

FXALGO NEWS

TOP STORIES

BNP Paribas re-engineers FX algos to deal with information leakage

BNP Paribas has significantly changed how its FX algo strategies are posted to the market in a bid to combat the growing issue of information leakage. In order to tackle this issue, the bank has moved away from using lit markets to a preference for dark orders, only using a very distinct policy when participating in external markets, increasing its participation on mid venues and investing heavily in expanding its internalisation capabilities, notably with the launch of a new BNP Internal eXchange (BIX) Peg mechanism. Asif Razaq, Global Head of FX Automated Client Execution, BNP Paribas, explains

that the move follows a growing issue with the liquidity landscape that algo orders are being exposed to. "There is now far more market data being consumed by a wider group of market participants than it ever has been in the past. The unfortunate downside this has created for FX algos is that when a lit order is posted to the market, this is instantly published in a market data update. This then has a knock-on effect on the execution performance of an algo. The new BIX Peg tooling helps further fill the space created by our shift away from placing lit orders in the market, meaning that internalisation has

now become a key differentiator in a top tier algo product." See page 4.



Asif Razaq

Goldman Sachs launches basket algos on MDPs

Goldman Sachs has rolled out its basket algo functionality to leading multi-dealer platforms (MDPs), enabling efficient algo execution of correlated pairs across FX spot, forwards and NDFs for a wide range of algo users. According to Dr Ralf Donner, Head of Marquee Execution



Dr Ralf Donner

Solutions at Goldman Sachs, the launch of basket algos on MDPs is a significant change for algo users seeking to trade correlated pairs efficiently, including systematic hedge funds and real money client segments, coupled with the workflow efficiencies offered by the MDP which allows users to batch trades directly from the EMS screen without additional steps. The move also coincides with the roll out of EFP functionality on currency futures, which merges futures and FX into a single product and is available via a leading MDP. Donner adds: "We are encountering some very interesting use cases and the basket algos on MDPs is definitely turning into a fascinating business opportunity. The market was obviously more than ready for this product." See page 3.

IN THIS ISSUE

p1: TOP STORIES

The latest industry stories

p6: REGULATORY ISSUES

What you need to know about the AI Act

p10: PROVIDER INSIGHTS

Bloomberg FXGO in the FX algo space

p12: INDUSTRY VIEWS

Customisation of FX execution algos

p19: EXPERT OPINION

Increasing algo volumes in Scandies

p20: MARKET WATCH

White Label and Algos as a Service

p26: ASK A PROVIDER

From information to knowledge

p30: TRADERS WORKSHOP

Lessons from low-impact FX execution

p32: INDUSTRY REPORT

Algo benefits from EMS consolidation

CitiVELOCITY

DYNAMIC ALGOS

Advanced market-making technology.
Deep liquidity. Ultimate control.



STILL INNOVATING

Contact your Digital FX salesperson to learn more

Goldman Sachs streamlines algo workflows with new MDP initiatives

As the first bank to offer portfolio algo execution in FX, Goldman Sachs continues to stand out in the market with its unique offering for clients seeking to execute a basket algo. Dr Ralf Donner, Head of Marquee Execution Solutions, shares why client uptake has been so strong since the offering was made available on multi-dealer platforms and looks at some of the unique use cases that demonstrate the benefits of the bank's unique approach.

The basket algo functionality has been available for some time now, but recently made the leap from the single dealer platform to a native integration with a selection of the multi-dealer platforms (MDPs). According to Donner, the offering itself is a uniquely native implementation, firstly in terms of the ease of use, with clients being to take the legs of the execution directly from the EMS screen on the MDP, batch them together and send them to the basket and then there is no further interaction needed. "Unlike some other more cumbersome workflows, our native implementation works off the bat," he adds.

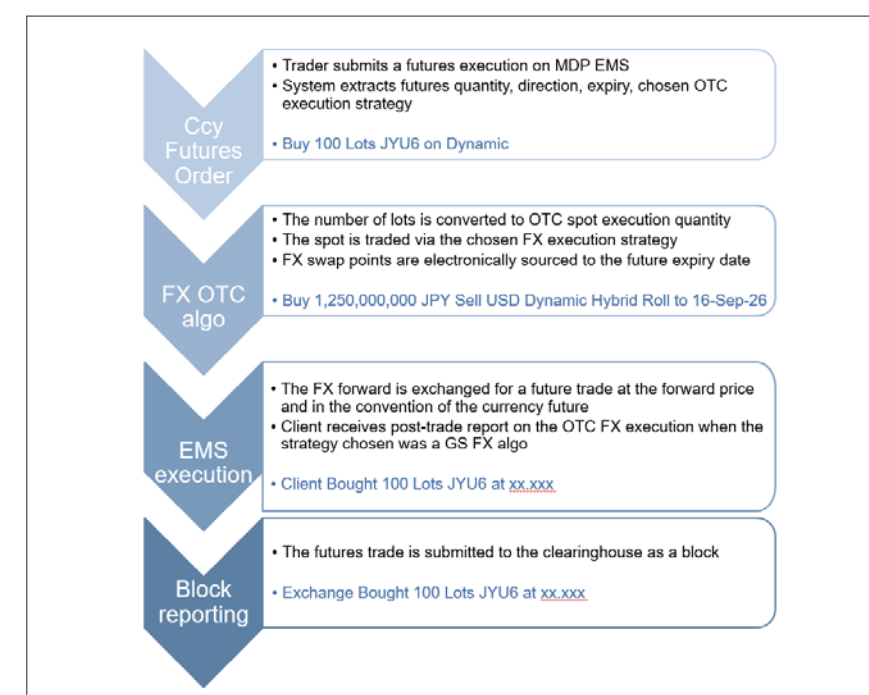
The reason why basket algo execution has been so successful also lies in the

ability for clients to trade correlated pairs, which is also proving to be the case on the MDPs, Ralf adds. "Several of the users, hedge funds and real money clients, are running a systematic strategy in spot FX and do not believe the orders should be independently executed, especially given that the strategy is on the basket. It makes more sense for them to send us the entire rebalance."

Other real money clients also have an issue of non-nettable accounts, where they must trade both legs of the order and simply cannot net the two, Donner adds. "However, our basket algo implementation ensures that if a buy order and a sell order are submitted to the basket algo, we will

simply net those for the user at no additional cost," he says. Further use cases have included more specialist algo executions, for example on China baskets or US Dollar Index (DXY) baskets. Donner adds that US based hedge funds have also started executing FX options portfolios or baskets using the basket algo on the MDPs. "This was driven by the need to efficiently delta hedge a basket of FX options, so it makes sense to use the basket algo to achieve this," he adds. "We are encountering some very interesting use cases and the basket algos on MDPs is definitely turning into a fascinating business opportunity. The market was obviously more than ready for this product."

At the same time, Goldman Sachs has been having particular success with its recently released EFP functionality on currency futures, which merges futures and FX into a single product. The product is available via a market leading MDP, with clients able to visit any landing page ie futures algo or cash equities algos, for example, and select how many lots to buy or sell of a currency future in the same workflow. "Customers can then opt to execute using our flagship dynamic hybrid algo, or Twap. We then convert that futures trade into an FX forward trade, then once it is done it is sent as a block trade to the exchange. So the order looks like a future all the way through, except that they are making use of OTC liquidity in the middle. This has been a very big growth business for us and once which may mark a significant shift in the nature of the FX algo market," Donner says.



A workflow diagram for EFP

BNP Paribas tackles issue of indirect market impact on FX algos

Liquidity access is becoming a key topic of interest for FX algo users, with the way market data is being consumed now having a notable impact on information leakage. Asif Razaq, Global Head of FX Automated Client Execution, BNP Paribas, explains why FX algo users are increasingly aware of this phenomenon and shares the unique approach BNP is taking to resolve the situation, including moving away from posting orders on lit venues and the launch of a new BIX Peg mechanism.

A primary benefit of FX algos when they were first introduced was the access they provided for clients to external markets, allowing them to post their interest across different venues to source liquidity and trade passively or aggressively across these different platforms. In response, banks typically adopted a hybrid approach, according to Razaq. “If the bank’s franchise liquidity was showing the best price, then the algo would trade internally with the bank, but if the market was cheaper, then the algo could look externally beyond the bank to source the liquidity,” he adds. “What we are now observing is that the market is becoming far more ‘electronically sensitive’, which means that these algo orders are now starting to leave digital signals, creating information leakage of


a client order. The minute we post an order onto an exchange and that order is transmitted through the market data we are finding that some sophisticated counterparts are reacting to the event.” While information leakage was initially attributed to HFTs, Razaq notes that there has also been a notable increase in regional banks looking at the market data before formalising where they think their price should be when they’re going to market make. “As a result, when they see a change in the order book, they will change their quote to reflect where they believe the market is trading,” he says. “That quote then gets transmitted back to these exchanges and the market shifts as a result. We call this phenomenon indirect market impact, or information leakage.” Now

passive algo orders are now likely to create this indirect market impact as a result, Razaq warns, which in turn means the algo creates more noise and the performance deteriorates.

“To solve this problem, we conducted A/B testing on posting orders which are fully visible to the market, versus posting orders only on dark venues. What we saw from this was that there was a big performance improvement by not posting orders to the market,” says Razaq. “In response, we shared these findings with our exchange partners and liquidity providers and explained that we will now need to stop posting lit orders to their venues because it is having an impact on our client orders.”

In collaboration with these vendors, BNP now utilise a new dark order type. According to Razaq, this new order type allows the bank to post algo orders to these venues as before, but this time the details of that order book are not transmitted to all counterparts, so no one in the market is aware of the order. “This means that the order still gets filled passively, but it is now no longer creating this indirect market impact,” he says.

RISE OF MID-MATCHING
“This phenomenon has become a bigger issue and clients are also aware of it,” adds Razaq. “We have explained to clients this change in our posting logic, where we no longer post in lit venues, or if we do, it is in only on a highly sporadic, randomised basis, so it doesn’t create a consistent signal.




BNP’s 4th generation algorithms are designed to service different levels of risk appetite. This is guided by finding the right balance of exposure between market risk and market impact

Our order placement philosophy has now shifted to placing the majority of our orders on dark platforms. After this shift, we ran a performance check and confirmed that it did indeed improve the performance of our algo orders.”

There is, however, a downside, Razaq explains. As the order is no longer being seen by the market, in turn this means that the fill ratios or fill speeds have significantly reduced. “Our next issue was that we then had to find a way to circumvent and fill that void, which is why we have since introduced more mid-matching capability into our platform,” he adds. “We have seen the

arrival of more mid-matching venues in recent years, to the point where we can now increase our participation in mid-matching venues for our algo orders.”

INTERNALISATION AND INNOVATION
On top of that, BNP also runs its own mid-matching platform, BNP Internal eXchange (BIX), which allows BNP clients to match orders off internally in a block match. Razaq explains that this cultivates a large block fill with zero market impact guaranteed on that trade, because that trade takes place internally on BNP’s platform. BIX was also recently expanded to not only




BIX enables client orders to match against BNP Paribas internal liquidity, providing fills without exposing the order to the wider market

accept client-to-client orders, but now any user (external and internal) using the algo platform can be part of that match. “The universe of BIX matching is starting to grow immensely. To aid that growth, we looked at innovative ideas on how to increase BNP’s market share. The more volume that trades through BNP’s algo platform, the more mid-matching fills we can offer.” This resulted in the launch last year of BNP’s gold label initiative, where the bank started allowing selected partner banks the ability to offer BNP algos to their client base. “All these new clients, will now utilise our algo platform, which in turn gives them access to BIX as well. By enriching the volumes that trade via BIX, we are also able to establish a new source of liquidity which cancels out the impact of not placing lit passive orders in the market anymore,” Razaq adds.

In addition to this, BNP has also introduced a new initiative, the BIX Peg product. Razaq explains that this tooling essentially introduces a partnership scheme with the bank’s market making desk.

“Through an electronic connection, we can now post a small component of the algo client’s order interest to BNP’s market making franchise, which provides client access to a unique pool of liquidity,” he explains.

This phenomenon has been evolving over the past few years but the impact was especially notable in the recent periods of heightened market volatility, explains Razaq, who adds that this accelerated the need to start migrating how BNP placed algo orders in the market. “We were also developing the BIX Peg tooling which now forms the foundation of our algo platform,” he says. “We have re-engineered our algos to move away from placing lit orders, which was the core of how algos used to source liquidity. BNP is a leader in the FX algo space and we responded to this shift in clients being more selective about internalisation rates and how we manage information leakage. Clients, especially more sophisticated clients, are becoming aware of these issues and are now starting to critique and pick the banks that can offer this advanced level of algo solutions execution.”



Reducing the information leakage associated with accessing external liquidity venues, for example in equities, is becoming an increasingly important consideration for institutional FX algo users

The AI Act is coming for your trading system: What you need to know

By Charles Glah, CFA UK (ASIP), Founder & Chief Product Officer, DeepAlgo Sovereign



Charles Glah

On October 19, 1987, portfolio insurance strategies did exactly what they were designed to do. They sold as markets fell. And as selling intensified, many strategies using similar dynamic-hedging logic generated further sell signals. By the close of trading, the Dow Jones Industrial Average had fallen 22.6%, the largest single-day percentage decline in its history.

Portfolio insurance was not the sole cause of Black Monday, but subsequent analysis identified it as one of the mechanisms that compounded selling into an already falling market. The algorithms had not necessarily malfunctioned. The deeper problem was the interaction between automated strategies operating simultaneously in stressed conditions.

Thirty-nine years later, the European Union is asking a related question. The EU AI Act represents the world's most comprehensive horizontal regulatory framework for artificial intelligence. Many of its provisions now apply,

although the timetable for the rules governing high-risk AI systems has recently been extended. Under changes adopted in July 2026, the requirements for stand-alone high-risk AI systems will apply from 2 December 2027, while those embedded in regulated products will follow from 2 August 2028.

For financial institutions running AI-driven execution, pricing, risk, surveillance or trading systems, the implications are significant. But firms first need to understand exactly where their systems sit within the regulatory framework.

WHERE TRADING SYSTEMS ACTUALLY SIT

The AI Act applies a risk-based framework to AI systems. High-risk applications face the most demanding requirements, including risk management, technical documentation, record-keeping, transparency, human oversight, accuracy, robustness and cybersecurity.

But there is an important distinction for trading firms.

AI-based algorithmic trading is not currently designated as a high-risk use case under the AI Act. ESMA made this explicit in its February 2026 Supervisory Briefing on Algorithmic Trading in the EU. It states that while an algorithmic trading system meeting the Act's definition of an AI system may be subject to relevant AI Act requirements, AI-based algorithmic trading is currently excluded from the scope of high-risk use cases. ESMA also notes, however, that the scope of high-risk use cases is subject to annual review.

That distinction matters. It means firms should not assume that an AI-driven execution algorithm automatically falls into the Act's high-risk category simply because it makes autonomous trading decisions.

Equally, firms should not conclude that the AI Act is irrelevant.

Where an algorithmic trading system meets the Act's definition of an AI system, relevant provisions may still apply depending on how that system is used. ESMA notes, for example, that AI used in algorithmic trading may potentially attract transparency obligations where it is intended to interact directly with natural persons.

More importantly, firms are already subject to extensive algorithmic-trading requirements under MiFID II and RTS 6.

The practical compliance challenge is therefore not simply:

Is my trading system high-risk under the AI Act?

It is:

Which elements of the AI Act apply to this system, how do they interact with our existing MiFID II obligations, and can we demonstrate effective governance of the system throughout its lifecycle?

THE EXISTING REGULATORY FRAMEWORK IS ALREADY DEMANDING

ESMA's February 2026 Supervisory Briefing is important because it makes clear that the arrival of AI does not replace the existing algorithmic-trading framework.

Investment firms remain responsible for their trading algorithms regardless of whether they are developed internally or supplied by a third party.

Systems must be appropriately governed, tested and monitored. Material changes require appropriate testing. Firms must maintain clear accountability and oversight, and outsourcing cannot transfer regulatory responsibility away from the investment firm.

ESMA goes further where AI is involved. It warns that a series of apparently minor model recalibrations could accumulate over time into a material change in model output without the resulting system being properly tested.

That is particularly relevant for machine-learning systems.

A conventional algorithm may behave differently because somebody changes its code. A machine-learning system may change behaviour because its inputs, parameters, training data or learned relationships change. The governance problem therefore becomes continuous rather than episodic.

Under Article 9 of RTS 6, investment firms already have to carry out an annual self-assessment and validation process covering their algorithmic trading systems, strategies, governance and compliance arrangements. ESMA says firms and supervisors should explicitly consider the use of AI within that process.

This is arguably the most immediate regulatory issue for many FX firms. The obligation exists now.

WHAT HAPPENS IF A SYSTEM IS CLASSIFIED AS HIGH-RISK?

Although AI-based algorithmic trading is not currently a designated high-risk use case, firms should still understand what the high-risk framework requires. The classification can evolve, and financial institutions may operate other AI systems that do fall within high-risk categories.

Several provisions are particularly relevant.



For financial institutions running AI-driven execution, pricing, risk, surveillance or trading systems, the implications are significant

ARTICLE 9 — RISK MANAGEMENT

A documented risk-management system must be established, implemented and maintained throughout the lifecycle of a high-risk AI system.

That includes identifying reasonably foreseeable risks, evaluating those risks, adopting appropriate mitigation measures and testing systems against defined performance metrics and thresholds.

This is not a one-off model-validation exercise.

For a financial institution, the practical implication is that governance has to follow the model into production. A model that was safe when approved may behave differently when market structure, liquidity, volatility, correlations or the data on which it relies begin to change.

ARTICLE 12 — RECORD-KEEPING

High-risk AI systems must include automatic logging capabilities enabling events to be recorded throughout the system's lifecycle.

The purpose is traceability. That does not mean that every individual trade must be accompanied by a human-readable essay explaining precisely why the AI made its decision. It does mean that firms need records sufficient to

understand and investigate how the system behaved, identify relevant risks and support post-market monitoring.

For trading firms, this has obvious parallels with existing expectations around algorithm logs, order records, parameter changes and testing documentation.

ARTICLE 13 — TRANSPARENCY AND PROVISION OF INFORMATION

High-risk systems must be sufficiently transparent to allow deployers to interpret their output and use it appropriately.

Providers must supply information about the system's characteristics, capabilities and limitations so that those responsible for deploying it understand how it should — and should not — be used.

This is sometimes described loosely as an "explainability requirement", but that phrase can be misleading. The Act does not prescribe one universal technical method for explaining every model decision. There is no requirement that every trading firm adopt SHAP, feature-attribution techniques or any other particular explainability methodology.

The appropriate level of interpretability depends on the system and its intended use. For a trading firm, the more useful questions are practical ones:



A documented risk-management system must be established, implemented and maintained throughout the lifecycle of a high-risk AI system

Can the people responsible for the system understand the factors that materially influence its behaviour?

Can they identify when it is operating outside the environment for which it was designed?

Can they tell when model confidence is deteriorating?

Can they reconstruct what happened after an abnormal trading event?

Those are much more meaningful questions than whether somebody can produce an elegant explanation of one isolated trade.

ARTICLE 14 — HUMAN OVERSIGHT

For high-risk AI systems, human oversight must be designed into the operating framework.

The Act requires that appropriate human overseers are able to understand the system's capabilities and limitations, monitor its operation, interpret its output and, where appropriate, disregard, override or reverse that output.

They must also be able to intervene in the operation of the system or stop it safely.

This does not mean that a human must approve every algorithmic trade. That would defeat much of the purpose of algorithmic execution.

It does mean that human oversight has to be effective rather than theoretical. A kill switch that exists somewhere in a manual but cannot practically be used when the system begins behaving abnormally is not meaningful control. That principle should already be familiar to MiFID-regulated algorithmic trading firms.

THE EXPLAINABILITY PROBLEM

There is a particular challenge when machine learning enters trading.

Traditional execution algorithms are normally constructed from rules that can be examined directly:

If volatility exceeds a threshold, reduce participation.

If spreads widen, become less aggressive.

If an order reaches a defined proportion of market volume, alter the execution trajectory.

The causal chain can generally be inspected.

Machine-learning models are different. A model may combine dozens or hundreds of variables and discover relationships that were never explicitly programmed by the people operating it.

That can produce better predictions. It can also produce a governance problem.

Suppose an AI-driven execution system suddenly becomes significantly more aggressive in EUR/USD at 10:17 on a volatile morning.

The question for a compliance officer is not simply whether the trade made or lost money. They may need to know whether the behaviour reflected volatility, liquidity, order-book imbalance, historical flow, model recalibration, a data-quality problem or an interaction nobody anticipated. That is why explainability in trading should not be treated as a fashionable branch of data science. It is increasingly becoming part of operational control.

THE REAL LESSON FROM 1987

Black Monday is useful precisely because the systems involved did not need artificial intelligence to generate systemic consequences. Portfolio insurance used relatively straightforward rules.

The danger came from feedback. Markets fell. Strategies responded by selling. That selling contributed to further market declines. Those declines generated further selling. AI introduces an additional dimension because some systems can adapt to changing data and discover relationships that are not explicitly coded by their designers. That does not make AI inherently dangerous. It does make governance harder.

ESMA's warning about successive small recalibrations accumulating into a material model change goes directly to this point. A system can evolve gradually until the model running today is meaningfully different from the one originally validated. The control framework therefore has to evolve with it.

WHAT FIRMS SHOULD BE DOING NOW

The first step is an inventory.

Firms should know which systems in their trading and risk infrastructure potentially meet the AI Act's definition of an AI system. That sounds obvious. It isn't.

Machine learning may sit inside execution algorithms, liquidity selection,

smart order routing, transaction-cost analysis, market surveillance, pricing engines, risk controls, client-flow classification and vendor products whose internal architecture the investment firm does not fully control.

The second step is classification.

Do not assume every AI system is high-risk. Determine what the system actually does, how it is used, whether any specific AI Act obligations apply and how those obligations interact with MiFID II and RTS 6.

The third is ownership.

Somebody must be accountable for what the system is doing. The model may have been built by a quant team. It may have been bought from a vendor. It may run in somebody else's cloud environment.

None of those facts removes the firm's responsibility for the algorithmic trading activity.


ESMA is explicit on outsourcing: firms should retain ongoing visibility and access to outsourced algorithms and must have the unconditional ability to monitor, suspend or terminate algorithmic trading when necessary. Regulatory responsibility remains with the investment firm.

The fourth is change control.

Machine-learning models create a temptation to regard continual recalibration as normal operation rather than model change. Regulators may take a different view when cumulative adjustments alter the behaviour of the system materially. Firms need thresholds for deciding when a change requires testing, validation or renewed approval.

The fifth is evidence.

It is no longer enough to say that a system is monitored. Firms need to be able to demonstrate how. What was tested? Who approved it? What parameters changed? What data was used? What happened when the model encountered conditions outside its normal range? Who could



For high-risk AI systems, human oversight must be designed into the operating framework

intervene? Was intervention technically possible? What records were retained?

Those questions should be answerable before a regulator asks them.

THE PENALTIES ARE REAL — BUT THEY NEED TO BE UNDERSTOOD CORRECTLY

The AI Act provides for substantial administrative penalties. The highest maximum — up to €35 million or, for undertakings, 7% of total worldwide annual turnover for the preceding financial year, whichever framework produces the applicable maximum — relates to infringements involving prohibited AI practices. Other infringements of obligations imposed on providers, deployers and other operators can attract penalties of up to €15 million or 3% of worldwide annual turnover, subject to the detailed provisions of the Act. Those figures are maximum penalties, not automatic fines for an algorithmic-trading compliance failure. But they demonstrate how seriously the European Union is treating AI governance.

THE DEADLINE HAS MOVED. THE PROBLEM HASN'T.

The regulatory timetable has changed.

The high-risk AI provisions that had originally been expected to apply from August 2026 have been postponed. Under the legislation adopted in July

2026, the principal requirements for stand-alone high-risk systems will apply from 2 December 2027.

For algorithmic trading firms, however, that should not be interpreted as an eighteen-month holiday.

AI-based algorithmic trading is not currently classified as a high-risk use case under the AI Act, but the classification is subject to review. More importantly, the MiFID II and RTS 6 obligations governing testing, governance, validation, oversight and control of algorithmic trading systems already apply.

The firms best prepared for whatever comes next will therefore not be those attempting to predict exactly where regulators eventually draw every classification boundary. They will be the firms that can demonstrate that they know which models they are running, what those models are allowed to do, how their behaviour is tested, who owns them, how changes are controlled, what evidence is retained and how a human can intervene when something goes wrong.

That is ultimately the issue. In 1987, the algorithms did exactly what they were designed to do.

The question the EU AI Act is asking is simpler:

Can you prove it?

Bloomberg FXGO and the extraordinary pace of change in FX algos

The unique position of Bloomberg FXGO in the FX algo space continues to ensure the platform is fast to recognise market changes and respond to shifts in client demand. Oleg Shevelenko, Product Head of Pricing and Execution for FXGO at Bloomberg, shares the stand-out areas of volume growth on the platform, how the electronification of algos is evolving and the importance of supporting clients to understand complex workflow challenges, including the automation reality gap.



What changes have you seen in FX algo use on FXGO and what is driving that demand?

Interestingly, we're seeing the growth in algo trading on FXGO across a range of client types including asset managers, hedge funds, and regional banks, and across instruments that would have been considered too complex or illiquid for algo execution not long ago. Buy-side

desks are managing more complexity with more currencies, more instruments, more regulatory obligations, except with the same or fewer people. Algos offer the consistency, repeatability, and governance that manual execution simply can't deliver at scale. That combination of operational pressure and improved platform capability is what's driving demand.

What factors are currently influencing the rate of algo growth in FX?

Three things are converging across the market.

First, the underlying infrastructure has matured: more venues are offering regulated electronic execution, liquidity in instruments like NDFs has deepened considerably, and the technology underpinning price streaming and order management has become significantly more robust. That creates the conditions for algo execution to work in markets where it previously couldn't. On FXGO, more than 250 banks stream spot liquidity across 100+ currency pairs,

alongside 16 banks providing streaming NDF liquidity across Asian and Latin American currencies and a growing number of providers streaming swaps. Using the expansion of streaming liquidity on FXGO as a proxy for the broader availability of electronic liquidity, we see a strong foundation for continued growth in algorithmic execution.

Second, instrument coverage has broadened industry-wide, with algo strategies now available across NDFs, swaps, and precious metals, allowing firms to standardise execution workflows across a much larger portion of their FX activity. On FXGO specifically, clients can now execute algorithmically across more than 190 currency pairs, including NDFs and precious metals, which means firms can standardise their execution workflows across a much broader portion of their FX activity.

Third, buy-side operating models are changing. Asset managers and hedge funds are under sustained pressure to do more with less with the same or smaller teams. Algos offer the consistency,

scalability, and governance that manual execution can't deliver at that kind of scale.

Is the rise more pronounced in particular types of algo product?

NDF algos have been a standout area, with volumes up year-on-year on FXGO. That's a market that was widely considered too fragmented and illiquid for algo execution until relatively recently. What's changed is a combination of better electronic liquidity, improved regulatory clarity, and a more sophisticated participant base.

APAC has been particularly strong, driven by the growth of NDF markets and a rapid shift toward electronic execution.

We've also seen strong demand for bulk routing to algos and the ability to route multiple orders to an algorithm simultaneously rather than one by one. This demand is particularly prevalent across asset managers, and it's a good example of how algo innovation is increasingly being driven by client workflow needs rather than technology for its own sake. We have exciting things in development here and in basket algo support with more to be announced soon.

Are there areas where algo use can be expected to rise?

NDFs will continue to be a major growth area as liquidity deepens and more venues develop NDF-specific strategies. FX swaps and precious metals are also gaining traction as clients look to bring more of their activity into a consistent electronic workflow. An orthogonal area of growth is workflow, where clients increasingly prefer to execute portfolios or baskets of orders using a common set of parameters and execution instructions, rather than executing each order individually. This enables clients to manage larger and more complex sets of orders more efficiently while applying a consistent execution strategy across the portfolio.

Can you explain where the electronification of algos actually stands today?

Algo execution in spot FX is well established and relatively mature.

Beyond spot, we're at an earlier stage; NDFs are developing quickly, but liquidity remains uneven across currency pairs and regions. FX swaps and certainly options are further behind still. The honest picture is that electronification is progressing, but unevenly, and the pace varies significantly by instrument, region, and client type. What's important to recognise is that electronification is not a binary state. It's a spectrum, and most institutions are somewhere in the middle, running electronic execution for part of their flow while relying on voice for the more complex or relationship-dependent trades.

What is the automation reality gap and why should clients be aware of this?

The automation reality gap is the distance between what firms say they can automate and what they actually automate in practice. The reality is that voice and electronic execution are complementary, not competing. Large-notional trades, complex structures, and relationship-sensitive flow will remain voice-driven for the foreseeable future. The risk is that firms set unrealistic automation targets, encounter friction, and lose confidence in the process.

Our focus is on supporting clients as they trade today, while identifying the surrounding workflows where automation can deliver meaningful progress.

What does FXGO do to additionally educate and support algo users?

Beyond the platform itself, we work closely with clients to build tailored roadmaps based on their specific needs and operating models. On the analytics side, we've invested in algo analytics hosting, which allows providers to deliver pre-trade and live-order performance data directly within the FXGO workflow, and we've integrated Bloomberg TCA so clients can compare algo execution against risk-transfer alternatives on a like-for-like basis. That gives clients a measurable, data-driven foundation for algo selection and continuous improvement.

What benefits can AI offer in improving FX workflows?

At this stage AI is primarily an efficiency tool, not yet a trading strategy. The real gains are in surfacing information faster and removing manual steps, not in generating alpha. And AI is only as good as the data underneath it. FXGO's advantage is that we sit on one of the largest pools of real-time executable FX data in the market with more than 20,000 in over 140 countries, and ADV exceeding \$900 billion. That data foundation is what makes AI tooling genuinely meaningful rather than decorative.

Are there additional applications of AI which can benefit FX algo execution?

Potential algo applications include using AI to improve liquidity and volume forecasting, dynamically adapt execution strategies to changing market conditions, optimize venue and liquidity-provider selection, and enhance transaction-cost and market-impact models. AI can also help clients select and parameterize execution strategies based on the characteristics of an order or portfolio, while post-trade analysis can identify execution patterns and feed those insights back into future strategy selection. The opportunity is therefore not necessarily to replace established execution algorithms, but to make them more adaptive, predictive and easier for clients to use. Across FXGO more broadly, we're continuing to invest in NLP across the platform with the goal of making the data that already exists on the Terminal actionable in ways that support faster, more informed execution decisions.

Is there anything coming up in the development pipeline you would be able to share?

The direction of travel across everything we're building is toward giving clients deeper programmatic access across more instruments and more stages of the trade lifecycle, reducing manual touchpoints at every step, from pre-trade price discovery through to post-trade reporting. The pace of change is genuinely extraordinary right now, and the firms that will benefit most are those investing in the data and workflow infrastructure today rather than waiting for the technology to mature further.



How much customisation do clients really want from their execution algos?

Looking back at the evolution of the FX algo offering, demand from clients for customisations and even bespoke algo strategies started to emerge as a key differentiator as early as 2020, notably around the time FX algo adoption started to significantly pick up in pace. Providers offering such services have since evolved and honed these services to meet the increasingly sophisticated requirements of more established algo clients, while others warn against the perils of introducing unnecessary complexity into a process that ultimately revolves around automation. Nicola Tavendale writes.

There has actually never been a true 'one-size-fits-all' solution in FX execution, claims Dr John Quayle, Head of Client Algo Execution at NatWest Markets. This is because, at its core, every execution decision involves a trade-off between minimising transaction costs on one hand and minimising execution time and the associated market risk on the other, he explains. "While 100 per cent certainty can be achieved through an immediate risk-transfer trade, that certainty comes at quite a high cost," Quayle adds. "Conversely, a well-designed passive execution algo can typically achieve better average pricing but may also introduce exposure to market movements while the order is being worked. The most common customisation discussion we have with clients at NatWest therefore revolves around the speed of the algo."

In practice, he argues that often this will translate as a discussion about where a client wishes to position themselves on the cost-versus-risk spectrum. So although some clients may be primarily focused on minimising market impact and are comfortable taking additional execution risk to achieve lower overall costs, other clients may tend to be more focused on speed, according to Quayle. He adds: "What has changed over recent years is that clients are



Dr John Quayle

"The most common customisation discussion we have with clients at NatWest revolves around the speed of the algo."

increasingly seeking execution strategies that can be calibrated to their specific objectives and risk appetite, which of course varies depending on the currency pair and time-zone. We do not see this demand for customisation being limited to a particular segment of the buy-side; whether the client is an asset manager, hedge fund or a corporate, the ability to tailor the balance between execution cost and market risk remains one of the most important aspects of FX algo design. Therefore, as liquidity providers continue to invest in analytics and gain a deeper understanding of their execution outcomes, we expect this trend towards greater flexibility and customisation to continue."

ATTITUDES AND EXPECTATIONS

In addition, client attitudes and expectations around FX algo customisation have evolved significantly in recent years. Buyside firms may well have become significantly more sophisticated, but the direction of travel is not simply towards more buttons and parameters, says Vittorio Nuti, Head of Algorithmic Trading, Deutsche Bank. He adds that clients are now increasingly asking for customisations based around improving execution outcomes, workflows and evidence, rather than seeking bespoke code for every trading style. "Historically, the discussion was often about basic controls such as speed, urgency, participation and benchmark tracking. The more recent client dialogue is about pre-trade analytics, real-time monitoring, post-trade TCA, API integration, governance and the ability to adjust execution dynamically as market conditions evolve," Nuti says. He believes that the main drivers behind this shift are factors such as best-execution scrutiny, more sophisticated execution desks, independent TCA from providers such as BestX, Curex and TradeFeedr, broader adoption of electronic workflows and an overall increase in client desire to explain outcomes internally.

In addition, Nuti shares that the most common requests from clients tend to be focus on key aspects of FX algo execution, such as aggressiveness, urgency, benchmark objective, timing, limit constraints, passive versus liquidity-seeking behaviour, use of internal liquidity, and the ability to intervene

during the order. He adds that clients also increasingly request customisations relating to their workflows and analytics: how orders are submitted, how live execution is monitored, how reports are delivered, and how performance can be compared against risk-transfer or independent TCA benchmarks. "In practice, many clients want to tailor how they interact with the algo more than they want to redesign the underlying execution engine," Nuti says.

SIMPLIFYING THE PROCESS

Meanwhile Pieter Oudshoorn, FX Sales for European Financial Institutions at NatWest Markets, adds that setting the most appropriate algo speed is often the most common focus for NatWest's algo clients. The parameters which will primarily affect this are firstly, the placement of child orders in the order book (joining the passive side of the market or going in at mid, for example) and secondly, the choice of liquidity pools, he explains. "In practice, however, most users specifically want to be 'passive' ie try to capture spread and join the passive side of the market," Oudshoorn says. "This in turn limits the flexibility from the order placement level, which leaves the choice of liquidity pools as the most impactful avenue for customisation." He explains that this is overall an especially rich topic and one which is evolving continually as new venues such as OneChronos enter the fray, and existing venues add new features and change their approach to curation. Oudshoorn adds: "At NatWest we are able to support the real-time management of liquidity pools, therefore this is a particular focus for client discussions and where we see clients regularly testing new combinations. Knowledge of which liquidity source performs well in which pair and which time-zone is crucial and is a key area where we can add value to the relationship."

Flexibility is also ultimately vital to the execution process – market conditions change, risk appetite evolves, and so the tools available do need to be able to easily accommodate this need, Oudshoorn explains. He argues that the pitfalls ultimately arise not from the scale or scope of flexibility offered, but rather from the challenge of gathering enough data for a given set of parameters to

Deutsche Bank
Investment Bank



FX. Unrestricted.

Aladdin
by BlackRock

HausFX
by Deutsche Bank

Contact us at haus.fx@db.com or visit db.com/hausfx

For institutional investors only. This does not constitute an offer to enter into any transaction. Transactions are subject to investment risk. The value of transactions may go down as well as up.
©2026 BlackRock, Inc. or its affiliates. All rights reserved. BLACKROCK, ALADDIN, eFRONT and PREQIN are trademarks of BlackRock, Inc. or its affiliates. All other marks are the property of their respective owners.
Details about Deutsche Bank's regulatory status are available on <https://www.db.com/about>



Vittorio Nuti

“In practice, many clients want to tailor how they interact with the algo more than they want to redesign the underlying execution engine,”

determine if they actually benefit the user or not. “The fact of taking market risk, that each order has a variable outcome as a result, means that getting a statistically significant set of data for a given set of parameters takes time,” says Oudshoorn. “If you always use the same settings, you will get there quickly, if you use five different settings on a regular basis, you need to be able to log the performance for each run (in itself not a given, depending on the analytics being used) and it will take you five

times as long to get there.” He adds that this is where third party services such as Tradefeedr can add value. While peer universe-style aggregations can help, Oudshoorn believes that these again tend to cover only a relatively small number of the customisation options available, which puts the onus back on the user to limit the variations used to improve the data. He adds: “The ability to quickly identify a small number of key default settings is key to successfully navigating this complexity, which is why we work very closely with our clients to thoroughly understand their execution requirements, leveraging our data and TCA services to map the most suitable parameter configurations.”

In addition, customisation requirements can differ between different client segments, with asset managers tending to care most about arrival-price performance, as well as other factors such as their portfolio workflows, benchmark tracking, governance and TCA evidence, according to Nuti. “Official-sector clients, on the other hand, typically place more emphasis on auditability, information leakage, transparency and clear suitability guidance, while quantitative and highly automated users are more likely to focus on factors such as API integration, parameter control and systematic workflow,” he adds. Furthermore, Nuti adds that regional banks tend to value workflow simplification, hedging automation and spread capture, as well as requiring access to institutional-grade execution technology but without the

need to build everything in-house. He explains that corporate and less frequent users will generally prefer to opt for simplicity, confidence and clear defaults rather than requiring a wide parameter set.

However, Nuti warns that such flexibility can become counterproductive if clients are asked to choose parameters they cannot reliably optimise, or when small parameter changes create complexity that is not justified by measurable improvement in outcomes. “For example, urgency, limits and timing choices can materially affect execution quality, but clients may not always have the market context to choose them optimally in real time,” he explains. “Too many choices can create decision fatigue, inconsistent usage and lower-quality execution decisions. A better model would be guided flexibility: keeping meaningful client controls but framing them around execution objectives, while also protecting the user with analytics, defaults and safeguards instead.”

FINDING THE RIGHT BALANCE

The most useful division of labour is instead when clients define the objective, and the provider helps determine the route, claims Nuti. He believes that clients should specify whether they care most about speed, market impact, benchmark tracking, discretion, completion certainty or leakage reduction, while providers should increasingly translate those objectives into settings using market intelligence, liquidity data, execution

analytics and pre-trade modelling. “This reduces the burden on the client and should make execution behaviour more consistent,” Nuti adds. “That does not mean removing client control, instead it means giving clients the ability to override or guide execution, while using provider expertise to recommend sensible defaults and dynamic adjustments.”

Quayle notes that clients now also need to choose the level of market risk that they are willing to take, and that is an aspect of the algo execution set-up which “simply cannot be automated”, he adds. However, he explains providers can instead make it easier for the user to select the combination of parameters that best suits their needs, which means that their expertise is vital. “The client determines the desired balance between execution cost and market risk, while the provider uses its expertise, analytics and market intelligence to optimise the underlying execution logic within those boundaries,” Quayle adds. “This highlights one aspect of customisation which is often overlooked – the need to easily set up customisable defaults.”

He explains that this is due to modern algos being able to support a very large number of parameters and potential combinations, but in reality, most clients only require a relatively small subset of those configurations to meet their day-to-day execution needs. In turn, the ability to create and save these customised defaults can significantly improve usability, operational efficiency and consistency of execution outcomes, Quayle says. “The industry is unlikely to be moving towards a fully automated model where providers make all execution decisions,” he adds. “Instead, the trend is more towards smarter automation within client-defined parameters; clients will set the risk appetite and execution objectives, while providers make it easier to implement those preferences and use their expertise to deliver the best possible outcome for the client.”

Nuti observes that leading providers also now appear to be moving toward a controlled-choice model: intelligent defaults, pre-trade guidance, real-time feedback, ability to intervene, and post-trade analytics to assess the result. “Deutsche Bank’s pre-trade

cone-and- outcome box tool is a good example of this direction,” he adds. “It helps clients understand not only expected execution duration, but also historical market-move ranges, forecast arrival-price performance, risk-transfer comparison, and the impact of switching between styles before or during execution. This allows the client to remain in control, but without the need to perform all the execution analytics manually.”

USE CASE FOR AI

Furthermore, AI and machine learning are also beginning to support the parts of algo execution where prediction and adaptation matter most – namely in the functions of liquidity forecasting, short-term alpha, market-move estimation, impact modelling, timing, and dynamic order placement, Nuti explains. He adds: “Investment in proprietary short-term forecasting is a large area for Deutsche Bank research. The objective is to use data and prediction to improve execution decisions, not simply to automate a static schedule.” A more advanced use case is adaptive execution that responds to changing liquidity and volatility, while still remaining inside client-defined objectives and safeguards, says Nuti.

He also argues that a natural next step is for execution systems to learn from a client’s historical trading behaviour. This would include areas such as preferred urgency, typical intervention patterns, tolerance for timing risk, benchmark choice, execution horizon and which outcomes the client actually rewards, says Nuti, while also allowing the system to recommend or apply client-specific execution profiles, instead of asking the user to choose between generic settings on every order. “The barrier is less about the technology and more about the explainability, governance and client consent elements. Any adaptive profile would need to be transparent, testable and easy for the client to challenge or override,” he adds.

This means that transparency also becomes more important as customisation and automation increase, according to Nuti. “If an algo is making more decisions on the client’s behalf, the client will increasingly expect to understand why those decisions were made and whether they were



Pieter Oudshoorn

“The value of customisation is only fully realised if clients understand not just what controls are available, but also how those controls influence the algorithm’s behaviour in practice.”

consistent with the stated objective.” He explains that algo clients increasingly want information on strategy logic, expected performance, limitations, liquidity sourcing, slicing, internalisation, execution path and to be able to compare this data with the pre-trade expectations. Nuti adds: “The FX Global Code framing is helpful here: algorithmic execution is a form of delegated execution, so providers need clear communication, documentation, governance and accountability.”

Oudshoorn agrees, adding that transparency is absolutely critical in building trust for the user when it comes to customisation. “The value of customisation is only fully realised if clients understand not just what controls are available, but also how those controls influence the algorithm’s behaviour in practice,” he adds. “Without that transparency, it becomes difficult for clients to predict outcomes, evaluate performance, or align an execution strategy with their objectives and risk appetite. At NatWest, we see strong demand for transparency around both the available customisation options



CUSTOMISATION REQUIREMENTS CAN DIFFER BETWEEN DIFFERENT CLIENT SEGMENTS



Transparency is absolutely critical in building trust for the user when it comes to customisation

and the practical implications of those settings.” For example, Oudshoorn notes that if a client were to choose a more passive execution profile, they would want to understand how aggressively the algo will seek liquidity, how completion risk may change, and how the strategy is likely to respond if market conditions deteriorate or liquidity becomes scarce. He adds: “Many clients will expect the algo provider to provide nothing less than 100 per cent of the information on execution decisions as and when required.”

For most people the ‘better outcome’ is reduced market impact and this is the most meaningful execution metric for the algo user, according to Oudshoorn. This is measured by looking at the all-in fill vs the spot mid at inception (implementation shortfall) and represents the actual cost of the execution in simple monetary terms, he adds. “A customised strategy is only genuinely better if it improves the overall execution frontier: either it delivers lower cost for the same risk, or lower risk for the same cost – this is the ‘speed vs impact’ equation again,” Oudshoorn adds. “Whether you are using a customised setting or a standard one doesn’t change this, and so the objective then is simply to be able to gather enough data to be able to make an informed decision. This comes back to the challenge of restricting the range of different settings used in order not to dilute the quality of the data too much – a common pitfall for those new to algos.”

DIRECTION OF TRAVEL

The use of analytics and TCA is also definitely helping users improve outcomes, Oudshoorn argues, with the realisation that the benefits of algo execution can now be measured, and that the relative performance of different providers can be compared is now widely accepted. He adds: “One of the biggest changes is that for many of the more active users meaningful segmentation of performance analysis can be achieved i.e. determining that a given setting which works well in G5 pairs, is not necessarily the same as that which gives the best outcome for CE4.” For example, a given venue is likely to have different levels of information leakage depending on the pair, Oudshoorn claims. “Therefore identifying this and setting up the appropriate customised defaults for the different use cases is one of the goals of high-quality TCA analysis,” he adds.

“Customisation should also be judged against the objective it was designed to improve,” says Nuti. Relevant metrics include arrival-price performance, for example, in addition to other factors such as implementation shortfall, risk-transfer comparison, market impact, fill quality, completion rate, internalisation rate and execution consistency. “The key is statistical evidence over a meaningful sample, not a single-trade anecdote,” he adds. “We always stress that algos should be evaluated on average and in aggregation because individual orders

can vary due to market risk.” As a result, Nuti claims that a customised strategy is genuinely valuable only if it improves the relevant metric without introducing offsetting costs such as slower completion, higher operational burden or poorer consistency.

Furthermore, TCA is making execution customisation more evidence-based, with clients increasingly using independent and internal analytics to compare algo results against risk-transfer benchmarks, peers, historical performance and alternative execution styles. “However, many risks and data gaps do exist which makes these types of evaluations complex,” Nuti warns. As buy-side execution analytics mature, firms are also likely to develop increasingly distinctive house execution styles, just as they already have distinctive investment philosophies and governance frameworks, Nuti believes. He adds: “Some firms will favour market-impact minimisation, others benchmark tracking, others internalisation capture, speed, certainty or reduced information leakage.” These preferences can be encoded into preferred execution profiles, workflow rules and governance-approved defaults, Nuti argues, and adds: “The likely future is not a unique algo engine for every client, but a common core execution platform with client-specific overlays that reflect trading philosophy, data feedback, operational workflow and governance expectations.”

Why algorithmic execution volumes in Scandies are steadily increasing

By Filippa Bång, Head of e-FX distribution at SEB



In FX, liquidity is only as real as the market impact it creates—particularly in Scandies, where depth is less resilient. This reflects the reality of a highly fragmented market, where liquidity is dispersed across multiple venues and inconsistent depth. Algorithmic execution provides a structured way to aggregate liquidity, reduce signalling risk, and optimise execution through adaptive routing and order placement. These benefits become even more critical in illiquid currency pairs, as limited depth increases sensitivity to order placement, routing logic, and market impact. We see algorithmic execution volumes in Scandies steadily increasing, as market participants turn to automated strategies to better navigate less resilient liquidity.

Liquidity conditions differ materially across currency groups. Major pairs tend to exhibit deeper and more resilient liquidity, while Scandie liquidity is typically thinner and more fragile. This also amplifies time-of-day effects in regional pairs, where

execution quality can deteriorate markedly outside local market hours, unlike in more liquid currency pairs where liquidity remains more consistent globally.

INTERNALISATION

In these environments, internalisation becomes a key driver of execution quality. For Scandies in particular, performance depends heavily on access to a strong client franchise with natural, two-sided flow. A broad and diversified client base, spanning corporates, institutions and retail, enables risk offsetting without relying on external liquidity. This diversity is essential for achieving stable and scalable internalisation.

It is also important to recognise that the definition of internalisation has evolved across the market, prompting increased focus on standardisation and best practices. Our approach is deliberately conservative, classifying flows as internalised only when offset against genuine client flow, actively sourcing liquidity from carefully curated, skew-safe client pools.

Adapting execution strategies to different market conditions is essential. The same execution parameters can appear passive in one environment and aggressive in another, reinforcing the need for dynamic configuration aligned with prevailing market conditions. This also extends to how skew is presented to clients, where robust leakage controls are required to prevent pricing adjustments from becoming unintended market signals.

Liquidity is fragmented and often deceptive due to a liquidity mirage;

where displayed depth fails to materialise upon execution; visible market depth can vanish the moment you try to execute a trade. This underscores the need to carefully evaluate where and how to trade, rather than relying strictly on displayed prices.

Ongoing assessment of liquidity providers is therefore critical, incorporating metrics such as fill ratios, execution consistency, and price decay. A venue offering a tighter top-of-book price may still underperform if it cannot deliver reliable size. In such cases, the cost of missed execution can outweigh the apparent pricing advantage, making consistency a more valuable attribute than spreads.

LIQUIDITY RELIABILITY

Across FX trading, and particularly in algorithmic execution, the focus is increasingly shifting towards long-term liquidity reliability. It is not only about how liquidity is priced, but how it behaves, how it recycles, how consistently it fills, and how execution quality holds across different volumes and market regimes.

For market participants willing to absorb time risk, algorithmic execution provides direct access to internal franchise flows combined with full transparency on execution outcomes. Effective liquidity management, including continuous optimisation of how flows interact, has enabled us to further improve internalisation levels, both within algorithmic execution and more broadly. Taken together, regional banks have a structural advantage through superior access to unique liquidity, even as international competition seeks to challenge it.



White label solutions and algos-as-a-service: More practical alternatives for FX providers?

By Paul Golden



Image by Shutterstock



Paul Golden

The economics of algo development have flipped. A decade ago, offering execution algos was a differentiator that justified the cost and manpower of an internal build - now it is table stakes.

Clients of regional banks and platforms expect the same strategy, controls and analytics they get from a global bank and they judge the outcome, not who wrote the code. Meanwhile the true cost of an institutional-grade build has become clearer.

“Most clients spend most of their time on a couple of algo variants,” says Jeff Leal, co-founder Qubealgo. “The prize is not an ever-longer menu; it is making the widely used strategies perform better while retaining the ability to offer something competitors cannot.”

Fellow co-founder, Martin Zinkin, reckons licensing execution technology is now viewed as evidence that a firm is putting its resources into areas where it differentiates, such as liquidity, relationships and service.

“Firms now understand just how complex institutional grade execution technology is to build and maintain and recognise that they may not have the resources to commit that time and money,” he says.

Two major shifts have driven this change. “The first is that the technology stopped being a black box,” says Zinkin. “The current generation of algos exposes how strategies are constructed and measures them with independent analytics. Secondly, buyers have matured. The key

questions that matter in any build-or-buy decision relate to future needs across products and asset classes, dependence on a vendor’s development and release cycle, long-run costs and the ability to integrate your own components without locking yourself in.”

He suggests that the environment such algorithms must live in is underestimated. “That is why the real cost of a build from scratch is in the testing and safety infrastructure - the ability to replay production data through identical business logic, back test against historical and synthetic markets, stress-test behaviour at many times peak load and prove that risk controls behave as intended under exactly those conditions.”

REDUCED TIME TO MARKET

Danielle Caravetta, director of global sales at Pragma by MarketAxess, says working with an expert in algorithmic trading platforms enables clients to get the look and feel of an in-house solution while significantly reducing time to market and ongoing maintenance. “There are obviously different ways that you can white label technology,” she says. “You can white label from a bank where they offer their technology and additional services outside of just the algorithm trading platform. Or you can outsource to an agency specialist like us who doesn’t do any trading and is not affiliated with any sources of liquidity.” When asked what capabilities clients expect from FX algos that were previously only available from global banks, Caravetta says that although customisation requests in FX are a fraction of what they are in equities, institutional and corporate clients will have very specific ways that they want a TWAP or a float-leg strategy to work and will expect their outsourced providers to support them.

“Post-trade analytics is a key requirement - everyone needs to have post-trade TCA so they can do their own analysis on how the algos are performing. If we are providing algos to a bank which is white labelling them to a corporate or institutional client, they need that post-trade TCA to provide to their client.” She refers to growing use of third party TCA providers and more metrics being used in post-trade TCA – so not only



Jeff Leal

“Institutions should weigh the provider’s team as heavily as the technology and ask who actually built the platform and where they learned the failure modes.”

looking at performance versus arrival or TWAP but also mark-out and reversion data.



Martin Zinkin

“Firms now understand just how complex institutional grade execution technology is to build and maintain and recognise that they may not have the resources to commit that time and money.”

“When choosing an algorithm provider, you need to work with someone who is going to be able to provide you with that level of transparency and that level of data.”

As for how providers of white labelled and outsourced FX algo trading services ensure transparency around the regulatory compliance and performance of their execution solutions, Caravetta observes that as long as they have the data and a surveillance process in place that can provide that data back to the client, that should meet most regulatory requirements.

“This goes back to the reasons why more firms will white label or private label,” she says. “Not only do you have to deal with the technology requirements - if you are building in-house then you are going to have more supervisory and regulation requirements.”

PARTNERING A PRAGMATIC CHOICE

As the cost of in-house development has risen and the ROI has narrowed, the industry has recognised that partnering with a proven, top tier algo provider is a pragmatic choice.

That is the view of Asif Razaq, global head of FX automated client execution at BNP Paribas, who says the challenges involved in developing institutional grade FX execution algos from scratch are mostly financial, technical and organisational. “Developing a competitive algo platform demands sizable investments in IT infrastructure, quantitative research talent and ongoing maintenance,” he explains. “Building the model yourself is expensive and a robust governance framework is required to monitor performance, conduct regular reviews, and satisfy regulatory expectations.”

In addition, the solution must be integrated with various multi-dealer platforms and provide a proprietary transaction cost analysis (TCA) capability. Some top tier providers also offer pre-trade and real time TCA services to further bolster their offering. Assembling all these components creates a substantial cost base and a lengthy development horizon, making it difficult

for new entrants to compete with established providers, observes Razaq, who adds that clients now anticipate a complete end to end execution experience.

“They expect real time visibility into how an algorithm is performing, the ability to run pre-trade cost simulations and analyse post trade performance,” he says. “Mid-trade adjustments are also expected, the platform should support a comprehensive list of currency pairs and analytical services surrounding the trade must also be easily accessible.”

Razaq observes that clients are increasingly moving toward consolidated execution desks that cover FX, futures, equities, rates and commodities within a single unit - a shift that creates demand for execution technology that can be applied uniformly across asset classes, delivering a consistent user experience and simplifying operational processes.

“Providers that have built a single, modular algo platform capable of handling multiple asset classes leveraging the same quant, IT, and trading expertise are gaining a competitive edge,” he says. “By offering the same familiar algorithms across FX, fixed income, metals and futures, firms enable clients to transfer their knowledge from one market to another, reducing learning curves and fostering loyalty.”

FACTORS THAT IMPACT CUSTOMISATION

The degree of customisation available depends on the vendor’s architecture and the resources they allocate to client specific development, explains Razaq. Some providers sell an off-the-shelf product with limited ability to customise, while others - particularly those that have developed their own in house algos - can create a large number of algo variants to suit particular mandates, risk tolerances or execution preferences.

“Transparency is built into the execution workflow,” he says. “Every slice of an order is logged with a timestamp, venue, price and quantity, producing a fully auditable evidence trail that is delivered to the client at the conclusion of the trade. This data allows clients to perform their own performance analysis or to engage third party analytics firms if they

prefer an independent view.”

Razaq says differentiation now hinges on the breadth and depth of the technology stack and on liquidity management, noting that providers that offer a complete suite - including pre-trade tools, real time TCA, intra trade amendments and comprehensive post trade analytics - stand out.

“Equally important is the ability to source and internalise liquidity, reducing the need to expose client orders to the external market,” he adds. “High internal matching rates improve algo performance and are increasingly sought after by clients. Institutions should therefore evaluate a provider’s transparency, regulatory robustness, liquidity sourcing capabilities, the extent of customisation and the provider’s commitment to ongoing innovation.”

Machine learning and artificial intelligence are already integral to modern algo design and Razaq believes their usage will only expand.

“These techniques enable algorithms to learn from market data in real time, adapt execution strategies on the fly and deliver superior performance relative to more static models,” he says.

“Vendors that have invested heavily in AI-driven modelling are able to offer algorithms that respond dynamically to changing liquidity conditions, news events or macroeconomic releases.”

Beyond execution, conversational AI allows clients to interact with the algo during a trade, asking questions such as ‘why is execution slowing?’ and receiving reasoned explanations or receiving proactive prompts like ‘non-farm payroll data is about to be released; would you like to accelerate execution?’

“Such dialogue-driven interfaces represent the next frontier in client focused algo services, turning execution algorithms into collaborative partners rather than black box machines,” adds Razaq.

DIFFERENTIATING FACTORS FOR PROVIDERS

On the question of how providers are differentiating themselves in an

increasingly competitive market and the attributes institutions should look for in an FX algo trading service provider, Caravetta suggests firstly working with a provider that has a track record of providing institutional algorithms that have been used by the largest players in the world.

“You want a provider that has robust technology because you don’t want outages when markets are volatile,” she says. “You should also look for strong quantitative expertise from a provider that really understands the market microstructure. We believe it is best to work with an agency provider because if you work with a principal provider, you are really only going to have access to their franchise liquidity.”

As for future developments, Caravetta notes that Pragma by MarketAxess already uses advanced AI in its algos to determine how to optimise routing, most importantly and that product demand will continue to grow.

“Right now, the biggest demand is still for spot algos but we also offer NDFs and firms will increasingly look to outsource routing and algos for NDFs and also maybe for swaps and forwards,” she adds.


“In terms of outsourcing in general, we will see more types of players outsource their execution. We have brought on a



Danielle Caravetta

“Not only do you have to deal with the technology requirements but if you are building in-house then you are going to have more supervisory and regulation requirements.”

number of clients who have historically built their algos in-house - whether that has been hedge funds or banks - and have realized that there is just no return on investment on the continual technology and regulatory spend. Partnering with a specialist allows them to reduce their ongoing cost and effort



The key to solving the puzzle of customisation is how the algos are built



Asif Razaq

“Developing a competitive algo platform demands sizable investments in IT infrastructure, quantitative research talent and ongoing maintenance.”

so that they can focus their resources on other areas that offer a bigger return on investment.”

According to Leal, clients now expect a complete set of strategies covering TWAP, VWAP, tracking, liquidity-seeking, opportunistic and full-amount styles, with real control over how each behaves. They also expect intelligent multi-venue

execution built on aggregated liquidity.

“Just as important, real-time, in-flight transparency, showing where child orders are working, what has filled and how execution is tracking against arrival, TWAP and VWAP benchmarks is now a baseline expectation,” he says. “The same goes for control: the ability to pause, amend or kill an order instantly with confidence that the system honours it.”

CLIENTS DEMAND TAILORED STRATEGIES

The expectation that raises the bar highest, though, is bespoke algo solutions. Sophisticated clients increasingly want strategies tailored to their flow, constraints and benchmarks and a provider simply cannot service that demand with an inflexible set of rented algos.

Firms that once ran an FX desk in isolation now face client demand for algorithmic execution in NDFs, precious metals, fixed income and digital assets, often through the same workflow and the same analytics. Building a separate stack for each asset class multiplies every cost and stitching together a different vendor per asset class is no better.

“Institutions increasingly want one execution architecture that spans asset classes and that has to be designed in,” says Leal. “Clients want the same

event engine, order framework, strategy building blocks and TCA across FX, fixed income and digital assets, with asset-specific modules layered on top for products, date conventions, risk and hedging.”


The key to solving the puzzle of customisation is how the algos are built. “Customisation comes in layers,” explains Zinkin. “At the surface, every strategy exposes parameters: urgency, participation, venue preferences, limit behaviour. Beneath that, providers can compose new strategies from building blocks, branded and tuned for their franchise.”

At the deepest layer, clients want full access to source code and can extend or replace components outright, retaining sole ownership of everything they build. Every customisation needs to be validated in replay against historical and synthetic data and stress-tested under production-like load before it goes near production.

Leal adds that at first glance all offerings look alike and that the real differentiation is in the detail. How deep does customisation go, parameters only or the ability to extend and compose new strategies? How much control does the client have over change, and on whose delivery timeline? What do the continuing costs look like? And is the client locked into a fixed set of algos, or can the service grow with their franchise?

“Institutions should also weigh the provider’s team as heavily as the technology and ask who actually built the platform and where they learned the failure modes,” he continues. “We built these systems at tier 1 and tier 2 institutions so our design reflects the pain points we lived with: integration friction, vendor queues, costly errors. Technology needs to fit the client’s choice of middleware, database and UI rather than imposing specific requirements.”

Finally, when looking at points of differentiation, Leal suggests firms should also ask questions about source access, client ownership of IP built on a platform, composable strategies and testing discipline.



Machine learning and artificial intelligence are already integral to modern algo design

THE BIGGEST **FX** CONFERENCE

Asset Managers | Pension Funds | Hedge Funds | Corporates

AMSTERDAM

September 15 - 17 | 2026




tradetech **FX**



Save 10% With Code: **EFOREX**
Free Buy Side Tickets Available

From information to knowledge in FX algo execution

By Aaron Shipper, Product Manager at JPMorgan Chase & Co.



Aaron Shipper

CHOOSING WISELY

With so many algo strategies available across providers, what should buy-side firms be thinking about when matching a strategy to a particular order?

The proliferation of algo strategies is a positive development -- but more choice doesn't automatically mean better outcomes. The question is whether firms have a structured way to navigate that choice.

It starts with the objective. Is this about minimising cost, managing risk, hitting a benchmark, or reducing footprint? Different objectives point to different strategies -- TWAP for time risk distribution, POV for market camouflage, aggressive strategies for urgency. The best selection frameworks are simple decision trees, not complex models.

Size relative to daily volume is a critical input. A \$200 million order in EUR/USD, where daily turnover runs into the hundreds of billions, is a very different execution challenge than the

same notional in USD/TRY or USD/ZAR, where it represents a much larger share of available liquidity. Cross pairs add another layer of complexity -- a EUR/JPY order is typically constructed from two legs, each with its own liquidity profile and spread dynamics, making the algo's ability to manage both components efficiently a key differentiator. The liquidity profile of the pair matters just as much as its size.

How should the characteristics of a specific trade – currency pair, size, urgency, time of day – shape the algo selection process in practice?

Every order has a fingerprint - and the more precisely the algo selection reflects that fingerprint, the better the outcome.

Currency pair determines liquidity depth, venue availability, spread

behaviour, and time-of-day patterns. Size matters not just in absolute terms but relative to average daily volume -- this determines how "loud" the order will be in the market. Urgency is worth interrogating: is it a genuine constraint or a preference? The distinction matters enormously for strategy choice.

Time of day shapes everything. London morning and Asia afternoon are effectively different markets with different liquidity profiles, spread dynamics, and participant behaviour. Expected holding period such as intraday versus multi-day affects which benchmark is appropriate and therefore which strategy makes sense. "One size fits all" selection remains the most common missed opportunity we see.

How important is a provider's ability to internalise flow, and what should buy-side firms understand about how internalisation affects their execution quality?

Internalisation or matching a client's order against the provider's own internal flow rather than routing to external venues can be one of the most significant factors in execution quality, yet it's often underappreciated in the selection process. When an algo can match flow internally, the benefits are tangible: minimal market footprint, reduced information leakage, and typically tighter pricing than what's available in the lit market.

End-to-end internalisation matters because it directly addresses adverse selection -- the risk that external counterparties are only willing to trade when conditions favour them rather than the client. When flow is matched internally within a provider's own ecosystem, both sides of the trade are natural, which removes the information asymmetry that drives adverse selection in the broader market. The more of an order that can be filled this way, the less exposed the client is to the signalling costs and predatory dynamics that come with external venue interaction.

How internalisation is accessed relative to external liquidity is equally important. The right approach is to access internal liquidity only when it offers a genuinely better outcome for the client - tighter



Time of day shapes everything

pricing, lower market impact, or reduced signalling -- rather than defaulting to it for the provider's convenience. At J.P. Morgan, internal liquidity is accessed when it is most advantageous to the client, competing on equal terms with external sources rather than being preferred automatically.

FEEDING THE MACHINE

What pre-trade information has become most valuable for informing algo selection and parameterisation, and how well is the buy-side community using it today?

The quality of an algo's execution often depends less on the algo itself and more on how well it's parameterised. The inputs matter as much as the engine.

Historical volume profiles are among the most valuable -- they reveal when liquidity typically appears for a given pair, and these patterns are reasonably consistent outside of event days. Spread behaviour follows similar rhythms: time-of-day, day-of-week, and event-driven widening all create predictable windows that should inform when and how aggressively to execute.

Volatility regime is another critical input. Whether the market is

trending, ranging, or stressed calls for fundamentally different parameterisation -- a set of parameters that works well in a calm market may produce poor results when volatility spikes.

Expected cost modelling brings these inputs together: given this order's size, pair, and current conditions, what should execution cost? The answer provides a baseline against which outcomes can be measured.

The gap we observe is that many firms select the right algo but accept default parameters. Pre-trade analytics exists to turn raw information into a specific execution plan -- but it's underutilised.

Once an algo is running, what real-time signals help distinguish genuinely adaptive execution from a strategy that's simply following a predetermined schedule?

"Adaptive" has become a popular label, but there are meaningful differences in how algos respond to changing conditions during execution. The signals worth watching are the ones that reveal whether the algo is genuinely reacting to the market or simply ticking through a schedule.

EVERY ORDER HAS
A FINGERPRINT



PEER COMPARISON ADDS FURTHER CONTEXT: HOW DOES THIS EXECUTION COMPARE TO SIMILAR ORDERS UNDER SIMILAR CONDITIONS?

Spread movement is one of the clearest indicators. Does the algo recognise when spreads are favourable and lean in, or does it maintain rigid pacing regardless? Similarly, volume patterns matter. An algo that adjusts participation when market volume is above or below average is responding to real conditions rather than following a clock.

Momentum signals add another layer. A genuinely adaptive algo modulates its aggression based on whether the market is moving for or against the order, rather than treating every interval the same. Fill rates offer a real-time feedback loop. If the algo is learning from its own recent success or failure and adjusting, that's a sign of genuine intelligence.

The buy-side can test this directly: run similar order profiles through different strategies and compare the behaviour, not just the outcomes. The patterns in how the algo trades often reveal more than the final number.

How can firms build a feedback loop where historical execution data meaningfully improves future decisions, rather than just sitting in a report?

The most valuable data any firm has is its own execution history. The question is whether it flows back into the decision process or simply accumulates in reports that nobody reopens.

Clients that can leverage the depth of data available from their liquidity providers have a real opportunity here. The raw material exists - detailed fill data, venue-level performance, timing analysis, cost attribution - but it needs to be systematically mined rather than passively received.

At J.P. Morgan, our FX Alternatives framework takes this a step further. It allows us to backtest client orders at different speed configurations, leveraging historical data and adjusting for market impact to produce what-if scenarios. Rather than debating in the

abstract whether a faster or slower execution would have been better, the framework provides a data-driven answer: here's what the cost profile would have looked like under different approaches, given the conditions that actually prevailed.

Provider assessment follows naturally. Comparing performance across providers for similar orders, controlling for conditions, gives a more objective basis for allocation than reputation or relationship alone. And automation helps -- systematic tagging, classification, and comparison can surface insights without manual effort. The firms that build this discipline create a compounding advantage over time.

KNOWING WHAT WORKED

Arrival price remains the most common benchmark. What are its limitations, and what additional dimensions give a fuller picture of execution quality?

Arrival price has earned its place as the default benchmark, but for a richer understanding of execution quality, it's worth considering what it captures -- and what it misses.

Arrival price measures from a single point in time, which may or may not represent the true decision point. For orders that take hours to execute, it conflates market drift with execution quality -- the price may have moved significantly for reasons entirely unrelated to the execution itself, and disentangling the two is difficult.

Additional dimensions worth examining include implementation shortfall decomposition, which separates delay cost, market impact, and timing luck into distinct components. Spread capture -- whether the algo earned or paid the spread -- reveals how effectively it used passive strategies. Volume participation shows how visible the order was in the market. And signalling cost measures whether the market moved adversely during execution in a way that suggests information leakage.

Peer comparison adds further context: how does this execution compare to similar orders under similar conditions? No single metric tells the whole story. A dashboard view that brings these dimensions together gives more useful insight than any one number in isolation.

How should TCA evolve to become a forward-looking decision tool rather than a backward-looking compliance exercise?

TCA started as a compliance and reporting requirement. The more interesting question is how it becomes a decision-making tool that actively improves future execution.

Backward-looking TCA answers "how did we do?" -- necessary but insufficient. Forward-looking TCA answers "what should we do differently?" -- and that's where the value compounds.

Segmentation is key. When you slice performance by pair, size bucket, time of day, strategy, and provider, patterns



We're moving toward real-time performance monitoring during execution

emerge that aggregate numbers hide. A firm might discover that its preferred strategy consistently underperforms for a specific pair during certain sessions, or that one provider delivers better outcomes for larger orders while another excels at smaller, more urgent flow.

This enables genuine hypothesis testing: "we believe a participation strategy works better than TWAP for EUR/USD orders over \$100M during London morning" -- does the data support it? If it does, that insight feeds back into selection frameworks and parameter defaults. If it doesn't, the assumption gets revised.

The maturity curve runs from reporting to analysis to insight to action. Most firms are somewhere in the first two stages. The competitive advantage lies in reaching the last two.

Looking ahead, how do you see the relationship between data, algo intelligence, and performance measurement evolving over the next few years?

The intersection of data, intelligence, and measurement is where the most interesting developments in FX algo execution are likely to emerge.

More data will become available -- deeper venue analytics, broader market microstructure signals, and richer internalisation metrics. Algo intelligence will continue to improve, with better real-time adaptation and

more sophisticated cost models that learn from an ever-growing dataset of execution outcomes.

While fully AI-driven execution remains ahead of us, AI is already being deployed in the areas surrounding algo execution -- integration with trading workflows, smarter order staging, and streaming pre-trade analytics that surface relevant context directly to execution desks. These tools help traders make more informed decisions about strategy selection and parameterisation without removing the human from the process. Over time, this kind of AI-assisted workflow is likely to become the norm rather than the exception.

Perhaps most significantly, measurement will become more integrated. Rather than TCA being a purely post-trade exercise, we're moving toward real-time performance monitoring during execution -giving both the algo and the human overseeing it the ability to assess and adjust while the order is still live.

The convergence is the key theme: pre-trade, execution, and post-trade are becoming a continuous loop rather than separate phases. For buy-side firms, the ones that invest in understanding and using data across this entire loop will extract more value from their algo programmes than those who treat each stage in isolation.

In a market where everyone has access to similar tools, the edge increasingly lies in how well you use them.

Lessons from low-impact FX execution

Mathijs Peeters, Chief Revenue Officer at Siege FX outlines how buy-side algos and bank e-FX desks can maximise outcomes from anonymous dark liquidity.



The core idea I want to talk about is that low-impact execution is not just where you trade. It is how you sequence liquidity, preserve anonymity, and manage signalling. After several years observing execution activity in the Siege MidPool, one conclusion has become clear: execution outcomes are determined less by venue access and more by the way users interact with it. Low-impact execution depends on using the right liquidity in the right sequence.



Resting time is not inactivity. It is an active execution decision

NOT ALL MID-POOLS ARE ANONYMOUS EVEN WHEN THEY ARE “DARK”

The term “mid-pool” is often used as if all such venues are similar. They are not. A mid-match is not automatically a low-impact match. Even when the order book is dark, if activity contributes to price formation or if identities can be inferred after execution the value of dark liquidity is compromised.

It is useful to distinguish between mid-pools in general (including dark ones) and a real anonymous MidPool, which combines three ingredients: a true dark pool of resting firm orders, a non-price forming mid-rate and participant anonymity before, during and after the trade. Siege's MidPool was built around those principles. But even the right venue design cannot fully compensate for a workflow that creates signaling elsewhere.

LESSON ONE: THINK IN TERMS OF WORKFLOW, NOT VENUE ACCESS

To make the most of accessing Siege MidPool when looking to improve execution quality, users must treat dark matching as part of a workflow, not as a standalone destination.

Their objective is not simply to obtain a match. It is to transfer risk without signalling it to the wider market. Sequence therefore matters. If a trader first shows interest in visible or semi-visible liquidity sources and only later comes to MidPool, some market impact has already been created.

The better question is often not “Where can I get a match?” but “In what order

should I expose this risk to minimize the information footprint of the entire risk amount?”

LESSON TWO: AGENCY EXECUTION ALGOS SHOULD START LARGE AND STAY LARGE

Agency algos often handle large institutional orders, frequently USD 100 million or more. Their objective is not simply completion. It is completion with minimum implementation shortfall and market impact.

For this flow, the greatest benefit accessing Siege's MidPool usually comes from placing a meaningful portion of the order into the pool at the start of the execution. While that order rests, the algorithm can execute calmly elsewhere, while genuine offsetting interest has time to find the order quietly.

Dark matching is probabilistic. A trader cannot force a natural match to appear but can maximise the conditions under which it becomes possible. Resting meaningful liquidity for sufficient time improves that probability.

The key is order replenishment. When a dark match occurs, sophisticated users replace the matched exposure with more of the parent order. Each anonymous match reduces the amount traded in lit markets. Each reduction in lit activity reduces signalling.

A large parent order should not be treated as disconnected small clips. It should be treated as a continuous opportunity to substitute visible execution with anonymous mid-rate interaction.

LESSON THREE: E-FX RISK DESKS HAVE A DIFFERENT PROBLEM

Bank e-FX desks face a different challenge. Instead of working a USD 100 million client order over time, they may need to reduce USD 10 million of risk quickly and cheaply without disrupting the market. To de-risk quickly, many desks try to maximise matching opportunities. On the surface this seems rational. If several mid-pools might provide a match, why not place small clips in all of them?

The problem is that this can create the very impact the desk is trying to avoid. Posting identical USD 1 million clips into several mid-pools simultaneously may look low risk, but if a match occurs on a venue where identities become visible after execution, that information will cause reactions elsewhere.

The desk may believe it is diversifying liquidity access but in reality it is diversifying information leakage. Small clips can still create signalling. Market impact is caused not only by size but by observable behaviour. Repeated small clips across venues can reveal direction, urgency and intent. The lesson for e-FX desks is that “small” does not automatically mean “low impact”

LESSON FOUR: SEQUENCE MATTERS MORE THAN VENUE COUNT

When evaluating low-impact execution, the order in which liquidity sources are accessed is often more important than the total number of venues being used. The first stage should be internalisation. The cleanest trade is usually the trade that never reaches the market. If risk can be offset internally without an external signal, that is normally the best outcome.

The second stage should be the anonymous mid-pool. One useful way to think about Siege's MidPool is as a multi-bank internaliser: participants seek natural offsetting interest from each other while interacting anonymously at a non-price forming mid-rate.

Only then should risk move towards venues more likely to generate signalling, price formation or post-trade leakage. Other mid-pools, bilateral



Market impact is caused not only by size but by observable behaviour

platforms and lit markets still have a role, but they should sit in the right place in the workflow.

Venue count is a poor proxy for execution quality. Three venues in the right order can be better than ten venues in the wrong order.

LESSON FIVE: LOW EXECUTION FEE IS NOT NECESSARILY LESS COSTLY

Low explicit fees can be misleading. Some bilateral mechanisms may charge only USD 1 or 2 per million and may occasionally allow participants to trade slightly better than mid through skewed pricing.

But explicit cost is only one part of execution cost. If a cheap venue creates information leakage that worsens the remaining execution, the apparent saving disappears quickly. A trader may save a few dollars per million on fees but lose far more if the residual order trades worse because the workflow created signalling.

The cheapest venue is not the venue with the lowest fee schedule. It is the venue that delivers the lowest total implementation cost after market impact, signalling and residual execution quality are considered.

LESSON SIX: PATIENCE REMAINS AN UNDERAPPRECIATED SOURCE OF ALPHA

Low-impact execution rewards patience. Resting time is not inactivity. It is an active execution decision.

For agency algos, if the trader has already decided to execute over time, there is limited opportunity cost in allowing a meaningful clip to rest anonymously at mid. The order is live, but it is not signalling.

For e-FX desks, patience is harder because the instinct is to reduce risk immediately. But the better question is not always “How quickly can I get rid of this risk?” It is “How much can I transfer without creating unnecessary market impact?” Urgency should be intentional, not automatic.

CONCLUSION

Future gains in execution quality will not come simply from adding more venues. They will come from better interaction with existing platforms.

To “Move Risk, Not Markets” the most successful agency algos maximize time spent seeking natural offsetting interest and replenish after successful matches. The most successful eFX desks sequence carefully: internalisation first, Siege MidPool second, replenish when matched and only use more visible liquidity sources afterwards.

The desire for lower fees can lead to much higher execution costs if measured properly across the full parent order measuring market movement vs fee per million USD.

Low-cost execution is not just about where you trade. It is about how you trade.

Trading Technologies sees FX algo benefits from EMS consolidation

Tomo Tokuyama, EVP, Managing Director, FX of TT, shares his response to the key findings in a recent study and the potential impact of EMS consolidation for FX algo use.

In a new survey looking at how buy-side firms approach consolidating FX with other asset classes, particularly listed derivatives, expanded algo capabilities was cited by 46 per cent of respondents as one of the top three most important benefits that would be achieved from consolidating EMS platforms.

The findings are detailed in a new white paper, *Bringing in FX: EMS Consolidation in a Complex Trading Environment*, created by Acuiti on behalf of Trading Technologies International (TT). Based on a survey of hedge funds, proprietary trading firms and asset managers, expanded algo capabilities was the third most cited expected benefit from a system consolidation, after a real-time view of risk (69%) and improved execution quality (52%).

According to the study, a trading operation that is able to bring greater standardisation across systems

also stands to improve algorithmic capabilities across the asset classes they trade, which the report adds is “an advantage for firms seeking to improve spreading tools for strategies like the basis trade”.

The research follows the recent expansion of the TT FX service, including the release of new functionality allowing traders to execute FX the way they trade futures using the same tools, screens and algo workflow, with bank algos now accessible from the same dropdown as futures algos.

Tomo Tokuyama, EVP, Managing Director, FX at TT, explains this addition of bank FX algos directly to TT FX has proven to be a significant selling point of the platform. “Multi-asset traders accessing bank execution algos for futures through TT can also access bank FX algos through the same platform and a consistent workflow,” he adds. “We



are also seeing growing interest from clients looking to build their own FX algos, which they can do using TT’s ADL (Algo Design Lab) product.”

Tokuyama adds that while the findings of the study were not surprising, they reflected many of the conversations around the benefits of EMS consolidation while helping to address some of the perceptions which firms still hold around switching to a more cost-effective system.

He explains that client feedback to the expansion of TT FX has been positive, with clients noting that it makes sense for users of TT for futures and options to use the same workflows to launch their FX algos as well. “The real value is giving traders a consistent execution workflow across asset classes, reducing fragmentation while expanding what they can automate,” Tokuyama says.



Got Questions?



Suggested Searches



what products do FXSpotStream support?



who was the fastest growing FX service in 2024?



does FXSpotStream charge clients?



is it true that FXSpotStream has 18 Tier 1 LPs available?



which vendor partners does FXSpotStream support?



does FXSpotStream offer a GUI?



does FSS have an Ultra Low Latency Network?



which algos do FXSpotStream support?

We’ve Got Answers

www.fxspotstream.com



FXSpotStream provides a multibank FX streaming and a matching service supporting steaming of pricing in FX Spot, Forwards, Swaps, NDF/NDS and Precious Metals Spot and Swaps. Clients can access a GUI or single API from co-location sites in New York, London and Tokyo and have the ability to communicate with all Liquidity Providing banks connected to the FXSpotStream Service.

FXAlgoNews is published 4 times a year in February, May, August and November

YOUR DETAILS (All information is kept confidential)

Name Job title

Company

Address

City Post/Zip Code

Country

Telephone Fax

E:mail

PRICE (tick the relevant box) *Your subscription will commence with the next issue unless otherwise requested.*

- ☐ Single annual digital subscription - £100
- ☐ Single annual digital + hard copy subscription (UK only) - £175
- ☐ Single annual digital + hard copy subscription (Outside UK) - £200

QUANTITY REQUIRED Discounts are offered for 3 or more subscriptions.

Please Fax a separate list of names & addresses. **TOTAL PAYMENT**

PAYMENT OPTIONS (please tick where relevant and follow the instructions)

CREDIT/DEBIT CARD ☐

Mastercard ☐ Visa ☐ American Express ☐ Debit card ☐

Card Number:

Expiry Date: /

CHEQUE ☐ make payable to **SJB Media Ltd** (please print name on the back of the cheque)

Post to Subs Dept: **SJB Media Ltd., Suite 153, 3 Edgar Buildings, George St., Bath, BA1 2FJ, UK**

BANK TRANSFER ☐ Please contact us on +44 1736 740 130 to arrange this.

NB. Payment must be received before subscription can commence. Please send me an invoice for my records ☐

This subscription is for me personally ☐ a company ☐

VAT number (Required for European Community based companies only)

Signed: Date:


Name in capitals:

Please e-mail this form back to:
Charles.Jago@fxalgonews.com
THANK YOU



BOOK OF THE MONTH

Python for Algorithmic Trading Cookbook



packtpub.com/en-us/product/python-for-algorithmic-trading-cookbook-9781806662036



FOR THE DIARY

QuantMinds: Strategies that set the standard

November 16-19, 2026: London



informaconnect.com/quantminds-international/?utm_source=chatgpt.com

Charles Jago
 Editor
 charles.jago@fxalgonews.com
 +44 1736 740 130

Nicola Tavendale
 News editor
 nicola@ntavendale.com
 +44 1736 740 130

Susan Rennie
 Managing Editor
 susie.rennie@fxalgonews.com
 +44 1208 821 802

Charles Harris
 Advertising sales
 charles.harris@fxalgonews.com
 +44 1736 740 130

David Fielder
 Subscriptions manager
 david.fielder@fxalgonews.com
 +44 1736 740 130

Tim Hendy
 Digital & Web services
 tim@thstudio.co.uk
 + 44 1209 217168

Matt Sanwell
 Design & Origination
 matt@designunltd.co.uk
 +44 7515 355960

Larry Levy
 Photography
 larrydlevy@gmail.com

Michael Best
 Events manager
 michael.best@fxalgonews.com
 +44 1736 740 130

SJB Media Ltd
 Suite 153, 3 Edgar Buildings
 George Street, Bath, BA1 2FJ
 United Kingdom
 Tel: + 44 (0)1208 82 18 02 (switchboard)
 Tel: + 44 (0)1736 74 01 30 (Sales & editorial)
 Fax: + 44 (0)1208 82 18 03

Printed by Headland Printers
 Published quarterly. ISSN 2056-9750




Although every effort has been made to ensure the accuracy of the information contained in this publication the publishers can accept no liabilities for inaccuracies that may appear. The views expressed in this publication are not necessarily those of the publisher. Please note, the publishers do not endorse or recommend any specific website featured in this newsletter. Readers are advised to check carefully that any website offering a specific FX trading product and service complies with all required regulatory conditions and obligations. The entire contents of FXALGONEWS are protected by copyright and all rights are reserved.



WEBSITE OF THE MONTH

FX Market Profile powered by Quant Analytics



cmegroup.com/trading/fx/cme-fx-market-profile-tool.html?itm_content=fx_2_mkt_profile_tool_122020



EDUCATION & TRAINING

Oxford Algorithmic Trading Programme

02 September 2026



sbs.ox.ac.uk/programmes/online-learning/oxford-algorithmic-trading-programme



Your eFX ALGO partner



Highly customisable Algo Strategies
to suit your requirements



Comprehensive market liquidity access
combined with CACIB internalisation



Fully fledged data set and analytics,
including BestX and now TradeFeedr



A special partnership with XTX
to add alpha to your trading style:
XTX Execution Algo

Contact your dedicated FX Sales to find out more

www.ca-cib.com



CRÉDIT AGRICOLE
CORPORATE & INVESTMENT BANK