocates a set of ‘easy to trade’ orders across a list of brokers. The identification of ‘easy to trade’ can be either done via a pre-trade model (estimated impact cost function) or by a set of rules such as percentage of daily volume and/or notional limits (all orders less than \$1 million are routed to the wheel for instance). At set frequencies – usually a month or a quarter – the performance across brokers is analysed and the wheel is reweighted such that the broker performing better receives more orders from



“Citi is enhancing its current historic data and analytics based approach by incorporating a forward looking simulation methodology, through the provision of ‘scorecards’. Citi believes that this hybrid method leads to better insight into performance measurement.”

James Baugh, Citi

the client over the next period. Note that this can lead to an imbalanced sample in the next period and therefore require statistically difficult comparisons across brokers, as some brokers have large samples and others significantly smaller ones. Having too many brokers on the wheel can potentially lead to orders being thinly spread and therefore sample sizes being too small. The optimal number of brokers is therefore client and flow specific. Finally, for fair comparison ideally a post trade normalisation of order performance across brokers is desirable. For instance, if Broker A received predominantly certain orders on news or event days when volumes and volatility are usually high; this needs to be factored into the analysis.

Brokers are continuously in the process of upgrading their algorithms based on internal

analysis as well as feedback from clients, however such redesigns are often time consuming and expensive in terms of resources. Therefore, fundamental changes in the trading platform occur only infrequently and usually most innovations are incremental in nature. Furthermore, short term changes to the strategies are unlikely to drive large fluctuations in execution outcomes for clients, given that significant care is usually taken by the broker to roll out only changes which are well tested and are expected to improve the outcomes on average. The volatility observed in rankings can often be greater than the infrequent changes to algorithms which leads one to suspect that the evaluation mechanism on the buy side can be improved by filtering out the noise in the data leading to more stable comparisons, from a broker perspective.

For medium to small size firms infrastructure is often restricted to consuming TCA provided by the broker themselves. Even to a completely impartial observer, it can be seen – due to significant differences in definitions, data capture and storage issues, and measurement errors – that performance measurement can vary widely across brokers making comparison problematic for the buy side. The subjective element of measurement is significant when comparing TCA across brokers.

In our experience even for those larger firms that are able to capture their own data, the numbers invariably do not match the brokers’ numbers and a significant amount of time and effort is subsequently required to reach some level of agreement and consensus, which is essential to ensure that the data is sufficiently clean before drawing any conclusions .

Whilst we have seen an increased interest in independent providers of TCA, which can apply a more uniform and consistent set of measures (benchmarks) against the data, they are not immune to these same challenges. Moreover, they lack the understanding and the nuances of broker specific execution strategies. Applying a one size fits all approach to data across brokers loses out on the broker specific differences that are key to achieving best execution. Without the right context as to why the algorithm behaved in a particular

way, the filtering of orders is based on rules which may not give informative outcomes.

Innovations at Citi

Citi is developing the use of advanced machine learning techniques to drive a more proactive approach to the TCA process. A number of these techniques, including the use of peer group analysis and clustering, are discussed in this article.

I. Peer analysis

A new feature selection approach using random forest models has been implemented to highlight how different variables contribute to the overall performance of flow from a particular underlying client. This tool is also helping to identify areas of underperformance via potential trader biases under different market conditions when the price may be moving in and out of favour, such as systematic modifications, limit price selections, cancellations, etc.

Actionable TCA of this nature is key in delivering client specific recommendations, which in turn is leading to better overall performance.

Machine learning based peer analysis has also been developed at Citi. Clusters are compared across different clients with similar objectives thereby identifying any systematic differences in order flow. Using clusters identified for each client,

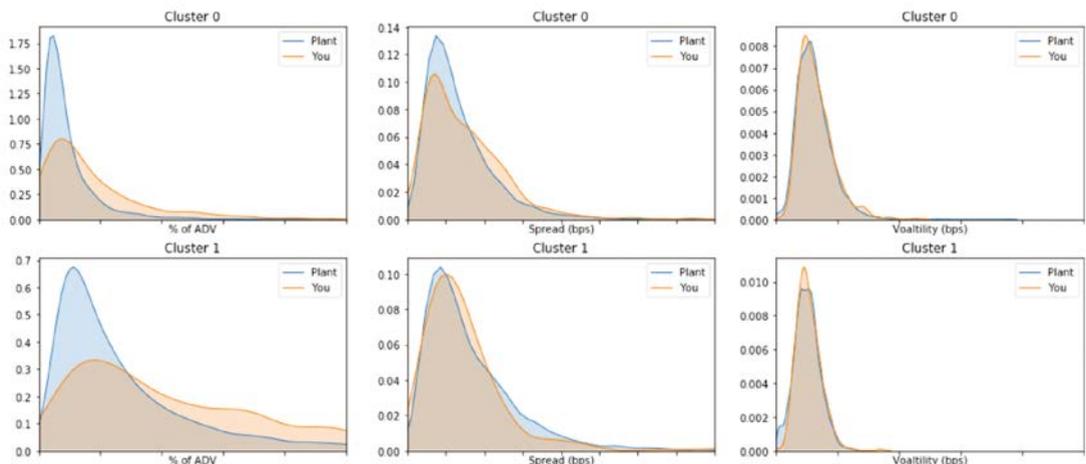
we can make recommendations on customisations and on strategies to optimise client performance. Exhibit 1 illustrates this approach.

II. Stock categorisation

Citi is also using machine learning methods to identify different clusters of orders. We use various stock level characteristics. The most standard being bid-offer spread, volatility and queue sizes. Other more microstructure level features are used to endogenously categorise each stock in the universe to a particular category. For instance one category may be high volatility, small queue size stocks. There are multiple applications of this approach both at the pre-trade and also post-trade phase.

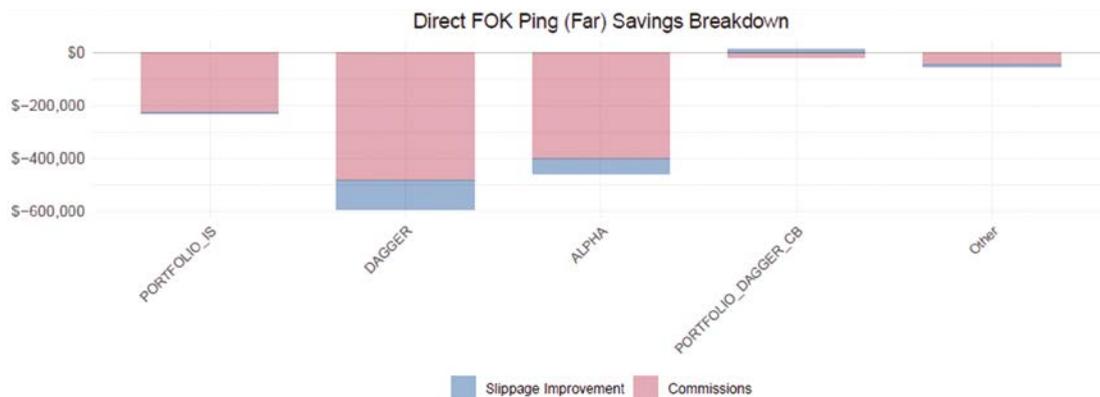
This model can be used pre-trade to determine the right strategy. For instance where a client is using VWAP algorithms for low volatility but wide spread stocks. The EAS team actively monitors the performance goals and can use these tools to come up with custom recommendations for different clients. In this case one recommendation may be to change the passive layering of the book. In essence security and order level clustering methods allow us to draw conclusions at the client level, which ultimately leads to better execution performance.

Exhibit 1



Source: Citi

Exhibit 2: Saving breakdown by strategy



Source: Citi

III. Forward looking

Citi is enhancing its current historic data and analytics based approach by incorporating a forward looking simulation methodology, through the provision of 'scorecards'. Citi believes that this hybrid method leads to better insight into performance measurement. Each scorecard highlights alternative strategies for liquidity provisioning. This innovative approach marries the client's objective with market conditions and available internal liquidity to come up with theoretical improvements in execution versus the client's benchmark.

The scorecard can be used by high or low touch trading desks and or central risk to adjust for commissions and make client specific recommendations depending on what the optimal strategy is for them to execute their flow.

Recommendations can be tailor made for clients with different flows. Exhibit 2 illustrates a sample scorecard approach, which shows cost savings / slippage versus commission rates for liquidity sourcing.

Optimising for TCA performance has been a key focus of the research and development here at Citi. Extensive study and back-testing has revealed that the price and size alone should not be the only consideration when interacting with different liquidity sources. Citi has employed machine-learning algorithms to monitor intra-trade

market activity and real time order performance to drive its algorithm behaviour. Real-time venue ranking also plays a crucial part in sourcing not only the best price and size, but also takes into account the quality of the liquidity. In regards to interacting with Systematic Internalisers, Citi's comprehensive real time venue analytics measures different characteristics such as short term toxicity, uniqueness of liquidity, reversion, good and bad fills, fill rate and quote and fill toxicity. For example, we not only look at the child-level fill performance to measure the quality of our fills, but prior to the child trade we establish the toxicity of the quote from each of the systematic internalisers.

This information is fed to the Smart Order Router to determine its routing decisions. Furthermore, the routing decision for taking and posting liquidity is also regulated for different strategy types based on the urgency level – a very aggressive strategy may still interact with toxic liquidity source to satisfy liquidity demand. These enhancements help us to tailor our execution strategies and improve performance.

IV. Case study

TCA driven enhancements to algorithms

Citi measures performance across a wide range of clients and uses the findings in designing better products that help clients achieve optimal outcomes – an example is given below. Citi is in

the process of rolling out significant enhancements to Dagger, its flagship liquidity-seeking algorithm (LSA) for equities. Earlier TCA reports gave a starting point of where the algorithm could be calibrated more effectively. This was combined with more research by the quantitative team to build a new algorithm which performs better against a wide spectrum of market conditions and stocks. Performance was evaluated using a standard A/B race on Citi principal order flow from Feb 2019 to May 2019. Trades numbered over ten thousand, with a value of over \$1bn. Results showed significant performance improvements. Some of the key benefits were:

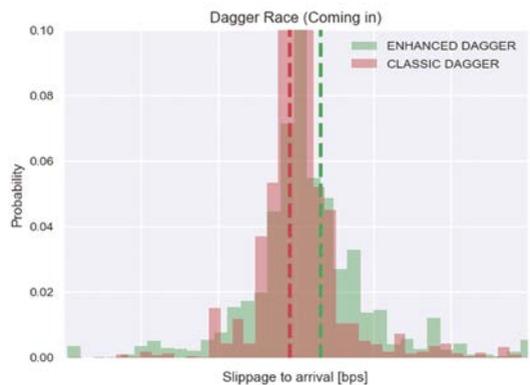
- Spread capture improved by 54%, with similar average fill rates;
- Spread crossing reduced by 42%;
- Similar average order duration and fill rates to the classic Dagger.

Key performance drivers were:

- Move away from heuristic based, reactive logic to a continuous optimal trajectory model based on machine learning analytics;
- Improved strategy sensitivity to possible adverse selection scenarios;
- Use real time analytics in addition to historical to rank venues;
- Machine learning algorithm for dynamically choosing passive allocation policy based on real time performance.

Performance was analysed on different aspects/slices of the data in terms of order as percentage of average daily volume (ADV), volatility, spread, market move and participation rates. The most significant improvement was observed in the category where the prices 'come in'. Exhibit 3 shows the distribution of the slippage for this category; note the significant improvement in performance. We also note that the performance remains consistent when the prices are moving away. Furthermore, the improved Dagger is more capable of sourcing liquidity passively and at mid. This is an example of where machine-learning components in the algorithm have helped to provide a better performance to the end client.

Exhibit 3



Source: Citi

Conclusion

In this note we have outlined some of the shortcomings of the post trade analytics process (of which TCA is an integral part), that in our experience clients have encountered frequently. Some of which can be fixed via more and better investments in infrastructure (better data quality), however others are more intractable (normalisation of orders by broker for instance) and require continuing research and analysis by our buyside colleagues. The Execution Advisory team here at Citi is a resource to assist with this process in every step of the way.

A number of innovations undertaken in this space have been outlined above. Our application of advanced machine learning techniques to TCA and using real-time analytics to minimise the implicit costs of trading is proving invaluable in providing the performance metrics our clients need to improve performance, meet best execution requirements and measure and rank broker performance.

Finally we gave an illustration of product design based on inputs obtained from TCA, which led to significant improvement for the client and better execution outcomes through the use of advanced machine learning techniques, calibrated using the findings from TCA. ■

For further information, please contact: emea.eas@citi.com



Do you know how your orders are routed?

By Henry Yegerman, Global Head of Sales, and Chris Sparrow, Head of Research, ISS LiquidMetrix

Do you know how your order is routed? Should you care? Lately, more and more emphasis is being placed on how orders are routed. Typically concerns are raised around whether brokers are routing in their own best interest or in their clients' best interest. Other considerations relate to whether certain trading venues are 'toxic' and lead to deterioration of performance.

Let's begin with a bit of background. In order to realise an investment opportunity, traders must first acquire a position. They can do this by buying or selling shares of stock in the 'market'. While this sounds quite simple and maybe obvious, the term 'market' hides a lot of complexity. The reason is

that the 'market' consists of many places to trade. While many years ago, there was typically only a single place to trade (often called an exchange), we now have many places to trade, some with names like MTFs, ATSS or ECNs. This means that we need to make choices about which of these venues to interact with when we go to purchase shares of stock. We also need to consider the ways in which we interact with the collection of trading venues we refer to as the 'market'.

While there are many possible questions to address, three key questions we would like to answer regarding our order are: When did we execute? Where did we execute? and How did we execute?

“Ultimately what we want to be able to do is to see the big picture of how our order interacted with the market. It therefore seems natural to visualise this interaction in a picture. We can develop such a picture by plotting the way we interacted along three dimensions.”

The portfolio manager generally knows why they want to execute, but the other three questions all relate to how the order was managed by a broker. The brokers use technologies such as algos that determine timing, and routers that determine where to direct the order and the way to execute.

One of the key things being managed in this approach is market impact. When we go to the market, we don't want to move the market too much as that can only hurt our performance. At the same time, the role of the market is to provide price discovery – the process of finding equilibrium prices that balance supply and demand. Therefore, we want to develop tools that can allow us to determine the answers to the questions of where, when and how we should trade so we can understand why we got the outcome we did.

Ultimately what we want to be able to do is to see the big picture of how our order interacted with the market. It therefore seems natural to visualise this interaction in a picture. We can develop such a picture by plotting the way we interacted along three dimensions.

The first dimension we care about is time – we want to answer the question ‘when did we trade?’ We want to be able to show how much volume we traded as a function of time. The second dimension we care about is the trading venue – we want to answer the question ‘where did we trade?’ The third dimension we care about is the way we interacted with the market – did we cross the spread and remove liquidity? did we post our order passively in the book providing liquidity? or did we execute in a dark pool? – we want to answer the question ‘how did we trade?’

Unfortunately, pictures are two dimensional, so how can we show these three dimensions in a single picture? Well, fortunately, there are colour pictures! So maybe we can answer all three questions by developing a colour image. We can show a time dimension along the horizontal dimension of our picture, and a venue dimension along the vertical dimension of our picture and then the colour can indicate how we traded.

We use the property that any colour is a mixture of red, green and blue. So, if we use red to represent the ‘active’ volume from crossing the spread, blue to represent the ‘passive’ volume and green to represent dark volume, then we can combine three ways of interacting with a venue into a colour which is itself a mixture of active, passive and dark volume. The intensity of the colour is proportional to the amount of volume while the actual colour is determined by the distribution of active/passive/dark. When the pixel is red, it means the volume in that bucket was all active. If the pixel is blue, the volume was all passive, while if green, then the volume was all dark. Other colours are mixtures of these three ‘primary’ colours.

This has real-life applications in evaluating the order routing of different brokers and can be used as an input into making Algo wheel strategies. Analysing execution data in this way allows us to clearly see the differences in order routing between brokers both under similar and different trading conditions. How differently do different algos trade? How similar are different brokers’ algos? How does my trading strategy influence my realised market impact?

Consider the differences between the routing in Figures 1 and 2. These 2 images both show a single order in AMD done withing 6 weeks of each other. The x-axis represents the time of day, and in this example shows 30-minute time buckets between the open and close of the market (9:30 – 16:00). Each row shows a different trading venue. The colour shows the way the order was traded by combing active, passive and dark fraction of volume as an RGB triplet, while the brightness represents the fraction of the overall order done at the given time on the given venue.

Fig 1: Broker A, symbol = AMD, date = 2019-05-03

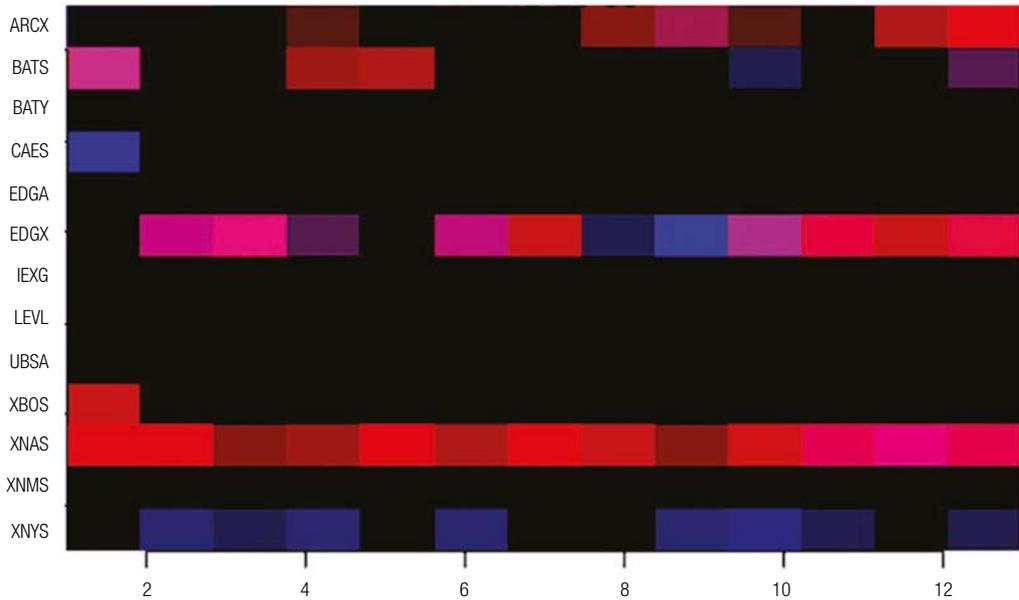
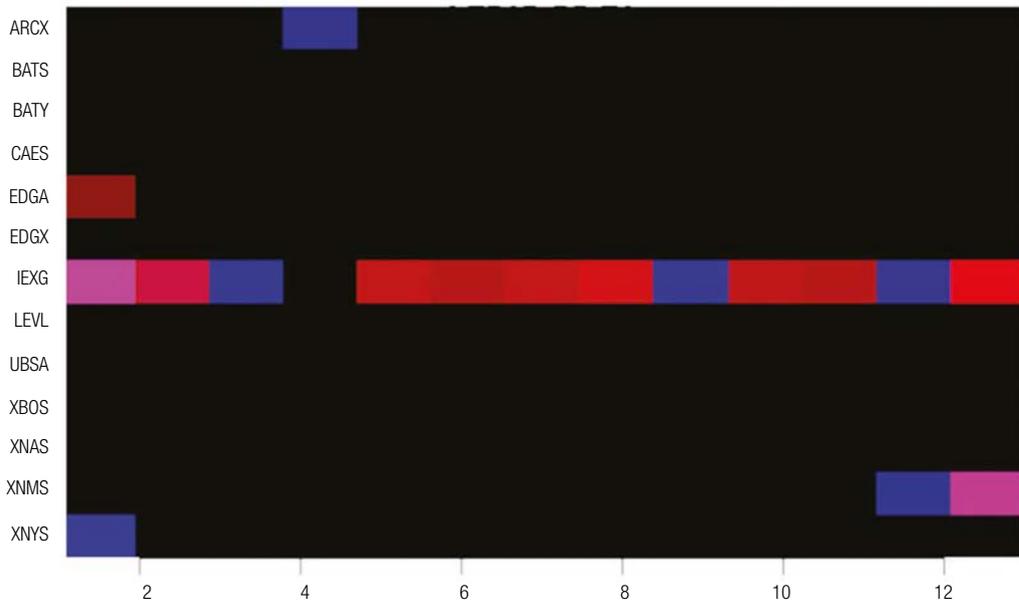


Fig 2: Broker B, symbol = AMD, date = 2019-06-21



We can see that Broker A mixes both aggressive and passive executions across multiple venues while Broker B largely executes on a single venue throughout most of the trading day.

This technique can easily be used across an

aggregate of thousands of orders to identify the broker's order routing tendencies.

In Figure 3 (low volatility), Broker B relies heavily on a single trading venue and is largely passive in execution style. However, in a high volatility



Fig 3: Broker B – low volatility

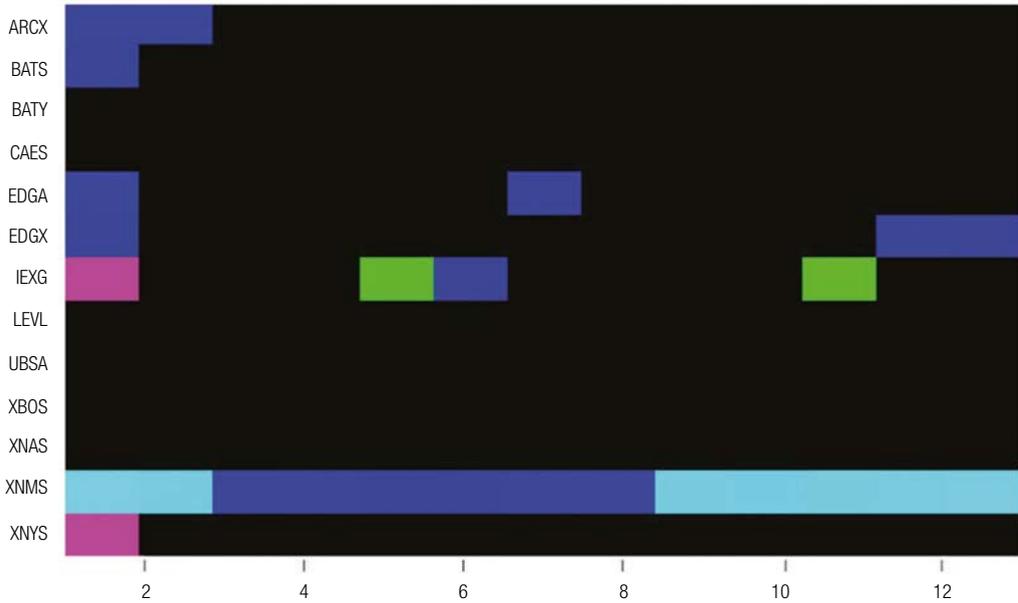
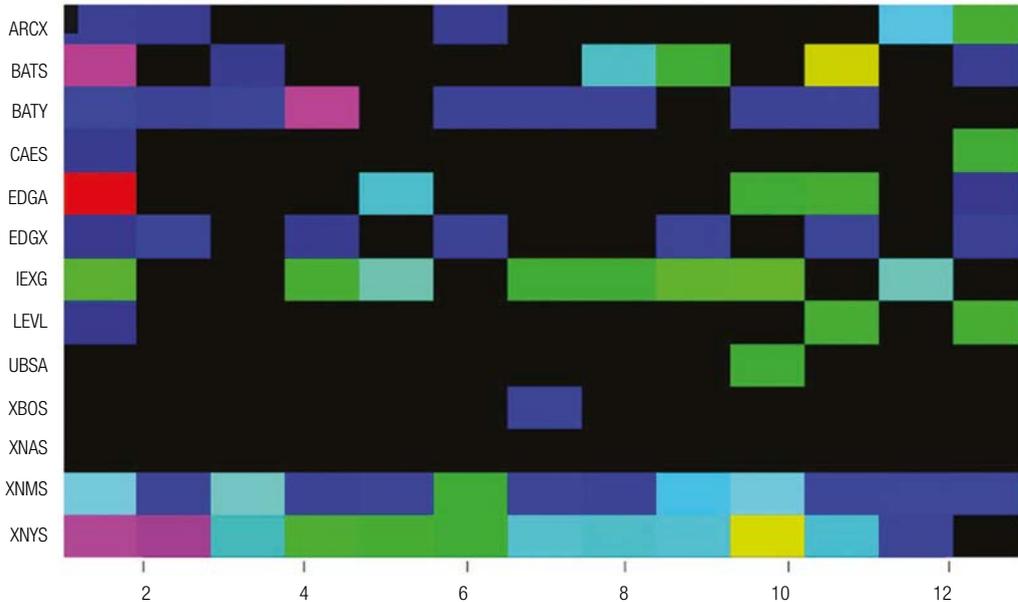


Fig 4: Broker B – high volatility



environment, order routing switches to using many venues and becomes somewhat more aggressive.

The use of color allows us to clearly see differences in venue selection, timing and the degree of active versus passive trading. More

importantly, the colors, brightness and size of each box can be converted into numerical equivalents which allow us to quantify the differences for best execution and specific applications such as Algo Wheel selection. ■

Fixed income – a best execution Legoland

The building blocks needed to achieve best execution depend not only upon qualitative and quantitative analysis of trading behaviour, but on the resources of sellside counterparties. Banca IMI's Umberto Menconi and Carlo Contino* discuss the importance of partnership in assuring execution quality.

What level of feedback should buy-side firms expect from their counterparties on execution quality?

Technological innovation, digital transformation and new regulatory obligations following the global financial crisis are dramatically reshaping the financial markets landscape, affecting every asset-class. Relationship dynamics, the level of engagement and the sellside and buy-side roles are shifting from the traditional bilateral workflow to a new paradigm. The buy-side is not only taking more ownership of its own trades but also rethinking its role in market transformation. It has also increased its appetite for sellside collaboration and investment in trading and risk management tools.

Changing dynamics have forced the sellside to evolve and adopt flexible models to satisfy the buy-side's liquidity demand requirements. They have to invest huge resources in technology innovation, in order to generate greater efficiencies as well as acquiring and sharing information and analytics in a more complex and fragmented market environment.

The biggest challenges remain efficient multi-market connectivity and real market data across the trade lifecycle: order creation, order placement, trade execution, TCA analysis, clearing, settlement and custody. The execution quality of a trading strategy, employed by buy-side execution desks to capture alpha or to limit portfolio risk, strongly

depends today on communication networks, collaboration tools, and continuous feedback. This is not only the case with buy-side execution desks, brokers and trading venues, but also among other financial ecosystem participants, such as custodian banks, clearing houses and central depositories. This is because best execution is more than just the transaction price. It also needs to embrace the entire trade workflow with a much broader horizon.

How does the execution protocol impact that?

Leverage limits, stricter capital requirements and balance sheet restrictions on proprietary trading are causing many sell-sides to rebalance towards agency broker models. They are helping clients to execute trades on trading venues rather than directly assuming portfolio risk. This market evolution had been pursued by Banca IMI from the beginning, with the launch of the Market Hub platform back in 2008. The transformation from a bilateral, voice-execution model towards a new communication and trading model, based on a network of trading platforms, electronic chats, systems connectivity, and sophisticated EMS/OMS tools necessitated a change in trading culture. This did not only include human resource hiring, but the adoption of different communication channels, language formats, trading protocols and a higher level of collaboration.



Umberto Menconi

To better execute orders, there is an increasing array of venues to use. This ranges from the conventional (MTFs, RMs, OTFs and ISSs) to the unconventional (dark pools) sources of liquidity, as well as different trading protocols (All-to-All, click-to-trade, CLOB, Auctions, Direct connectivity, RFMs, Portfolio Trade, anonymous RFQs-to-all). However, despite the growth, the traditional electronic RFQ (request for quote) model is still central.

We are still a long way from real market structure changes, and the new market ecosystem integration remains a 'work-in-progress'. However,

there are several drivers for a stronger fixed income trading ecosystem including standardised and integrated market connectivity solutions and communication protocols to reduce complexity. There are also shared technological infrastructures among all participants as well as different trading relationship levels, with respect to the recasting of players' roles, and putting the different protocols in place. Moreover, there has been broader implementation of machine learning and artificial intelligence solutions for the data 'tsunami' created after the introduction of MiFID II. Last but not least



Carlo Contino

is the increasing demand for readable smart data and value analytics.

When supporting high- and low-touch executions, what should dealers be cognizant of in supporting their client's performance?

In fixed income it is possible for the buy-side to fully automate some of their execution, thanks to more data availability, axes and price streaming activity. Buy-side execution desks see this as an opportunity as well as a game changer. A high-touch and low-touch execution strategy is increasing automation

in execution of small size tickets, while leaving more time for traders to concentrate on larger tickets on a bilateral bespoke basis as well as the design of new investment ideas.

To perform high-touch and low-touch execution, the buy-side needs to modify its best execution rules accordingly, based on both a quantitative and qualitative approach. It has to reflect their perception of what is low-touch or high-touch as well as ensure trading decisions are addressed to the right counterparties. It also has to make the best use of low-touch efficiency in small size

execution and of high-touch risk trades. Firms have to measure execution performance through new metrics tools, even though we are still in the early stages for TCA in fixed income. TCA is still controversial due to a lack of reliable transparent data compared to some other asset classes.

To support buy-side execution performance, the sell-side should constantly monitor their execution quality and offer customised and efficient trading solutions for high-touch and low-touch business, market connectivity, readable market data and comprehensive reporting services.

What sort of feedback and interaction can better support an understanding of dealer/client objectives?

On one hand, the sell-side has been spending plenty of time focusing on data consumption and, more importantly, wondering how they can transform that data to ensure the buy-side can use it and that it adds more value for both sides. On the other hand, the buy-side has focused more resources on workflow automation and collaboration enhancement with their sell-side partners. What is clear today is the need for collaboration with the goal of establishing a relationship based on value.

In this new and complex environment, a pivotal role can be played by fintech vendors when bridging the gap between buy-side and sell-side, and better supporting interaction. The key moving forward for both buy- and sell-side firms is to have a more open and connected networking infrastructure. The end objective is to create a best-of-breed solution that brings the buy-side and sell-side together.

What are the challenges in benchmarking performance?

Since market data can address the adverse effects of a more fragmented trading environment, demand for real market data has increased. However, today, the cost of acquiring real market data from the various data sources is expensive and there is also a lack of standards. This results in a significant hidden cost to consume and manage data, and represents one of the main challenges

in fixed income performance benchmarking. It reduces the potential of using it as a measure of trading performance, improving execution and applying more complex benchmarking. The equity market can be used as a reference, but fixed income execution needs more caution due to its unique characteristics.

How does a bank then optimise performance based on feedback from multiple clients?

The style in which the buy-side and sell-side interact has changed significantly in the last few months. Investment banks can no longer offer everything to everyone. The buzzwords of today are, network compatibility, partnership, interoperability and execution model flexibility. Therefore, Banca IMI has set up a client service desk, which is fully dedicated to improving the client experience – across asset classes, as well as the full operative process – that is in close proximity to the sales and trading desks.

This service strives to better serve and increase the specialisation of our services for our clients. Customer experience offers a key opportunity to build investor loyalty and to create a competitive advantage for both the buy-side and sell-side. To achieve an edge in the current financial market, the sell-side needs to move forward with an increased reliance on innovative technology in order to deliver returns for the buy-side, reduce market complexity and real market data consumption cost. Firms also need to be able to integrate their trading systems and data tools into the client workflow.

By moving from trading-centric towards a customer-centric, unified value proposition, fintech companies are bridging the value chains of various industries to create an 'ecosystem' that should reduce customers' costs, be more user friendly and provide them with new experiences. In this collaborative environment, a data-centric business model and client service feedback are gaining ground as the real new assets and value-adds for the sell-side/buy-side relationship. ■

**Umberto Menconi is head of Digital Markets Structures, and Carlo Contino is an analyst in Digital Markets Structures, at Market Hub, Banca IMI, part of the Intesa Sanpaolo Group.*