NORGES BANK
INVESTMENT MANAGEMENT

01 | 2015

SOURCING LIQUIDITY IN FRAGMENTED MARKETS

ASSET MANAGER PERSPECTIVE

ISSN 2387-6255

As a large participant in global asset markets, Norges Bank Investment Management has to ensure that our liquidity sourcing strategies reflect ever-changing market conditions. In particular, off-exchange trading volume has evolved and increased in recent years, requiring new strategies to successfully access this liquidity potential. Much of this volume occurs in trading venues that have limited pre-trade transparency – often referred to as 'unlit' or 'dark' venues. These venues have emerged to address two distinct developments in global asset markets – the increased 'institutionalization' of asset management on the one hand, and the advent of computer-based trading and the emergence of new forms of liquidity providers such as high-frequency traders on the other[1]. Based on our own investment and trading experiences, we evaluate whether trading venues with limited pre-trade transparency contribute to well-functioning markets, and present a wish list of venue characteristics.

1 *As background reading on a related topic, see "High Frequency Trading – An Asset Manager's Perspective", Norges Bank Investment Management, Discussion Note #1 (2013).*

Date 17/04/2015

The Asset Manager Perspective series articulates Norges Bank Investment Management's views and reflections on issues topical for the financial industry. They are not meant to be definitive, rather they are intended as timely contributions for the benefit of all market participants. The series is written by employees, and is informed by our investment research and our experience as a large, long-term asset manager.

Contact information:
amp@nbim.no
www.nbim.no

Introduction

Equity trading has become increasingly fragmented in recent years - more than 30% of US trading volume now occurs outside the recognized exchanges, for example. This is often referred to as 'dark volume', and the trading venues as 'dark pools'. We believe that this is a somewhat unfortunate generalization. While non-exchange trading venues are characterized by limited pre-trade transparency about the intent to trade, they differ substantially in their organization structures, their matching protocols and the way in which they are used. Assessing their contribution to equity markets thus requires a more nuanced analysis. In this paper, we will provide a taxonomy of such trading venues along these lines.

The limited pre-trade transparency differentiates these venues from exchanges, where limit orders advertise the intent to trade and hence provide optionality indiscriminately to all other market participants. In this note, we will refer to trading venues with limited pre-trade transparency as 'liquidity pools', differentiating them from exchanges.

The rise of such trading venues and their comparative lack of transparency has meant that they are often seen as one of the potential problem areas of modern market microstructure. In particular, they may lead to an efficiency drag on the price discovery process, if information about the intent to trade is not disseminated broadly.

As a large participant in asset markets globally, Norges Bank Investment Management has a more differentiated view of the role that liquidity pools can play in market clearing and an efficient price discovery process. There is a wide range of trading venues that operate with limited or no pre-trade transparency. These venues are utilised at different stages of the investment process, either directly by the investor, or by a broker employed as an agent. The impact of such liquidity pools on market quality needs to be analysed based on these different uses.

Liquidity Pools and Well-Functioning Markets

Norges Bank Investment Managment assesses the impact of specific market structure elements through our 'Well-Functioning Markets' framework[2]. Do liquidity pools contribute to well-functioning markets? For present purposes, we define this as supporting a market structure that maximises natural liquidity (long-term, natural buyers and sellers can find each other with high probability) while minimizing cost (rent extraction by intermediaries such as high-frequency traders, exchanges and broker/dealers should not be excessive). Liquidity pools have several characteristics that have the potential to help achieve these objectives:

- they can efficiently facilitate direct block trading between institutional investors,

2 See "Well-Functioning Financial markets", Norges Bank Investment Management, Discussion Note #13 (2012).

- they can serve as competitive checks on exchange monopoly power, and
- they can be tailored to specific market participant requirements, and innovate rapidly.

These benefits have to be weighed against the potential efficiency drag on the price discovery process. Institutional investors will typically prefer sourcing liquidity from pools with limited pre-trade transparency. Trading in transparent, 'lit' venues such as exchanges can thus be an indication that liquidity is not available in liquidity pools. This means that the residual volume that is traded on 'lit' venues becomes more informative. This can increase instantaneous volatility and the cost of price discovery.

In addition to the (intended) lack of pre-trade transparency, many liquidity pools are also less than fully transparent about the way that they operate – for example, the nature of the matching engine, the type of potential counterparties, and the degree to which pre-trade transparency exists for some counterparties through IOIs. We view this as a predominantly a negotiation issue between investors and brokers as their agents on the one hand and venue operators on the other. We have seen substantial improvements in operational transparency over the last few months (for example, through the publication of the ATS filings in the US). Nevertheless, this lack of transparency can impact the well-functioning of markets, and give rise to excessive rents being extracted by some intermediaries.

Liquidity pools thus have the potential to contribute meaningfully to well-functioning markets, but this is crucially driven by implementation details and operational transparency. To assess whether they succeed in this, we next introduce a straight-forward segmentation of the universe of liquidity pools. We show that the type of liquidity pool significantly impacts their potential contribution to well-functioning markets.

Classifying Liquidity Pools and their Impact on Market Quality

Trading venues that operate with limited pre-trade transparency can be classified along several different dimensions. Butler (2007) segments 24 US liquidity pools into 16 different types, for example[3]. This classification is based on several characteristics, such as the pricing mechanism, the nature of the order flow, and the type of counterparties in the pool. Understanding these characteristics is important for the order routing direction we give our brokers, for example – we deliberately exclude a number of trading venues based on these and other characteristics.

In order to assess the impact of liquidity pools on market quality, we find it more useful to classify them according to the stage of the investment process in which a venue is used. Some venues – often the ones that have been established the longest – focus on direct block crossing between investors, and typically appear early in an investor's execution plan. These pools show large average trade sizes but low fill rates. Other liquidity pools appear later in an investor's execution plan, and are often accessible only after the investor has delegated the execution to a broker. These pools – including most of

3 See Butler, G., "Liquidity Aggregation: What Institutional Investors Need to Know", The Journal of Trading 2.2 (2007).

those operated by broker/dealers, exchanges' hidden order books, as well as independents and HFT ping destinations – have much smaller trade sizes (comparable to those in lit exchanges) and higher fill rates.

Chart 1 shows the typical rank ordering of trading venues for an institutional buy-side trader. Everything else equal, the trader will prefer to execute the trade in a block-crossing venue with another buyside firm, because of the larger trade size and (typically) the absence of direct market impact cost. If the trade urgency is high, such that the low expected fill rates in such a trading venue are not acceptable, the trader will move to other execution strategies that combine a higher probability of order fill with higher expected total cost. This will typically include sell-side agency execution (whether on a cash trading desk or through the use of algorithms) or the utilization of a market maker's risk carrying capacity.

Chart 1: Typical Rank Ordering of Trading Venues for a Buy-Side Trader

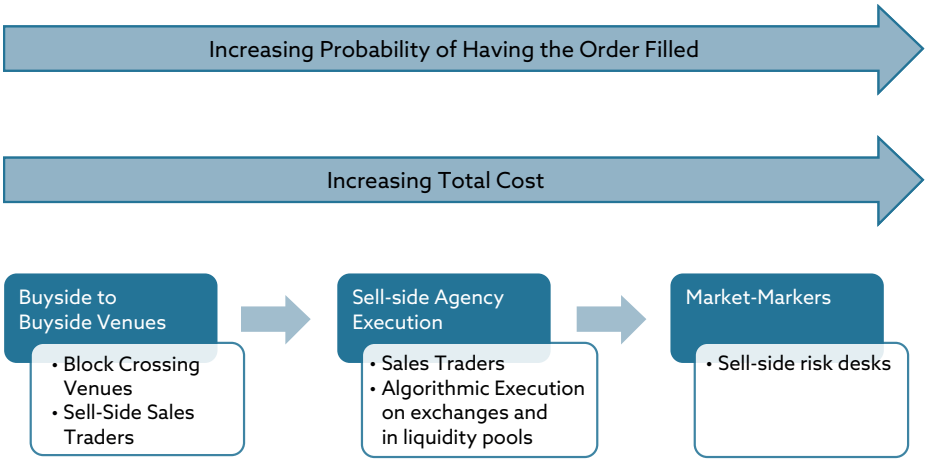

Block Crossing Venues

Block crossing venues are the modern-day equivalent of 'upstairs trading', a concept which is at least as old as exchanges. Their utility depends on the imbalance between the investor's desired order size and the typical trade size on the exchange. The greater the imbalance, the more attractive upstairs trading becomes, notwithstanding the lower fill probabilities.

This imbalance between investor's desired order size and typical trade size on exchanges has increased in recent years. On the one hand, increased execution automation through the introduction of trading algorithms has tended to decrease the typical trade size on exchanges, as a 'parent' order is sliced into sequential 'child orders'. On the other hand, the desired order size has increased as the 'institutionalization' of asset management has continued to progress[4]. While this has gone hand-in-hand with an increase in more passive investment strategies, the net result has been fewer, but larger orders. In many cases, these orders will be driven by cash flow considerations (such as

4 In the US, for example, more than 80% of a typical large cap stock's shares outstanding are held in investment vehicles that are professionally managed. Other markets tilt even more heavily towards institutional management. See Rydqvist, K., J. Spizman and I. Strebulaev, "The Evolution of Aggregate Stock Ownership", CFS Working Paper, No 2011/18 (2011).

inflows/outflows into mutual funds), rather than by new information about the expected excess return opportunities in individual securities.

These large orders will in most cases ultimately be filled by similarly large contra-orders from other institutional managers, potentially after passing through the hands of one or more liquidity-providing intermediaries[5]. Given the level of institutionalization reached, it is unlikely that the order can be filled through a series of small retail orders. This makes block crossing venues increasingly attractive destinations, and we expect to see continued growth and gradually increasing fill probabilities.

The uninformed cash flow considerations driving many of these large orders mean that the efficiency of price discovery is not impacted if trading occurs in block crossing venues. At the same time, the rent extraction potential for intermediaries is minimised. We would thus argue that trading in such block crossing venues is a net contributor to well-functioning markets.

How much trading should we expect in block crossing venues? Ready (2013) shows that these venues account for between 5 and 8% of US large cap institutional flow[6]. The remainder of institutional flow is either block crossing through other avenues (primarily manual trading by broker/dealer cash traders) or through algorithmic execution, often delegated to broker/dealers. There are several factors driving this split in trading volume. Higher market volatility, for example, will tend to decrease the proportion of trading volume done in block crossing venues, since the opportunity cost of waiting for a match increases with volatility. Similarly, exogenous restrictions on trader patience – such as upper limits on cash holdings in mutual funds – will tilt trader behaviour towards greater certainty in execution, and away from block crossing venues.

In addition, institutional characteristics of these types of venues are critical. Fragmentation is a real concern – the probability of finding a match reduces geometrically with the number of venues. Competition amongst these venues will help to limit the effect of fragmentation by introducing a rank preference ordering amongst institutional traders. Services are part of this competition, such as blotter scrapers. These allow for the identification of possible matches, potentially before the trader is actively working the order.

Other services relate to the information stream and the order types available at the venue. While we have some concerns about the possibility of information leakage from conditional orders, for example, we believe that these can be structured in such a way that the impact of the leakage can be controlled while still meaningfully increasing the probability of finding a match. This can be achieved, for example, by including a minimum automatic execution quantity, with the conditionality attached to the possible upsizing of an order.

5 Institutional trading volume accounts for between 25-30% of total trading volume in US large-cap stocks. The rest (minus retail's contribution) is accounted for by shorter-term liquidity providers, including HFTs. This means that an order potentially passes through a number of intermediaries before the institutional contra-order is found at some point in the future – what the CFTC has characterised as 'hot potato volume'. See Crow, C. and S. Emrich, "'Real' Trading Volume", Morgan Stanley Quantitative and Derivative Strategies (2012) and Ready, Mark, "Determinants of Volume in Dark-Pool Crossing Venues", University of Wisconsin-Madison (2013).

6 See Ready, Mark, "Determinants of Volume in Dark Pool Crossing Networks", University of Wisconsin-Madison (2013).

We are seeing considerable innovation in this space, adapting the offerings of block crossing venues to the needs of today's large market participants. We would expect, therefore, that the proportion of volume going through such block crossing venues to continue to increase.

Non-Block Liquidity Pools

If the liquidity available in a block-crossing venue at a point in time is insufficient, the buyside trader may decide to source liquidity from other venues, depending on the urgency of the order. These include lit exchanges as well as liquidity pools that do not typically transact in block sizes. Most broker/dealer operated liquidity pools are of this type, as are independents, exchange-operated hidden books, as well as ping destinations that can be characterised as a form of automated market making.

There is considerable nominal fragmentation in this space, with a large and growing set of available trading venues. However, unlike for the case of block-crossing venues, the *effective* fragmentation is considerably smaller than the nominal fragmentation. At this stage of a trade, typically a broker as agent is tasked with sourcing liquidity, generally through use of an algorithm. These algorithms typically have access to a wide range of trading venues through the broker/dealer's order router, including multiple liquidity pools. Subject to best execution obligations, the venue routing decision will be driven by economics. This includes explicit access costs for the broker, as well as competitive considerations. These competitive considerations can lead to fragmentation, where brokers may try to 'internalise' their client flow. However, in our experience, most broker algorithms will route to other trading venues, limiting the effective liquidity fragmentation.

Sourcing of liquidity at this stage will lead to market impact, necessary to attract liquidity. To ensure the market impact paid is reasonable, algorithms will break up a large 'parent order' into small 'child orders' that are executed sequentially. Pacing of these child orders, as well as the logic used to route them to different venues, are the key variables used to control the market impact.

The splitting of parent orders into sequential child orders leads to executed trade sizes on lit exchanges and non-block crossing liquidity pools that are typically considerably smaller than those on block crossing venues. While the venues that are accessed at this stage of execution generally have similar executed trade sizes, they differ substantially in probability of fill (from near-certainty for market orders sent to exchanges, to very low for most HFT ping destinations), 'toxicity' (information leakage) and cost.

The competitive environment for trading venues has the potential to be effective in limiting the rent extraction by exchanges and liquidity providers. However, it can also lead to the emergence of trading venues that have a cost/benefit trade-off that is unattractive to the investor, since the incentives and costs of order placer and broker are not aligned. This means that the investor has to direct the broker not only on trading strategy benchmarks, but also on permissible venues. For example, we do not believe that the liquidity from HFT ping destinations is worth the information leakage costs. For other

venues, as well as exchange-operated hidden books, the picture is more nuanced and requires constant monitoring of execution quality.

We believe monitoring has to be both qualitative and quantitative. On the qualitative side, transparency around the operating procedures of both lit exchanges and liquidity pools, in particular around order types and matching logic, are of paramount importance. We have seen the publication of Reg ATS forms for many of the liquidity pools, as well as greater clarity around order types and matching logic on exchange hidden books. On the quantitative side, monitoring has to be an ongoing process, involving the analysis of broker executions on both a macro level (such as comparisons to trading benchmarks and expected market impact) and a micro level (such as an analysis of the venues that orders are routed to and executed on, compared to aggregate volume distribution).

Sourcing Liquidity from Liquidity Pools – An Asset Manager's Perspective

As a large participant in global equity markets, Norges Bank Investment Managment utilises a wide range of trading strategies. We actively use block crossing venues as one of our preferred methods of execution, but we also delegate execution to brokers. The brokers' algorithms are calibrated to a number of different trading strategies, including VWAP, aggressive liquidity seeking, and passive/stealthy opportunistic trading. Some brokers' algorithms are restricted to trading on lit exchanges, while others can utilise liquidity pools and exchanges' hidden order books. Our algorithmic choices and their parameterisation are driven by our past trading experience, and conditioned on current market conditions.

Table 1 shows summary statistics of the execution venues used by some of the broker algorithms most frequently used by our US Equities trading desk since Jan 2014. We show two VWAP algorithms, as well as five representative algorithms which are calibrated from aggressive to patient. All of these algorithms are permitted to execute in liquidity pools – indeed, the most patient algorithm in our sample exclusively sources liquidity from such venues.

Table 1: US Execution Destinations for Broker Algorithms

Execution Destination	Aggressive				Patient	VWAP	
	1	2	3	4	5	6	7
Lit Exchanges	98.1%	49.6%	91.9%	31.2%	0.0%	83.0%	99.9%
Inverted Exchanges	6.8%	0.0%	16.8%	1.1%	0.0%	5.6%	7.4%
Liquidity Pools	1.9%	50.4%	8.1%	68.8%	100.0%	17.0%	0.1%
Broker's Own Pool	0.9%	50.2%	3.9%	8.0%	16.3%	16.6%	0.0%
# of liquidity pools accessed	18	8	11	15	4	4	3

Overall, these seven trading strategies accessed 25 different pools, in addition to the 13 (at the time) exchanges. We generally see brokers favouring their own pools – nearly 100% of the liquidity pool volume share (50.2% of 50.4%) in Algorithm #2's case, to a low of around 11.6 % for Algorithm #4. Even though the brokers preferred their own pools, they also routed to other

liquidity pools – Algorithm #1, the most aggressive, accessed 18 out of the 25 liquidity pools overall, while the most patient Algorithm #5 just accessed 4 liquidity pools. This behaviour is as we would expect. More importantly, it shows that effective fragmentation of liquidity is not much of an issue for algorithmic execution strategies. The order routing technologies embedded in broker algorithms due to best execution obligations ensure that the algorithms access many pools of liquidity. This does not mean that the 'toxicity' of trading venues does not matter – best execution in the time domain is a considerably more challenging task.

The limited effective fragmentation experienced by algorithmic execution strategies with order routers is exemplified in the differences in execution venues for two of our VWAP algorithm providers. The algorithms show comparable performance in our post-trade analysis versus the VWAP benchmark, yet have very different execution venue characteristics. While Algorithm #7 almost exclusively executes on lit exchanges, Algorithm #6 executes 17% of trading volume in liquidity pools, most of that in the broker's own pools.

Based on our experience with broker execution algorithms, and our on-going post-trade review of execution performance, we believe that liquidity fragmentation is not much of an issue for algorithmic trading. However, we are concerned about the 'toxicity' and the potential for information leakage of several liquidity pools. We have introduced a white list of permitted trading venues. The list explicitly excludes HFT ping destinations, and actively takes into account client tiering systems that brokers commonly offer for their own liquidity pools.

Future Developments in Liquidity Pools – A Wish List

We believe that block crossing venues as well as other liquidity pools play an important role in ensuring well-functioning markets. Broadly, they aid in limiting the rent extraction ability of intermediaries. Block crossing venues are effective in bringing together natural institutional trading intentions, and will become increasingly important in a largely institutionalized asset market. Other types of liquidity pools serve as effective competitive pressure to limit rent extraction by intermediaries such as broker/dealers, exchanges and HFT liquidity providers. Price discovery is not significantly hampered, in our view, as de facto linkages across trading venues are established through brokers' order routers, as well as the activity of high frequency arbitrageurs[7].

There is currently a global discussion on the role and impact of liquidity pools. New regulations in Europe propose limiting the trading volume that can be executed in non-block crossing liquidity pools to 8% of overall volume under a Reference Price Waiver (block crossing venues, in general, will be able to operate with Large-in-Scale waivers). In the US, several exchanges have suggested limiting their access fees, which would change broker economics and may increase the attractiveness of lit exchanges to order routers. This demonstrates the effectiveness of the competitive pressure that liquidity pools can bring to the market. We believe that these discussions have to

7 See, for example, Comerton-Forde, C. and T.J. Putniņš, "Dark Trading and Price Discovery", Journal of Financial Economics, forthcoming (2015).

be seen in the broader context of overall exchange costs – especially including data fees.

While liquidity pool competition is effective in limiting rent extraction, it is at least as important that the role of exchanges as the arbiter of the price discovery process is not impeded. The rise of liquidity pools over the last few years has made trades that occur on exchange more informative – since exchanges will typically appear later in the investor's preference order, a trade on an exchange is an indication of the absence of liquidity in other trading venues. Unsurprisingly, this has led to a deterioration in liquidity quality on exchanges – we typically see less displayed depth on limit order books, and greater instantaneous volatility as limit orders are cancelled and re-set. HFTs are often blamed for this, and are seen as exploiting the larger order placers coming to an exchange. While there may be some truth to this – the technological arms race in equity trading certainly would suggest so – we believe that much of the change in the structure of exchanges' limit order books are reflective of greater informativeness of trades coming to the exchange.

This suggests an alternative approach to assessing the impact of liquidity pools on the price discovery process, based on limit order book depth and the ratio of instantaneous volatility to daily volatility. We support investigations of these effects, such as the SEC small-cap pilot project currently under review.

At the same time, we believe that greater transparency around the operating procedures of both liquidity pools and exchanges is needed. For liquidity pools, the publication of Reg ATS forms by many of our US brokers is a useful first step. Perhaps even more pressing is transparency around the order types and their matching priority for exchanges' hidden order books. We are in favour of further transparency about operating procedures and available order types, particularly if they differ by client. If we do not feel we have sufficient transparency in a given trading venue, it will not be on our whitelist of permitted trading venues for our brokers.

Lastly, the broker order routing process is critical. Subject to best execution obligations, which are primarily concerned with the comparison of execution prices to contemporaneously available quotes, a broker's routing decision will be determined by economic considerations. A broker's economic decision to route to one destination in preference to another might impact subsequent execution quality through information leakage. Recognizing this, we believe that skewing the broker's objective function – through the imposition of price benchmarks, as well as through active limitations on the set of permitted execution venues – is a critical fiduciary duty of investment managers.

Block crossing venues, which generally will rank higher in a buyside trader's preference ordering, play an important role in a market that is dominated by ever-increasing institutionalization. We believe that these trading venues should have even greater prominence, and are actively working on establishing and strengthening such venues. For block crossing venues, liquidity fragmentation is a real concern. Block crossing venues are potentially subject to rapid changes in relative market share, driven by the fee structure as well

as ancillary services offered by the venue. This can increase the search cost for buyside traders, to the detriment of investors. We therefore support the development of utility-like block crossing venues, which would serve to limit the possibility of rapid changes in relative market share, and to increase the fill probability.

Intermediation and rent extraction are a necessary part of asset markets. However, there is a natural tendency for this rent extraction to reach excessive levels. Market evolution serves to keep this tendency in check. We view the emergence of liquidity pools as an example of such beneficial evolution. However, they in turn introduce novel avenues for rent extraction, primarily through insufficient transparency. Asset owners and managers need to show continued vigilance and a proactive research-based approach to analysing and adjusting potential excesses.