

# On the Limits of Markouts and Venue Curation

Traders on both the buy-side and sell-side focus much attention and energy on trading venues and order types. There is a notion, implicitly encouraged by many trading venues, that best execution can be achieved through **venue curation**—routing to good venues and shunning bad ones. But the reality is more complex: venues offer tradeoffs, and best execution can only be achieved by using venues and order types selectively and intelligently, based on stock characteristics, order characteristics, and dynamic signals.

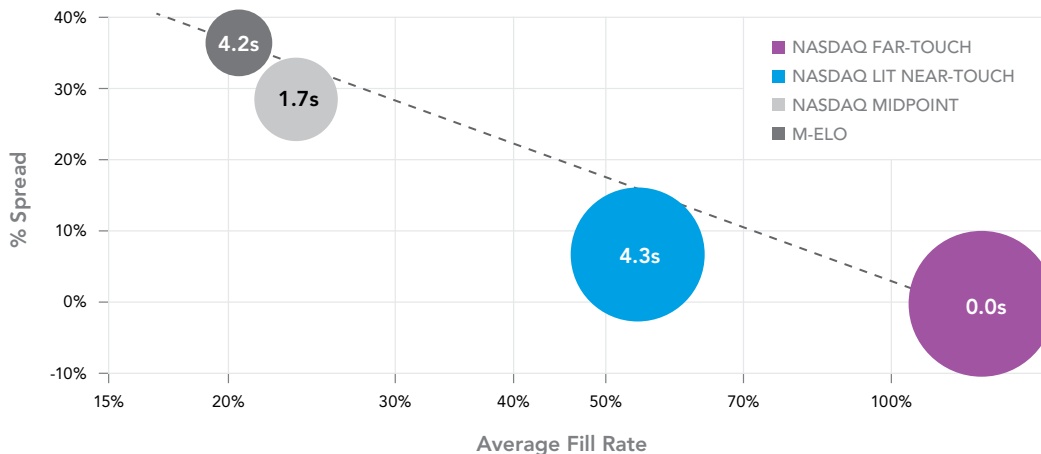
In competing for order flow, trading venues tend toward simplistic marketing claims to imply that their venue or order type is better than others. Displayed venues often tout their volumes and fill rates. Dark venues and hidden or discretionary order type pitches often focus on fill quality, usually measured by markout. Block-oriented venues talk

about trade size. The reality is that venues and order types offer **tradeoffs** between execution quality and fill rate. There is no best venue or order type that should be used across all stock characteristics, trading goals, and market conditions.

Nasdaq Economic Research has highlighted this tradeoff across several Nasdaq order types, reproduced below.<sup>1</sup>

This two-dimensional picture highlights the key limitation of markout: it only measures what happens after a fill—it doesn't measure the likelihood or cost of **not** getting filled. In general, being pickier—avoiding bad fills—improves markout, but comes with a cost: avoiding a fill now may force you to trade later at a

<sup>1</sup> <https://www.nasdaq.com/articles/what-markouts-are-and-why-they-dont-always-matter-2020-07-23>



**FIGURE 1**  
Trade-off between markout and fill rate for NASDAQ order types.<sup>1</sup>

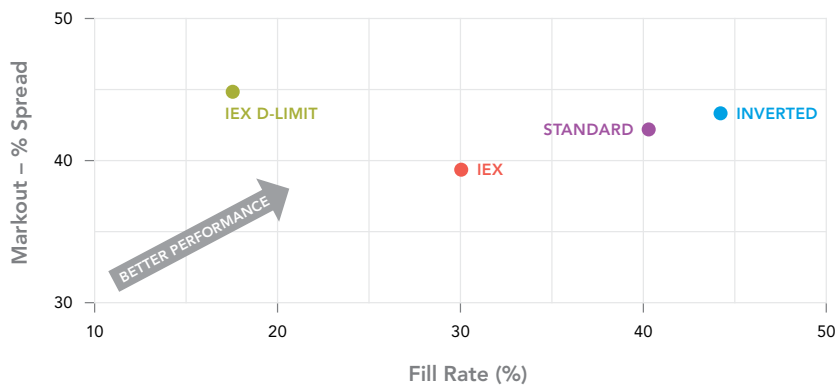


FIGURE 2

### Markout and Fill Rate of Near-Touch Limit Orders

Markout and fill rate of IEX regular limit and IEX D-Limit vs. regular limit orders on two of the largest maker-taker (Standard) and taker-maker (Inverted) exchanges. Markouts are measured 1 second after execution relative to the far touch and include fees. D-Limit appears to have better markout but lower probability of fill. This doesn't tell us anything definitive about D-Limit, but illustrates the first problem with the IEX analysis.

worse price—whether to stay on a VWAP schedule dictated by the trader, or just to eventually complete the order. This isn't rocket science—most traders are familiar with these tradeoffs in the context of using optimistic limit prices, where the benefit of holding out for a better price has to be weighed against the risk you might “miss the market.”

It's generally the balance of these two competing effects, and whether they match the specific trading situation, that determine whether a venue or tactic is good or bad. Using a picky order type when trading an urgent order can degrade performance. Simple use of markout does not show this obvious fact.

## A Markout Case Study: IEX D-Limit

Last fall IEX introduced the D-Limit order type, continuing their track record of innovation. Simply put, D-Limit orders are designed to pull from the market and avoid execution when it looks like the market is likely to trade through them. This is a reasonable idea: avoiding being run over is generally a good thing. IEX posted a blog highlighting this benefit by showing that D-Limit orders have better markout than other venues' ordinary passive orders.<sup>2</sup>

This is fine as far as it goes, but common sense tells us that IEX's D-Limit might come at some opportunity cost in missed liquidity—not just because of the order pulling, but also because IEX as a venue only represents about 3% of market volume, most of it at mid-point, and without the large rebate to attract takers as

some inverted venues have. If those opportunity costs are too high, use of this “smart” order type could hurt performance.

IEX followed up with another blog to answer this question, as they put it “are D-Limit's performance and price improvement worth its lower fill rate?”<sup>3</sup> Unfortunately there were two big limitations in their analysis.

The first problem is that D-Limit was compared not against posting a limit order on say Nasdaq or NYSE, but posting a standard limit order on IEX. This leaves the possibility that while IEX D-Limit is preferable to a regular IEX limit order, it is still inferior to options available on other venues with a higher probability of fill. This is a serious problem because the market running away before you get filled is exactly the kind of opportunity cost you should worry about when posting a picky order on a small venue.

The second problem is the analysis looked only at D-Limit orders that were pulled and where the market then didn't tick up above the original limit price within a second. But this leaves out many possible bad outcomes of posting an order, including all the cases where the order wasn't executed and the market ticked away to above the original posting price. A proper accounting of the cost of any trading tactic must be a probability-weighted average of all the possible outcomes—the bad as well as the good. IEX presumably reasoned that such bad outcomes would be equal between regular IEX and IEX D-Limit orders—but this returns us to the first problem: how many bad outcomes are there when posting IEX D-Limit vs. say posting a regular limit order on Nasdaq, and how do these weigh against the benefits?

2 <https://medium.com/boxes-and-lines/d-limit-one-of-these-things-is-not-like-the-others-3c6b438d3a6e>

3 <https://medium.com/boxes-and-lines/d-limit-performance-the-fill-rates-race-4dcd26661a98>

Despite these flaws, we give IEX kudos for their transparency—other venues, especially ATSEs, for example, provide questionable performance metrics—but IEX sets a positive example in publishing the details of the analysis behind their claims. The underlying problem is that **venues just don't have the data that would allow them to evaluate execution quality fairly**—so analyses provided by venues that focus on markout should be viewed with this caveat in mind.

## Markout is not a Venue Characteristic

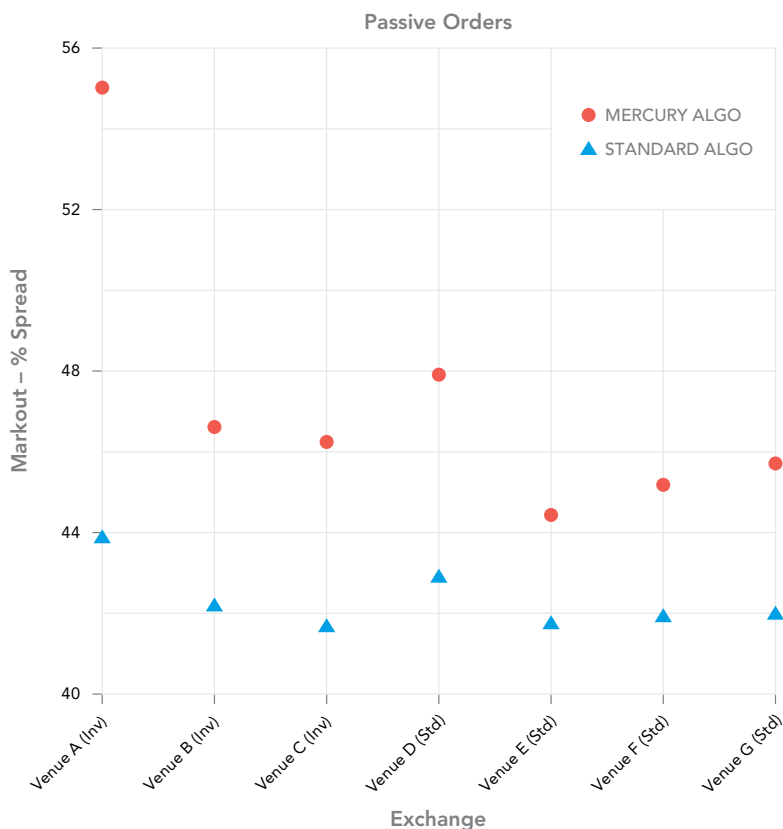
Markout *can* be a useful tool for traders who see the parent order and understand why and when each child order is sent (and canceled).

To use markout properly, it is important to understand that even when you have all the data, **markout is not mainly a characteristic of a trading venue or order type**, but also of the selective and biased way that an algorithm or trader places orders.

This may seem abstract or hypothetical, so we illustrate it using Pragma data. Figure 3 shows markout for passive posted orders across a set of venues from two different algorithms, “Standard”, Pragma’s highly regarded VWAP algorithm and, “Mercury”, Pragma’s new AI-based algorithm, which was in beta testing in 2020. The data set includes over 488,000 parent orders traded over 2020. Overall, the Mercury algorithm had significantly better overall performance as measured by VWAP shortfall and similar metrics.

Strikingly, the **same venues have different markouts** for the two algorithms. This clearly demonstrates that markouts are not an inherent characteristic of venues or order types; they depend at least as much on the times that traders choose to send and cancel their orders.<sup>4</sup>

4 These orders are from the same traders on the same days, with each parent order assigned to one algorithm or the other based on a fair coin toss. As a result, there is no systematic bias that would account for different markouts at the same venue—verified by closely matched distributions of stock and order characteristics (not shown). The orders were traded from the same servers, so technical aspects like latency are also the same between the groups.



**FIGURE 3**  
Same Venue, Same Order Type, Different Algo

Markouts for two algorithms posting similar near-touch limit orders to the same set of venues. The venues are 7 of the largest exchanges, including both standard maker-taker and inverted venues. Markouts are measured 1 second after execution relative to the far touch and include fees.

## Algorithmic Logic is Primary

Intelligent use of dynamic market signals and intelligent order routing are critical to achieving best execution.

The Mercury algo makes better trade-offs between different order types to reduce overall markout across order types by about 3% of spread through the use of more sophisticated, dynamic market signals.<sup>5</sup>

This demonstrates that **the primary driver of**

5 If you are about to complain about our sloppy use of markouts without adequate qualifying context, congratulations! You’ve been paying attention. In this particular case, as mentioned, the performance benefits of Mercury reflected in the markouts are also borne out in parent-order level shortfall improvement and we discuss some of our methodology in, [Measuring Execution Quality](#) which we published in Fall 2020.

**execution quality is algorithmic logic, not venue curation.**<sup>6</sup> Venues and order types offer a palette of options that a well-designed algorithm can use to good effect—but it's not the case that an ordinary algorithm can achieve strong execution quality simply by using the right venue or order type.

Why can't an exchange or ATS solve a bigger part of the trading problem? The core issue is that any venue—whether it's IEX, Nasdaq, IntelligentCross, MEMX, or Level—has to simplify their order type to the point that it's practical for brokers to use. In practice, this means choosing some point on the performance vs. opportunity cost tradeoff. Users of the order type then have to accept that tradeoff whether it's part of a 5 second aggressive sequence or a 6 hour ultra-passive VWAP. The more complexity a venue exposes to work around this limitation—say by offering additional parameters to control the order type, or varying the behavior by symbol—the more difficult it becomes for a broker to understand and use it effectively.

Sophisticated execution algorithms will use similar signals and tactics, but carefully customized to the stock characteristics, the specific demands of a customer's order, and the dynamics of what's happening in the market.

---

<sup>6</sup> We're speaking generally; for short-lived liquidity seeking algorithms, maximal liquidity access is important, and there is relatively little opportunity for sophisticated use of signals.

## Conclusion / Summary

Careful analysis of the performance characteristics of venues is critical to achieving the highest level of execution quality—but the analysis itself, and the algorithm using those venues must be sophisticated enough to make sure each venue is used at the right time and in the right way.

The bar for best execution keeps getting higher. Pragma's next generation AI based algorithms significantly improve shortfall by making more sophisticated use of dynamic signals in the context of stock characteristics and liquidity demands of the specific parent order. A simple qualitative venue curation approach to execution quality that doesn't have that kind of flexibility can no longer deliver state of the art performance.

Finally, wise traders will take markout statistics and claims about performance comparisons in general with a huge grain of salt. They will inquire how opportunity costs were evaluated, and how it was determined that "all else was equal" in the order flow in the two groups being compared before believing any marketing claims.