

Execution in Spot FX

Sanjith Ponnamaneni¹, Chris Jones², Ed Tricker³

Abstract

In this paper we highlight some of the mechanics, nuances, and pitfalls unique to OTC FX trading and propose that building a bespoke FX Aggregator is a way to overcome some of these challenges. We also highlight some of the considerations that are needed in such an endeavor.

¹ Director - Smart Execution

² Managing Director - Portfolio Solutions & Research

³ CIO - Quantitative Strategies

1. Introduction

For most CTA and Systematic Macro investment managers, Foreign Exchange (FX) is a core market, along with equity index futures, bonds/rate futures, and commodity futures. However, FX is unlike these other asset classes in as much as it is typically traded through OTC forwards rather than exchange-traded futures due to the relative depth of OTC FX market compared to the less voluminous FX futures market. This does, however, present a challenge given the fragmented and nuanced nature of the OTC FX markets, and specific expertise and systems are needed to overcome these challenges.

This paper explores the nature and nuances of FX markets, such as fragmentation by various types of liquidity provider and types of quote. Furthermore, we show how a carefully constructed FX quote aggregation system can overcome many challenges presented by market fragmentation and enable efficient execution.

2. FX Spot Vs Futures

Futures are exchange-traded and therefore have some advantages such as access to the exchange limit order book and protections from market manipulators. The FX market, on the other hand, is highly fragmented with no central exchange. Price discovery is based on the relationships with various market participants and the quotes received from each counterpart. This makes execution in FX more challenging, but with the potential for relatively low cost of trading due to tight spreads and high liquidity that helps make it a compelling choice over FX Futures.

3. The FX Market Place

Trading in FX markets is estimated at \$6.6 trillion per day with Spot FX at about \$2.2 trillion/day¹. However, liquidity in Spot FX is highly fragmented. At the top (by volume) are the two primary markets – EBS (G10 pairs) and Thomson Reuters (focused on Commonwealth pairs). There are several electronic marketplaces (ECNs) with HOTSPOT and FASTMATCH among the more widely used. ECNs have become more popular due to

high access fees and quote restrictions in primary markets. In addition, there are multiple aggregators (e.g. Bloomberg) which gather quotes and liquidity from other sources such as banks and non-bank institutions (e.g., XTX). Banks also have access to the interbank market, where they can transact significant volumes at tight spreads. Figure 1 shows the distribution of FX volumes across different channels.

4. Liquidity Provider Mechanics

A *liquidity provider* (LP) in spot FX is very similar to a market maker on a futures or equities exchange. They make markets by providing a two-way quote and profit primarily by capturing spread. They typically hold the position for some time, looking for an offsetting trade from other clients or offloading it into the market (if the trade is risk-increasing or loss-making for them).

A LP quotes a bid and ask price individually to each client, which is derived as follows:

1. The LP establishes a mid-market price using prices from primary market, ECN, and interbank quotes. These quotes are fed into a proprietary model which provides a mid-market estimate.
2. Half spread is applied to either side of the mid-market price, and resulting quotes are sent to the client. The spread used depends on many factors (currency pair, volatility, individual client profitability – this last factor is identified by a LP using *mark out* analysis and discussed in more detail in later sections). Customers providing the most profits can usually expect to receive the tightest spreads.
3. Based on the current position of the LP, the mid-market price is adjusted to encourage or discourage flow on a particular side (this is known as *skew*). For example, if a LP is long in EUR/USD and wants to reduce their position, they will reduce their offer price (thus encouraging buy orders from clients).

¹BIS Quarterly Review, International banking and financial market developments, Dec 2019. Available online at https://www.bis.org/publ/qrtpdf/r_qt1912.pdf.

How were FX trades executed in April 2019?

Percentage shares in total turnover

Graph 2

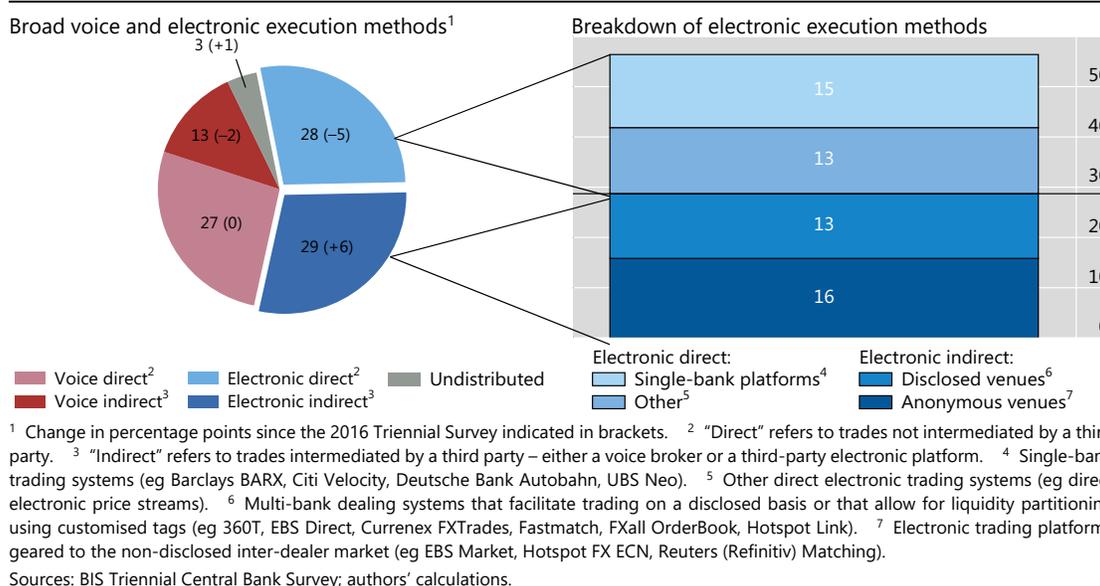


Figure 1. Spot FX Volume Distribution. Original source: BIS Quarterly Review, International banking and financial market developments, Dec 2019.

4.1 Types of quotes

LPs provide two types of quotes:

- Sweepable** Sweepable quotes allow you to hit multiple price levels or multiple quotes at the same price level. They are used typically by discretionary traders who want to build up (or reduce) large positions in a short period of time. LPs quote a wider spread because they are less sure of the risk they are taking on.
- Full Amount** When receiving full amount quotes, a client is limited to trading against a single quote at any given time. Large systematic firms may favor the smaller quote sizes and tighter spreads afforded by full-amount quotes since they typically execute over a longer time horizon using VWAP/TWAP type algorithms.

In addition, when going through an aggregator (for example, Bloomberg):

- Disclosed quotes** Both sender and receiver know identity of each other. Sender can customize their quotes based on receiver, and receiver can choose whether to trade against a particular sender. Systematic firms can expect to receive tight spreads on disclosed quotes because this scenario is similar to maintaining direct connectivity to underlying LP's in the aggregator.
- Non-Disclosed/Anonymous quotes** Sender and receiver do not know the identity of each other. This will help the receiver mask their trading activity. However, because sender does not know the identity of the receiver, they cannot customize the quotes for each receiver. For systematic firms, this could lead to higher spreads.

Example

If quotes in Table 1 are sweepable quotes, and the client wants to buy 3M of this pair, they can hit both the 1M quote @ 1.2610 and 2M quote @ 1.2615 at the same time. If these are Full Amount quotes, client needs to hit the 5M quote @ 1.2620 and request only 3M to be filled against this quote.

Bid Size	Bid	Ask	Ask Size
1,000,000	1.2605	1.2610	1,000,00
2,000,000	1.2600	1.2615	2,000,00
5,000,000	1.2595	1.2620	5,000,00
10,000,000	1.2590	1.2625	10,000,00

Table 1. Sample Quotes for EUR/USD

4.2 Last Look

Liquidity providers typically provide a non-binding quote and decide whether to accept an order against this quote if and when they receive an order. This mechanism is called *last look* and allows the LP to perform risk and credit checks. It also allows them to accept only those orders which are profitable for them. In addition, they can hold an order for a short time before making the decision. As one would expect, in many cases, this is not advantageous for the party sending the order because if the LP rejects, it may not be possible to refill the order at the same price from another LP.

4.3 Mark-Outs

Liquidity providers are constantly monitoring the profitability of each client. One tool they use in their analysis is known as a *mark-out*. A mark-out plot looks at the price an order is executed and market mid-price over many intervals following the execution. Figure 2 illustrates three different scenarios. The red line shows the LP losing the spread (captured at trade time) rapidly. LPs will typically mark this flow as bad for them and start quoting higher spreads for clients showing this profile to account for the rapid decay. Such a scenario can arise when a client has many LPs in a pool (indicating adverse selection) or uses high-frequency alpha signals as part of their trading. Both scenarios are bad for the LP. On the other hand, the client showing the green profile where the LP retains the spread captured at trade, will typically receive the tightest spreads from LP.

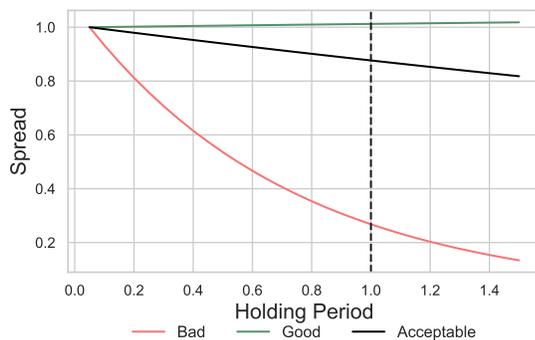


Figure 2. Markout (% of Spread retained vs Time) from LP's perspective.

5. An FX Aggregator

In the previous sections, we introduced the challenges posed by systematically trading FX markets. Nuances such as last look, skew and mark-out analysis add complexity to the problem of optimal execution.

One way to overcome some of these challenges is through the construction of a carefully assembled *FX aggregator*. This is a system (often with associated algorithms and front end) that draws in quotes from multiple sources and venues to streamline execution.

An FX aggregator allows you to connect to multiple venues, aggregate quotes, and view a consolidated book from which you can choose the best quote to trade against. Below we highlight the considerations needed when constructing an FX aggregator such as the number and nature of participants in the pool along with their hedging styles.

Who Should Be Part of the Pool?

It is preferable to choose counterparties with similar pricing and hedging models as it helps ensure that you minimize market impact from your execution. If one of the LPs has an aggressive hedging model and is offloading risk much faster than others, this information will show up on the mark-out profile, and soon other counterparties will start quoting higher spreads.

Exposure to ECNs and primary markets can help during market dislocations, place limit orders, and generate alpha signals. However, trading on ECNs and primary markets carries the risk of higher market impact, as any action on these markets is immediately visible to everyone.

The Role of ECNs/Primary Markets in an FX Aggregator

When a firm trades with a bank, the bank typically holds this position for some time (i.e., has a long holding period) while waiting for some offsetting flow from another client. This provides for a cushion in terms of market impact. Conversely, when trading directly on ECNs or primary markets, every action on the order book is published to all the participants. Therefore each action carries some amount of market impact with it. At the same time, primary markets and ECNs are a large source of liquidity, especially during times of high volatility in markets, when banks might pull back their quotes. They also provide an opportunity to post limit orders and save on spread costs.

When choosing the counterparties to trade against, it is tempting to include as many as you can. However, this can be counterproductive to optimal execution.

Adverse Selection

With many LPs in a pool, there is a chance that one of them has failed to update their quotes fast enough. When a client executes against this quote, the LP is in trouble because the market has already moved away. As a result, the LP is sitting on an instantaneous loss. If this happens consistently, the LP will start widening their quotes and start hedging this flow immediately. This causes indirect market impact and is undesirable for both LP and the client. This scenario is most likely to occur when there are many LPs in a pool and when some LPs have a considerably different pricing model compared to the majority.

The Number of Counterparties in the Aggregator

It is often the case that the quality of counterparties in the aggregator is more important than the number of LPs in the pool. As the number of LPs increases:

1. Each LP receives less flow
2. Adverse selection kicks in, especially with the LPs receiving less than approximately 5% of the flow. These LPs were chosen to trade with because they could not get away during high market movement. When they look at these trades, they will then immediately offload the position and create market impact resulting in a losing scenario for both sides.
3. More of the flow is being directed to LPs because of their pricing inaccuracy, and this results in higher rejection rates

By our observation, broad industry consensus is that an optimal pool contains approximately 8 LPs per currency pair. This ensures that all LPs are getting a reasonable amount of flow while reducing the chances for adverse selection among LPs.

5.1 Analyzing Counterparties in an Aggregator

It is crucial to constantly monitor the flow going to each counterparty in your aggregator and engage with them to understand how each manages the order flow. Firms typically look at metrics such as order fill rates, % of flow going to each LP, last look reject times, and cost of each reject. In addition, mark-out profiles can show how the counterparty is managing the flow. A mark-out with an increasing slope after execution could mean the counterparty is immediately hedging the flow, thereby causing market impact. A mark-out with a decreasing slope could mean the counterparty is leading the market and is providing aggressive offers when the market is about to go up (and vice versa).

6. Conclusion and Outlook

The first part of this paper described the challenges of trading OTC FX, such as fragmentation, lack of standardization, and nuances around pricing. The second part explored how the development of an FX aggregator can alleviate some of these issues. We have highlighted some of the considerations and trade-offs needed in constructing such a tool: 8-10 liquidity providers per currency pair, ideally with a degree of homogeneity of pricing and hedging models, provides an appropriate mix. Ongoing analysis of the flow going to these counterparties is always advisable.

Legal Disclaimer

THIS DOCUMENT IS NOT A PRIVATE OFFERING MEMORANDUM AND DOES NOT CONSTITUTE AN OFFER TO SELL, NOR IS IT A SOLICITATION OF AN OFFER TO BUY, ANY SECURITY. THE VIEWS EXPRESSED HEREIN ARE EXCLUSIVELY THOSE OF THE AUTHORS AND DO NOT NECESSARILY REPRESENT THE VIEWS OF GRAHAM CAPITAL MANAGEMENT. THE INFORMATION CONTAINED HEREIN IS NOT INTENDED TO PROVIDE ACCOUNTING, LEGAL, OR TAX ADVICE AND SHOULD NOT BE RELIED ON FOR INVESTMENT DECISION MAKING.