

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.