XTX MARKETS

PATIENT ALGOS AND ‘REGRET’: THE PROBLEM WITH PASSIVE

MAY 2021

THE TRADE-OFF

It is well understood that the more patient you are willing to be, the cheaper on average you can trade.

In return for this you must run some market risk: the market may drift further if you take an hour to patiently work an order than if you push the same order through in ten minutes.

Execution specialists who aim to reduce costs across a large portfolio of algo orders - accepting market risk on individual orders - will therefore often opt to trade patiently.

This is especially true when the trader knows her order is purely transactional in nature and non directional - hedging an M&A transaction or corporate proceeds abroad for example.

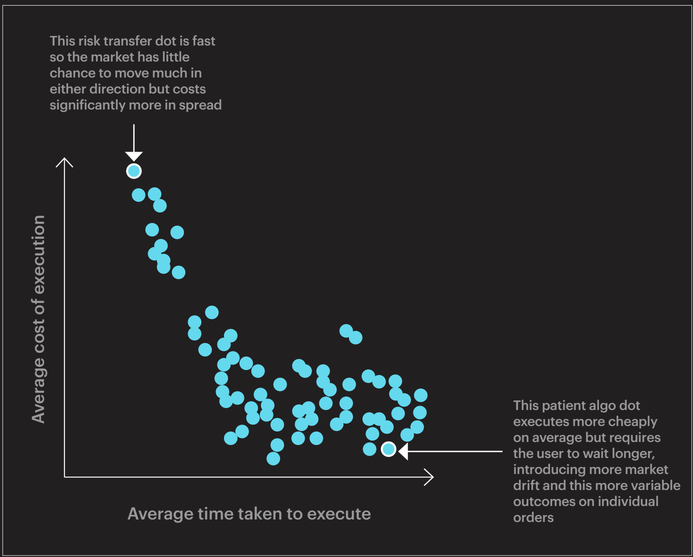

Illustrative image (not real data) showing the trade-off between speed of trading and cost of trading.

FLOATING OR PEGGED ALGORITHMS

A popular type of algorithm exists for these patient orders and its basic operation is to attach itself to the bid (if buying) or offer (if selling), often using the primary market as a reference.

As the bid or offer moves up or down, so the algorithm floats with it. The algorithm receives fills when other people cross its spread. It will never cross the spread itself - it waits patiently to obtain fills. There are all sorts of variations in terms of the exact implementation but this is the core idea.

Aggregated market data (Bid/Ask)

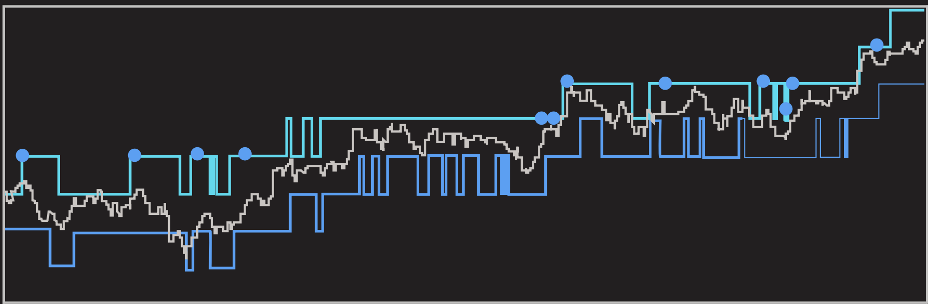

Illustrative image (not real data) showing a sell float algorithm.

HOW DO YOU KNOW IF A FILL IS PASSIVE?

On the primary markets in FX - EBS and Reuters - fills are tagged as 'maker' (passive) or 'taker' (aggressive). This is one definition, however it is quite limited as not all fills will occur on the primary market.

A nice generic way to see if a fill is passive is to compare it to a neutral - ideally independent - mid-price at the time of trade.

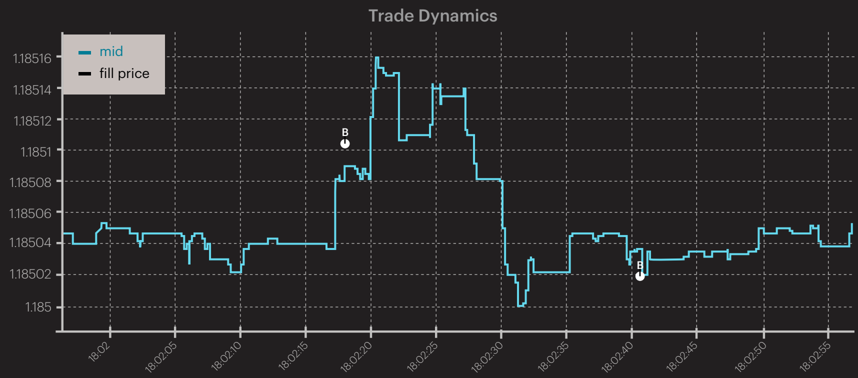

Illustrative image (not real data) showing child buy orders plotted against a neutral mid

Let us say the independent mid-price happens to be 11. Any algo buy fills that come in below 11, we might consider passive. Any algo buy fills above 11, we might consider aggressive.

Using this simple method we can even quantify how passive or aggressive each fill is.

For example a buy fill at 9 is two pips better than mid ... a buy fill at 10.9 is only 0.1 pips better than mid. Clearly, more spread was captured on the first fill.

ADVERSE SELECTION

It seems like patient execution is all about getting filled passively. However there's a problem: adverse selection.

It is easy to get filled passively if the market is about to move against you. This is called being 'picked off' or adverse selection.

For example, imagine a market that is 10 / 12. You are filled on the bid at 10. Five milliseconds later the price has dropped to 5 / 7. That is a passive fill but a pretty horrible one. It would have been better to have been patient and pulled the 10 bid so you could trade more cheaply in the 5 / 7 market that followed.

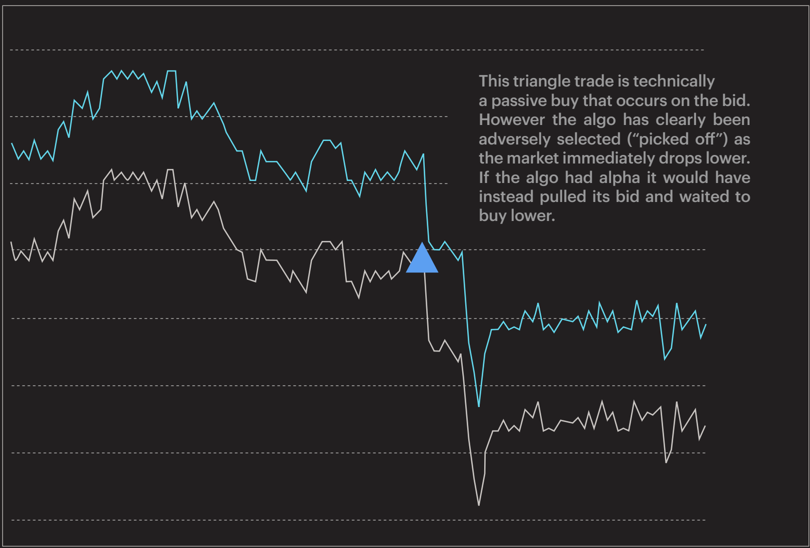

Illustrative image (not real data) showing a passive buy that has been adversely selected.

Imagine the algo is buying and pegged to the primary bid. However, the algo is always slow to react and adjust its bid when the price changes.

This is classic adverse selection:

• Whenever the market ticks up the algo will not buy: someone faster to react will place their bid above the algo's bid and take those fills;

• Yet whenever the market ticks down the algo will get filled as it is slow to move lower ... but will buy at the old, higher bid price.

EVALUATING PASSIVE FILLS IN INDEPENDENT TCA

So how do you tell if it was a high quality passive fill or a pick off?

PASSIVE OR NOT?

The first thing is to look at independent analysis in Transaction Cost Analysis platforms. These platforms all have independent mids. This is interesting because immediately you can see one interesting statistic: How many fills were genuinely passive at time of trade? That is, how many buys occurred below the independent mid at the time?

The next thing to look at is how many fills remain passive vs the independent mid a second later. Remember the 10 / 12 market that drops immediately to 5 / 7. In this case we compare 10 (fill) to 6 (mid of 5 / 7). We can instantly spot that this was an order that was picked off and did not genuinely capture spread.

Algo	Passive % at time of trade	Passive % after one second
Algo 1	78.4	16.7

Illustrative results (not real data) showing a hypothetical algo's child order results

You bought just before the price dropped! In the above illustrative example we can see that the passive fills are heavily adversely selected. Only 16.7% appear passive a second later.

ADVERSE SELECTION ('REGRET')

To capture this more precisely, you could measure the adverse selection by determining the spread paid or earned ($/m) vs an independent mid one second post fill. Here you are just asking how much - if at all - did the independent mid move from time of trade to one second later? If it does not, i.e. you bought at 10 and a second later the market is still at 10/12, then it is a genuinely passive fill.

SPREAD CAPTURE

You can see all of this in a single picture.

Averaged across all child fills, you would want to see the following for each algorithm in each currency pair. These are similar to the market impact charts you may have seen but inverted so they are from the algo user's perspective.

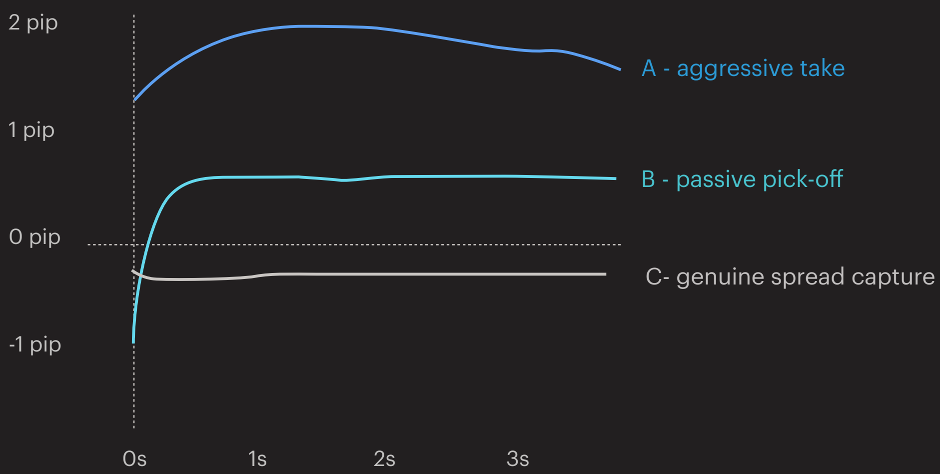

Illustrative image (not real data) plotting the child fill curve for three hypothetical algorithms

What this chart does is plot the average child fill for each algorithm against an independent mid at various points in time. At time of trade (t= 0s) for example we can see Algo B earned 1 pip versus the neutral mid. It captured spread. After 1s however, this picture is transformed: its fills are offside by 0.5 versus the neutral mid one second later.

Let’s take an example where we buy at 10 in a 10 / 12 market.

Spread paid is a measure (in bps or $ per mio or pips) of the fill rate relative to a neutral mid at time of trade. Here we bought at 10 and the mid was 11. So we earned 1 pip.

However, if one second after the trade the market drops to 8.5 / 10.5. that then gives a mid of 9.5. Comparing the buy at 10 against this we get a spread paid of 0.5 pips. We got picked off!

That is represented in B - passive pick-off in the chart above.

BEYOND CHILD FILLS: SIGNALLING RISK

We have looked through some useful ways to see if a passive fill genuinely captured spread for you or was merely picked off.

The skill of the algo operator will determine this: a float that is never picked off might produce wonderful results in terms of child fill quality whilst a float that is often picked off might suffer poor performance, despite both nominally being 100% passive.

However, there are several other factors that affect the overall result. For one thing, float orders suffer from the fact that you buy more easily in a falling market than a rising market - classic adverse selection. In our view though, the key factor is information leakage.

By showing a bid in the market, other participants may learn there is a large order active. Even if the bid is not filled it may signal the presence of a large buy order. This effect is harder to capture but real.

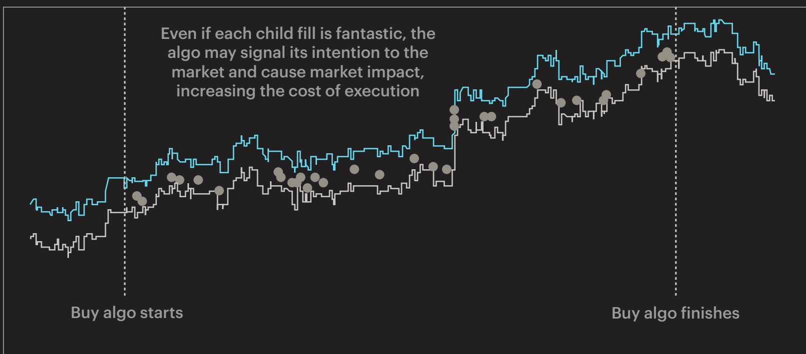

Illustrative image (not real data) of a hypothetical buy order showing that, even if child fills are high quality, information leakage may cause the overall fill to be poor.

The best way to observe this all-in effect is to measure the all-in slippage to arrival price on each parent order. That is, the independent mid-price when you set the order live compared to the all-in fill. This is of course the overall cost of execution to the trader's firm.

The difficulty is that these results are noisy. You need hundreds of runs to evaluate each algorithm in a single currency pair to +/- $100/m accuracy. Comparing 4-5 algorithms in several pairs would therefore require thousands of orders, which any individual client may not have.

The solution is peer universe analytics.

PEER UNIVERSE ANALYTICS

Almost all of the popular independent FX analytics firms offer this product today.

The idea is that many clients opt-in to a shared universe of metadata. No client can see anyone else's orders or sensitive details. This remains private. However, all opted in clients can see aggregated results. For example, they might look at the implementation shortfall of all algorithms in GBPUSD of around 50m for the month of March.

When all the results are aggregated, the noise is reduced and the good algorithms float to the top of the results and poor algorithms to the bottom. Even better - a client may see if a particular algorithm performs well without having to try it for themselves and pay away performance whilst learning. If an algorithm improves that will be visible, too.

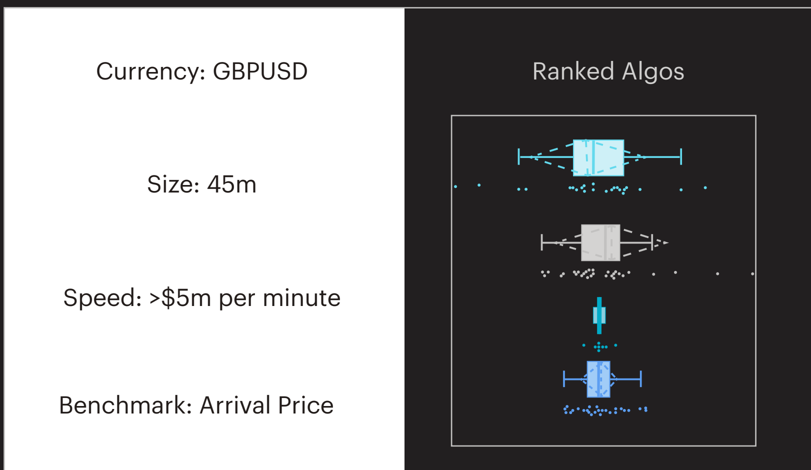

Illustrative image (not real data) of a peer universe comparison of algorithms.

CONCLUSIONS

Patient execution remains a logical choice for traders with large orders who wish to reduce their cost of trading. However, there are four things worth bearing in mind:

- Floating or pegged orders may prove an effective way to achieve this but must be measured because 'passive' does not necessarily mean a high quality 'patient' fill.

- Independent analytics may report on the true passive spread capture: the $/m spread paid or earned vs an independent mid in the seconds post fill. This will help separate the high quality passive fills vs those that were merely picked off.

- Beyond child fills, to get an idea of the overall performance of each algorithm at a parent order level it is necessary to capture other effects such as information leakage.

- Given the difficulty in obtaining a sufficient sample size, one way to achieve this is for clients to opt-in to 'peer universe' analytics offered by independent analysis companies.

For patient traders who wish to reduce execution costs, everything is visible in the data if you know where to look. You are likely to find the results extremely interesting.

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