



IS IT TIME FOR AGGREGATION  
TO GET A BIT SMARTER?

SEPTEMBER 2021

# IS IT TIME FOR AGGREGATION TO GET A BIT SMARTER?

On algo desks the concept of a Smart Order Router (SOR) is generally well known. A smart order router considers multiple factors – for example price, historical impact, likelihood of fill – and then routes orders to venues accordingly.

Aggregators are common in FX, with many technology providers having appeared over the last decade. However, an aggregator is not necessarily a smart order router.

We will look at one approach for improving the output of an aggregator, mainly from the perspective of a typical regional bank, and discuss ways in which these aggregators could add more value for their clients.

	Bid	Offer	Spread
Best	30.1	30.5	0.4
LP1	30	31	1.0
LP2	29.5	30.5	1.0
LP3	30.1	32	1.9

*Figure is purely illustrative and not based on real data. We can see that LP3 has the best bid but the worst offer. Aggregation allows the trader to benefit from skews across multiple liquidity providers.*

# HOW REGIONAL BANKS USE AGGREGATORS TODAY

The typical aggregator will take in, for example, five LP (liquidity provider) prices and order them from best to worst. This allows the client to compete each order and obtain the best price available to them at any given time. This may improve execution outcomes and unsurprisingly aggregation has become popular amongst regional banks.

A typical regional bank may have one aggregator but two separate users: the electronic desk and the voice desk. It is possible that an LP may stream different pricing to each desk, tailored to the nature of each counterparty's flow.

The way most banks work is that they will quote and win a corporate RFQ (request for quote) e.g. 25m EURUSD. They will then slowly work their way out of this. They might, for example, sell 1m every 30 seconds.

Their hope is that the spread they charge on the 25m is more than the spread they pay on the 1m hedging clips and the market does not move before they have hedged.

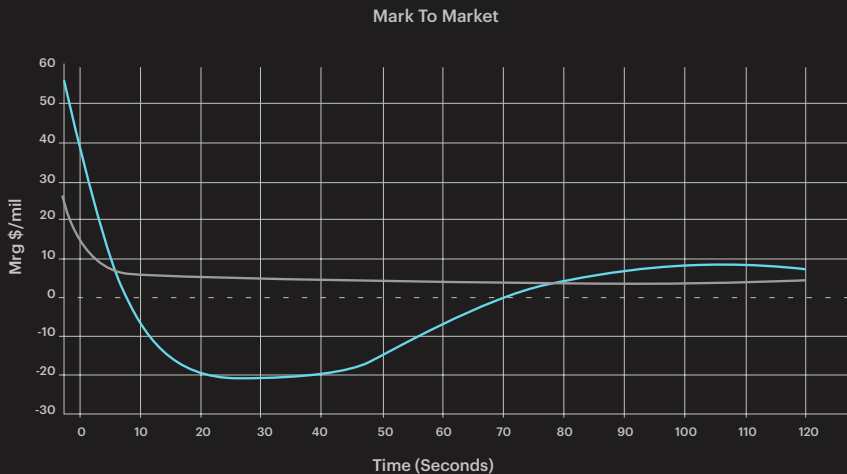
Electronic desks and voice desks both commonly use aggregators in roughly this way, even if the former is automated and the latter will use more discretion.

# WHO WINS THE TRADE?

As we discussed before, the trade is won by whoever shows the best price. This sounds good. However, it is potentially expensive for the bank.

Imagine you are a bank with 25m to hedge. The LP with the best price is 0.1 pip better than the next LP. However, you observe that this LP causes a lot more impact (0.3 pips compared to 0.1 pips) on average in the 30 seconds post trade before you send your next child order.

Whilst you may save 0.1 pip on the initial 1m order, you will then lose 0.2 pips (0.3 - 0.1pip) on the 24m balance that you have left to trade. The loss is almost 50x the saving.



*Figure is purely illustrative and not based on real data.  
The grey venue exhibits far lower post trade impact than the blue venue.*

You might also consider the impact of rejections.

Imagine a voice desk wishes to risk transfer a 20m GBPUSD ticket just before data comes out. LP A is 0.1 pips tighter than LP B, but LP A only has a fill rate of 75% compared to LP B with 100%. It would be sub-optimal for the aggregator to just pick the best price and ignore the potential cost of being rejected and having to re-trade at a worse price.

# ALL-IN AGGREGATION LOGIC

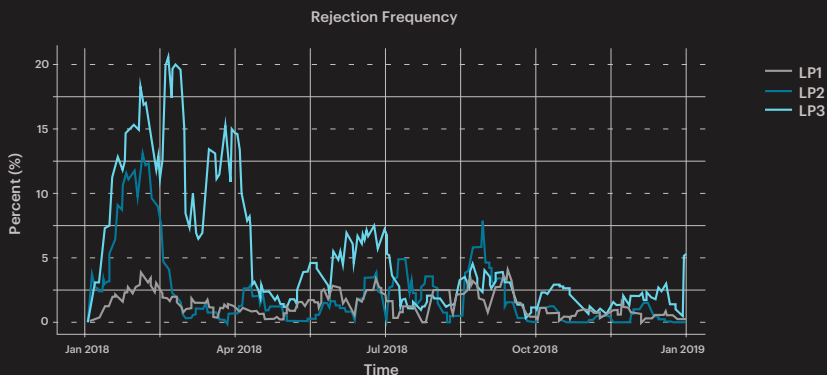
A handful of sophisticated regional banks route based on all-in logic. What this means is that they weight the displayed prices by market impact and likelihood of fill. Let's look at a simple example:

Over the last 10,000 EURUSD trades LP A has:

- a reject rate of 10%;
- a cost of rejects<sup>1</sup> of 0.4 pips;
- average market impact<sup>2</sup> at 30 seconds of 0.2 pips.

If LP A showed a bid of 1.20506 we would then subtract  $(10\% * 0.4)$  and 0.2. That is  $50.6 - 0.04 - 0.2 = 50.36$ .

The aggregator would then route on the basis of this all-in score (50.36 not 50.6). The LP with the best all-in score should win the trade as the all-in cost to the regional bank and client of the aggregator will have been taken into account. This means that the best displayed price will not necessarily be guaranteed to win the trade.



Source of data: Tradefeedr – CMC Markets, originally in *Aggregation and Liquidity Management in the FX Market (XTX; 2018)*.  
The certainty of fill rate clearly varies by LP. This information should be incorporated into routing decisions.

This is clearly a more sensible way to route flow as it takes into account all of the potential costs to the client who uses the aggregator, rather than only one of them.

<sup>1</sup> We will define cost of rejects as the average mid-to-mid move between the timestamp of a rejected order being sent to an LP and the timestamp of an order being rejected by that LP.

<sup>2</sup> We will define average market impact as the average mid-to-mid move between the timestamp of an accepted order being sent to that LP and the timestamp 30 seconds later.

# AGGREGATOR VENDORS

This feels like an interesting area to explore for aggregation software vendors. It is a crowded market and this approach might be a nice way for a vendor to stand out from the crowd.

Almost all aggregators have tied price logic – deciding which LP wins the trade when their prices are the same – but generally this is an incomplete solution. The reality is that, outside top of book G3 prices, most prices are not tied. Furthermore, the tied price logic doesn't sufficiently capture the cost of market impact, which as described above may be the dominant cost for regional banks hedging.

Applying all-in routing logic is actually pretty trivial for a software vendor, provided that they have the capability to generate and store these metrics. There are then some interesting UX (user experience) decisions: What price should be displayed to GUI (graphical user interface) users? Should it be the all-in score or the displayed price or the displayed price of the LP with the best all-in score? In reality, this is something that is quite easy to solve in consultation with clients.

The software vendors also have to be careful to amass a minimum sample size of orders before applying this logic: a reject rate of 50% on two trades is not very meaningful and no decisions should be made on the basis of a tiny sample size.

# CONCLUSIONS

Applying all-in logic or asking your vendor to do so has the potential to result in meaningful improvements in hedging costs.

Why not discuss the above with them and ask if it is feasible?

You may also wish to get feedback from your LPs and see what perspectives they have to offer on the topic.

# IMPORTANT NOTICE

This Document is issued by XTX Markets Limited (“XTX”), which is authorised and regulated by the Financial Conduct Authority of the United Kingdom (the “FCA”), with FCA FRN: 711945. XTX is a private limited company incorporated in England & Wales, with company number 09415174. XTX’s registered office and principal place of business is R7, 14-18 Handyside Street, London, N1C 4DN, United Kingdom.

This Document is issued by XTX only to and/or is directed only at persons who are eligible counterparties for the purposes of the FCA Rules. This Document must not be relied or acted upon by any other persons (including, without limitation, persons who are retail clients or professional clients for the purposes of the FCA Rules). The distribution of this Document 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 is for information purposes only. The Document does not 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, 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.

Any reference in this Document to a counterparty trading with XTX (or XTX trading with counterparties) is a reference to XTX providing liquidity quotes through certain specific platforms, and all trades will be settled or cleared with or through XTX’s prime broker(s) or CCPs, and not with XTX directly.



**[XTX]**  
**MARKETS**