

Date: September 20, 2017

To: Global Foreign Exchange Committee ('GFXC')

Cc: FXWG Chairman Guy Debelle
MPG chairman David Puth
GFXC Workstream Leads, Chris Salmon and Simon Potter

Feedback on Last Look practices in the Foreign Exchange Market from LMAX Exchange

Information about the respondent:

LMAX Exchange is authorised and regulated by the FCA as a Multilateral Trading Facility (MTF). LMAX Exchange operates a central limit order book (CLOB) with streaming, no 'last look' liquidity only supplied by institutional market makers, banks and non-banks.

Trading on LMAX Exchange is governed by the [LMAX Exchange Rulebook](#) which does not permit 'pre-hedging' and 'last look' practices, thus ensuring fully transparent and fair execution for all its clients and market makers.

LMAX Exchange is known in the FX industry for its long-held conviction that the 'last look' practice is open to abuse and it should not exist in a transparent, fair FX marketplace.

www.lmax.com

Response to Specific Consultation Questions:

[As the first market participant to commit to the FX Global Code](#), LMAX Exchange welcomes the opportunity to provide feedback on Last Look practices in the Foreign Exchange Market.

Though the FX Global Code is a positive starting point for restoring trust in the FX industry and creating globally consistent guidance, LMAX Exchange doesn't believe that the Code goes far enough on restricting or banning the potential market abuse that can result from the use of 'pre-hedging' and 'last look', in its wording of Principle 17.

Questions: Principle 17 of the Code states that *"During the last look window, trading activity that utilises the information from the Client's trade request, including any related hedging activity, is likely inconsistent with good market practice because it may signal to other Market Participants the Client's trading intent, skewing market prices against the Client, which (1) is not likely to benefit the Client, and (2) in the event that the Market Participant rejects the Client's request to trade, constitutes use of Confidential Information in a manner not specified by the Client"*.

LMAX Exchange agrees that any trading activity, utilising the Client's order information, during the 'Last Look' window does not benefit the client and at the very least, constitutes use of confidential information in a manner not specified by the client. Furthermore, we are not aware of any situation or scenario where pre-hedging during the 'Last Look' window can be beneficial to the client or where clients benefit from skewed market prices against their orders, caused by information leakage during the 'Last look' window.

The current wording in Principle 17 effectively legitimises pre-hedging, which could stand accused as front-running during the 'Last Look' window. Legitimate working of a client order will be indistinguishable from unethical front-running for pure profit-making utilising privileged Client's order

information. Front-running is considered 'unethical practice' in capital markets, defined as *'unethical practice whereby someone with advance knowledge of a specific market order in, say, shares, bonds or a currency from a client steps in ahead and buys for their own account. When the client's usually much larger order is executed and drives up the price, the private purchase can be sold at a profit'*. Thus, if front-running is acknowledged as 'unethical' across all asset classes, why isn't there a stronger stance in the Global Code on pre-hedging activity during the 'last look' window?

Thus, as a direct response to the consultation questions, LMAX Exchange recommends removing **'likely'** from **'likely inconsistent with good market practice'** in the wording of Principle 17, referring to 'any hedging activity during the last look window utilising the information from the Client's trade request.

Longer-term, LMAX Exchange believes that the Code should ban 'last look' at least on anonymous multi-dealer trading venues; it can be argued that the practice may still have its place in the disclosed bi-lateral trading relationships (i.e., bank to specific client), if both counterparties prefer to trade with 'last look'. Banning 'last look' will avoid any potential for market abuse, that has already been evidenced by recent scandals and legal investigations for the misuse of 'last look' by some of the most reputable, global financial institutions. LMAX Exchange believes that the need for 'last look' has become obsolete; the technological advancements and availability of real-time streaming market data, enabling instantaneous price checks, have entirely eliminated the need for 'last look' as a risk management tool. The practice of 'last look' that doesn't exist in any other asset class, erodes trust in FX trading at the time when the industry needs to reinstate much-needed transparency and fairness in FX markets.

LMAX Exchange viewpoint on the practice of 'last look':

1) Advances in trading technology have replaced the need for 'last look' with more superior risk management tools:

The practice of 'last look' is a legacy business solution to an historic technology problem. As is the case now, market makers reasonably wanted to protect themselves against sudden fluctuations in the market (as described in Principle 17 of the FX Global Code of Conduct¹) when they didn't have robust technology to stream prices without introducing enormous amounts of new market risk into their businesses. At the time, the obvious patch for the shortcoming in technology and the continued desire to responsibly manage risk while at the same time increase the breadth of the business' market making reach, was to introduce what is now known as 'last look'. Today, the same financial institutions have invested greatly in people, technology platforms, and electronic trading is no longer a small offshoot of a bank's core FX business it is the core FX business. This attraction of resources and advances in technology platforms across the marketplace, also calls for the evolution of market standards and best practices.

Our view is clear at LMAX Exchange. We can process over one hundred million orders per day, cancel and replace orders in sub-one-hundred microseconds, perform seven million real time risk calculations per second, and conduct pre-trade credit checks instantaneously. Technology has moved on dramatically, and the same needs and rationalisations for 'last look' are no longer the same.

In today's environment, we feel it creates an uneven playing field biased against clients, whether they are cognisant of it or not, and equally as important, against financial institutions who are trying to create a new market order with greater transparency and equality, yet are forced to

¹ Principle 17, p.21 '*..last look is a risk control mechanism used in order to verify validity and/or price. The validity check should be intended to confirm that the transaction details contained in the request to trade are appropriate from an operational perspective and there is sufficient available credit to enter into the transaction contemplated by the trade request. The price check should be intended to confirm whether the price at which the trade request was made remains consistent with the current price that would be available to the Client'*

compete with those who are not yet willing or mandated to do so. It can be a very daunting task if not every market participant operates to the same level of standards.

LMAX Exchange is not a proponent of 'last look' for numerous reasons, including the possible abuses and optionality that it introduces into pricing. LMAX Exchange, along with a small handful of other venues, is demonstrable proof that trading without 'last look' is a wholly viable option for both market makers and the clients who take their liquidity. Moreover, by doing so, these participants disassociate or distance themselves from a practice which is open to abuse, while at the same time improving both the transparency and quality of execution in the market place. This is not a transition that takes place overnight, but sooner, rather than later, market participants will have to accept the mature, transparent nature of the foreign exchange market and its place and status with other mature asset classes and their markets and realise the practice of 'last look' is no longer defensible.

2) 'Last look' creates a disorderly market and liquidity mirage in the anonymous multi-dealer execution environment:

Though LMAX Exchange believes the need to use 'last look' has become obsolete, it can be argued that the practice may still have its place in the disclosed bi-lateral trading relationships (i.e., bank to specific client), if both counterparties prefer to trade with 'last look'.

The situation is different on anonymous multi-dealer platforms, where clients are trading on anonymous quotes streaming from multiple LPs, each using 'last look' according to their own discretion. Furthermore, 'last look' on multi-dealer platforms allows LPs to quote more venues than they are willing to fill, causing a liquidity mirage and increased fragmentation, in turn leading to disorderly markets.

3) The 'last look' practice significantly diminishes the trader's control over execution quality and costs, thus creating opportunities for market abuse and undermining any regulatory initiatives aiming to impose stricter controls over execution factors:

Discretion over the LP's use of 'last look' and its consequences on the trade execution quality have been evidenced in our recent white paper ['TCA and fair execution. The metrics that the FX industry must use'](#). The paper, containing the analysis of the independent data set of over 7 million trades (both firm and 'last look' liquidity), revealed the significant level of discretion used by 'last look' liquidity providers regarding:

- The length of 'last look' window or discretionary hold time before order execution – constituting one of the most significant hidden trading costs (e.g., 100 milli-seconds of hold time can cost the client up to \$25/million²)
- The bias in passing the underlying market behaviour on limit-orders to clients – our analysis demonstrated that clients trading on 'last look' liquidity were not getting the full price improvement, costing them up to \$40/million³ in unrealised value from price improvement

Such discretion highlights the potential for market abuse of 'last look' – for detailed analysis, please refer to Appendix I, containing selected chapters from the LMAX Exchange white paper ['TCA and fair execution. The metrics that the FX industry must use'](#).

Furthermore, the level of discretion that LPs have over the order during the 'last look' window makes it impossible for trading institutions to have sufficient control over execution quality, in turn undermining any regulatory efforts that impose stricter levels of responsibility on traders for achieving best execution.

² ['TCA and fair execution. The metrics that the FX industry must use'](#), Part II - Quantifying the cost of hold time, p.60-63

³ ['TCA and fair execution. The metrics that the FX industry must use'](#), Part II - Quantifying the value of price improvement, p.50-55

To this extent, we believe that in the context of 'last look' liquidity MiFID II best execution standards are unattainable for the buy-side. Only when trading on firm liquidity, the buy-side participants are able to 'take all sufficient steps' to obtain best possible results for execution factors such as price, costs, speed, likelihood of execution and settlement, size and nature of the trade⁴.

4) Ongoing scandals and investigations related to the abuse of 'last look' provide further evidence that the practice is misused and in turn deters trust in the FX marketplace:

Asymmetrical application of 'last look', failure to pass fully and transparently price improvement to clients as well as consistent use of pre-hedging during 'last look' window are examples of 'last look' abuse in the recent scandals and ongoing investigations:

- Barclays Bank fined \$150m for abuse of 'Last Look' by the NYDFS (November 2015)
- Legal claim by NFA brought against FXCM and Effex Capital (February 2017)
- Legal claims filed by Alpari (US) against 6 banks for abuse of 'Last Look' (July 2017)

5) Not banning or taking a stricter stance on 'last look' diminishes the Code's intentions:

The Code's intentions are stated as follows: *'to promote a robust, fair, liquid, open, and appropriately transparent market in which a diverse set of Market Participants, supported by resilient infrastructure, are able to confidently and effectively transact at competitive prices that reflect available market information and in a manner that conforms to acceptable standards of behaviour'*.

Our view is that **robustness** of the market is diminished by 'last look' and the discretionary nature of the duration of 'last look window'; **fairness** is impaired by not banning pre-hedging during the 'last look' window and not requiring Market Participants to pass on price improvement to Clients; **liquidity of the market** is affected by liquidity mirage and fragmentation, enabled by 'last look'; **openness** is deterred by Market Participants having sole discretion over the use of controversial practices.

- **The use of 'last look' at a sole discretion of a Market Participant creates conflict of interest with the Client and takes away from the Client control over execution quality and trading costs.**

The Code rightfully states that during the 'last look' window the Client is left with 'potential market risk in the event the trade request is not accepted'. Since this is a substantial risk, the Code needs to elaborate on the benefits for the Client to trade with 'last look' and make it obligatory for the Market Participant to pass on all the benefits that may arise from 'last look' execution to the Client. To this extent, the Code needs to raise awareness of different types of liquidity available, firm and 'last look', and provide examples and guidance to help Clients make informed choices when selecting the appropriate liquidity for their trading strategies.

- **Standardised application of price improvement vs price slippage to Client limit orders needs to be enforced on 'last look' liquidity.**

The Code needs to take a strong stance that Market Participants are required to pass full price improvement, resulting from market fluctuations during 'last look' window to the Client, in the same way as they treat price slippage. By not giving full price improvement on a 'last look' stream, though fully disclosed, the Market Participant is disadvantaging the client, which contravenes Principle 8.

- **Monitoring and enforcing 'correct use' of information from the Client's trade request during the 'last look' window is close to impossible on multi-dealer platforms.**

⁴ MiFID II Best Execution obligation

Information about the Client's order, in advance of 100% execution, is valuable and open to abuse. Even if the information is not used for hedging activity, it can be used for future pricing by Market Participants (i.e., determining, spreads, fill ratios, hold time).

Principle 17 proposes, as a good market practice, that Market Participants engage in a dialogue regarding the handling of their trade requests with Clients. This proposal is constructive and relatively easy to implement and enforce in bi-lateral disclosed trading relationships. Unfortunately, on multi-dealer platforms where Clients are trading on 'last look' liquidity streaming from many different LPs (each with different disclosures about the use of 'last look') and where information is passed around in milliseconds / or even micro seconds, it becomes very difficult, if not impossible, for Clients to monitor whether their trade requests are treated in accordance with disclosures and confidentiality.

Furthermore, the encouragement of the dialogue between the Client and the Market Participant is too reliant on the following assumptions about the FX marketplace:

- Lack of conflict of interest between Market Participants and the Clients;
- Open access for every Client, whether big or small, to have the same dialogue with every Market Participant, who may be pricing their orders;
- Full awareness and understanding by Clients of all available execution alternatives and the associated trading costs for each alternative.

Unfortunately, the FX marketplace doesn't operate in accordance with the above assumptions:

- Until Market Participants are required to pass full price improvement, when using 'last look', to the Client, the inherent conflict of interest exists between the two parties;
- Unless there is standardisation of the use of 'last look' by LPs for each specific multi-dealer venue, it's impossible for all clients to have similar information access to LPs' disclosures and to have the same access to each LP to discuss individual disclosures;
- Finally, in the context of the OTC-traded FX marketplace that operates without centralised pricing benchmarks, the only way clients can make informed choices about the liquidity source for their execution strategy is to conduct FX Transaction Cost Analysis (TCA). Compared to equities, FX TCA methodology is still in its infancy - it doesn't address differences between execution costs on firm vs 'last look' liquidity and it hasn't reached sufficiently wide adoption by the marketplace to become a useful blueprint for Clients in their dialogue with Market Participants.

Effectively, Principle 17 places the complete responsibility on the Client for understanding 'last look' and 'pre-hedging' related disclosures (driven by the sole discretion of the Market Participant), without removing the conflict of interest between the two parties and without equipping the Client with capabilities and tools to have an informed dialogue with the Market Participant. As a result, Principle 17 reinstates sole control over execution with Market Participants and doesn't contribute to levelling out the playing field between Market Participants and the Clients.

Detailed Recommendations:

Short-term:

- Remove '**likely**' from '**likely inconsistent with good market practice**' in the wording of Principle 17, referring to 'any hedging activity during the last look window utilising the information from the Client's trade request';
- Remove discretion from Market Participants on how they treat price changes, resulting from market fluctuations during 'last look' window, and require them to pass the full price improvement on limit orders to Clients, in exactly the same way price slippage is treated;
- Raise the level of awareness of execution alternatives for the Clients and promote a standard set of metrics that Clients can use to calculate total trading costs for each alternative. For this, the FX TCA methodology needs to be developed to capture a comprehensive set of metrics

that measures execution quality across both firm and 'last look' liquidity. The development of such FX TCA methodology needs to be an industry-wide initiative, intended to equip Client with the ability to make informed choices and have dialogues with Market Participants about the use of 'last look' on a level playing field. LMAX Exchange has made some headway in developing FX TCA methodology able to capture the nuances both liquidity pools (*see the latest TCA white paper*). LMAX Exchange would welcome cooperation from industry participants to develop the methodology further and would be happy to contribute to any industry-wide activity targeted at educating the marketplace about execution alternatives.

Longer-term:

- LMAX Exchange believes that once 'pre-hedging' is forbidden during the 'last look' window and Market Participants have to pass full price improvement to Clients, the practice of 'last look' will cease to exist by itself, as in today's electronic FX marketplace which operates in micro-seconds, Market Participants have much more superior risk management tools than 'last look';
- However, if the conflicts of interest between Market Participants and Clients, inherent in current wording of Principle 17, are not addressed, it would be a far more efficient use of industry/regulatory resource to ban 'last look'. Eliminating potential for abuse whilst insuring fair and transparent treatment is to not allow either 'pre-hedging' or 'last look'.

APPENDIX I

In the latest white paper [TCA and fair execution. The metrics that the FX industry must use](#), LMAX Exchange analysed execution quality on firm vs last look liquidity, using the independent Third Party Aggregator data from over 7 million trades (sent to 7 'last look' and firm LPs) during 2016. Among other findings, the analysis revealed discretionary nature used by 'last look' liquidity providers in passing on price improvement and the duration of hold time or 'last look' window applied to the orders. Such discretion highlights potential for market abuse of 'last look', below are the extracts from the white paper related to Price Variation and Hold time analyses.

PRICE VARIATION - SLIPPAGE AND PRICE IMPROVEMENT

(p.23 in the full [LMAX Exchange TCA and fair execution white paper](#))

Price variation is a trader's view of the difference between a desired or expected price and the actual execution price achieved by an order. While attention is often focused on slippage (i.e. execution at a worse than expected price) when using market orders we should expect to experience both slippage and improvement. Traders using price constrained orders (limit or PQ) may have been conditioned to expect neither; limit orders cannot slip and many traders do not even consider measuring price improvement.

Measurement of slippage or improvement requires information which may only be available in the trader's own logs. We cannot rely on orders to carry the price which prompted the decision to trade – market orders do not carry a price at all and the price on a limit order is not necessarily the same value as the decision price – making this metric potentially both opaque and highly subjective. However, the order placement behaviour of the TPA is far more predictable, allowing us to measure the impact of price variation consistently and objectively across LPs.

When the TPA receives a customer order, it waits until the market data it receives from the LPs indicates that the order can be filled, meeting all price or size criteria specified. Once suitable market conditions are identified the TPA selects one or more LPs, captures the current best price on the relevant side of the market and sends some or all of the order to the selected LPs for execution as a 'leg'. We have calculated slippage or price improvement per leg by looking at the difference between the logged market price at the time the decision to trade was made and the actual fill price received. This approach removes much of the individual variation from the data, treating the TPA as a single customer trading with each of the LPs and requesting the current price available for immediate execution.

We have excluded numbers from infrequently traded currency pairs (any instrument with less than 100,000 trades over the 12 month period of the data set). The remaining sample set consists of trades in EURUSD, GBPUSD, USDJPY, AUDUSD, GBPJPY, USDCAD, EURJPY, EURGBP, NZDUSD, USDCHF, EURCHF, EURAUD, AUDJPY and AUDCAD, which together represent 91% of all successful trades.

We have reported slippage and improvement using the FX conventions of 'pips', i.e. the 4th decimal place of the price other than for currency pairs priced in JPY where the 2nd decimal place is used. This introduces some comparability issues across currency pairs and over time, for example 1 pip GBPUSD is a smaller proportional slippage than 1 pip AUDUSD, and a 1st January GBPUSD pip is a smaller proportional slippage than a 1st November GBPUSD pip due to the depreciation of GBP over the year. However as all pip values fall within a range close to 0.01% of traded price (between 0.006% and 0.016% at the extremes) we have erred on the side of using familiar units over something abstract but more mathematically accurate such as basis points.

Market orders

Table 6 shows the proportion of market orders receiving fills where prices showed slippage, were as expected or showed improvement.

Venue	Slippage	As expected	Improvement	Ratio of slippage to improvement
Bank 3	0.00%	100.00%	0.00%	0
Non Bank 2	1.14%	98.18%	0.68%	1.68
Non Bank 1	19.40%	70.04%	10.56%	1.84
LMAX Exchange	4.36%	93.54%	2.10%	2.08
Non Bank 3	0.64%	99.15%	0.21%	3.05
Bank 2	3.47%	95.65%	0.88%	3.94
Bank 1	7.16%	92.15%	0.69%	10.38

Table 6: TPA market order price variation statistics

Chart 1 shows the percentage of orders that experienced slippage or improvement at 0.1 pip intervals. Negative numbers indicate slippage (a worse price than expected) while positive numbers indicate price improvement (a better price than expected).

In addition to the skew and shape of the distribution, it is important to note the scale is limited to +/- 5 pips for illustrative purposes. In many cases the maximum improvement observed is less than 5 pips (denoted by the green marker) whereas the maximum slippage observed is in most cases more than 5 pips away from the zero point (indicated by the red marker). In the TPA data, only LMAX Exchange exceeds 5 pips of price improvement.

The price variation of market orders falls into two distinct categories. There are those venues which show both slippage and improvement at an approximately 2:1 ratio and those for which the slippage is dominant with little or no price improvement.

As LMAX Exchange operates a firm central limit order book offering best execution in price-time priority, we might expect a more neutral result. The skew towards slippage suggests that behaviour in this data set is linked to market direction, demonstrating a propensity towards buying in a rising market and selling in a falling market. This leads to a natural bias towards slippage and away from improvement. If we take LMAX Exchange behaviour as an approximation of the pure market, then this ratio becomes an interesting metric for market order price variation. This allows us to distinguish between those venues which are passing the underlying market price behaviour straight through to the customer against those which show a higher bias towards slippage.

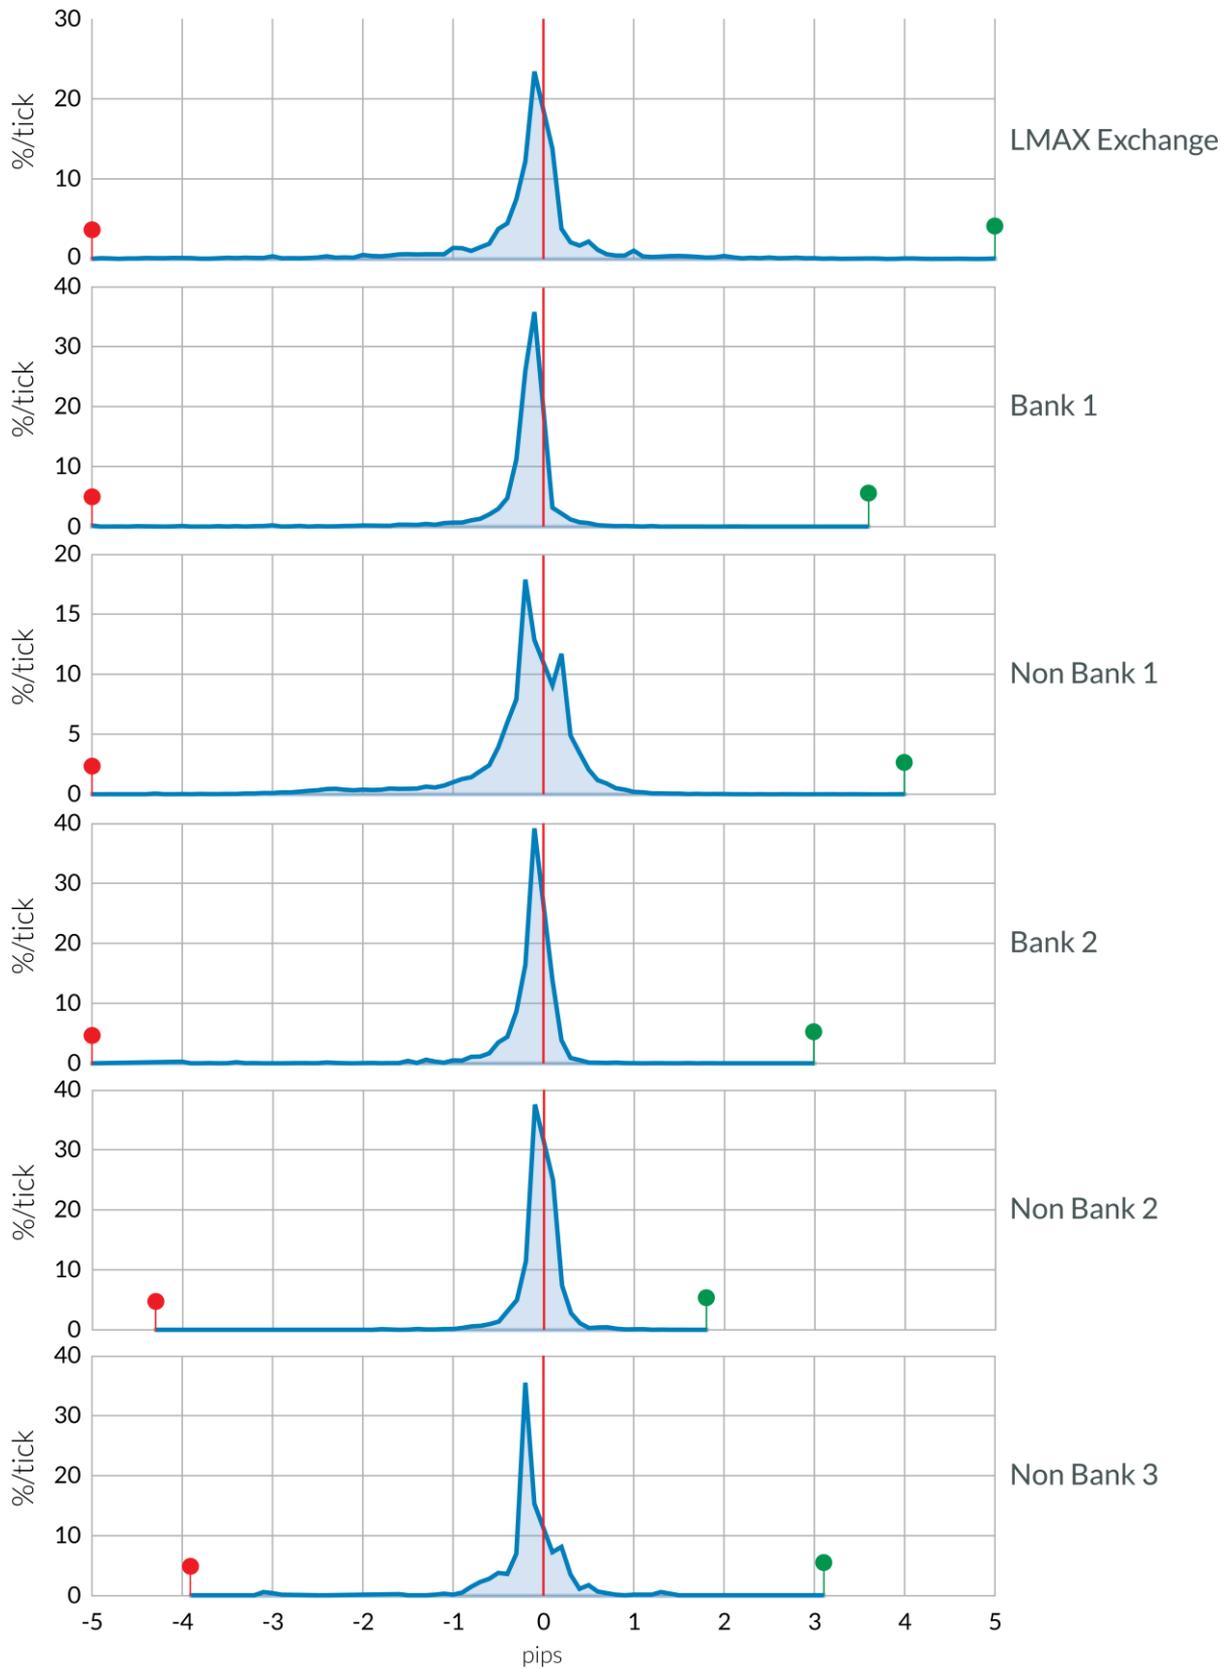


Chart 1: Market order slippage by venue

Limit/PQ orders

The situation for order types with price constraints is more interesting. These order types prohibit slippage, and the TPA sets its limit price to the same value it uses as a reference level to calculate slippage or improvement for market orders, so naively we might expect that the price variation for such orders would have a similar incidence and distribution to the price improvement side of the market order charts shown above.

With the exception of LMAX Exchange, this is not the case. Table 7 shows the proportion of limit or PQ orders receiving price improvement by venue, alongside the market order price improvement from the same venue for comparison.

Venue	Order type	Improvement	Market order improvement
LMAX Exchange	Limit	6.358%	2.10%
Bank 1	PQ	0.001%	0.69%
Non Bank 2	PQ	0.000%	0.68%
Non Bank 3	Limit	0.000%	0.21%
Non Bank 1	Limit	0.000%	10.56%
Bank 2	PQ	0.000%	0.88%
Bank 3	Limit	0.000%	0.00%

Table 7: TPA limit/PQ order price improvement statistics

Only LMAX Exchange exhibits a significant level of price improvement for limit orders. Improvement is either negligible or entirely absent for limit orders executed on all other venues. As a further illustration of the mechanism driving limit order price improvement, chart 2 shows the distribution of the level of improvement received by both limit and market orders on LMAX Exchange, showing the percentage of orders that received improvement at 0.1 pip intervals

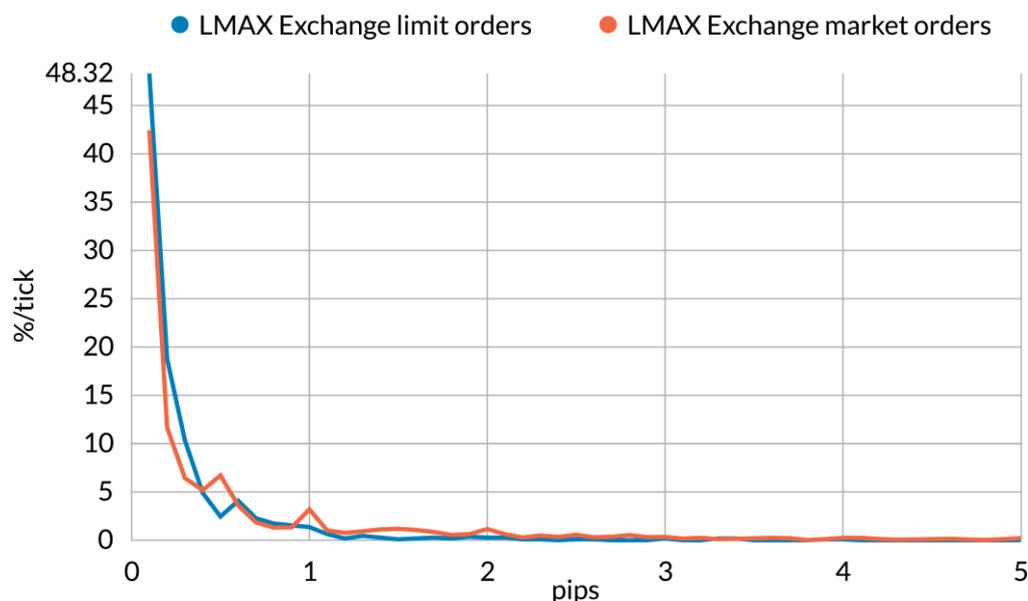


Chart 2: Price improvement for LMAX Exchange market and limit orders

The consistent distribution of price improvement observed for both order types is a key characteristic of firm liquidity. Limit prices only constrain the worst execution price for an order. When better prices are available, limit and market orders behave identically.

In contrast, the very different price improvement behaviour observed for market and limit orders on last look liquidity demonstrates a fundamentally different approach to filling limit orders in which LPs exercise their option to fill almost every order at its limit price, even though the evidence of fills on market orders indicates that a better price should be available for some proportion of the time.

HOLD TIME AND EXECUTION LATENCY

(p.29 in the full [LMAX Exchange TCA and fair execution white paper](#))

Execution latency is the time taken between an order being transmitted from the trader's system and the receipt of a response. Hold time is the commonly used name for discretionary latency where the execution of an inbound order from a trader is deliberately delayed pending a decision to fill or reject by the liquidity provider's systems. This period of time is also referred to as the last look window.

Hold time/discretionary latency is just one component of execution latency, so we must first look at other causes of latency before we can assign hold times to each venue in order to compare this aspect of the execution quality of last look and firm liquidity.

We will divide execution latency into the following components:

Systematic. The time required to complete the necessary operations to execute the trade, including network round trip time, transit through any pre-trade risk control system, matching engine cycle time and any other systematic delay applied across all customers of the LP;

Tail. Each cause of systematic latency will also have a characteristic jitter with causes at network, operating system or application level. In addition, platform capacity constraints ranging from microbursts to sustained higher traffic rates during market announcements can lead to queueing and congestion giving a familiar long tail latency distribution;

Discretionary. Any time added where the order is held prior to executing a trade. LPs may apply or vary hold time based on their assessment of a customer's market impact, the current market conditions or their own appetite to trade in a given direction.

Each of these components is subject to variation over time. Systematic latencies may be affected by hardware or software upgrades which may change the LP's latency profile. Tail latencies may likewise be affected by capacity upgrades or constraints. Lastly hold time may be adjusted by LPs in response to a change in market conditions, strategy, policy or simply based on developing insight into a customer's trading behaviour.

While we are primarily concerned with discretionary latency in the direct comparison of firm and last look liquidity, information regarding the non-discretionary causes of latency is also valuable in its own right, as this can be used to make order routing decisions as well as for TCA purposes. For example, if the latency of a particular LP degrades badly during busy times, this information may be used to augment best price or volume criteria in selecting an execution venue.

Chart 3 shows the execution times for rejects and fills for a particularly interesting last look LP in the TPA data set providing a clear example of each of these different types of latency. The execution time is recorded to the nearest millisecond and the frequency of occurrence is shown on a logarithmic scale. The chart spans the whole year of 2016.

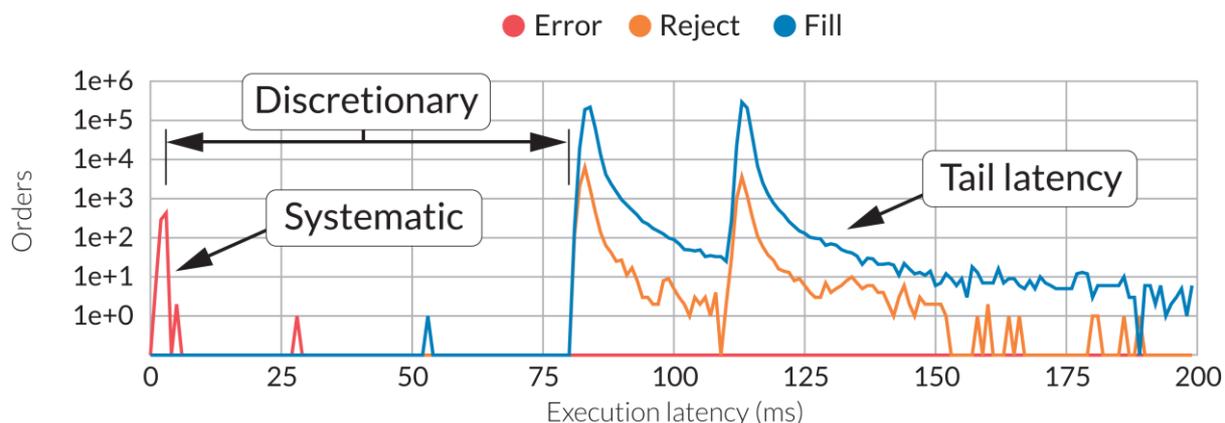


Chart 3: Detailed execution times for Bank 3

An analysis of this kind would normally use supplemental information gathered by the trading infrastructure to determine some parts of the systematic latency. For example the base network latency can be estimated by using session level FIX messages – heartbeats or test requests – which are typically processed at the edge of the LP’s trading platform. Unfortunately that level of data was not available to us in the TPA data set, and we were then forced to determine the systematic latency from the execution time profiles available. Fortunately there are some markers in the data that can help us.

For the LP in chart 3, there is an interesting pattern in that fills and non-error rejects indicate that the minimum response time is around 81-82ms. However, when we looked at rejects due to errors – as defined earlier – a response time of 2-3ms is evident. 99.7% of these errors were caused by a reject at the pre-trade risk control level, rather than a programming or FIX session level error. This is then an error from within the platform – not an immediate reject at the edge.

With the moderate assumption that the next logical step within the platform would be matching the order against available liquidity, we can then assign a systematic latency of at least 2-3ms. The discretionary latency or hold time would then be 80ms for this LP. It is unlikely that an order would transit the network and pre-trade risk control systems within 3ms and then take a further 80ms to be placed unless there was a hold time in play.

The execution latencies for all of 2016 for each LP in our set are shown in chart 4 (below), which plots the millisecond latencies for fills, errors and rejects against the number of orders experiencing that level of latency. There are several features which stand out and bear further investigation:

- Histograms for the same class of event (e.g. fills) which display multiple peaks in the latency histogram;
- LPs where the peaks for fills, errors and rejects occur at different modal latencies;
- Long tails to execution latency distributions.

Our first task is to investigate each of the features above so that we can determine a characteristic systematic latency and hold time for each LP. We will investigate the first 200 ms of latency in detail. In some cases the latency distributions extend beyond this, however, latencies much beyond 200ms are usually a very small proportion of trades and our goal here is to derive the base characteristics of hold time for each LP.

Defining the systematic latency as being the mode of the first peak in the execution time histogram (whether from fills, rejects or errors) and the hold times as being the difference between the systematic latency and the mode of the second peak, we can produce the following table of systematic latencies and hold times. Rejects and fills are examined separately as their latency histograms may differ as in the example above.

Venue	Systematic (ms)	Fill hold time (ms)	Reject hold time (ms)
LMAX Exchange	1	0	0
Bank 1	4	5	1
Non Bank 1	1	90	90
Bank 2	1	9	5
Bank 3	4	80	79
Non Bank 2	1	0	0
Non Bank 3	1	0	0

Table 9: First glance modal hold times by LP

A quick comparison of table 9, which attributes very similar latency profiles to Non Bank 2, Non Bank 3 and LMAX Exchange, and chart 4 (p. 32), which shows a very different visual signature for each, indicates that our initial scorecard is not telling the whole story. Below are relevant extracts related to the analysis of use of Price Variation and Hold Time

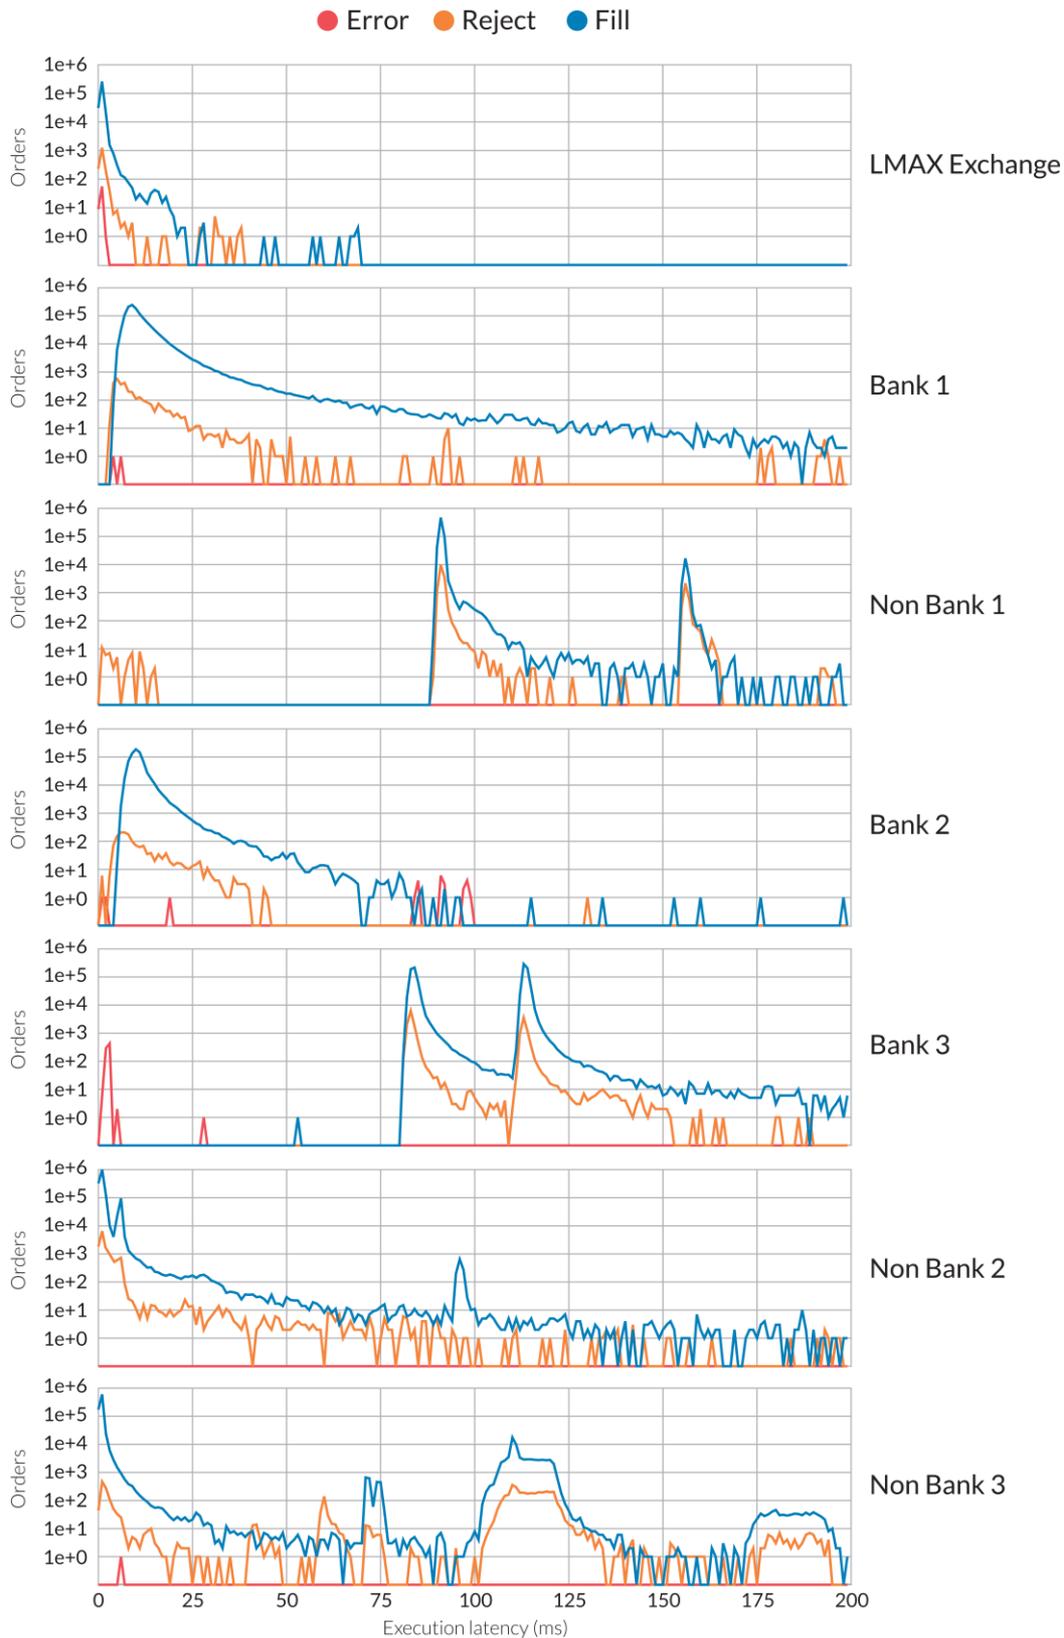


Chart 4: Execution times by venue for 2016

The multiple peaks and wide variation in the tail distribution of the latency histograms displayed by last look LPs require further investigation, and are suggestive of arbitrary changes to discretionary latency which, by definition, do not occur on firm liquidity.

APPENDIX II

Examples of law suits related to abuse of 'last look'

<p>Barclays Bank fined \$150m for abuse of Last Look by the NYDFS http://www.dfs.ny.gov/about/press/pr1511181.htm</p>	<p>Nov'15</p>
<p>Legal claim by NFA brought against FXCM and Effex Capital https://www.nfa.futures.org/BasicNet/Case.aspx?entityid=0308179&case=17BCC00001&contrib=NFA</p>	<p>Feb'17</p>
<p>Legal claims filed by Alpari (US) against 6 banks for abuse of 'Last Look' http://www.profit-loss.com/articles/news/banks/six-banks-face-lawsuits-over-last-look</p>	<p>Filed in Jul'17</p>