FX Passive Order Analysis

VERSION: 1.0

DATE: 7/27/22

Introduction

Venue selection is an important part of trading in FX markets where a variety of ECNs and liquidity providers compete for client flow. An FX algo strategy that attempts to capture spread by placing passive orders must consider the quality of liquidity available at a particular venue and how much quantity to commit to that venue in order to maximize the probability of getting a fill at a favorable price. The strategy must also balance trading at the "primary" ECNs, EBS and Reuters, that impose 1M order minimums and secondary ECNs that allow trading in smaller sizes.

In this study we analyze over 500,000 passive order chains placed by clients via FX TWAP and VWAP strategies for a variety of directly traded currency pairs over a period of two years from February 2020 to January 2022. By measuring differences in arrival slippage, we observe that trading in smaller sizes results in price improvement on the order of 7% of spread over 1M orders. For the clients that prefer trading in larger sizes, we also note that for 1M orders, trading on secondary ECNs outperforms primary ECNs for both EBS and Reuters currency pairs.

Methodology

An FX strategy that follows a pre-determined trajectory such as TWAP or VWAP may attempt to execute passively by placing a passive order at the reference near-touch price on an ECN. If the market moves away, the passive order will re-peg by modifying the limit price or canceling and sending an updated order until it either gets filled or the strategy has to cancel it and cross the spread in order to keep up with the trajectory. We consider the full sequence of events from the initial placement through all the re-pegs to the final fill or cancel as a single **passive order chain** and measure its price performance relative to mid at arrival, i.e. the time of original order placement. Since not all such passive chains result in a full fill, as a cleanup cost, we use the far-touch price at time of cancel for any quantity that was not filled. In order to compare performance across pairs, we normalize slippage by the average reference bid-ask spread. In particular, the performance metric is defined as follows:

$$Slippage = (1\ for\ buy, -1\ for\ sell) \times \frac{ChainFillPrice - ArrivalMidpointPrice}{AvgBidAskSpread}$$

Note, smaller values of slippage indicate better performance.

Order quantity and venue are chosen randomly by the strategy prior to order placement, which enables a fair, apples-to-apples comparison of chains across different venues and order sizes. When comparing the performance of primary ECNs to secondary ECNs, only orders for 1M, which are eligible to trade on both types of venues, are considered.

Passive Chain Size Analysis

Table 1 shows a high-level summary of the passive order chain data, broken down by order size of 1M versus less than 1M in base currency. Pairs that traditionally trade at EBS, e.g. EUR/USD, USD/JPY, USD/CHF, and a few others, are considered separately from all other currency pairs that traditionally trade at Reuters.

In order to control for differences in parent order difficulty and pair characteristics between the chains we compare, we also include measures of average spread in basis points and the order's estimated participation rate, the effective percent of ECN trades (POT)[1].

Pair Primary	Size	Num. of Chains	Fill Rate	Avg Spread (bps)	Parent Rate (POT)
EBS	<1M	210,527	0.78	0.93	4%
EBS	1M	54,104	0.84	0.87	11%
Reuters	<1M	345,931	0.67	3.00	4%
Reuters	1M	42,797	0.83	1.83	12%

Table 1: Summary of passive chain data grouped by primary exchange and order size.

We note that Reuters pairs have higher average spread than the mostly liquid pairs like EUR/USD and USD/JPY that trade on EBS. Even though the slippage metric is spread-normalized, we need to control for spread as a proxy for other possible differences in currency pair trading characteristics. The average POT is also sufficiently different between the groups we compare, <1M and 1M, so we need to control for POT in our performance comparisons as well, since chains belonging to orders with higher POT have less time to execute passively before having to cross the spread.

Figure 1 compares the performance of passive order chains for 1M versus less than 1M in base currency, broken down into POT and spread buckets, so that each bucket compares passive order chains from parent orders with similar difficulty. The results shown combine EBS and Reuters pairs, however, note that the observations are consistent for both EBS and Reuters pairs. For each bucket, the figure shows two histograms of the bootstrapped mean of the chain slippage[2], one corresponding to chain quantity under 1M (red) and the other for chains of 1M (blue). The center of each histogram corresponds with the average slippage for that group and the width of the distribution represents the uncertainty in that average. If the two distributions overlap notably, the results indicate no statistically significant difference between the two groups. We look for separation between the two distributions, the blue and the red, as an indication of the statistically significant differences.

[1] Percent of trades, POT, compares the parent order quantity to the ECN volume that is estimated to trade during the order duration based on historical trade data.

[2] Bootstrapping the mean refers to generating many random subsets of data in each bucket for each group and calculating the quantity-weighted mean slippage for each subset. The distributions of values obtained from all the subsets is plotted as histograms shown here. The distribution of the bootstrapped mean helps to evaluate the statistical significance of results. Two times the standard deviation of this bootstrapped mean corresponds with the 95% confidence interval.

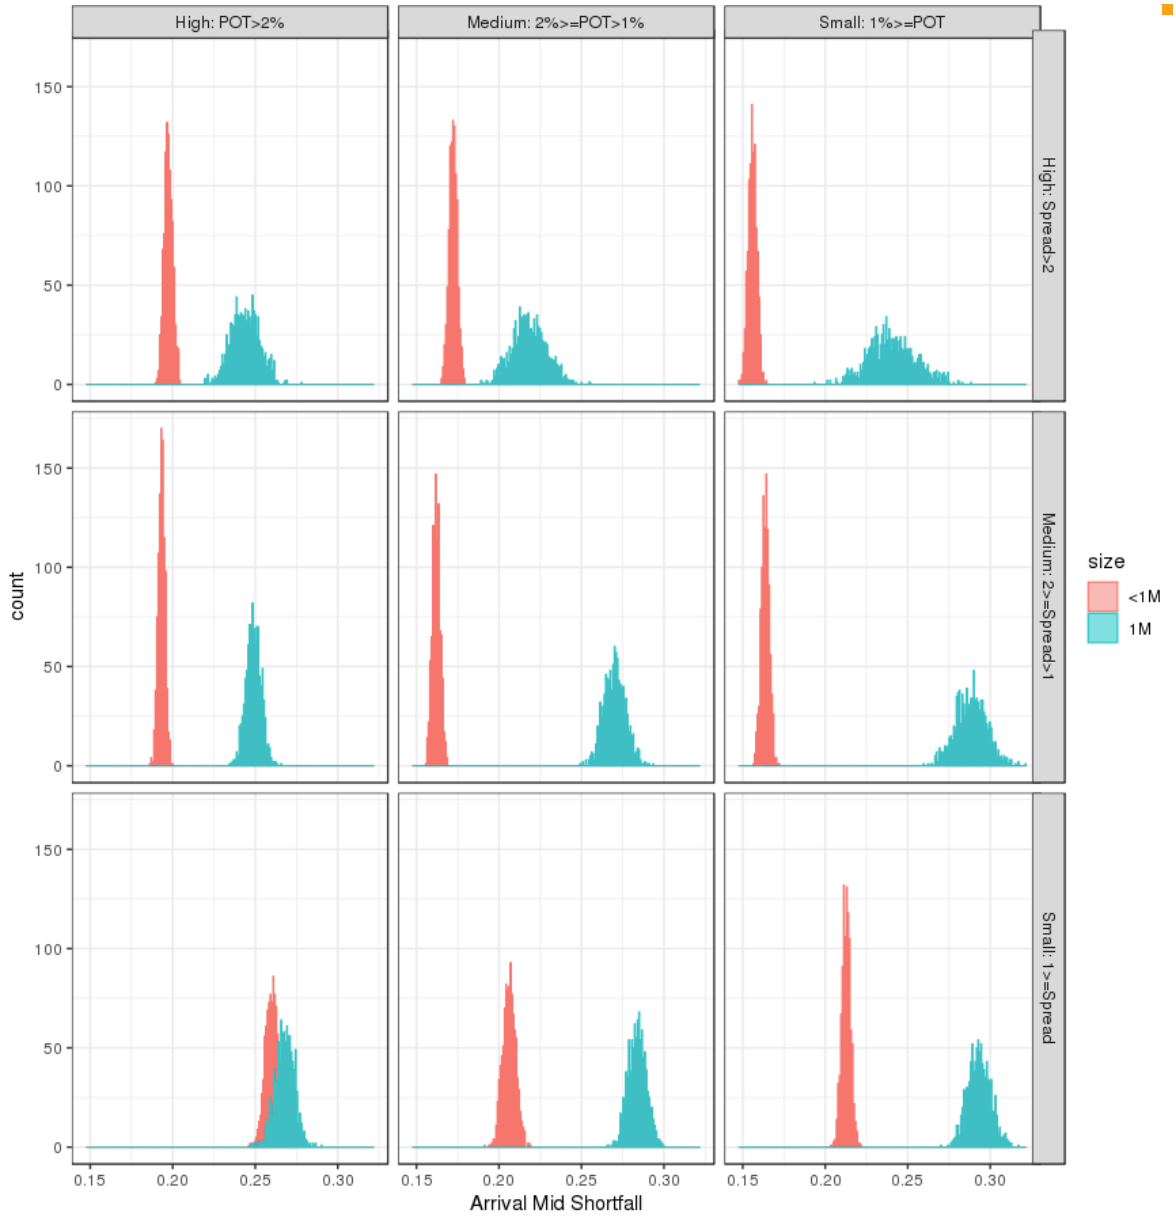

Figure 1: Histograms of the bootstrapped mean of passive chain slippage grouped by spread, rate and order size.

We observe that smaller orders outperform 1M orders, as indicated by the red histograms appearing consistently to the left of the blue histograms, and the results are statistically significant as there is little to no overlap between the blue and red histograms. This result is consistent with the observations from our previous studies and a general hypothesis that trading in smaller sizes results in lower adverse selection leading to better overall outcomes.

Primary Vs. Secondary ECNs

When trading in larger sizes of 1M or more, we can evaluate whether resting at the primary ECNs or secondary ECNs results in better performance. Table 2 shows a high-level summary of the passive order chains for 1M across primary and secondary ECNs. The differences in average spread and POT, though not as pronounced as in the comparison of small and large chains from Table 1, still require us to compare in buckets of like parent order difficulty.

Pair Primary	Venue	Num. of Chains	Fill Rate	Avg Spread (bps)	Parent Rate (POT)
EBS	Secondary ECN	239,049	0.79	0.92	5%
EBS	Primary ECN	25,582	0.85	0.88	4%
Reuters	Secondary ECN	365,089	0.68	2.94	4%
Reuters	Primary ECN	23,639	0.83	1.81	7%

Table 2: Summary of passive chain data comparing primary and secondary ECNs for orders of 1M for EBS and Reuters pairs.

Figure 2 compares the performance of passive order chains for 1M posted at the pair's primary ECN versus at secondary ECNs in buckets of comparable POT and spread.

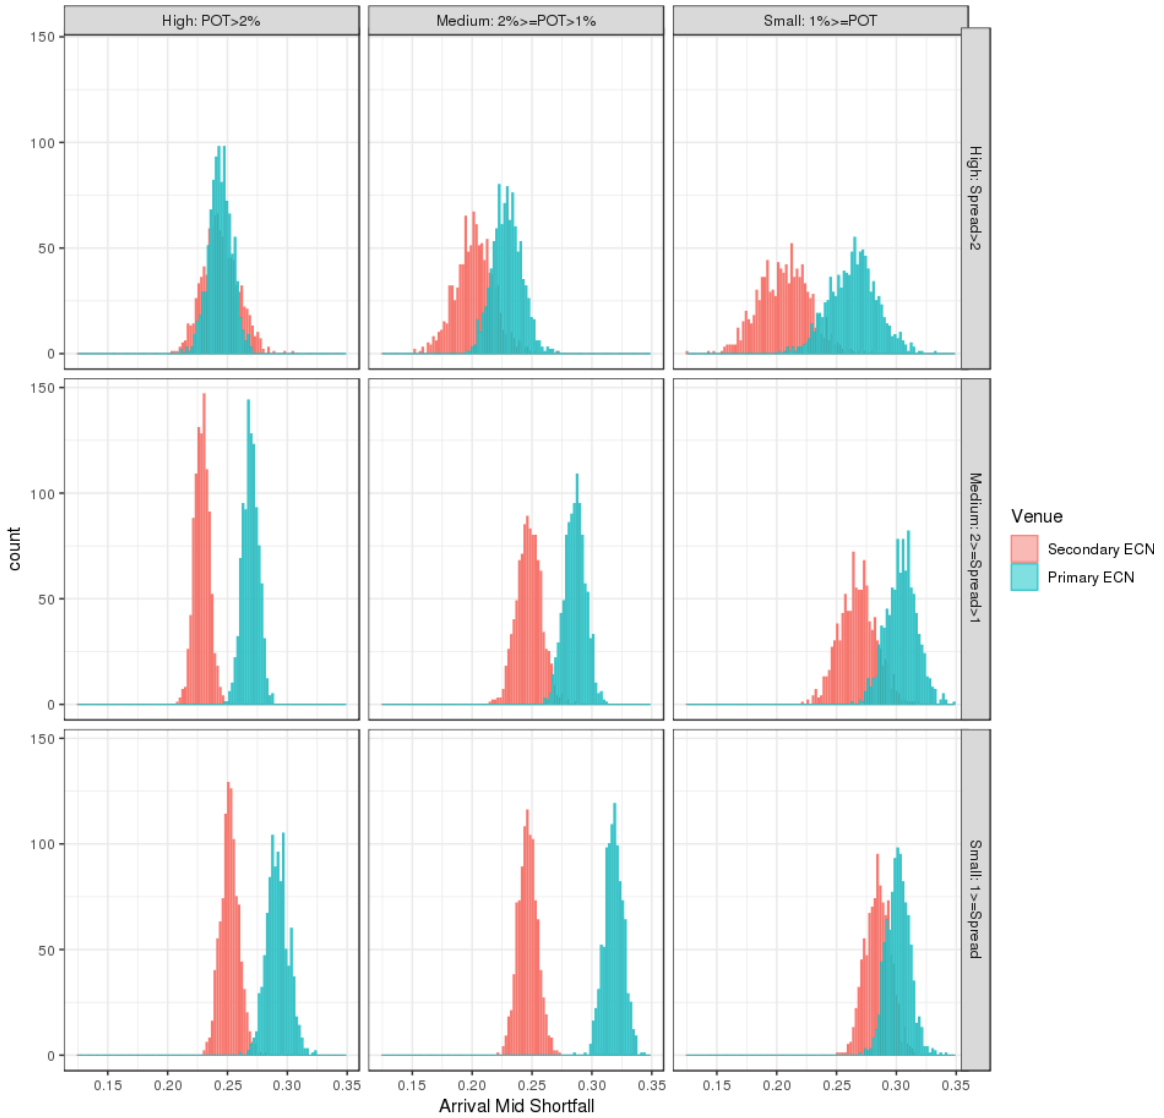

Figure 2: Histograms comparing of the bootstrapped mean of passive chain slippage for primary and secondary ECNs, grouped by spread, parent rate (POT).

We observe that posting on secondary ECNs generally outperforms posting on each pair's primary ECN, as indicated by the red histograms positioned largely to the left of the blue histograms, and the results are statistically significant as indicated by the separation in most of the distributions of the bootstrapped means for each group. We note that while orders in both groups have an initial quantity of 1M, orders placed on secondary ECNs are eligible to get partial fills, and that may have contributed to the improved price performance.

Conclusion

We have analyzed the performance of passive chains over the period of two years, 2020-2022, focusing on order characteristics such as size and venue selection. We observe, with statistical significance, that posting in sizes smaller than 1M results in lower slippage than posting in size of 1M. Further, when there is a need to trade in larger sizes, posting on secondary ECNs leads to less slippage than posting at primary ECNs. It is not surprising that as liquidity generally moves away from primary ECNs, secondary ECNs play a more important role. Which secondary ECNs offer better performance for different pairs is an area of interest for future study.