

The Fifteen Year ITCH

October 18th 2010

Back in May, a small US broker dealer published a white paper highlighting that information published in relation to some exchange or MTF dark pools could allow participants to identify the direction and longevity of (supposedly dark) resting orders. In Europe, another broker-dealer brought the issue to the attention of their buy-side clients, who demanded that problem be addressed, forcing the affected venues to amend their data feeds (which they did within days). All water under the bridge, or so I thought.

But recently there have been further accusations that exchanges continue to “deliberately sell confidential order data to HFT firms”, with some suggesting that there’s a grand conspiracy amongst exchanges, regulators and HFT firms to defraud institutional investors. These new allegations are being levelled at exchanges in relation to their public lit order books. I think they’re wide of the mark and reveal a lack of appreciation for how public data feeds work.

Here’s a quote from my previous (in fact, first) blog entry:

“When exchanges first started offering electronic order entry disseminating a public data feed, participants wanted to be able to identify their own orders and executions in the public data feed. This allowed participants to know their queue position in the order book, and to display this on a client front-end. It allowed them to perform better transaction cost analytics – by identifying which executions on the ‘tape’ were theirs. It allowed them to measure the latency of the public market data against their own Execution Reports. And it allowed multiple OMS and EMS systems within the firm to identify their own orders & executions without having to feed each system with drop-copies of the order entry/execution feed.”

I repeat this to illustrate that the inclusion of OrderIDs in exchanges’ public data feeds was a response to participant demand. Dark pools aside, participants do expect exchanges and MTFs to provide this information.

First, brokers and market data vendors need to build and maintain a copy of the order book for each instrument – converting individual order-add/amend/delete and trade messages disseminated by the market into a depth-of-book representation that can be used for trading decisions. The OrderID assigned to each order is essential to facilitating this process.

Second, investors, brokers and markets all need to be able to relate individual trades back to the orders they belong to. This is essential for order management (e.g. to know the cumulative traded quantity and residual quantity for an order), for transaction cost analytics and for regulatory compliance amongst other things. The way this is typically achieved is by making the OrderID an attribute of each Trade. So it’s easy to find and sum all the Trades linked to a particular OrderID. This in turn makes it important that an OrderID is persisted throughout its lifetime – as changing the OrderID half way through would cause the ‘loss’ of the related trades, and consequently over-trade errors (which I know to be true from personal experience – a lowlight of my days as a trader).

But, I also said then:

“The specs make it relatively easy to identify iceberg/reserve orders as soon as the visible peak is first refreshed, and also to identify pegged orders as soon as they are modified by the market. “

This observation seems to form the basis of the ongoing allegations. So I want to delve a little deeper and explain why changing the way these data feeds work would represent a huge cost to the industry for little or no benefit:

How does it work exactly?

- When an Order is first received, an OrderID is assigned to it, and both communicated back to the participant (in the acknowledgement message) and disseminated in the public market data (allowing the participant to see where they stand in the book).
- From that point forward, the OrderID is used to communicate any events affecting the order:
 - Each execution reported back to a participant carries the OrderID to which the trade belongs. If it was a visible order that traded, a single “Order Executed” message in the public data feed tells recipients that there has been a trade against the specified order, and hence that the remaining quantity in the book should be reduced. If it was the visible portion of an iceberg/reserve order, most communicate that there was an execution against the order, but that it remains alive with a new display quantity (typically with a loss of time priority).
 - For each order amendment or cancellation, the public data-feed refers to the OrderID and communicates which attributes have been amended.
 - And when the exchange automatically adjusts the price of a pegged order, the data feed refers to the OrderID and communicates the new price. (I believe some exchanges go so far as to label pegged orders explicitly, although I confess I don’t see a good reason to do so).

The persistence and dissemination of the OrderID in the public data feed is key to enabling participants to trade and manage their orders effectively, but also makes the market more transparent than some participants may have appreciated.

Exchange and HFT detractors argue that even if we arrived at this situation innocently, it’s still wrong, and that institutional traders had no idea that their information was being ‘compromised’ in this fashion. Since nothing has changed since the now widely-used ITCH protocol arrived on the scene over 15 years ago (with its specification public ever since), it’s somewhat surprising that this is news to some market professionals. But, timing aside, how serious are the concerns about information leakage regarding iceberg and pegged orders, and should markets be changing their public data feeds to assuage the critics?

Ceasing to publish any OrderIDs in the public data would render market data useless and break most OMS systems. But, in relation to iceberg and pegged orders, could exchanges switch from publishing amendments to the existing OrderID to instead sending a cancellation of the order and then the addition of a replacement order (with a different OrderID)? This sounds alluring, except that:

1. Without significant design changes, this would break the link between trades and the OrderID, and as I explained above, that's not a good idea. Avoiding this would require very substantial investment by data vendors, brokers, OMS vendors etc – for which there is very little appetite.
2. This would double the volume of market data being disseminated in relation to order amendments and iceberg executions (two messages instead of one).
3. And most importantly - it wouldn't materially reduce information leakage. Even if replacement OrderIDs were used in the scenarios above, iceberg executions and pegged order amendments would still glaringly obvious to any consumer of the data feed from the immediacy with which they followed executions or peg reference-price changes. In short, we'd make making market data less efficient and forcing an overhaul of OMS systems for no good reason.

So what should an institution do?

Frankly, they probably shouldn't worry about it too much. If they're using a sophisticated broker, then the broker will have already developed their algorithms to mitigate the risk of information leakage. They could insist that brokers don't use the iceberg functionality of exchanges – but such a decision would come with a cost of less participation in the marketable liquidity passing through the exchange. They could insist that their brokers don't use exchange pegging functionality – but I expect they'd find that most brokers already don't (in fact, due to lack of demand, Turquoise didn't implement pegging functionality in its new Millennium Exchange platform). Or they could insist that brokers cancel and replace orders rather than amending them (and again would find that this is common practice for many big brokers already).

The simple fact is that 'lit' markets are *very* transparent, exactly as regulators wanted them to be. Meanwhile, all the empirical data suggests that this improved transparency, the competition amongst various markets, and the supplanting of traditional market making by the electronic variety have coincided with an ongoing reduction in the total transaction costs experienced by institutional and retail investors alike. And whilst correlation doesn't prove causation, this positive trend reduces the force of arguments that this level transparency is harmful to institutional investors.