


## The Blind Men and the Elephant



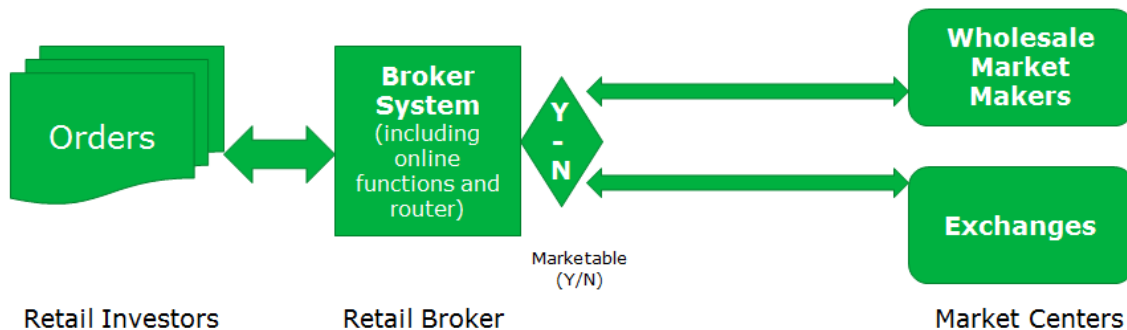
In the parable "the blind men and the elephant" an elephant is surrounded by several blind men, all of whom use their sense of touch to describe what the elephant is like. In the story, each of them examines a specific part of the elephant, and, therefore, describes the creature quite differently. In most versions, the men end up in an argument over what the elephant looks like, as they all have a different view based on what part they were touching. The lesson, of course, is that each of the observers had a unique perspective, and only by combining all of them, would they collectively be able to accurately describe the elephant.

The quest to evaluate trading quality is remarkably similar; combining several key perspectives of the process is the only true way to gain an accurate picture, particularly for institutional investors. From the time a portfolio manager makes a decision to purchase or sell stock, many steps take place. Their order traverses their trading desk, a routing broker dealer and, quite often, multiple market centers consisting of exchanges, alternative trading systems and market makers. It is important to measure the performance of each of these participants in the proper context in order to gain an accurate understanding of the full process.

In the case of retail investors, the SEC set up a framework in 2001 (Rule 605) that has done a credible job of spurring improved execution quality, but there are two caveats to keep in mind. First, it took many years before the full benefit of the disclosures helped retail investors, since the brokers who routed those orders were not covered by the rule. (Rule 605 only applies to the market centers which actually execute these orders - not to the retail brokerage houses that route their client orders to those market centers). Over the ensuing decade, however, the "best execution" policies of the large retail brokers changed to measure execution quality of the orders they route, using Rule 605 as a starting point for the analysis.

The second point is that retail brokerage is quite straightforward. The vast majority of orders are routed without modification or subdivision to either market makers or to exchanges as “held” orders. The only decision that is often made by the routing system of the retail broker is to determine marketability (i.e. if they are aggressive enough to be executable at the best bid or offer). If they are marketable, then they are most often sent to wholesale market makers, who generally deliver execution quality inside of the best bid or offer. If they are not marketable, then they are sent to a stock exchange to be displayed, pursuant to the limit order display rule. Since most orders are routed in their entirety, both the market centers and retail brokers have materially similar views on the execution quality of these orders. In short, routing and executing brokers have the same perspective.

## Retail “Held” Order Flow

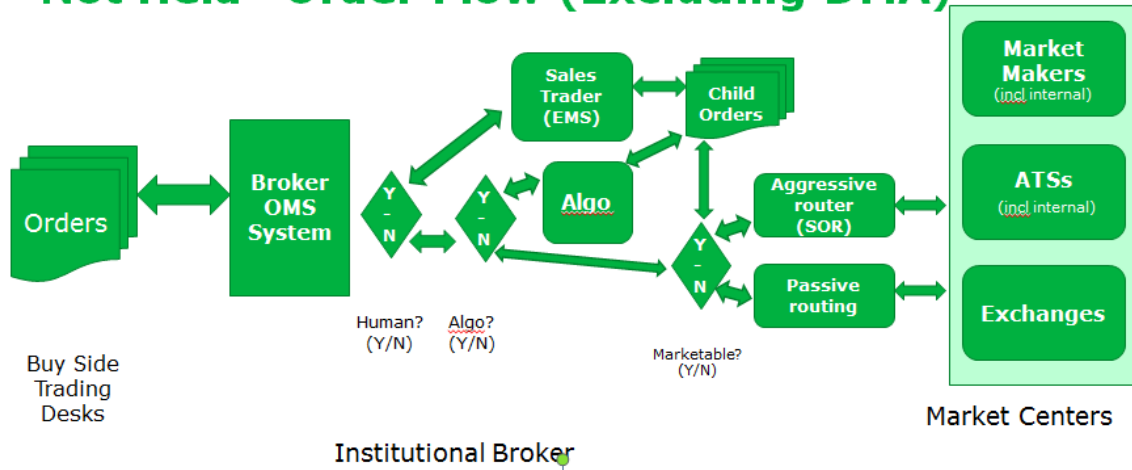


Institutional orders are much more complex; routing brokers, along with their asset manager clients, make several other decisions. These orders are routed to brokers’ OMS systems and there are several decision layers. First is whether it is high touch or not. If so, it gets routed to a sales trader on a desk for further action (who can subsequently take any of the later possible actions including splitting the order up into “child” orders).

Second, there is the decision to either send the order to an algorithm to be worked over a defined time period, or send immediately to the market. In both of those cases, the client provides information regarding the amount of discretion the broker may use, the timeframe of the trade and the benchmarks that the broker is accountable for meeting. (Note that the case where a client directs the broker to send an order to a specific market with no discretion, called DMA or Direct Market Access, is excluded from this discussion. This is because brokers have no decisions to be judged on, so metrics on DMA orders, while important for market centers, are not relevant for brokers.)

Third, if either the sales trader or algorithm decides to send an order to the market at a specific price, the system determines if that order is marketable or not (called a passive order). If it is not marketable, it can be routed directly by the algorithm or sales trader to one or more market centers either as a displayed or hidden order, or a combination of both. If it is marketable, most firms utilize a smart order router that attempts to fill the order as efficiently as possible using either sequential or parallel routing strategies.

## “Not Held” Order Flow (Excluding DMA)



While much more complex than retail, it is not true that institutional orders are too complicated to measure. For many years I have heard industry participants claim that "determining best execution for institutional trading is more art than science" or state how hard it is to measure. But the reality is that institutional flows have some key similarities to their retail counterparts.

An examination of the institutional order flowchart shows that, despite all of the decisions and variations, all orders ultimately end up being executed by a series of aggressive and passive orders sent to market centers. Therefore, as long as the resulting "child" orders are grouped into statistically comparable categories, valid comparisons can be made. The key, like with the blind men and the elephant, is to consider the perspective for each order and aggregate them within the appropriate one.

This means that, in order to measure a routing broker, it is important to report separately based on both the types of orders received and sent. An example of this was described in the appendix to the IHS Markit comment letter on the SEC Order Disclosure rule<sup>i</sup>. The following matrix is an example of such reporting as for each routing broker it shows an overview of routing activity by venue and in total, execution quality and volume statistics broken out into categories that correspond to the key decisions in the institutional diagram (high touch, algorithm, or immediately routable). The report shows statistics per market center categorized by marketability of the child orders sent, the time in force for those orders (since immediate or cancel orders are quite different to day orders) and by order size. This particular report also includes both agency and principal negotiated block trades, as those are considered critical to the "high touch" business.

Not Held Order Report – Venue Specific Data												
				Totals	1 - 499	500 - 1,999	2,000 - 4,999	5,000 - 9,999	10,000 - 24,999	25,000 - 99,999	100,000 +	
Venue 1 (etc. including aggregate statistics)			* Execution Quality Stats → (including PI, Fill Rates, EQ, Realized Spreads)									
	Marketable	DAY	High Touch Orders									
			Electronic Algorithm Orders									
			Immediately Routable Orders									
		IOC	High Touch Orders									
			Electronic Algorithm Orders									
			Immediately Routable Orders									
	Non-Marketable	DAY – displayed	High Touch Orders									
			Electronic Algorithm Orders									
			Immediately Routable Orders									
		DAY – non-displayed	High Touch Orders									
			Electronic Algorithm Orders									
			Immediately Routable Orders									
		IOC	High Touch Orders									
			Electronic Algorithm Orders									
			Immediately Routable Orders									
		Negotiated Block Trades – as principal			High Touch Orders							
					Electronic Algorithm Orders							
					Immediately Routable Orders							
	Negotiated Block Trades as agent (including use of conditional orders)			High Touch Orders								
				Electronic Algorithm Orders								
				Immediately Routable Orders								

If reporting like this was readily available from all broker dealers, it would provide a good starting point for institutions to measure the quality of order routing. As I will discuss in future posts, there is a need to augment some of the statistics provided, but the structure proposed here would be a major improvement. To illustrate how this structure would work, consider three prototypical “bad behaviors” that would trigger “red flags” using this type of analysis. In all of these cases, further detailed analysis would be justified.

**Overly aggressive manual trading** – This behavior would show up as a combination of poor execution quality metrics for marketable day orders and a high percentage of those orders in the “high touch” category. Analysis would indicate that when a firm’s sales traders are directly involved in trading orders, that they often pay through the spread from their EMS. Note that a high percentage of negotiated block trades would not be nearly as problematic, as those are generally agreed to directly by clients.

**Over-reliance on probing dark pools** – This behavior would show up as a combination of extremely large IOC order volumes with very low fill rates within any of the order sources. It is not necessarily indicative of a problem, but such data should trigger an evaluation of the firm’s smart order router to compare its fill rates at the BBO to available liquidity. In cases where firms overuse IOC orders in order to minimize their own transaction costs, the SOR will often “miss” available liquidity, resulting in the client incurring higher transaction costs.

**Over-reliance upon posting displayed orders** – This behavior would show up as a combination of extremely large order volumes in the non-marketable day category with low fill rates, particularly in venues that pay large rebates. This type of trading can be indicative of trading strategies which either create a lot of market impact or incur large opportunity costs (or both) and should be further analyzed by aggregating the market movement from the time unfilled orders are posted to when they are cancelled.

So, what does all of this have to do with the blind men in the parable? The answer is that, while execution quality statistics within the subcategories recommended will provide useful information, understanding the perspective of the routing firm provides necessary context for evaluation. As an example, if a broker’s “high touch” routing of marketable orders showed relatively high transaction costs, but this was less significant overall than the volume and excess liquidity provided by that firm’s “negotiated” trades, then clients should probably continue trading with the broker. While the report might spark a conversation about how the broker could make improvements, the overall results would justify their value. Similarly, if a routing broker has high costs for routing of marketable orders in its “immediately marketable” orders, it could indicate either a poor router or that their clients are overly aggressive. There are a myriad of similar examples, which is why we believe that this framework is a necessary start towards a true contextual analysis of institutional trading.

---

<sup>i</sup> <https://www.sec.gov/comments/s7-14-16/s71416-13.pdf>