

AGENCY ALGORITHMS AND ELECTRONIC LIQUIDITY PROVIDERS

AUGUST 2025



Liquidity & Market Structure Commentary
From XTX Markets' Distribution team

INTRODUCTION

Agency brokers dedicate a lot of effort and expertise^[1] to optimizing algorithms for their buy-side clients. This includes deciding when and how to interact with Single Dealer Platforms (SDPs) such as the one XTX Markets operates. This note will explore those interactions and share our experience of what works well in practice.

The objective is to give readers, many of whom may not be directly responsible for managing SDP interactions, some context on the main design choices facing an algorithm provider, as well as the trade-offs associated with each decision.

Heads of Trading at buy-side firms often ask us for our views on these topics. The following sections capture what we consider to be the critical points.

[1] In our experience, both in Europe and the US, there is a dedicated expert at the sell-side firm who will work with each liquidity provider to give feedback and try to optimize the interaction over time. Beyond having expertise, one of the big value-adds brokers can provide to individual buy-side clients is the huge amount of data they accumulate. As we will discuss later in the note, there are certain evaluations that require a sample size that may be beyond what is achievable for many individual buy-side firms.

SHOULD ONE TRADE WITH SDPS AT ALL WHEN USING BROKER ALGORITHMS?

SDPs are not one homogenous group. They can be markedly different in terms of origin story, pricing style, product set, firm culture and so on. Accordingly, it does not make sense to talk about SDPs 'en masse' and we would strongly advocate for performing analysis at the individual SDP-level based on objective execution data. For example, via a sophisticated broker, a trader may have sufficient control to opt into trading with SDP 1 and SDP 2 but not SDP 3.


Hitesh Mittal at BestX makes a [compelling case](#) for A/B testing the inclusion of SDPs at parent order level. We agree. In terms of A/B test implementations we would stress the earlier point on evaluating each SDP on its individual merits rather than as an averaged group of entirely different liquidity sources. Given parent order shortfall results are noisy, this type of parent order analysis is more likely to be conducted at broker level to obtain statistical significance.

Mark-outs are another useful tool. Ultimately everything will appear in the parent-level result but it can take a lot of time to perform those evaluations, as many orders are required before the noise yields to signal. Mark-outs take less time to converge and may provide clues meanwhile as to the nature of a liquidity provider. It is important to make like-for-like comparisons and care must be taken when performing this analysis^[2]. Commonly evaluators will focus on two metrics:

- Spread paid. One way to test if price improvement (to NBB0) was meaningful is to evaluate it not only at the time of trade but shortly afterwards. For example, if buying in a 10/20 market, a mid-fill of 15 is attractive as it has saved five units vs the offer. However, if one second later the market has become 05/15 that 'saving' looks a lot less impressive. A commonly used measure is therefore to mark fills to NBB0 mid at both time of trade and compared to a few seconds post-trade.

[2] It is critical when producing these comparisons that they are as like-for-like as possible. This means not simply comparing two venues naively as the results may prove misleading. It should mean evaluating the markouts of each venue on a 'control' subset of randomly routed flow. We discuss this concept in more detail in the 'scorecard and tied prices' section later in this note. Consider conceptually two SDPs that are identical but one is at the top of the waterfall and one at the end. The SDP that sits in first place will mechanically have more benign markouts, if naively compared with no control, so using markouts to justify its place at the top of the waterfall creates a problematic loop and a randomized control is necessary to objectively compare the characteristics of both liquidity providers.

- Post-trade impact. Market impact is a big contributor to the cost of executing larger orders. If you have 100,000 shares to buy it is not attractive to buy the first 1,000 cheaply yet find that the external market moves much higher while you still need to purchase the remaining 99,000 shares. Plotting the post-trade market impact characteristics of each venue, based on thousands of child fills, yields some insight into the nature of the liquidity available and the risk-warehousing characteristics of each liquidity provider.



This is a purely illustrative image. It depicts two liquidity providers, one of whom (A) has more observable post trade impact than the other (B). It is critical when producing these comparisons that the comparisons are as like-for-like as practically possible. This means not simply comparing two venues naively as the results may prove misleading. It should mean evaluating the markouts of each venue on a 'control' subset of randomly routed flow.


To our knowledge, all brokers enable each client to opt out of any particular venue to which they have access. However, some brokers require clients to actively opt into each venue while others perform their own analysis and aim to set 'sensible defaults' with the option to customize them if needed. In our experience ~40% of brokers operate an 'opt out' model whilst ~60% operate an 'opt in' model in the US. In Europe the same brokers have iterated over time and now largely operate on an 'opt out' model where many – but not all – systematic internalizers are eligible venues, based on long-term observed performance, unless a client specifically wishes to opt out.

Reasonably, one could disagree on the most appropriate model. However, it should be uncontroversial to state that the average buy-side client has far less data and fewer quant resources dedicated to this kind of analysis than their broker. This is especially true when performing parent-level A/B tests of individual SDPs. Accordingly, we would argue brokers are well placed to evaluate the data and decide – likely sharing the data and reasons for their conclusion – which defaults will produce the best outcomes for the users of their execution platform. Provided this process is fully transparent, any buy-side firms that feel strongly would retain the ability to opt out.

INDICATIONS OF INTEREST

Most brokers that interact with our SDP do so by consuming continuously streamed Indications of Interest (IOIs). This means that they know at any given time for any given symbol we are available to buy or sell and the size. We understand IOIs to be a commonly available feature of all leading SDPs. Fill rates when trading via an IOI should approach 100%^[3] – some technical rejects will exist where there is some latency in responding to an updating IOI but brokers monitor this assiduously.

Accordingly, when an order arrives at the broker, they can consider our liquidity against other options and decide whether to send us an order or not. If we have no IOI available on the side and stock that interests them, they have no need to send us an order because they already know we are not available. This technology removes any potential conflict of interest in terms of information leakage on unfilled orders.



This is a purely illustrative image. It shows a broker who is aggregating multiple IOI streams from liquidity providers. The broker can then decide whether to send an order to any of these liquidity providers, based on their IOI availability.

In some cases, brokers (or downstream vendors) may not be able to consume or process IOIs. With approximately 10,000 symbols the message count can quickly become enormous. In the case where a 'blind ping' is a technological necessity we would strongly argue for using a hosted room^[4] to mitigate any risk of information leakage.


[3] This assumes a technically competent vendor that correctly processes IOI messages and responds in a timely manner. At the time of writing, a representative bulge bracket algo provider has 99.7% volume-weighted fill rate YTD when trading with XTX's SDP.

[4] See our retail-focused overview note here: <https://www.xtxmarkets.com/assets/clients/whitepaper-rooms.pdf>

In this case the broker may send an IOC to an ATS hosted room, whose matching engine has a lit view of the resting orders from liquidity providers and knows whether an incoming IOC can be matched. If no suitable liquidity provider orders are in the book, the ATS can cancel-back to the broker (unfilled) without making any liquidity provider aware of the order's existence^[5].

Many broker algorithms have sub-strategies which 'rest in the dark'. For example, they may wait patiently at mid without placing in any lit book. These stages can account for a surprisingly large quantity of fills for patient orders. A conditional order placed onto a hosted room may replicate this functionality. Alternatively, the broker may synthetically rest the order in its own system by monitoring incoming IOIs and hitting one if it appears suitable.

This latter feature is novel and not supported by all brokers, meaning that it is possible for an SDP to stream mid interest to sell that it is unable to match with a resting order to buy at mid. While all development work must be justified, our experience is that capturing such liquidity can provide a boost to mid fills and overall performance.



[5] Implementation details of each ATS are subject to change and users should verify the logic with ATS representatives.
The information provided is an accurate overview based on the author's understanding of these products at the time of writing.

SCORECARDS AND TIED PRICES

If a broker has an order to sell 500 shares at mid and finds themselves in the happy situation of having three liquidity sources willing to buy at mid, each with 1,000+ shares available, they have a choice to make. Which should they trade with?

This decision can have a huge effect on liquidity providers and the liquidity they choose to provide to that broker in future. Empirically the fills when there is lots of competition for the order tend to be the highest quality fills for liquidity providers.

Liquidity providers make decisions on the liquidity they wish to provide to each broker client– see later section – based on the average of the past flow received. So, if a broker always awards ties to venue A instead of venue B, it is likely that the average for venue B will worsen. The result may be that they will show that broker less competitive liquidity in the future than they otherwise could or would have.

A good tie-breaking logic will seek to incentivize liquidity providers to add value to the book for future orders. For example, the most common logic we see today is $\text{price} > \text{size}$. If selling 500 shares and three LPs are available at mid for 1,000+ shares the broker will simply hit the largest. This incentivizes the liquidity providers to increase size. While not useful on this particular order, perhaps some future larger order will benefit from that incentive.

We would argue that this can be improved upon. In practice what we've seen work extremely well is a scorecard approach combined with randomization. A scorecard will take into account a blend of execution factors that the broker considers important: size; presence in illiquid stocks; ratio of mid: full spread fills; post trade impact; fill ratios etc. These are combined into an overall score.


TOTAL SCORE TABLE

LP	SIZE	IMPACT	PI	FILL RATE	TOTAL SCORE
A	7	8	9	9	33
B	6	7	8	9	30
C	9	5	6	8	28

This is a purely illustrative image. It depicts a simple multi-factor scorecard.
In reality each broker will use different inputs and weight each input differently.

Note that to get a higher score the liquidity provider is incentivized to show more mid, have lower market impact, show more size, provide liquidity in illiquid stocks, have an extremely high fill ratio – all useful things. Then an element of randomization is introduced. This is critical as it allows unbiased ongoing evaluation of each liquidity provider. Finally, the randomization is weighted in some form based on the scorecard overall result for each liquidity provider.

RANDOMIZED ALLOCATION TITLED BY SCORE




This is a purely illustrative image. It depicts how a stronger or weaker scorecard may slightly bias the allocation for each liquidity provider away from purely random allocation. Whilst a better scorecard will result in more price ties the randomization ensures that A does not always beat B deterministically, allowing for competition and like-for-like analysis post trade.

DOES EACH BROKER GET THE SAME LIQUIDITY FROM EACH SDP?


No. When brokers connect to all-to-all Central Limit Orderbooks (CLOBs) the liquidity each broker can access on a particular CLOB is identical. This is because a CLOB is not allowed to segment and provide different prices to different users. All brokers receive the same size and bid/offer price from any given CLOB venue, it will not vary based on the broker you choose to access it.

SDPs can – and do – tailor their liquidity to each of their broker clients based on the nature of flow that they receive from them. The differences can be meaningful. For example, below is a plot of mid presence across the top 5,000 stocks that XTX's SDP offers for a well optimized US broker. It shows that XTX is available at mid approximately 70% of the time for this broker.



This plot was produced by XTX using analysis of internal data on a production agency broker stream in the period 14th July 2025 – 18th July 2025. It shows that for most stocks in the 5,000-stock universe a one-sided mid IOI was available around 70% of the time.

Many brokers receive no mid at all. Size is another common variable that can be adjusted at stream level. Below we compare production for touch streams across brokers and the difference between the largest size shown and smallest size shown is almost 30x.



These plots are produced by XTX using analysis of internal data on production streams for five representative agency brokers. They show for the median stock on each stream in the period 14th July 2025 – 18th July 2025 how frequently each stream was present at mid and the relative size available indexed to stream A. For example, stream E was available at approximately 30x the size of stream A. This highlights the large difference in available liquidity even when accessing the same SDP via different channels.

Whilst two brokers may have access to the same venue, they may obtain markedly different liquidity and results based on how they decide to interact with it^[6].

Brokers do not typically only have a single stream from each SDP. There are many trade-offs on individual SDP streams – generally across the variables of presence, size, and price (mid or touch or somewhere in between). For example, a mid-stream will typically be less present than a touch stream. A small stream will be more present than a large stream and so on. These trade-offs are discussed collaboratively with each broker and SDP that iterate together to find a productive long-term balance.

The most technologically sophisticated brokers may thus have four or more separate SDP streams from a single SDP to benefit from highly specific per-stream liquidity characteristics to satisfy their underlying algorithm objectives. For example, a small-sized mid-stream that maximizes mid presence alongside a mid-stream with less presence but in much larger size. This allows the broker to tailor what exactly they want rather than having a generic single stream from each SDP which must average down on size/presence/price to the lowest common denominator.

[6] Note that, for risk-warehousing SDPs, it may not be beneficial to interact at all with high participation/urgency algorithms or multi-day executions. This is because these executions themselves are likely to have market impact and, if the SDP suffers large losses on these interactions, it will show reduced size or mid presence to the detriment of other more patient orders whose characteristics merit exceptional price or size improvement to lit venues. This is one area where the judgment (and technological flexibility) of each broker can make a difference.

STREAM CURATION


As discussed above, in a well optimized set-up brokers will often have a variety of streams and use them for different purposes. Below is an example of how a broker might incorporate four streams from a single SDP:

STREAM	PRICE LEVEL	SIZE	USE CASE
A	MID	SMALL	CHILD ORDERS ON PATIENT ALGORITHMS
B	TOUCH	SMALL	CHILD ORDERS ON ALGORITHMS "CATCHING UP"
C	MID	MEDIUM	PARENT LEVEL ORDERS WHERE THE IOI CAN FILL THE REMAINING BALANCE IN FULL
D	TOUCH	LARGE	PARENT LEVEL ORDERS WHERE THE IOI CAN FILL THE REMAINING BALANCE IN FULL

This is an illustrative table, which represents hypothetical use cases for various SDP streams at a broker.

If the broker can support a menu of streams, such as above, it has the option to map various combinations of streams as it deems appropriate to underlying algorithms and stages within each algorithm. For example, SDP child mid-streams might be enabled for all types of algorithm – liquidity seeking, dark aggregation, scheduled – but within those categories only parent orders that execute below a certain level of urgency.

Indeed, not all flow may be suitable for SDP interaction. Extremely high participation orders or multi-day orders are likely to result in meaningful losses for risk-holding SDPs who will then have less incentive to show large sizes or mid liquidity for future orders. Such orders might be more suited to anonymous all-to-all venues and bypass SDPs entirely.



THE ROLE OF THE BUY-SIDE IN SHAPING CONVENTIONS

When we trade via broker algorithms, transparency flows one way: the buy-side knows which SDP it accessed, but SDPs see only the broker. We believe this “broker-in-the-middle” model is a natural fit. Even our so-called “bilateral” trades with European buy-side desks are routed through agency brokers, who intermediate each transaction^[7].

Agency brokers offer three key advantages:

1. Neutral Aggregation – they combine liquidity from many SDPs and lit or dark venues, fostering genuine order-by-order competition.
2. Execution Expertise – beyond routing, they offer consulting, analytics, and ongoing platform optimization.
3. Alignment of Interests – as fiduciaries, they share objective performance data and help clients calibrate venue defaults, opt-in choices, and tie-breaking logic.

We also encourage open dialogue between buy-side firms and brokers. By understanding a broker’s design choices – stream curation, tie-breaking scorecards, opt-in logic – clients can give more precise feedback and tailor strategies. After all, as our data from the same SDP (XTX) proves, execution outcomes can vary significantly across brokers (size shown, mid-presence, fill rates)^[8].

BROKER	MID	TOUCH	STREAM CURATION	RESTING PHASE	TIE PRIORITY	IOI	SDP	SDP CHOICE
A	YES	YES	LIMITED	NO	PRICE > SIZE	YES	CLIENT OPT-IN	GRANULAR
B	NO	YES	COMPREHENSIVE	NO	RANDOMIZED SCORECARD	YES	CLIENT OPT-OUT	ALL OR NONE
C	YES	YES	COMPREHENSIVE	YES	RANDOMIZED SCORECARD	YES	CLIENT OPT-IN	GRANULAR

This is an illustrative table, which a buy-side firm might use to understand and compare how each of their brokers interacts with SDPs.

No two brokers will have configured their SDP interaction logic identically. But by discussing these topics and understanding the choices each broker makes, buy- and sell-side can co-evolve industry norms and improve execution quality for everyone.

[7] See for example <https://flextrade.com/resources/a-new-era-in-bilateral-liquidity/>.

[8] See plots on earlier pages.

DISCLAIMER

This Document is issued by XTX Execution Services LLC ("XTX"). XTX is registered with the Securities and Exchange Commission (SEC) as a brokerdealer (SEC#: 8-70532, CRD#: 309266) and its Designated Examining Authority is FINRA. The registered office of XTX is 64th Floor, 50 Hudson Yards, New York, NY 10001, USA.

The distribution of this Document is restricted by XTX and may be further restricted by law. No action has been or will be taken by XTX to permit the possession or distribution of the Document in any jurisdiction where action for that purpose may be required. Accordingly, the Document may not be used in any jurisdiction except under circumstances that will result in compliance with any applicable laws and regulations. Persons to whom the Document is communicated should inform themselves about and observe any such restrictions. This Document may not under any circumstances be copied, distributed, published or reproduced, in whole or in part, without the prior written consent of XTX.

This Document is a marketing communication and is not and should not be construed as investment research or a research report. This Document is for information purposes only. Nothing in this Document constitutes investment, financial, tax, legal or other advice, nor does this Document constitute an offer to transact in, or the solicitation of an offer to transact in, securities, derivatives, FX transactions, or other financial instruments, in any jurisdiction.

Although the information in this Document is believed to be materially correct at the date that this Document has been communicated by XTX, no representation or warranty is given as to the accuracy of any of the information provided. Certain information included in this Document is based on information obtained from sources considered to be reliable. However, any projections or analysis provided to assist the recipient of this Document in evaluating the matters described herein may be based on subjective assessments and assumptions and may use one among alternative methodologies that produce different results. Accordingly, any projections or analysis should not be viewed as factual and should not be relied upon as an accurate prediction of future results. Furthermore, to the extent permitted by law, neither XTX nor its employees, directors, officers, shareholders or service providers assumes any liability or responsibility nor owes any duty of care for any consequences of any person acting or refraining to act in reliance on the information contained in this Document or for any decision based on it. Past performance cannot be relied on as a guide to future performance. XTX does not undertake any obligation generally to update or revise any information contained in this Document, except as may be required by law.