Characteristics of different liquidity pools and their effect on execution

*Julio Concha, Eran Fishler, Maarten Bar and Peter Fraenkel**

Over a few years, the number of ATS's has risen from a handful to several dozen, each with its own unique characteristics and rules. In many ways, the distinction between dark pools and lit pools has broken down, with large volumes of hidden orders trading in displayed markets, orders being internalised in broker pools on their way to the open market, and a complex and generally obscure set of routing relationships and IOIs among all the venues and services. Effective navigation of today's fragmented liquidity landscape consequently requires ever more effort and investment in quantitative monitoring and research into countermeasures. In this chapter, we highlight a few examples of the surprising and disturbing effects we have discovered, and a few of the countermeasures we have developed to protect execution quality while still providing access to as much liquidity as clients demand.

The growth of dark pools presents opportunities to the buy-side trader by offering increased liquidity and the promise of crossing large quantities quickly, with little market impact. However, the sheer number of available options now poses a major challenge. Consider, for example, the questions faced by a trader trying to buy 10,000 shares of an illiquid stock. Should they place parallel orders in all destinations? What type of orders (limit, hidden, IOC, etc.)? How much should they allocate to each destination, and how long should they allow an unexecuted order rest? What about throttling the order to avoid price impact? Is there a risk of information leakage? The answers often hinge upon seemingly minor differences between destinations. In order to make effective use of dark pools, it is critical to understand their individual features in detail.

Probably the most important feature of a dark pool is its

*
Julio Concha,
quantitative engineer

Eran Fishler,
director of research

Maarten Bar,
quantitative engineer

Peter Fraenkel,
CTO

Pragma Securities

Execution strategies

order-handling procedure, e.g., whether IOIs are used, to which counterparties the order is exposed, etc. However, this procedure is typically confidential. Moreover, the dark pool operators themselves may not be aware of the effects of their policies. Thus, the best way to gain insight is through detailed empirical analysis of orders and executions.

Such an analysis can lead to surprising results. For example, we found that one dark pool lets orders rest for 30 seconds and then exposes them to a second, unknown pool. This policy can result in information leakage. As another example, we have discovered that destinations differ widely in terms of the actual execution price of a mid-quote order.

Another important aspect of dark pools is the type of liquidity they contain. Common wisdom posits that large fills are the result of large, information-less orders. This would imply that large crosses have no price impact whatsoever. Our research contradicts this rosy picture. In fact, it is the type of liquidity in the pool that is crucial, rather than the average fill size. We demonstrate, through a careful analysis of order executions, that adverse selection increases with fill size, unless strong anti-gaming measures are in place. Even more importantly, in some major pools, large fills are accompanied by significant price reversions after the fill, providing strong indication of gaming. After discussing an extreme case with a pool operator, we learned that that pool had recently added a high-frequency trader as a liquidity provider.

Finally, we address the issue of information leakage. We start from the observation that information leaked by one destination can manifest itself in a different one, rendering the task of detecting leakage highly complex. Our methodology overcomes this problem and in fact helped us detect specific leaky dark pools.

Mechanics of dark pool order handling

Each pool develops its own set of policies for matching and order handling. These different rules, often confidential, influence the potential for information leakage and overall performance of the pool. In this section, we concentrate on two major destination characteristics: time to first fill and price improvement.

Time to first fill

Informed investors trade in reaction to alpha signals and require speedy execution. On the other hand, destinations vary widely in terms of the time needed to fill an order. This is

Execution strategies

FIGURE 1: TIME ELAPSED BETWEEN ORDER PLACEMENT AND FIRST FILL AT A BROKER/DEALER-SPONSORED DARK POOL

After 30 seconds, orders appear to be exposed to a second liquidity pool. Orders that remain unexecuted after 30 seconds should be canceled. Inset: A typical dark pool does not exhibit the second spike.

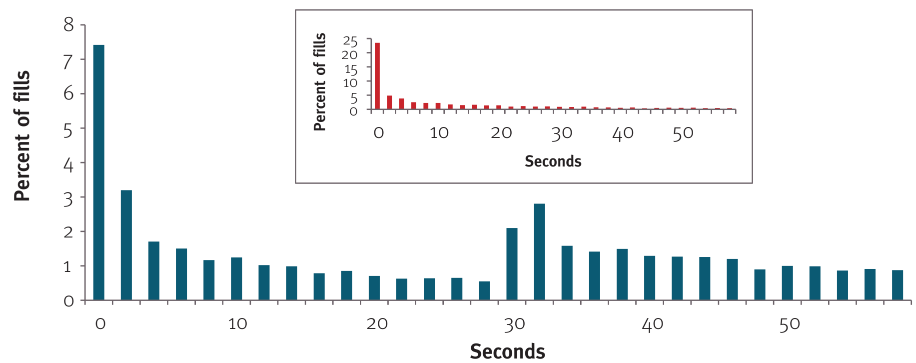

Source: Pragma Securities

not only a matter of liquidity, but also of the internal operation of the destination itself. An obvious example is periodic versus continuous operation, e.g., POSIT versus POSIT Now.

We consider here the time in seconds between order placement at a given destination and the first (perhaps partial) execution of that order. Figure 1 depicts a histogram of this metric for a broker-dealer sponsored dark pool. The large spike near zero corresponds to fills against resting orders, and is common to all dark pools. However, the second spike, at 30 seconds, is quite unique. This spike indicates either that our unexecuted order is routed to a second pool, or that the destination broadcasts IOIs.

The implications are twofold. First, the delay comes at a cost in terms of alpha extraction. Second, this Flash-like behavior provides ample opportunity for information leakage. A smart allocation algorithm could, for example, discourage gamers by randomly canceling live orders to this destination before 30 seconds elapse.

Price improvement (or not)

Most dark pools claim to cross orders at the mid-quote, providing the dual benefit of better prices (by not crossing the spread) and no market impact (by not exposing a directional imbalance). In what follows, we demonstrate that, contrary to common wisdom, mid-quote orders may not execute at the mid-quote.

Figure 2 depicts a histogram of bid/ask shortfall for three different liquidity pools, in units of bid/ask spread. Assuming a buy order, the bid/ask shortfall equals zero if the order was executed at the mid-quote, 1/2 if executed at the ask, -1/2 if executed at the bid, -1/4 if executed halfway between the mid and the bid, etc. From the figure, we can observe that in destination 1, the overwhelming majority of trades (76%) take place at the mid. Small deviations do happen, due to timing discrepancies. These deviations are rare and roughly symmetric, so the average shortfall is zero.

For destination 2, we observe a different pattern. A large spike at -1/4 indicates price improvement, consistent with the fact that, in this destination, dark orders interact with displayed market flow. Thus, the allocation algorithm can trade off price improvement against the risk of information leakage associated with open market destinations (see below).

Finally, in destination 3, we again observe that the majority of the executions occur at the mid-point. However, about 20% of the fills occur at the other side, e.g., at the ask (or worse) for buys, even though a mid-price order was submitted. This cannot be attributed to timing discrepancies. After we alerted the pool manager and explained our analysis, it surfaced that they constructed their own NBBO from individual data feeds. These feeds included marketable Flash and Bolt orders that resulted in a seemingly locked market and an erroneous NBBO, a fact that had gone unnoticed. The problem has since been corrected.

Fill size analysis

Dark pools exhibit varying execution sizes, in part because different venues attract different market participants and kinds of order flow, e.g., retail, prop desks, institutional. In exchange-based dark pools, for example, hidden orders interact with regular orders, which tend to be smaller. Other destinations have large minimum order quantities and attract only buy-side participants with block-size orders.

Recent studies have argued that larger crosses are beneficial. The rationale is that interaction with displayed order flow causes market impact and leaks information about the existence of a residual order. In what follows, we demonstrate that large crosses can in fact have substantial market impact.

Execution strategies

FIGURE 2: HISTOGRAM OF EXECUTION SHORTFALLS RELATIVE TO THE MID-QUOTE AT THE TIME OF THE CROSS FOR A MID-QUOTE ORDER, FOR THREE DESTINATIONS

A shortfall of 0.5 means we paid half the spread. Destination 2 provides systematic price improvement. Destination 3 should be avoided due to the abnormal number of executions that cross the spread.

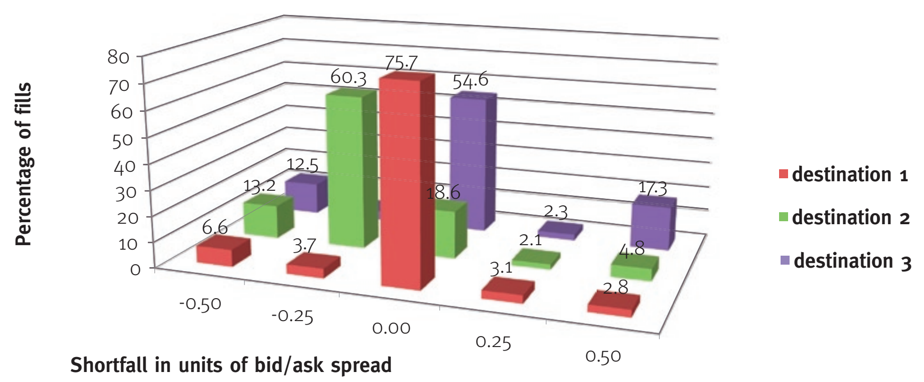

Source: Pragma Securities

In order to detect the market impact of large crosses, we need a more detailed analysis than a naive comparison of Implementation Shortfall. For example, a destination with high latent liquidity tends to provide fills soon after an order begins to trade. These early fills are not subject to the accumulated market impact present in later fills, and will introduce a favorable bias in the average shortfall of that destination. Instead, we analyse price moves before and after a fill on a fine-grained time scale.

Let t_o be the time of a specific fill, $p_m(t)$ the mid-quote at time t, and sf(t) the shortfall at time t relative to that fill. The time t could be a few seconds or minutes before or after t_o. Mathematically represented:

$$sf(t) = \frac{p_m(t) - p_m(t_o)}{p_m(t_o)} \text{ x side.}$$

In Figure 3, we depict sf(t) for two different destinations and two fill size groups. For each destination, we calculated the average returns for all fills, and again for large fills only. We define 'large fills' as those in the top 20% by value relative to a large data set containing executions at many destinations. This ensures a fair comparison.

Execution strategies

FIGURE 3: AVERAGE PRICE MOVES AROUND A FILL EXECUTED AT ONE OF TWO DESTINATIONS, A AND B.
Larger fills at B are preceded by steeper unfavorable price moves. Inset: price reversion right after a large Destination B fill. Destination B should be avoided.

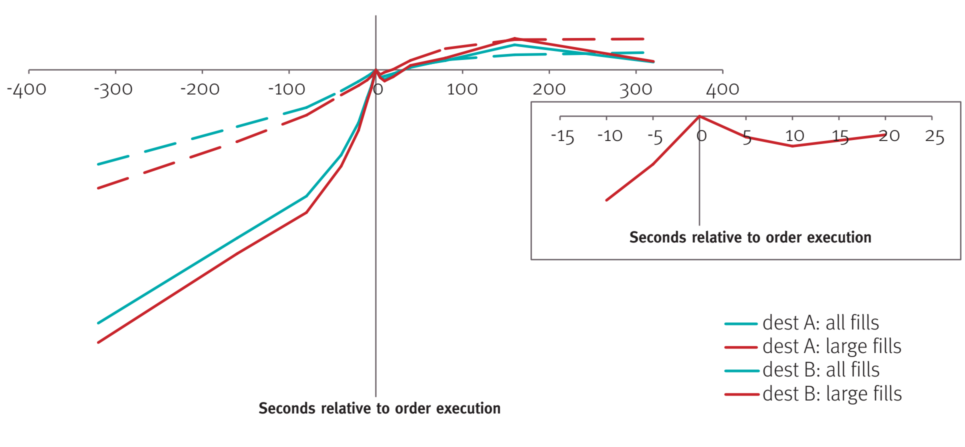

Source: Pragma Securities

The first striking fact is the differences between the two destinations. Fills at destination B are preceded by steep adverse price moves. Even more striking is that the larger the order size, the steeper the price moves. One possible explanation for the latter is simple adverse selection: we are more likely to get large fills after quick moves that make the price more attractive to counter-parties. However, this does not happen in destination A. A better explanation is information leakage followed by gaming: the price is being manipulated just prior to the fill. The price reversion during the 10 seconds after the fill provides convincing support for the latter theory. The fact that this pool draws large orders appears to attract traders engaged in gaming.

The conclusion is that the large fills available in this pool come at a price, contradicting the notion that large crosses are 'impact-less'. In fact, a sensible allocation strategy would send only small orders to this destination, or avoid it altogether.

Quantifying information leakage

Information leakage can manifest itself in a variety of ways: a dark pool can leave a detectable trail on the tape, pools can send IOIs to select

FIGURE 4: AVERAGE IMPLEMENTATION SHORTFALL FOR ORDERS EXPOSED TO A DESTINATION, RELATIVE TO ORDERS NOT EXPOSED TO THE SAME DESTINATION

Destinations 10-15 should be used carefully.

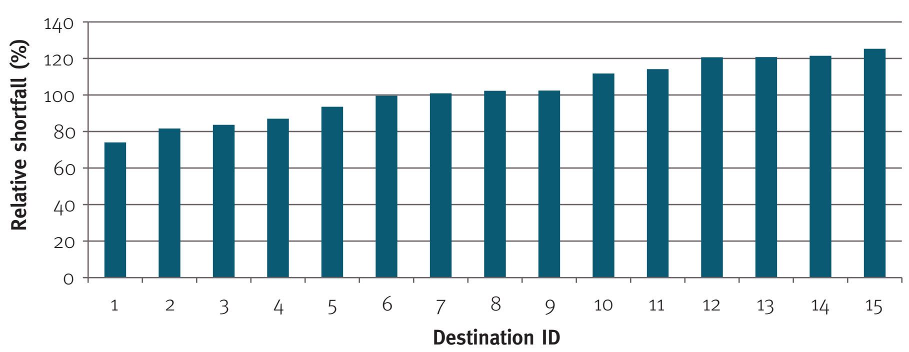

Source: Pragma Securities

market participants, high-frequency traders can send small probing orders, etc. These heterogeneous mechanisms make it difficult to identify sources of information leakage. More importantly, large orders are often worked in multiple destinations in parallel. Even if only one of these destinations leaks information, the Implementation Shortfall of the whole order will be affected. Note in particular that unfavourable executions need not occur at the same destination that leaks information. Thus, pinpointing the source of the leakage from the fills alone is impossible.

In order to detect information leakage and correctly identify its source, one needs to compare two separate groups of similar orders: one that was exposed to all destinations, and another that was exposed only to a smaller subset. If the larger group exhibits significantly higher shortfalls, then it contains leaky destinations.

To implement this method, we considered orders that were variously exposed to 15 destinations. For each destination, we computed the average shortfall of orders that were exposed to it, and of orders that were not. We expect the quotient of these two averages to be about 1. However, as seen in Figure 4, it can be significantly different from 1. Exposing an order to destinations 1-5 improves shortfalls by 5% to

Execution strategies

25%, possibly due to the additional liquidity. In contrast, exposing an order to destinations 10-15 increases shortfalls by 10% to 25%, indicating that these destinations may leak information. A smart allocation algorithm should utilise these destinations carefully, e.g., by using large minimum cross sizes and/or low trading rates. Destinations exhibiting large persistent deviations from 1 should simply be avoided.

Conclusion

Liquidity pools are numerous, diverse and ever-changing in terms of their order types, order-handling procedures, constituents, etc. Treating an execution service as a benign black box – whether an individual dark pool or an algorithm that acts as a gateway to an unknown network of other pools and venues – risks a loss of information control and execution quality.

Our approach is to monitor pools quantitatively on an ongoing basis. This is a complex process that requires sophisticated analysis techniques as well as comprehensive historical data. Yet this analysis is critical for the design of effective countermeasures, including early-warning systems, smart allocation algorithms and anti-gaming capabilities. These mechanisms, in turn, provide a measure of security that gives buy-side traders access to the liquidity available in dark pools, while avoiding potential pitfalls.

Finally, we note that this type of analysis does not apply only to dark pools. We have carried out similar studies regarding open-market and other destinations relevant to non-aggregation algorithms, such as VWAP, POV and IS. In this context, metrics such as latency and market impact are critical. Just as with dark pools, we have found that an in-depth understanding of market microstructure is an invaluable tool in the quest for optimal execution.

pragma