

Taking a more holistic approach to the business of FX Liquidity Management



Nicholas Pratt

Liquidity distributing banks are looking to build more holistic liquidity management systems that are able to cater for a faster, more diverse, more complex and more competitive marketplace. Nicholas Pratt examines some of the challenges involved and the ways they can be resolved.

Harder, Better, Faster, Stronger is the title of a popular song produced by French techno music producers Daft Punk. In addition to selling thousands of copies, it has also become a musical tribute to the ability of technology to make the world a more efficient place. This clearly evident in the following lyrical refrain:

*“Work it harder, make it better
Do it faster, makes us stronger
More than ever, hour after hour
Work is never over”*

This chorus of sorts neatly encapsulates how technology developments have the ability to create far greater processing efficiency yet also generate far more complexity, producing an equilibrium of sorts where the end result is that ‘work is never over’.

Such a thought is evident in the FX industry where algorithmic technology has created a whole range of new trading patterns, participants, platforms, prices and practices, all of which create new challenges. And the one part of the FX trading process where these challenges are possibly felt more than any other is within liquidity management, and particularly for those distributing banks that sit between the 18 or so major market-making institutions that provide the bulk of the market’s liquidity and the rest of the market, from corporates to retail investors.

On the one hand they are receiving more prices from more sources – hedge funds as well as major sell-side banks – and are then distributing them to an ever-more demanding and diverse customer base. Added to all this is the fact that everything has to be done at a space quicker than the speed of light. Some people in the market witnessed a similar phenomenon in other asset classes, such as equities and futures. First of all, pit-based trading was replaced with screen-based trading and then this quickly gave way to the greater use of algorithms and APIs which then created a massive change in the size, frequency and volume of transactions.

Lack of aggregation technology

“When the same thing happens in FX, a lot of firms servicing the buy side will not be ready,” says Kevin Ashby, chief executive of Velsys, a UK-based developer of FX liquidity management software. More and more hedge funds are trading aggressively but most sell-side banks don’t have the low latency technology to match them. In this age of low latency, if you’re 50 millisecond slower than the next liquidity provider, you will get pipped. If you’re a bank aggregating liquidity for mostly corporate clients then latency is not an issue but if you’re dealing with institutional investors or fund managers looking to make money on small price movements in the market, then you will be judged on speed. If there is a divergence on speed between 16 or 17 liquidity providers, then the laggards will either lose business, be taken advantage of or be left with no choice but to start turning clients off”

Liquidity aggregation systems have to be fast and have to have dashboards that make the users look at the end-to-end latency and the dynamics of each channel they are trying to squeeze data down, says Ashby. The systems also need to have comprehensive liquidity architecture in place that comprises a full range of

metrics such as smart order routing, auto-hedging rules, various data sampling techniques and filters that take out volatility, says Ashby. “It is all about using the liquidity more intelligently and such is the speed and size of the market these days, you can no longer have humans doing this.”

Ashby cites the latest BIS figures on the size of the FX market that show volumes between 2004 and 2010 more than doubled, partly due to the growing involvement of non-bank participants. However during the same period the average trade size halved meaning that the number of trades must have quadrupled. In the equities market the number and volume of trades has also increased enormously; algorithmically driven trading is accounting for an increasing proportion of the market; liquidity has





Kevin Ashby

“If you’re a bank aggregating liquidity for mostly corporate clients then latency is not an issue but if you’re dealing with institutional investors or fund managers looking to make money on small price movements in the market, then you will be judged on speed.”

fragmented and smart order routing is becoming a necessity. Ashby sees no reason why the FX market will not follow the same path.

“We are already dealing with so many sources of liquidity and so much pricing data that if you are not using some serious computing in order to do it quickly, you will get lost. The whole world is changing and you have to deal with liquidity intelligently but you cannot rely on the pricing engines to do all of this because they are already doing too much as it is.”

Complex Event Processing

Complex event processing (CEP) is often heralded as the potential saviour for FX dealing banks hit by the deluge of data that has resulted from the increased trading volumes. But Ashby says there are a number of issues with CEP that are not always considered. “For a start there is a substantial cost involved. There are a

number of firms that will offer you simple aggregation technology with all the supply chain management that comes with it. But once firms decide to move on from that and look to bring in CEP software, there is a severe increase in the level of complexity. Not only are you bringing in new software but you also need to hire a team of quantitative analysts to decipher all the data. It is a huge step and not all firms are ready to take it.”

“There are also two sides to using CEP – there is the liquidity coming in from hundreds of data sources that needs to be aggregated and managed to create an internal price. And then there is the pricing data that is distributed to clients. We are already seeing banks struggle to use CEP for both. I think CEP technology has to be used for what it is good at and that means allowing a proprietary trading desk or buy side firm to build and test trading algorithms. Liquidity management, market connectivity and distribution are becoming more and more commoditised, which points towards off the shelf products. Similar to what Patsystems, Trading Technologies and GL Trade did in the futures industry during the last decade”

The talk of low latency is often categorised as an arms race and a battle of resources and spending power but, says Ashby, it is not solely a question of computing power and resources but more about the intelligent techniques that are being used in the programming to reduce the speed of processing. “If you have hedge funds with a team of quants that are coming up with new trading ideas every week and trading in microseconds, you cannot spend days to respond to their bids. And you cannot spend two years developing a CEP engine that tries to respond to all of these bids. What you can do though is have a dashboard that monitors all of these activities and turns off the flow that could be considered toxic – after all, flow is only toxic when you don’t like it. The winning banks will be the ones that have a dashboard rather than a rear-view mirror and have infra-red lights that are able to see through the fog and what is happening in the distance. And the key to liquidity management is intelligent analysis and an understanding of supply chain management.”

Quality of liquidity sources

One of the challenges in the liquidity management space is how to add qualitative analysis to what is essentially a quantitative process – how do you rank the quality of various sources of liquidity? In theory you could use an enhanced pricing engine that

employs complex event processing (CEP) technology to assign a weighting factor to particular liquidity providers, in particular currency pairs in addition to the price of any sourced liquidity, to decide which liquidity provider to execute against. While this may work in theory, we are yet to see a fertile market for such an approach, however there are a rising number of application developers providing exactly this mix of qualitative and quantitative information.

“We are primarily focused on liquidity redistributors that deal with 18 or 19 of the top banks and then feed that liquidity down to their clients,” says Yaacov Heidingsfeld, founder and chief executive of TraderTools. “We gather and provide a wealth of statistics to both the providers and the distributors that tells them who were the most successful trading partners over the course of a month in both specific currency pairs and in all currency pairs – i.e., who offered the best price and who was able to fill the highest percentage of both orders. By using this

information and this technology, both liquidity providers and liquidity redistributors are able to monitor and improve their relationships.”

The developments in technology has forced banks, to some degree, to deal with non-traditional participants such as high-frequency hedge funds and multi-asset traders who are able to place electronic orders through ECNs and compete with the traditional FX banks for business. The new participants are using CEP engines to enter the market and while they do not have the long-established relationships with liquidity takers that the major FX banks have, they are able to offer highly competitive pricing in an increasingly transparent world. And they are being included in the aggregation engines of the liquidity providers.

Two consequences of the increasing number of non-traditional FX participants are an increased focus on technology and greater adoption of the kind of trading tools that have been more commonly used in the equities market, for example. “The technology is constantly pushing the industry to new speed limits. It is not so much a question of latency but if you are using algorithms, the faster you can generate orders, the more successful you will be,” says Heidingsfeld.

New trading models

One of the biggest changes in the FX marketplace that has resulted from the increased use of algorithms, is the fact that not all participants are posting two-way prices, says Heidingsfeld. A lot of the new algo-driven trading models are creating orders that are either *buy* or *sell*, but not both, and posting them on multi-party anonymous platforms. This is a fundamental change from the traditional trading model where banks post *bid* and *offer* prices – a change that has sparked a debate between the traditional and non-traditional participants in the FX market, says Heidingsfeld. He goes on, “The traditional players think that all participants should be forced to issue two-way prices and the non-traditional players disagree, arguing that just because this is the way things have always been done, does not make it right. They argue that they are providing extra liquidity, even if it is one-way, thereby satisfying market demand. If that were not the case, there would not be anyone executing against their offers.”

But from a technology perspective, liquidity management systems have to cater to both models and this is where there can be a knock-on effect in



Yaacov Heidingsfeld

“The liquidity mirage applies to any institution that does not understand that when they make the same price available on multiple venues they are creating the potential to be traded against in a much larger way than originally intended.”

the shape of the liquidity mirage. “A lot of these algorithmically-driven execution systems generate a lot of orders in the same areas. The liquidity mirage applies to any institution that does not understand that when they make the same price available on multiple venues, they create the potential to be traded against, in a much larger way than originally intended. That is why I need to have faster algorithmic ability – because if my price is traded against, I can adjust my positions and prices accordingly,” notes Heidingsfeld.

This is not to say that the order-driven market described above and created by the new participants in the FX market will ever exceed the quote-driven market operated by the major FX market-making banks. But it is indicative of the trend for trading practices and models to be imported from one asset class to another. One such trend is the growing penchant among traders for multi-asset trading and including FX within these strategies.

Heidingsfeld does not see this trend as having a major impact on the FX market. “It is good to have a multi-asset approach to viewing P&L and managing risk but, in terms of trading, I think most participants will continue to position themselves in what they’re good at,” says Heidingsfeld. “So I don’t think liquidity management will become more holistic, it will become more specialised. While the larger players may be able to be all things to all people, their smaller counterparties will be driven towards specific liquidity providers. I think the FX market is big enough to cater to all of these different approaches – quote-driven or price-driven; single-asset class or multi-asset class – without affecting the price dynamics. The most important thing is that liquidity begets liquidity.”

Convergence

According to Carl Martin, group technology director at Eurobase International the big challenge in liquidity management revolves around the blurred boundaries between different participants and what used to clearly constitute the buy-side and the sell-side. “Previously in FX the buy-side would try to get the best price and hit it while the sell-side would just aim to have a competitive price but now things are becoming much more convergent, liquidity is more dispersed and the trading relationship between buy and sell-side is more circular. The best bid and offer is still important but it is also becoming important to protect your liquidity providers.”

As a wide range of participants enter the FX market, there are far more trading strategies and objectives at play, says Martin. “Some may be taking a long view of the market, whereas others may be looking to get out as soon as they can, and others may be simply looking to balance their books. A liquidity management system has to be able to interpret all of this and to put all of this liquidity into context and to do it all on an automated basis.”

Martin refers to this as the ‘liquidity switch’, where feeds are constantly analysed and liquidity providers are judged against one another. Ultimately FX firms have to make a judgement call or an educated guess but through the use of systems and mathematics it is possible to reduce the amount of guesswork needed.

The need for liquidity management systems to keep up with the different strategies and behaviour is indicative



Carl Martin

“The aggressive, alpha generating algorithms are out and the intelligent, cognisant algorithms are in because they are able to absorb much more information aside from simply price.”

—TraderTools' Trade & Earn Program:—

Trade \$1m Earn \$60

At TraderTools, we don't just say that aggregation improves profits, we prove it, with analytics software capable of quantifying ROI from FX liquidity aggregation. According to statistics on actual customer flow in 2010, our aggregation solution improved profits by an average of \$60 for every \$1 million traded, after costs. We systematically collect this business intelligence together with our customers to analyze flow and optimize FX trading.

If you're looking at aggregation software for FX, make sure your vendor can back up claims of profit improvement with actual numbers, the way we do. For more information about how you can start earning \$60 more for every million you trade, contact Mark Mayerfeld, EVP, International Sales at tel. 1-212-782-3550 or email earn@tradertools.com.

TraderTools
The eFXperts

www.tradertools.com



of a general trend that has become evident in the post-financial crisis – that liquidity-taking banks are now being more sensitive to their liquidity providers overall qualities rather than simply judging them on price alone and on a deal by deal basis, says Martin. “I think the financial crisis caused everyone to take a deep breath and reassess exactly what they are trying to do when it comes to FX. There was a realisation that relationships are important.”

Relationships matter

The growth of algorithmic-driven trading models has enabled dealers to quickly analyse multiple prices and execute accordingly but, on occasion, this has involved neglecting the longer-term and more qualitative aspects of the relationship between liquidity providers and liquidity takers, says Martin. “If you put price aside from a moment, there are other things you want from your liquidity provider. For example, you want them to back you once you make a deal with your customer. Ultimately these banks distributing bids and offers are importers and exporters and the one thing that loses them customers is when a displayed price is no longer available because the market has moved against them.”

The re-evaluation of the relationship in FX liquidity management has led to some give and take on both sides – the liquidity takers have agreed to use a longer time horizon when judging their liquidity providers rather than on a deal-by-deal basis and to not simply go with the lowest price on offer. In return, the liquidity providers have agreed to back their clients trading activity even when the occasional deal may lose them money. And this has had an effect on the types of technology used in liquidity management. “There is now a greater level of sophistication involved in the pricing models and the liquidity management systems,” says Martin. “The aggressive, alpha generating algorithms are out and the intelligent, cognisant algorithms are in because they are able to absorb much more information aside from simply price.”

Martin has also noticed the increased presence of multi-asset traders in the FX market and the effect this is having on the development of trading technology and tools. “In the past we had clients come to us in different silos based on asset class but those classifications are becoming more blurred. In FX the corporates at the bottom of the food chain are multi-asset by definition because they are entering the

FX market in order to facilitate another transaction. And at the more sophisticated end of the market, the various triangulations strategies are joining together different asset classes and this is now affecting liquidity providers in the FX market,” says Martin.

It is also affecting the vendors who supply the liquidity management systems, says Martin because those that only operate in a single asset class will find it difficult to adapt to a multi-asset class environment should the participants increasingly seek to use one system to cover multiple asset classes rather than the asset-class specific approach that has prevailed up to now. “I think if you are just an FX platform, then you have a lot to worry about. To develop a new version of the technology and bring it to market is relatively trivial but to develop the ability to add a whole new asset class to your platform is a massive job.”

Challenges with upgrading

For Harry Gozlan, founder and chief executive of SmartTrade, the biggest challenge for FX participants in the liquidity management space lies in upgrading their systems and the complexity involved when migrating from their current system, which may be five or more years old to a more powerful new system equipped with all the latest trading tools and metrics. “They are not starting with a blank sheet. The difficulty is that there may be some of the pieces already in place such as auto-hedging rules or pricing engine but to get everything harmonised and ready at the same time is a huge undertaking,” says Gozlan.

“The only possibility is to implement the solution one component at a time and ensure that you take a very open approach to the development because there are often different protocols involved. This gives banks much more flexibility for the future which is important because much of this technology is still in an evolutionary stage. This is also true of the securities market where a new state-of-the-art product arrives and you need new components to work with it, be that software or hardware.”

So what kind of principles are involved in the development of the latest liquidity management systems? According to Gozlan, the method behind the metrics is much as it always was. “In the classical sense it is all still about extracting maximum value with minimum risk and all of the various metrics – real-time positioning, auto-hedging rules and real-time

profit and loss – are important because you have to be able to react to market prices. But it is a more volatile market currently so the way that banks are able to distribute prices to their clients is important because it must be suited to their trading styles.”

The adaptation to different customer behaviour is seen in the frequency of hits that a bank is willing to accept from each customer, how much information is shared by the clients (as is the case in a quote-driven trading model), the throttling aspects, the provision of a last look facility and many other things that were previously done on the phone but now must be done through the use of intelligent technology, says Gozlan.

“I don’t think it is a problem of a fundamental change in customer behaviour but the fact that there are more different types of customer and the fact that everything is being done faster and in a more competitive environment. Previously when everything was voice-based you could give everyone the same price and then we went to a tiered approach of gold, silver and bronze for the different customer types. But now, if you want to extract maximum value from the market, you need hundreds of different tiers in your distribution engine.

“Real-time hedging must be based on clever analysis of the risk and the price you give to the client has to be able to adapt to their changing behaviour and the changing prices in the market. And all of this has to be done fast enough or else you will not end up with the trades that you want. But you cannot have all of this capability acting on a high frequency in one machine. The most important point in building a liquidity management system is that it must be properly designed.”

Part of this intelligent design is the correct use of technology, says Gozlan, particularly something like CEP. “Complex event processing is a great technology for developing a clever calculation engine but it is not a trading technology. The idea is to have lots of price feeds going into the CEP engine, then embed some rules within it that enable it to calculate the different margins involved for each client before sending prices out. But CEP should not be used as a trading platform or to execute trades because it is essentially an analysis tool.”

Holistic approach

The other key to developing a more holistic approach to liquidity management, says Gozlan, is to make



Harry Gozlan

“Liquidity management is a whole job in itself, so it is not something you can approach with an improvised attitude.”

greater use of outsourcing and external hosting services on offer. This makes it easier for banks to set up a centralised approach within their own set-up and rely on external providers to ensure all the necessary connections to external sources of liquidity. “It enables those banks that are in the very highest tier to solve all of their networking issues without having to rely on a large and well-resourced internal IT department. And if these externally hosted components are able to sit closer to the connectivity points and if they can be controlled by very precise executions, then the liquidity management will be more efficient than if everything was done through an internal network,” he says, “The larger banks may still be able to cover all of this internally but for the smaller banks, such an approach is almost mandatory. Liquidity management is a whole job in itself, so it is not something you can approach with an improvised attitude,” Gozlan concludes.