FX Trading: Broker Panel

Author

Damien Lodge
Head of Currency Trading

When assessing the strategy of broker panel selection, (the counterparties selected to compete for our business), it is essential to understand the micro-structure of a market and how liquidity is distributed by liquidity providers and accessed by clients. The market structure is nuanced and due to the interactions of market participants, a larger panel of brokers does not automatically deliver better results. Execution strategy is complex and takes into account many dynamic factors that evolve in real time and we believe that best results are achieved by careful selection of liquidity providers, innovative technology and strong relationships with counterparties.

JPMAM strategic solution

In 2013, J.P. Morgan Asset Management (JPMAM) Global Fixed Income Currency & Commodities (GFICC) Team made the strategic decision to build a proprietary connectivity solution to access liquidity after identifying shortcomings within the market structure. We found that real liquidity was confined to a smaller subset of tier 1 market makers and that liquidity was simply being recycled by other participants. JPMAM is a fiduciary and executes all foreign exchange (FX) business on an agency basis, meaning that all trading cost savings are passed on for the direct benefit of our clients. We recognised that we could drive pricing lower by removing third party brokerage costs and building stronger relationships with our strategic partners to deliver the best client outcomes.

To do this, we built proprietary, direct, bilateral connectivity to a panel of brokers who were selected for their ability to provide liquidity, their technology expertise and a willingness to be partners in the solution. In return, JPMAM does not charge our liquidity providers a brokerage fee, enforcing the cleanest and most transparent price possible.

Our FX trading panel consists of 10 liquidity providers, who we refer to as our 'standard counterparty list' and contains primary liquidity providers and key strategic relationships. In addition, we have bilateral connections to another seven liquidity providers who we can also use to support client requests and access specialist liquidity, if required.

Counterparty performance is reviewed regularly and modifications can be made to our standard list to reflect changes in liquidity providers' capabilities. Although we have a larger panel available, we will curate the size of the panel for price requests based on our analysis to achieve the best outcome for our clients.

J.P.Morgan
ASSET MANAGEMENT

JPMAM's key advantages

- JPMAM executes on an agency basis, so all efficiency gains directly benefit the clients.

- Optimal pricing that is clean without embedded charges.

- As one of the largest asset managers, JPMAM is able to command the best price tiering (most competitive pricing) based on our global presence across asset classes. This is particularly useful for balance sheet intensive FX trades such as forwards and swaps, which have the largest impact on efficiently managing multi-currency portfolios. Price tiering is determined by a liquidity provider's willingness to allocate its balance sheet to a client, which reflects the profitability of the business. At JPMAM, our global diversity, size, relationship and total contribution to counterparty return on equity (ROE) ensures we receive the best tier of pricing.

- The JPMAM execution protocol is to trade in competition using Request for Quotes (RFQs) to broken dates (non-standard market date conventions for both spot and forward rates) without disclosing the buy or sell intention. This process leverages our liquidity panel expertise in aggregating the fragmented market structure and providing the best available price.

- We maintain ownership of our trading data which supports our internal trading analytics and best execution processes.

- We have bilateral trading agreement documentation (ISDA's) and are not reliant on credit intermediation which reduces other explicit costs incurred in models such as prime brokerage.

- We have the flexibility to build bespoke solutions for our workflows, such as post trade account allocation, benchmark and algo execution protocols.

Market structure

The primary or interbank market consists of wholesale market makers, who will make markets on the traditional primary venues such as EBS, Reuters (Refinitiv) and CME which are known as Central Limit Order Books (CLOB). A CLOB is an exchange style execution method where rules govern the matching of all bids and offers according to price and time priority.

Price is governed by a minimum tic increment and notional size. It allows all users to trade with each other as opposed to being intermediated by a dealer or third party. In addition, pricing is also distributed via secondary markets or Electronic Communication Networks (ECN's) such as Hotspot, Currenex, Fastmatch as examples. The difference being that secondary markets do not operate as CLOBs, rather they are 'managed liquidity pools'. A managed liquidity pool is a venue (intermediary) where liquidity is curated to match takers to other takers and makers without specific rules which results in participants views of liquidity all being different. This is the first challenge, "all liquidity is not equal".

It is important to highlight that the market is far smaller than is generally perceived. If we look at a study of liquidity provider match codes on primary venues, it is clear that the majority of the market is made up of a small number of participants. The best way to view primary market liquidity is a spine from which the secondary and other markets are priced and hence acts as the price boundary.

Market participant contribution to primary market liquidity

(Exhibit 1) shows the contribution of liquidity providers (anonymised) to market liquidity based on volume and trade count (left hand side). Average trade size is represented on the right hand side with 100% being $1 million, which is the minimum increment. What we are highlighting is that the 80% of volume blue line intersects the black cumulative volume line at approximately five liquidity providers (20% of the providers surveyed). Hence 80% of volume is provided by 20% of liquidity providers.

Exhibit 1: Primary Markets
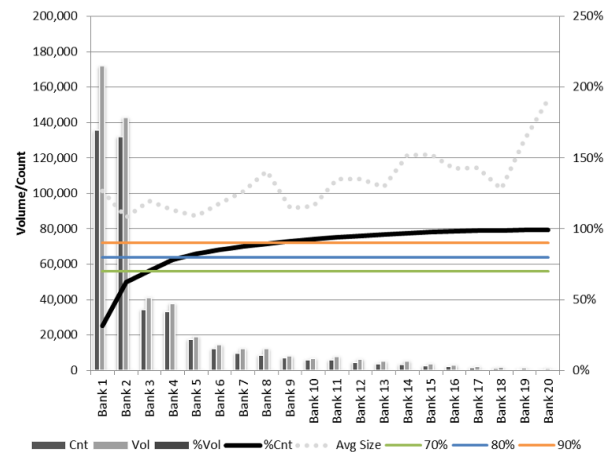
Source: Goldman Sachs, market study over 2018/19 data showing banks who are providing the liquidity when the surveying bank was a liquidity taker (aggressed the price).

Liquidity is complicated

Consumers of this liquidity, which includes smaller, non-market making banks, will redistribute the pricing via secondary markets. Most secondary markets are just a primary liquidity price replicated in a managed pool. The result being a multiplier effect on liquidity distributed to multiple platforms. For example, secondary venues receive different types of liquidity feeds from which they curate different pools.

Typically, a liquidity provider would provide different liquidity streams, as an example, full amount, firm and not firm, which can then be broken into a RFQ or a request for stream (RFS). These streams are then replicated across different geographical server co-locations. If we take an example to highlight how many liquidity streams are being priced off the primary market; let's assume there are 10 liquidity providers distributing prices to 10 secondary venues, each with three types of liquidity streams, two request protocols and two geographical regions. This implies we have 10 * 10 * 3 * 2 * 2 = 1200 liquidity streams hinged off the primary venue boundary price.

Types of liquidity streams impact how the secondary market pools are curated. For example, a non-firm stream means the liquidity provider has the option of last look to protect against adverse price movement, a firm price is similar to that of exchange liquidity, hence different streams will have different spreads to reflect the risk involved. **This redistribution of liquidity is one of the major challenges to market impact** as market makers lose control of who the end client is and hence the behaviour of the price taker.

(**Exhibit 2**) shows a representation of the connectivity. In green we have an additional facility which is a mid-pool and is most easily described as an equity-like dark pool. The mid-pool allows participants to post an interest to match at a mid-price as defined by the primary venue CLOBs. The three primary venues in blue and in pink are some examples of secondary venue connections. The three on far right are secondary venues but have slightly differing rules in how they curate their liquidity.

Exhibit 2: Example of market connectivity

For illustrative purposes only.
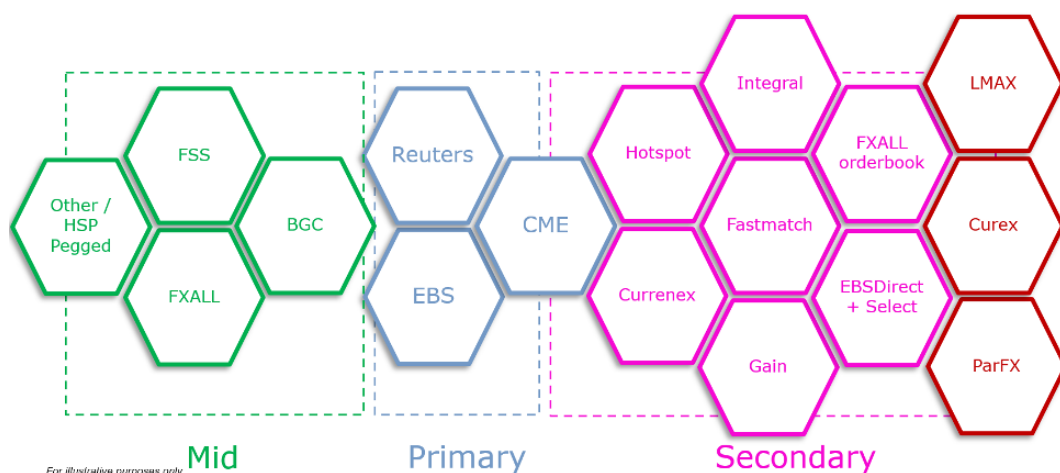
Source: Goldman Sachs, 2023.

Finally, the larger liquidity providers will also run large e-books which manage and internalise direct flow from their client's franchises. In many cases, the size of these books in volume terms are larger than some of the dominant secondary venues. This is an important factor when selecting counterparties for your panel as internalised risk is less transparent to the broader market and will likely have less market impact.

In conclusion, primary/interbank market liquidity comes from a small number of providers and is hinged on a price boundary which comes from primary CLOB venues. A liquidity provider's ability to provide the best pricing largely depends on their interaction with the secondary market liquidity pools, and the size of their own e-book.

How do clients access this liquidity? – The third party venue

The challenge for the liquidity providers and banks is distributing their pricing and penetrating end user client bases such as asset managers, hedge funds and corporate clients. The majority of clients do not have the technology expertise or the capacity to be able to access multiple sources of liquidity and hence traditionally used the single dealer platforms of their brokers. As the focus on best execution became sharper, the opportunity arose for new venues to enter the market and provide a solution that allowed client connectivity to multiple liquidity providers, but also facilitated a simplified avenue for liquidity providers to reach a larger client base.

Essentially, these platforms sit between the liquidity providers aggregated liquidity and the end client. These platforms include FXAll, FXConnect, Flextrade, 360T and Bid FX along with numerous others. The single stand out features of the platforms, now referred to as Multi-dealer Trading Facilities or MTFs, is that clients need only a single connection to access multiple liquidity providers, and similarly banks' connectivity complexity to clients is reduced. This relationship is a brokerage model, where the platform onboards clients and provides the service for free, but the liquidity provider is charged a fee by the platform to provide pricing to the end user.

This has multiple implications; it is a great model for the platform venues, because client onboarding is an easy sell as there is no charge, and as the onboarded client base grew so did the platforms' leverage over the liquidity providers.

As a result, client pricing has become controlled by the venue charging model, as the liquidity provider's pricing is tiered depending on the venue brokerage schedule which impacts transaction costs for the end client.

In conclusion, the advent of the third-party venue has added another layer of complexity to accessing liquidity, as liquidity providers tier pricing based on brokerage fees, which are not transparent, and results in costs being passed onto the client. Direct relationships with counterparties are diminished and controlling market impact becomes more difficult.

How many liquidity providers do we need?

We have established that:

- The 80-20 rule (80% of the volume is provided by 20% of liquidity providers), applies to the primary market liquidity (Exhibit 1).

- Secondary market liquidity is essentially derived from the spine of the primary CLOB boundaries.

- Due to the complexity and number of liquidity connections, market impact is difficult to control.

So, what is the optimum number of brokers required in an RFQ to minimise spread cost and market impact? This is a complicated question as there are many variables to consider within the data set. To support our belief that panel size and market impact are inversely correlated, and that an intersection offers an optimal solution, we refer to research conducted by BestX, JPMAM's third party Transaction Cost Analysis (TCA) provider. The total asset under management for BestX® clients is approximately $47 trillion as of 1 February 2023 and the opt-in pool tracks over $110 trillion of trading volume which BestX uses to derive trading behaviour studies.

"Putting multiple counterparties in competition is rationally perceived as a way to reduce transaction costs. However, the process of establishing and maintaining relationships with dealers is costly, and thus it is not always economical to have more counterparties. Moreover, a higher panel size may induce more information leakage. Therefore, there is a balance between more competition and potential information leakage to arrive at an optimal panel size for a given trade." — BestX.

Relationship between panel size, cost and market impact

(Exhibit 3) shows a side-by-side comparison between the effect of cost (left chart) and market impact (right chart), against the number of brokers included. The size of the bar reflects the range and hence confidence interval. The below is transaction cost by panel size for G10 trades, in 0.01 basis points.

Exhibit 3: The number of brokers impact on cost

(a) BestX Total Cost by panel size

(b) 30 second revaluation by panel size
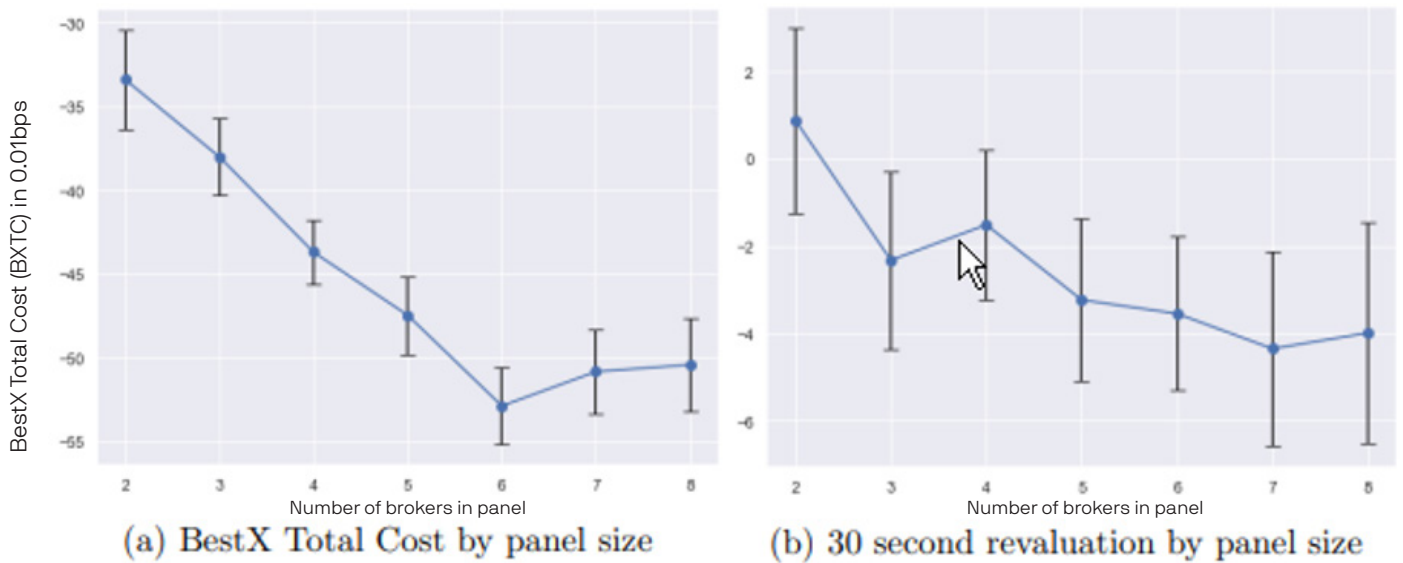
Source: BestX, as of 19 January 2022.

As noted by BestX and illustrated in the charts above "When the panel size is less than six, more competition is lowering the overall transaction cost despite inducing more information leakage. However, where panel size exceeds six, the overall transaction cost is not reduced by a larger panel size, and information leakage is still amplified."

The results are intuitive. Transaction cost savings diminish as the panel size increases, yet market impact will continue to increase. Based on JPMAM trade data and our trading experience, we would hypothesise that the continuous increase in market impact results from the addition of non-core liquidity providers to the panel. As discussed earlier, we established that there are a limited number of core liquidity providers, and a large portion of liquidity is actually recycled via secondary venues. Assuming within the BestX data pool that the participants have access to the core five to seven major liquidity providers, adding additional non-core liquidity providers to the panel would imply a higher probability that the liquidity is being recycled, hence having a larger impact. As we will discuss in the next section, this impact is important when managing larger trade sizes and measuring execution performance from arrival time of the first tranche.

We have conducted similar analysis ourselves but focused on the participants in the panel (**Exhibit 4**). We measure the winning spread (the bid-offer spread

shown by the winning broker) versus the theoretical spread (the combined best bid and best offer). We can construct this from our competing quote data collected from the non-directional RFQ.

The difference between the winning spread and the theoretical spread represents the potential saving of having competition within the panel. We observed that there is no improvement in potential saving (a wider gap between winning and theoretical) when comparing a panel of our top five liquidity providers versus all available brokers. This indicates non-core liquidity providers add no additional benefit to the panel.

This is in line with expectations, as all participants have access to the primary market price but due to the way the secondary venue liquidity pools are managed, they may not have equivalent access to secondary liquidity. Even ignoring this concept of "all liquidity is not equal", we would still expect our top liquidity providers to outperform as they have large e-books – JPMAM receives top tier pricing – and due to our execution strategy, liquidity providers are more comfortable in showing price skew to our requests. We benefit from skew because our liquidity providers know we don't redistribute pricing and as a result are willing to show axed prices which reflect the risk they hold internally on e-books.

JPMAM analysis of spread improvement dependent on panel composition (EURUSD)
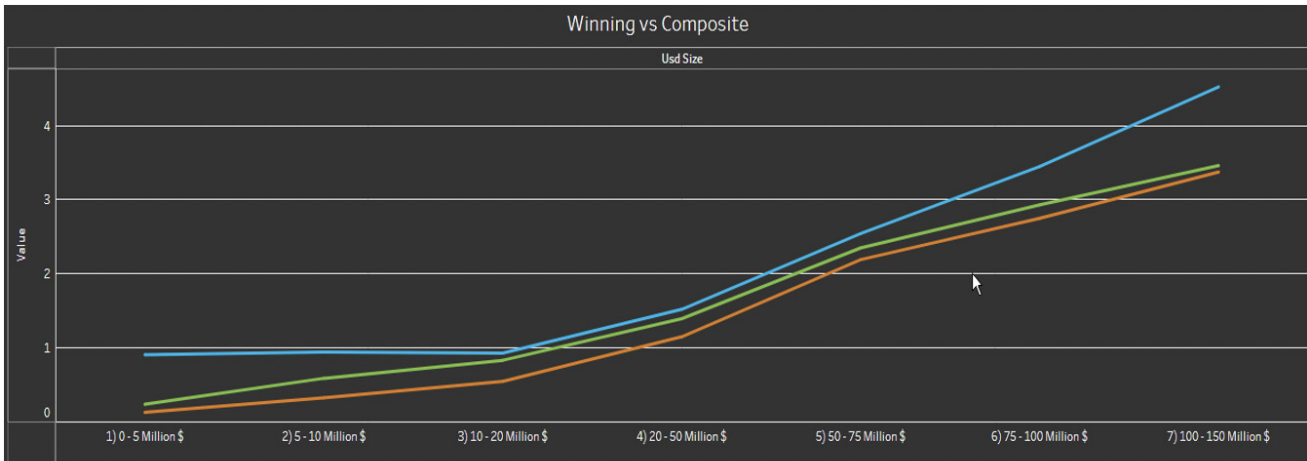
Exhibit 4: a) Spread saving for top 5 liquidity providers

Orange - Theoretical spread: Green - Winning spread: Blue - Static spread matrix supplied by brokers

Usd Size	Ccy Pair	Min. Bp Winning Spread	Avg. Bp Winning Spread	Max. Bp Winning Spread	Min. Winning Spread	Avg. Winning Spread	Max. Winning Spread	Avg. Composite Spread(in Bp)	Avg. Composite Spread(actual)	Usd Net(in MM)	Number of Records
1) 0 - 5 Million $	EUR/USD	0.00000000	0.26623409	24.47528310	0.00000000	0.00002798	0.00273000	0.14229698	0.00001495	10,167	16,863
2) 5 - 10 Million $	EUR/USD	0.10214573	0.58257422	3.89632144	0.00001000	0.00006107	0.00042800	0.32078349	0.00003360	9,008	1,373
3) 10 - 20 Million $	EUR/USD	0.13104440	0.80560721	3.99764034	0.00001460	0.00008506	0.00040260	0.54250468	0.00005725	11,937	886
4) 20 - 50 Million $	EUR/USD	0.33658759	1.35950794	5.81487125	0.00003720	0.00014384	0.00062070	1.12465178	0.00011901	15,750	514
5) 50 - 75 Million $	EUR/USD	0.83594858	2.32298192	11.37202301	0.00009450	0.00024512	0.00113100	2.14730010	0.00022646	8,237	136
6) 75 - 100 Million $	EUR/USD	1.37497787	2.89228921	7.40171378	0.00015350	0.00030719	0.00078000	2.70806364	0.00028762	5,733	67
7) 100 - 150 Million $	EUR/USD	1.52056713	3.59973046	10.85620724	0.00017220	0.00038521	0.00115900	3.57761096	0.00038292	8,324	69

Exhibit 4: b) Spread saving for top 15 liquidity providers

Orange - Theoretical spread: Green - Winning spread: Blue - Static spread matrix supplied by brokers

Usd Size	Ccy Pair	Min. Bp Winning Spread	Avg. Bp Winning Spread	Max. Bp Winning Spread	Min. Winning Spread	Avg. Winning Spread	Max. Winning Spread	Avg. Composite Spread(in Bp)	Avg. Composite Spread(actual)	Usd Net(in MM)	Number of Records
1) 0 - 5 Million $	EUR/USD	0.00000000	0.23239834	24.47528310	0.00000000	0.00002445	0.00273000	0.12201449	0.00001285	14,253	26,084
2) 5 - 10 Million $	EUR/USD	0.08793992	0.58321340	3.89632144	0.00001000	0.00006144	0.00042800	0.32055588	0.00003377	13,553	2,058
3) 10 - 20 Million $	EUR/USD	0.08860692	0.82646672	4.50773675	0.00001000	0.00008761	0.00048020	0.54221654	0.00005746	16,913	1,262
4) 20 - 50 Million $	EUR/USD	0.33658759	1.39396212	5.81487125	0.00003720	0.00014809	0.00062070	1.15110620	0.00012233	23,668	780
5) 50 - 75 Million $	EUR/USD	0.83594858	2.34171486	11.37202301	0.00009450	0.00024785	0.00113100	2.18409724	0.00023103	9,876	163
6) 75 - 100 Million $	EUR/USD	1.37497787	2.92187393	7.40171378	0.00015350	0.00031104	0.00078000	2.74187843	0.00029185	6,336	74
7) 100 - 150 Million $	EUR/USD	1.52056713	3.66007663	10.85620724	0.00017220	0.00038958	0.00115900	3.57452075	0.00038072	9,034	75

Source: J.P. Morgan Asset Management DNA trading data tool, data as of 01/01/2022 to 08/02/2023.

We measure skew in our performance analysis as an indicator of the value a counterparty adds to our panel. We do this by calculating the USD value and basis point difference between the winning price of a counterparty against the next best price within the quote. Our results show that our core liquidity providers have the largest contribution, likely due to the size of their e-book and willingness to show skew, in comparison to second tier providers who are more likely to be quoting and clearing in small sizes, priced off of the CLOB spine boundary.

In conclusion, it is not only panel size but also the quality of the liquidity providers you include in the panel. The quality of the liquidity is largely derived from the mutual relationship which drives how you access the market and hence the tiered price you receive. Price tiering is critical within your panel and maintaining this tiering is paramount.

Transaction cost analysis – oranges and lemons when making comparisons?

Keeping in mind the complexity of the market structure and the role price and liquidity distribution plays on market impact, the process of TCA is not necessarily a like-for-like comparison between asset managers. To make a true comparison, it is fundamental that the outcome being measured is clearly defined. For example, are we measuring:

● the spread from arrival mid of a trade, or

● the cost of the total execution based on the arrival time for the first tranche?

An excellent real world example is the TCA analysis of an algorithmic trade (algo). An algo trade is characterised as a large notional value being broken down and transacted as multiple smaller tranches (normally systematically, targeting a specific result).

Explicitly an algo execution is a trade-off between reducing the bid to offer spread (implied cost) of each individual tranche by reducing size versus an increase in risk of adverse price movement resulting from market impact, time delta and unpredictable forward volatility.

Analysis of our own trade universe highlights that the spread saving by size reduction is consistent and predictable, however, the total cost of the trade based on the arrival price at the beginning of the algo is less consistent and vulnerable to unknown changes in market conditions.

Algorithmic trade execution highlighting market impact and cost comparisons

(Exhibit 5) is a graphical representation of an actual algo execution to sell a very large notional EURUSD trade. You can see the arrival price as the blue line and is 1.100268 which is the starting point for measuring the total cost of the execution. The pale blue band represents the price boundary throughout the time period i.e. the bid/offer, and the coloured dots represent each individual execution slice. Visually you can see in terms of each slice the execution is within the bid offer and have been executed on very tight spread.

The overall cost, which is measured as 1.100268 starting price minus 1.099403 average price, is equal to -8.65bps. The broker also shows us performance vs a risk transfer price which is the price they would have shown us to trade the full amount in one trade. This price would have been 1.57bps better than what was achieved by the algo slices. This difference highlights the time delta and volatility risk. In this case, the decision to use the algo was discussed and taken in line with the investor's objectives.

**Exhibit 5: Screen shot of the live monitoring of an algo trade**
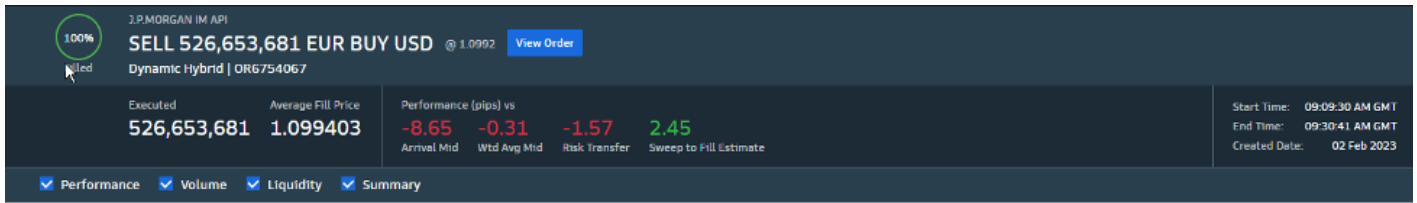
Source: J.P. Morgan Asset Management as of 01/01/22 – 06/02/2023.
Source: J.P. Morgan Asset Management DNA trading data tool, data as of 01/01/2022 to 08/02/2023.

However, using the BestX analysis (**Exhibit 6**), our spread to mid cost on the individual slices is 0.13bps. We can also benchmark on the individual slice performance vs expected cost (i.e. the spread saving on each individual slice executed within the algo). On this measure, we outperformed by 0.53bps for a USD saving of $30,000.

Exhibit 6: Third Party TCA vendor analysis of the spread to mid cost of each individual slice within the algo trade.
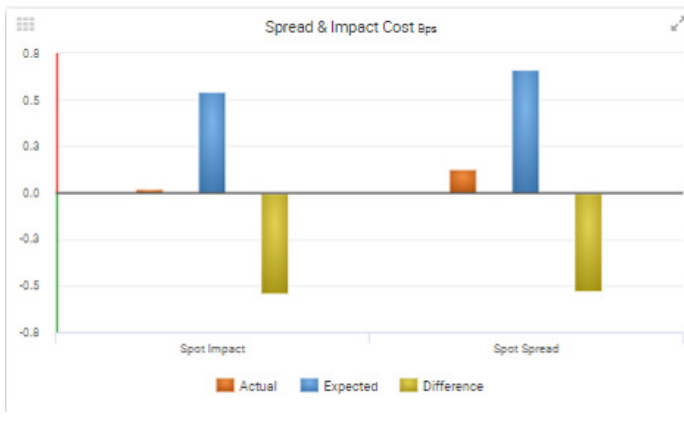
Source: BestX, as of 02/02/2023

This highlights that performance is dependent on understanding the measurement and desired outcome. This example can be extrapolated to comparing asset managers and their reported TCA costs. The TCA result can be highly distorted if Manager A reports a cost based on an average of small trade sizes vs Manager B who measures cost over the full execution of the trade. Manager A would have a better result in terms of spread cost, but Manager B would outperform based on the complete execution from arrival time. There are multiple other factors which can also bias this performance when looking at a simple measure based on spread to mid, such as size, timing, manager execution strategy and portfolio objectives.

A more detailed analysis highlights the importance of liquidity providers who are primary liquidity providers and have large e-books which allow them to control market impact. The blanked-out box in (**Exhibit 7**) represents the counterparty e-book and ability to internalise the flow. In this example, only 20.5% of liquidity was accessed via secondary market venues and 65% of the flow was matched passively either internally or on the primary market.

Exhibit 7: Analysis of venue where the slices were executed indicating the counterparty ability to internalise flow

VENUE NAME	MIC CODE	QUANTITY (EUR)	QUANTITY (USD)	% OF FILL	TAKING	PROVIDING
TOTAL		526,653,681.15	579,004,395.7	100.00%	27.19%	72.81%
∨ INTERNAL		172,503,681.15	189,665,332.5	32.75%	1.91%	30.85%
		171,670,000.00	188,748,741.76	32.60%	1.75%	30.85%
		833,681.15	916,590.74	0.16%	0.16%	0%
∨ PRIMARY		211,000,000.00	231,964,396.00	40.06%	4.75%	35.32%
EBS MARKET	EBSC	210,000,000.00	230,865,128.00	39.87%	4.56%	35.32%
REUTERS	RTSL	1,000,000.00	1,099,268.00	0.19%	0.19%	0%
∨ SECONDARY		108,150,000.00	118,896,507.2	20.54%	20.54%	0%
FASTMATCH	FAST	13,000,000.00	14,291,819.5	2.47%	2.47%	0%
HOTSPOT	HSFX	20,650,000.00	22,701,956.7	3.92%	3.92%	0%
FX SPOT STREAM	FSSA	12,000,000.00	13,192,811.00	2.28%	2.28%	0%
GAIN GTX	GTXS	60,500,000.00	66,510,874.00	11.49%	11.49%	0%
HOTSPOT NO LAST LOOK	HSFX	2,000,000.00	2,199,046.00	0.38%	0.38%	0%
∨ MIDPOOL		35,000,000.00	38,478,160.00	6.65%	0%	6.65%
MIDBOOK 4	MID4	10,000,000.00	10,994,019.00	1.90%	0%	1.90%
MIDBOOK 1	MID1	16,000,000.00	17,590,009.00	3.04%	0%	3.04%
MIDBOOK 2	MID2	9,000,000.00	9,894,132.00	1.71%	0%	1.71%

Source: Goldman Sachs Marquee – J.P. Morgan Asset Management live trade. Data for trade date 02/02/2023.

To conclude, we have shown the complexity of comparing manager TCA performance like-for-like without having a detailed understanding of the measurement process. In addition, it is important to highlight that the bid/offer spread is not an explicit cost to the portfolio and hence is less transparent. When comparing the performance of different managers, the most effective measure is the overall portfolio return as this will capture both the transaction costs as well as the added value generated by the investment process. Frictional execution costs are an element of overall performance.

Counterparty concentration

There is also a trade-off between optimal pricing and distribution of risk across counterparties. It is important to keep in mind the full impact on price of such a distribution of risk.

Consider the example where, as a manager, JPMAM is rolling a client's FX exposure, which is large and is likely to have market impact. It may be the case that JPMAM is mandated to distribute the risk across multiple liquidity providers as opposed to being able to transact the full amount, in competition, at a single price.

Assuming we allocate across five different liquidity providers, we would ask a price for 20% of the total notional in competition and the most economic price would win. The second tranche is then competed between the remaining four brokers and again the most economic price would win the trade. However, Broker 2 would now not be showing as competitive a price as in the first request and, in addition, the market is likely to have repriced to a new level as a result of the first execution. This process continues across the remaining brokers until the final trade is completed, each execution being at a wider spread and potentially a worse price.

We know the last execution will be at the worst spread, but we also need to consider the price delta from the first tranche. It is this price delta that is less transparent when conducting TCA on a trade-by-trade basis. In addition, there are other equally important but less measurable costs. Due to the multiple tranches causing a repricing in the market, this has a direct impact on the ability of the winner of the first trade to cover the risk.

There are multiple elements to consider:

- Liquidity providers are providing a price on a smaller parcel of risk, hence the price is not truly reflective of market risk;

- The winner of the first trade is likely to be a counterparty who is providing a better tiering of pricing, and this counterparty is being impacted the most by potential adverse repricing;

- This trading strategy has a marked impact on your relationship with the liquidity provider and is likely to lead to less competitive pricing in the future.

Conclusion

JPMAM, as the agent, aims to minimise execution costs and deliver the best outcomes for our clients. We have looked at the nuance in measuring execution cost, considering both price and impact, the complexity of primary and secondary market liquidity, and the importance of being able to access liquidity without incurring additional costs charged by third party venues.

Our strategic counterparty selection is the key to achieving the best execution outcomes and we have evidenced optimal panel size based on our own, and third-party analysis. Pricing is driven by our global relationships and premier tiering, facilitated by our proprietary OMS/EMS direct connectivity solutions.

All liquidity is not equal and JPMAM commands the highest quality liquidity.

J.P.Morgan
ASSET MANAGEMENT