Incentivizing Trading Behavior Through Market Design

Daniel Aisen, Co-founder and Head of Quantitative Strategy, IEX
December 2017

Introduction

This paper provides an overview of the broad range of trading practices by various participants in the U.S. equity market. It explores the benefit and harm to the overall market of various types of behaviors and discusses specific ways in which IEX has attempted to promote productive trading strategies and discourage harmful practices through thoughtful pricing, system architecture, and market design.

We discuss several harmful speed-based trading strategies but focus primarily on short-term "pick off" strategies in which a market participant trades with a slower counterparty at an imminently stale price, such that the counterparty experiences near-instant remorse. From the counterparty's perspective, this scenario—also known as *adverse selection*—is akin to purchasing an airplane ticket only to find the exact same ticket available at a cheaper price immediately after the purchase is locked-in.

In particular, this paper highlights the upcoming introduction of IEX's CQI Removal Fee, an execution fee designed to isolate and disincentivize aggressive trading strategies attempting to pick off susceptible orders at soon-to-be-stale prices when the market is in transition.

Participants in the U.S. Stock Market

There are several different types of participants in the U.S. equity market, each with distinct needs and objectives. This paper will predominantly focus on participants and trading mechanics involved in institutional trading.[1]

In the institutional world, most trading activity is split between exchanges and alternative trading systems (i.e. dark pools). With the exception of IEX, the vast majority of trading on each of the exchanges results from executions against displayed orders at the NBBO. Conversely, nearly all dark pool trading results from interaction between hidden orders, often at the midpoint of the NBBO, but also sometimes at the bid or offer.[2]

[1] On the retail trading side, most of the interaction occurs off-exchange via direct bilateral relationships between retail brokers (acting on behalf of retail investors) and proprietary trading firms known as retail wholesalers. Retail traders tend to have relatively small-sized trading needs without substantial short-term alpha and as a result, are the most attractive counterparty for a market maker. As such, retail market makers compete aggressively to obtain first-look access to such order flow by offering a combination of payment for order flow (PFOF) and price improvement to retail brokers. It is worth noting that passive retail orders can make their way to the exchanges as retail brokers can earn special outsized rebates when adding liquidity on certain markets, most notably Cboe EDGX. However, by-and-large retail trading and institutional trading are segregated in the U.S. stock market, tied together only by the rules of Regulation NMS, which dictates that all trades must occur in line with the prevailing quotes at the national securities exchanges.

[2] Note that the Canadian equity markets do not generally allow off-exchange trading to occur at the bid or offer. In the U.S., this concept is known as "Trade-At," and the incumbent U.S. exchanges have been lobbying for a Trade-At rule to be introduced in our market as well. In fact, the exchanges were recently successful in incorporating a Trade-At bucket in the ongoing Tick Size Pilot that has been underway for the past year, although to date the pilot program has not yet yielded compelling evidence for or against the practice. The exchanges' central argument is that dark pools free-ride off the price discovery that occurs on the exchanges and that preventing off-exchange trades at such prices would result in more centralized liquidity and more efficient markets. We find this argument to be thinly-veiled and self-serving, but there is something to be said about consolidating liquidity in fewer places. We will discuss the deleterious effects of fragmentation later in this paper, but on this particular issue we do not take a hard stance.

There are a wide range of participants in the institutional equity trading ecosystem, but for simplicity's sake, we narrow our description to four classes: investors (i.e. asset managers), proprietary trading firms, brokers, and venues. There is significant overlap between these participants; for example, many proprietary trading firms have broker-dealer licenses, many brokers operate dark pools and market making desks, as well as asset management divisions, and just about all exchanges have affiliated routing brokers. Still, we find it reasonable to think about the participants in any given trading scenario under these categorizations.

Broadly speaking, we generally think of investors and proprietary trading firms as the two types of end-customers on either side of a trade. One major difference is that investors are typically only on one side of a security at any given time (i.e. only buying or only selling), whereas proprietary trading shops are often agnostic as to the security itself and thus willing to buy or sell anytime if the price is right. Additionally, investors tend to take larger positions and have longer holding periods. Finally, proprietary trading firms tend to have a far more in-depth understanding of market structure dynamics as well as much faster technology at their fingertips than do most institutional investors.

Institutional Investor Trading Strategies

Asset managers tend to devote most of their resources toward making investment decisions, not the actual execution of individual trades. Most institutional traders lean heavily on the tools and expertise of their brokers to effectively navigate the markets. They may make high-level decisions, such as which broker to trade through and how aggressively to take on a position, but they do not typically drill down to venue or order type selection on a trade-by-trade basis.

Brokers typically offer a wide range of tools and tactics to their clients, from basic algorithms like TWAP (time-weighted average price) and DMA (direct market access), to highly customizable dynamic quantitative strategies, often with colorful names and descriptions. Depending on its objectives, each of these algorithms makes several decisions during an execution including, among other things, which venues and order types to utilize, how to break up and space out the submission of individual child orders, and the price at which to trade. Things get even more complicated if a broker integrates block-crossing venues, conditional order types, and internalizers into their strategies.

An institutional broker's decision points when sending an individual order to a single market boil down to three main choices: post or remove; displayed or hidden; and passive, midpoint, or aggressive.

There are many more granular choices a broker can make when sending an order to a trading venue, but these choices provide a reasonable overview of the standard low-level actions a broker takes. Our goal has always been to attract a mix of these various broker order flows in addition to healthy participation from proprietary trading firms. Most trading venues primarily solicit either displayed order flow (exchanges) or hidden order flow (dark pools), but usually not both. IEX is somewhat unique in the U.S. market in that we have attracted a mix of both hidden and displayed volume.

Proprietary Trading Firm Strategies

Over the past several years, high frequency trading has been vilified in the media, portrayed as a predatory practice in which sophisticated algorithmic traders scalp pennies off regular investors.

At the same time, HFT practitioners have argued that they provide an irrefutable benefit, making markets more efficient by providing liquidity. Spreads have never been narrower, and retail traders have never had it better – both very true statements.

The word "HFT" has been used as an umbrella term to describe all automated proprietary trading strategies, but this is an unproductive generalization. We believe that much, if not most, proprietary algorithmic trading activity in the U.S. equity market contributes a substantial positive service to the

industry. However, there are also certain strategies that extract profits from the market without providing any real value. Throughout the rest of the paper, we define the specific strategies we consider predatory, explain their potential harm, and quantify their prevalence.

One simple way to distinguish between proprietary trading strategies is by examining whether the strategy predominantly adds or removes liquidity. This doesn't tell the full story, but it does serve as a helpful gauge. While not a hard-and-fast rule, we generally consider proprietary strategies that predominantly add liquidity to provide a net-benefit to the market, whereas predominantly removing strategies are often—but not always—net-harmful.

Most proprietary trading strategies boil down to connecting natural buyers and sellers, who presumably would not usually find each other if left to their own devices. Whether it's a standard market making strategy providing instantaneous liquidity to a buyer or a seller in a security, or an index arbitrage strategy connecting the market for an ETF to its underlying securities, this basic notion stands.[3]

Even a complex statistical arbitrage strategy that identifies price dislocations between correlated securities connects buyers and sellers across related, though not directly interchangeable, securities.

Before we go further, here is a breakdown of how we think about various proprietary trading strategies:

Strategy	**Description**	**Tag Line**	**Examples**	**Add / Remove**	**Net Impact**
Market making	Always willing to buy or sell a security with the hope of capturing the spread and/or rebates for a profit.	Connect buyers and sellers trading at different times	Posting two-sided quotes	Add	Positive
Fragmented liquidity arbitrage	Connecting buyers and sellers who otherwise might not find each other	Connect buyers and sellers trading in different places	Index arbitrage; dark pool arbitrage; cross border arbitrage	Remove	Neutral-to-positive
Structural arbitrage / latency arbitrage	Taking advantage of fleeting opportunities created by nuances in market structure	Disintermediate between multiple market participants who would likely otherwise interact on their own	Stale quote arbitrage; crumbling quote arbitrage	Remove	Harmful
Order detection	Identifying a large order in the market, buying or selling	Disintermediate between one larger trader and the rest of the	Pinging for the presence of hidden orders;	Remove	Harmful

[3] The same goes for an auction-based strategy in which a trader enters Opening Auction Imbalance Only (IO) orders on Nasdaq at 4am, and then upon receiving Nasdaq's first imbalance message at 9:28am, identifies securities whose auctions have severe enough imbalances and races to take on positions in the pre-market in the same direction. Their IO orders will presumably offset the imbalance when the auction takes place, closing out their newly acquired position for a profit. In this case, the proprietary trader is connecting pre-market buyers or sellers with participants in the opening auction.

	alongside it, and then flipping back the position	market who would likely otherwise interact on their own	electronic front running		

Investors generally benefit from market makers. Market makers consistently provide liquidity and almost always at the best prices. The more market makers out there, the better the market is for investors.[4] While investors prefer trading with other investors whenever possible, market makers do contribute a valuable service filling in the gaps.

Market makers also benefit from investors—without investors, there would be no spreads to capture, and market makers would make no money. As discussed earlier, market makers prefer interacting with small investors and spaced-out broker algorithms like VWAP. Smaller orders spaced out over time are easy to absorb and capture spread on without taking too much risk. On the other hand, market makers can be harmed by large, aggressive investor orders (e.g. aggressive liquidity seeking algorithms or efficient smart order routers) because they can potentially get run over. Getting run over is the worst possible outcome, and a market maker considers this subset of investor order flow toxic. Smaller investors may complain about large investors running them over too, but this isn't really a market structure dynamic—it is just unfortunate timing on the part of the small investor.

We use the term “fragmented liquidity arbitrage” to encompass strategies that connect buyers and sellers seeking to trade the same or similar securities, but unable to find each other due to market fragmentation. For the most part, these strategies do provide a value to the market. It is worth noting that unnecessary fragmentation is a boon to these strategies. In an ideal world, brokers would be perfectly efficient at sourcing liquidity, and there would be fewer such arbitrage opportunities. Overall, we consider this category of strategies to be bloated but neutral-to-positive.

We use the term “structural arbitrage” to describe strategies that make money by exploiting structural inefficiencies, for example trading at a stale price after the NBBO has ticked or gaming the opening auction. Structural arbitrage strategies provide no positive value to the market and are toxic to investors and market makers alike.

Order detection strategies are similarly toxic to everyone they interact with and similarly contribute nothing positive to the market. The most abusive retail wholesaler practices—in which a wholesaler receives a large retail order, causes a substantial price dislocation in the broader market, and then sticks the retail investor with a terrible price—would fall into this category as well.

It is important to emphasize that many of these strategies overlap. For example, a standard market making strategy may incorporate elements of an order detection strategy to better protect itself from getting run over. If a market maker sees a large order sweeping through the market, headed straight for it, it might adjust its quotes to be more passive. This is common sense, and a market maker cannot be blamed for attempting self-preservation. Said market maker might theoretically even use this knowledge to proactively take on a position alongside the large order ripping through the entire market. A market maker who does so blurs the line between providing a significant service to the market and detracting from it.

Additionally, it is important to highlight that none of the strategies described thus far are explicitly prohibited under the law. Unlike illegal manipulative strategies like spoofing, every strategy listed above is

[4] Now, one could argue that the ubiquity of sophisticated market makers at the best prices crowds out investors, and that were these price levels not already saturated, investors might have greater opportunity to trade passively in favorable market conditions. This is a sound argument, although it is hard to imagine an investor effectively competing at all on this front without major advancements in the tools available to them. This line of thinking does lead to interesting potential book priority models, such as prioritizing investor interest at a price level over market maker orders, but this is a topic for another day.

allowed under the existing regulation. Our position is that it would be impractical and ill-advised for regulators to attempt to outlaw even the trading strategies that we consider to be net harmful. We live in a capitalist society, and the typical operators of these strategies are not customer-facing and thus have no fiduciary obligations. We believe there are conflicts of interest in the markets that sometimes enhance these opportunities, and those conflicts should be eliminated, but clever traders who take advantage of these opportunities do not deserve our ire. Further, even if we did believe these practices *should* be illegal, it would be better to simply close the loopholes that make them possible, rather than rely on the regulators to police them. Why hire guards and install security cameras if you could just as easily lock your savings in an impenetrable vault?

One interesting point is that traditional investors are not the only participant on the losing side of the net-harmful strategies listed above. Some of these strategies often profit at the expense of other proprietary traders. For example, take the case of stale quote arbitrage, in which an opportunistic trader recognizes a change in the NBBO and immediately attempts to initiate favorable trades at the old price before the various market centers become aware of the NBBO change. In this case, suppose another proprietary trading firm is making a market at the now-stale NBBO in a dark pool. Even if the market maker recognizes the NBBO change and attempts to update or cancel its order, the opportunistic trader may beat it to the punch. Still, it is probably safe to say that proprietary market makers tend to be in a better position to protect themselves from harmful strategies than traditional investors.

Fragmentation is unequivocally bad for investors and unequivocally good for the latter three categories of proprietary trading strategies listed above. For market makers, the impact of market fragmentation is less cut and dry, as it certainly makes things more complicated, but can also reduce the risk of getting run over (especially if the market maker employs some of these more predatory tactics). Unfortunately, it is very difficult to allow the introduction of new competition and innovation in the market without increasing fragmentation.

Ultimately, it is the broker's job to navigate the market on behalf of its investor clients, and the broker is probably in the best position to help an investor avoid order detection strategies and unnecessary interaction with fragmented liquidity arbitrage strategies. Unfortunately, there is not much a broker can do to avoid structural arbitrage, as these opportunities exist within the fabric of stock trading infrastructure. We believe that only the exchange operators are in a position to realistically make a dent in preventing these scenarios, although they are by no means typically economically incentivized to do so. The remainder of this paper explores several of the steps we at IEX have taken to hamstring these practices and touches upon some of the ideas we have considered but chosen not to pursue.

Using Fees and Rebates to Incentivize Desired Behavior

One interesting trend in equity trading over the past 20 years has been the proliferation of fee/rebate pricing schedules. If adding strategies are so much more desirable than removing strategies, why not pay rebates to adders and charge fees to removers? A trading venue choosing a pricing schedule like this undoubtedly leads to more adders and fewer removers than there would be otherwise. The main problem is that this pricing scheme creates a conflict of interest for customer-facing brokers, particularly given that, for the most part, exchange fees and rebates are not passed through from brokers to their clients. This means a broker who has the opportunity to get its client done by crossing the spread at the prevailing market price has an economic incentive to instead rest the order passively to collect a rebate, risking potentially missing the client's price altogether.[5]

[5] One thought experiment is whether an exchange could get more creative around pricing by offering proprietary trading firms high rebates for adding and charging them disproportionately for removing—to create the right incentives for them—while providing flat pricing to customer-facing brokers to avoid exacerbating this conflict of interest.

One common criticism of IEX is that we have not done enough to encourage lit trading on our market. Currently, only about 25% of volume on IEX is the result of a displayed order being hit/taken, compared to approximately 85% of volume on the other US equities exchanges.

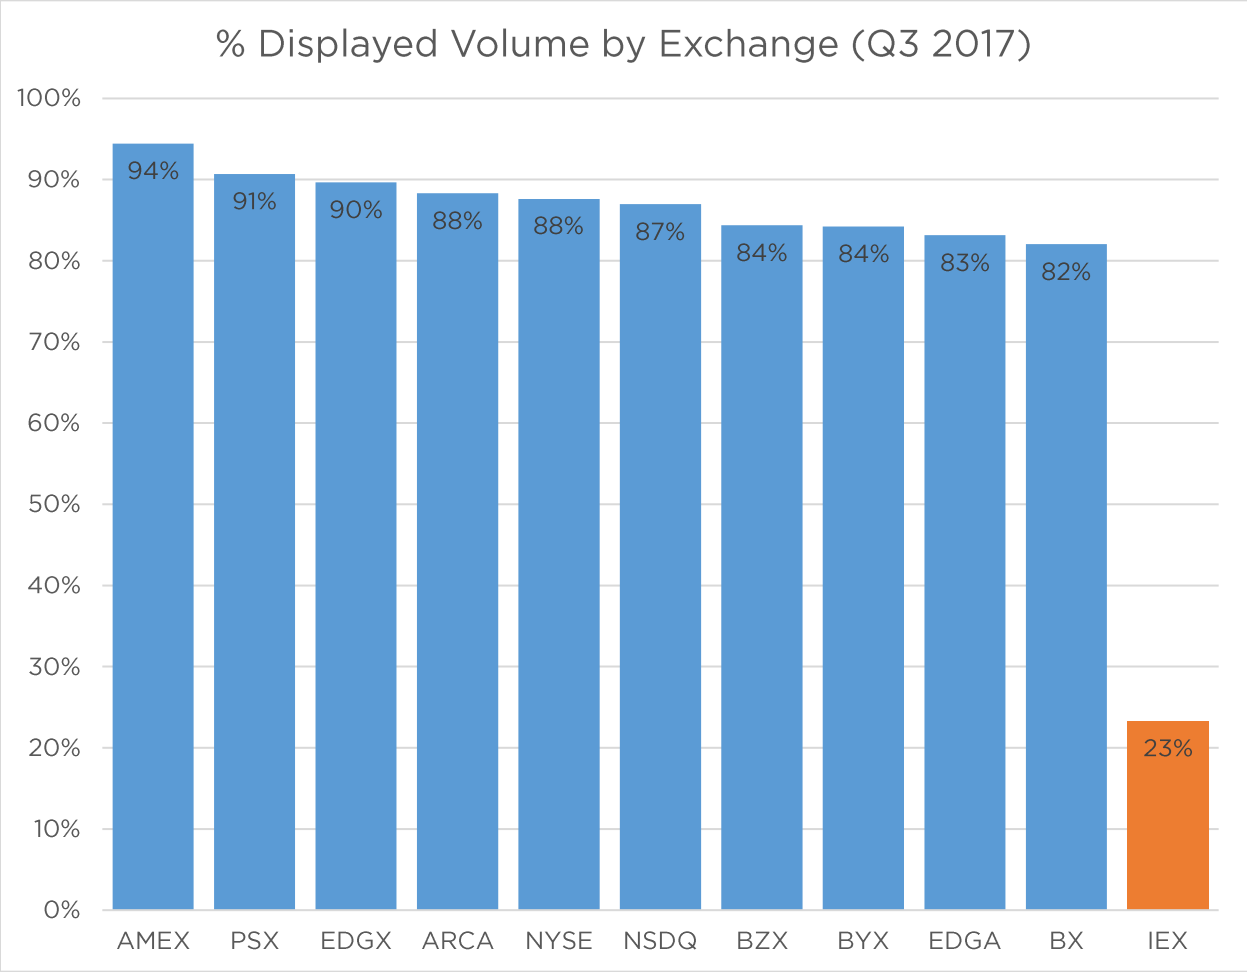

Average Daily Lit Volume %, Single Stocks; Source: SEC MIDAS and IEX

Part of this trend is due to the fact that IEX does not pay rebates. All things being equal, a market maker will choose a rebate-paying exchange over a non-rebate-paying exchange. Additionally, while we have introduced several novel pieces of functionality designed to protect our members from various forms of structural arbitrage, none of these features has directly benefited resting displayed orders.

Preventing Structural Arbitrage and Promoting Displayed Trading

Since our founding, preventing structural arbitrage has been a high priority at IEX. While we believe it is important to reduce the impact of these strategies on both displayed and non-displayed orders, to date, the more unique features we have introduced are disproportionately relevant to resting hidden orders.

First of all, there is the speed bump: the 350-microsecond, 38-mile long coil of fiber optic cable that all order messages must traverse before entering the IEX trading system. The speed bump is designed to ensure that IEX's view of the NBBO is more up-to-date than whatever view the market participant had when they decided to send their order to IEX. While this helps prevent hidden orders from executing

immediately after the NBBO moves through them, displayed orders on IEX contribute to the NBBO, so it is virtually impossible for the market to move through them without them executing.

Second, there is the Signal—the quantitative model IEX uses to predict an imminent quote change and prevent discretionary peg and primary peg orders from executing at unfavorable prices. This feature only benefits these two order types, both of which are fully hidden.

To promote lit trading, IEX has not historically charged a fee to add or remove displayed shares; both parties have received a free execution. As a result, IEX has found itself in a favorable position in the standard economics-based, inter-market effective queue (i.e. most routers allocate shares to remove quotes on IEX before the standard maker-taker exchanges which charge a roughly 30-mil fee to remove). This leads to a shorter time-to-fill and a lower incidence of adverse selection for adders on IEX, just like on the inverted exchanges, which charge a fee to add and pay a rebate to removers.[6]

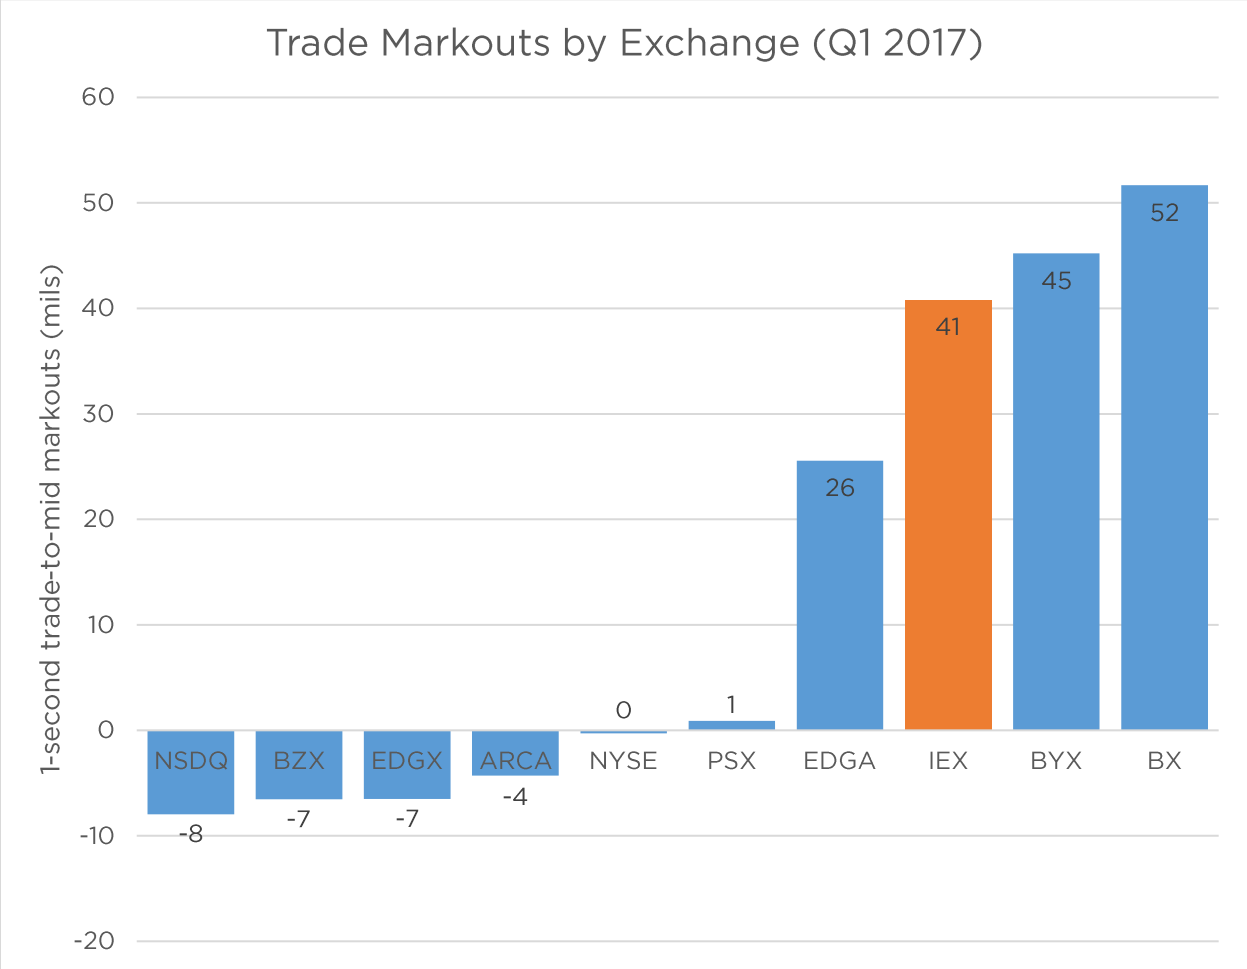

1-second trade-to-mid markouts, adder's perspective, executions at the NBBO, non-locked/crossed markets, Source: A Comparison of Execution Quality across U.S. Stock Exchanges*; NYSE TAQ data*

[6] On January 1, 2018, IEX will begin charging 3-mils to each side for lit executions, but we do not anticipate this price increase to make much of a difference from an execution performance perspective as it will not significantly alter IEX's positioning in the economics-based, inter-market effective queue.

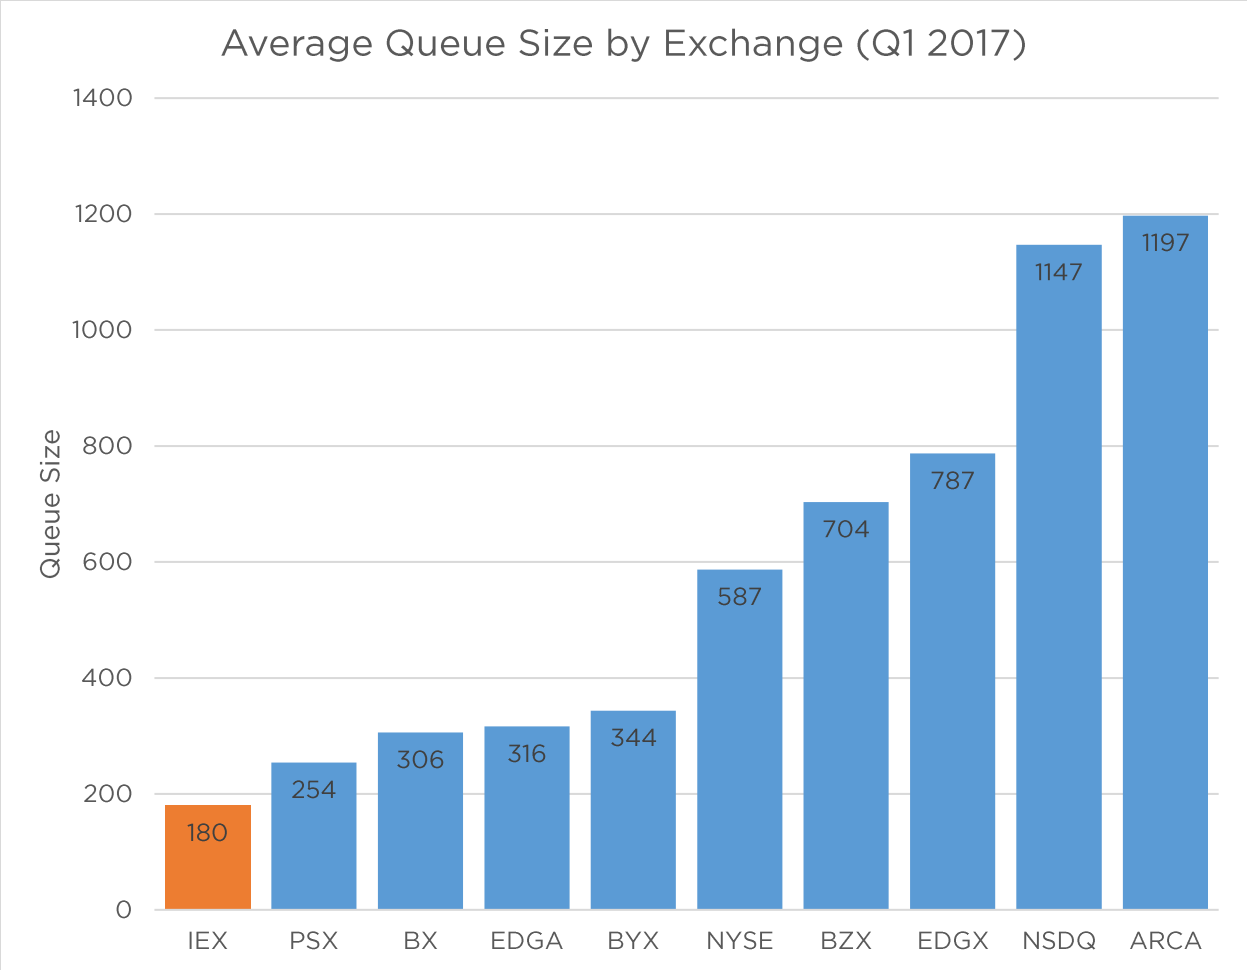

Time-weighted average queue size at the NBBO, Source A Comparison of Execution Quality across U.S. Stock Exchanges*; NYSE TAQ data*

An adder can expect that in the aggregate, they will obtain quicker and better fills on IEX than on the standard maker-taker exchanges, but the fact remains that we simply refuse to pay for order flow with rebates. We believe our value proposition should be enough to attract displayed orders and compensate for the lack of rebates, but so far at least, lit order flow remains a relatively small subset of the IEX ecosystem.

Over the years, we have considered many ways to improve the experience of resting displayed orders on our market including applying crumbling quote protection directly to lit orders or even halting our market altogether when an NBBO change is likely imminent. We have also followed intently the progress of various interesting speed bump proposals on other markets geared toward promoting lit trading, such as the CHX one-way speed bump proposal and the TSX Alpha model in Canada. Each of these ideas has led to vigorous debate and discussion, but ultimately we have shied away from these ideas because we have felt that each one would either be too disruptive to normal continuous trading or likely to have unintended secondary consequences, such as a deleterious effect on liquidity for investors attempting to access the quote.

The challenge has always been to:

1. Discourage opportunistic traders from picking off displayed orders when the market is in transition

2. Without giving market makers the opportunity to fade their quotes when a large investor sweeps the street
3. While also avoiding contributing to the conflict of interest for client-facing brokers

We believe we have found an answer: we have decided—rather than paying a rebate to directly *incentivize* displayed orders—to charge a narrowly-tailored fee to *disincentivize* opportunistic order flow that "picks off" displayed orders. Our goal is to reduce the incidence of this net-harmful order flow, in turn further improving the experience of displayed liquidity adders on IEX.

Before we dive into more detail on this new fee, let us take a brief aside to understand what "good fills" and "bad fills" are from the perspective of the liquidity adder.

Good Fills vs. Bad Fills

In a simple world, a market maker is willing to buy on the bid and sell on the offer all the time hoping to capture the spread. Good fills happen when a small buyer or seller enters the market and trades with the market maker without running them over. Bad fills happen when a large trader (i.e. a *freight train*) enters the market and trades with *all* the market makers, pushing the price of the stock through them. This is also known as getting run over or getting *adversely selected*.

It turns out there are actually two distinct types of scenarios in which adverse selection occurs: when the market is stable and when the market is already in transition.

On a sub-second basis, the market is stable the vast majority of the time. In other words, in almost every instance where an investor (or an algorithm trading on behalf of an investor) takes an action, nothing is happening in the market immediately beforehand. When an investor sweeps the entire inside quote, if they do so in an effective manner, market makers get run over. This occurs in the course of normal market activity.

On the other hand, if an investor attempts to sweep the street with a particularly sloppy router (e.g., one that pings multiple dark pools prior to sweeping the exchanges), market makers may see them coming and have a chance to get out of the way.

Immediately after an effective router sweeps the street, or while an ineffective router is in the process of sweeping, the market will react to this influx of new supply or demand and adjust accordingly. During these periods, we consider the market to be *in transition*. Other examples of events that may cause the market to react include a price move in a correlated security or a news announcement. Note that when we describe a market as in transition here, we are talking on the scale of single-digit milliseconds or less.

When the market is in transition, speedy market makers often attempt to cancel their quotes to avoid getting run over. Additionally, other fast participants rush in to pick off any sitting ducks, for example slow participants who are unable to react on such short timescales or market makers that are ever-so-slightly slower than they are. In this example, we call these participants "pick off traders."

For example, during a period from August 2015 to January 2016, when the IEX router experienced high latency to NYSE, and correspondingly low fill rates, we found that half of the liquidity we missed was the result of market makers canceling their quotes and, surprisingly, half was the result of other fast participants swooping in and picking those quotes off before we got there.[7] To us, this was a somewhat shocking result; prior to running this analysis, we were under the impression that most fading liquidity was due to speedy market makers observing a large counterparty headed their way and adjusting their quotes to protect themselves. This analysis showed that just as often other fast participants were proactively

[7] https://www.sec.gov/comments/10-222/10222-395.pdf

trading to get in front of a market-moving sweep. It highlighted the notion that sophisticated market makers also get picked off by other fast traders, which raises an interesting point:

When an institution sweeps the street in a suboptimal way that leaks its intentions, it kicks off a race between two other parties:

1. The market maker that sees the institution coming and attempts to cancel its quote; and
2. Potentially multiple other speedy traders who observe this scenario playing out and attempt to pick off the quote ahead of the institution.

The party that arrives to the exchange first wins the race. But what is most interesting is that there could theoretically be many fast traders in the latter category, and if any one of them wins, the market maker gets stuck with a bad trade. This means that in this fading liquidity scenario, when an institution fails to access a quote, it creates a many-versus-one race, in which multiple fast traders compete with one market maker to get to the exchange first. If there were say three such pick off traders, we might expect the market maker to get picked off three times for every one time it managed to successfully cancel its quote. In reality, the market maker can probably afford to be a little more “trigger happy” with its cancel request, as the cost of a false positive for the market maker (giving up potentially favorable queue position) is likely smaller than the cost of a false positive to a pick off trader (an actual loss of capital). In our experience, these competing factors seem to balance each other out, resulting in a pick off trader winning the race roughly half the time.

This scenario is just one example of a fast trader swooping in and causing unnecessary disintermediation. We consider these types of ultra-short-term pick off traders harmful to the market. In contrast with other types of arbitrage traders who connect disparate buyers and sellers, these traders connect buyers and sellers that can find each other on their own. As such, they provide zero value to the market and simply extract profit at the expense of other market participants.

The crumbling quote signal is designed to identify when the market is in transition and protect resting orders from these pick off traders. When an effective institutional router sweeps IEX, the Signal does not realize the market is going to transition until afterwards.

For example, see the profile of orders received on IEX from self-classified agency broker-dealers within 1ms of the signal firing:

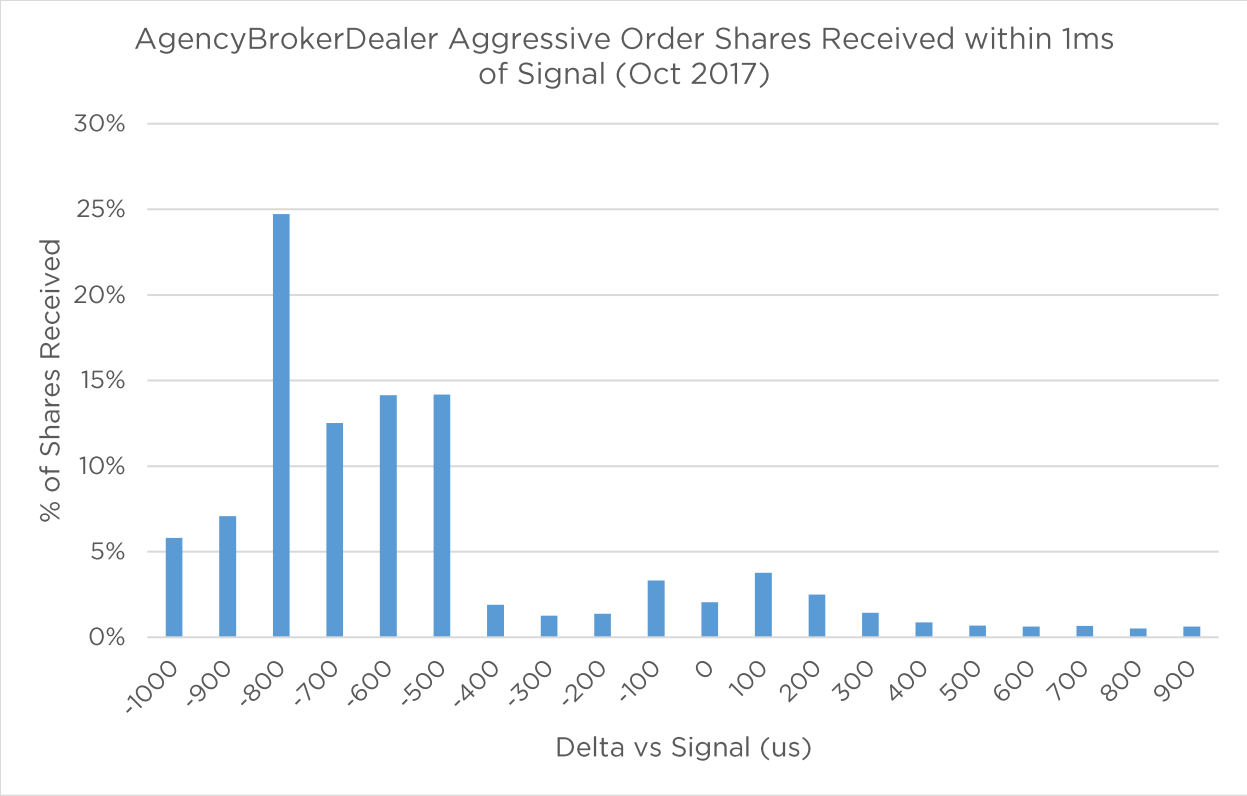

Agency B/D spread-crossing order flow received on IEX within 1ms of the crumbling quote signal firing, Source: IEX

As you can see in the chart above, most aggressive order activity from institutional agency brokers around Signal fire occurs on the negative side of the spectrum, well before the Signal has fired. Only 14% of such order flow arrives *after* we have observed a crumbling quote. This suggests that these firms are typically the *initiator* of a market event. They take action in an otherwise stable market, *causing* the market to transition. In contrast, see the aggregate of all self-described proprietary trading firm order activity on IEX around the signal firing:

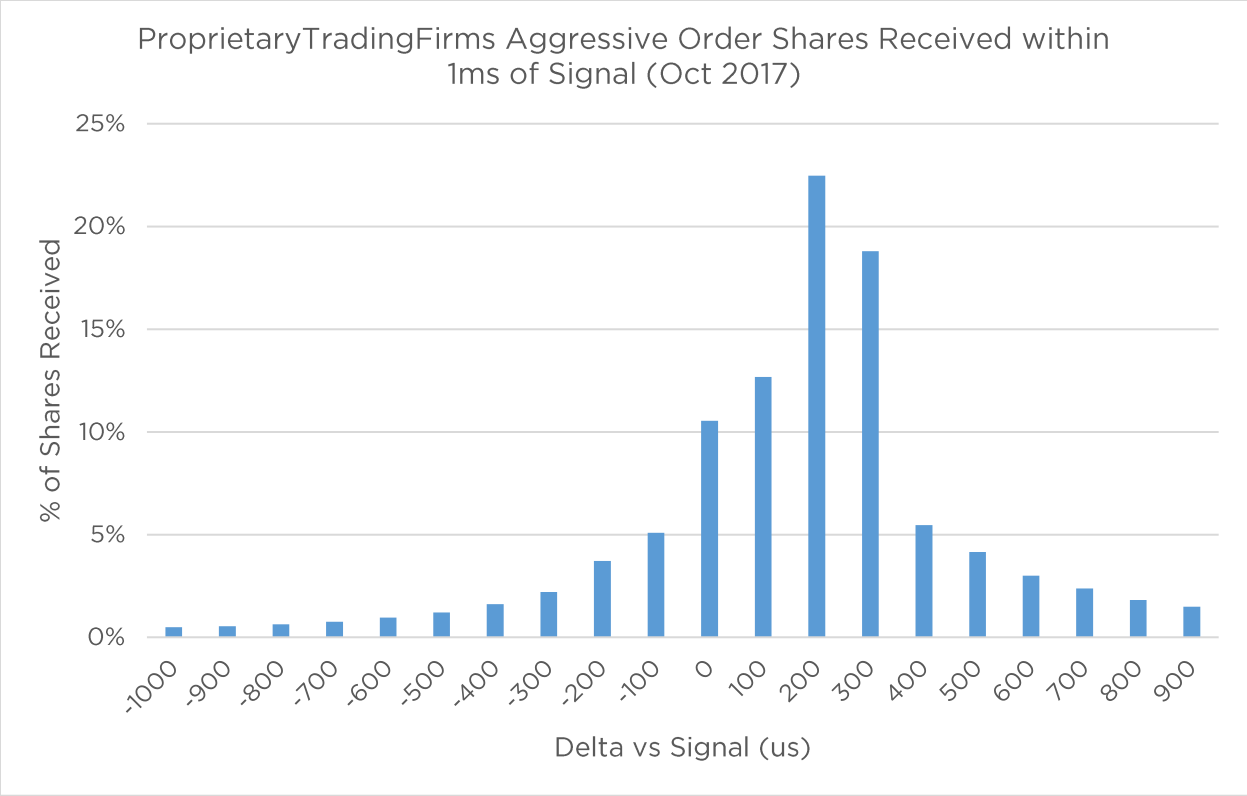

Proprietary trading firm spread-crossing order flow received on IEX within 1ms of the crumbling quote signal firing, Source: IEX

A whopping 83% of aggressive orders received from these firms in the 1ms window around a Signal fire, occur after we have observed the quote to be crumbling. This suggests that order activity from proprietary trading firms typically occurs *in response* to a market event, while the market is already in transition.[8]

This data demonstrates that the signal is quite effective at identifying when the market is already in transition, and can therefore be used to distinguish between the two types of adverse selection. **As such, we are launching a new 30-mil fee for removers while the signal is on to disincentivize reactive pick off flow and reduce adverse selection for all resting orders, including displayed orders.** We discuss the details of this fee in the next section.

The Crumbling Quote Remove Fee

More specifically, for Members (denoted by MPID) who remove more than 1 million shares during the Signal over the course of a month, and for whom this activity represents more than 5% of their total volume on IEX, all such executions above this threshold will be charged a 30-mil fee. The standard fee on IEX is 9-mils for removing hidden liquidity and, effective January 1, 2018, 3-mils for removing displayed liquidity, so when this fee is applied, it will be a substantial increase from the status quo. We will only charge the fee on executions above the 5% threshold because our goal is not to penalize members for incidentally removing during these times. Based on historical analysis of October 2017 trading data, only a small subset of our members would have exceeded the threshold in a meaningful way, and we believe

[8] Additionally, note that the largest columns in the above chart are in the 200-300 and 300-400 microsecond post-signal buckets. Accounting for IEX's 350us speed bump, these spikes suggest that most proprietary trading firms identify that the market is in transition at roughly the same speed as the IEX Signal, or perhaps even slightly faster (but not fast enough to compensate for the speed bump).

that these Members are deliberately operating or enabling an opportunistic strategy designed to pick off resting orders while the market appears to be in transition. Our hope is that this fee will serve as a powerful deterrent to this strategy, and that these Members will refrain from running this strategy on IEX once the fee goes into effect. Additionally, prior to the fee launching, we are disseminating a code on execution reports identifying trades that would have counted toward the fee to give members ample time to adjust their strategies accordingly.

While we expect our trading volumes will likely take a hit in the short term, we believe that the result will be a substantially improved experience for resting orders on IEX, and we will be much better for it in the long run.[9]

In the Signal, we already had an effective means of distinguishing reactive order flow from investor-initiated orders. By simply adjusting our fee during those moments to the maximum allowable amount, we hope to create a powerful disincentive to this particular opportunistic strategy while still having a firm, immediately accessible quote, and without adversely affecting other types of removers.

Expected Impact of the Fee on Displayed Orders

So how do we expect this new disincentive to impact the trading experience of displayed orders? To explore this, we examine the trading experience of self-classified non-proprietary trading firms when adding displayed liquidity on IEX in the below three charts:

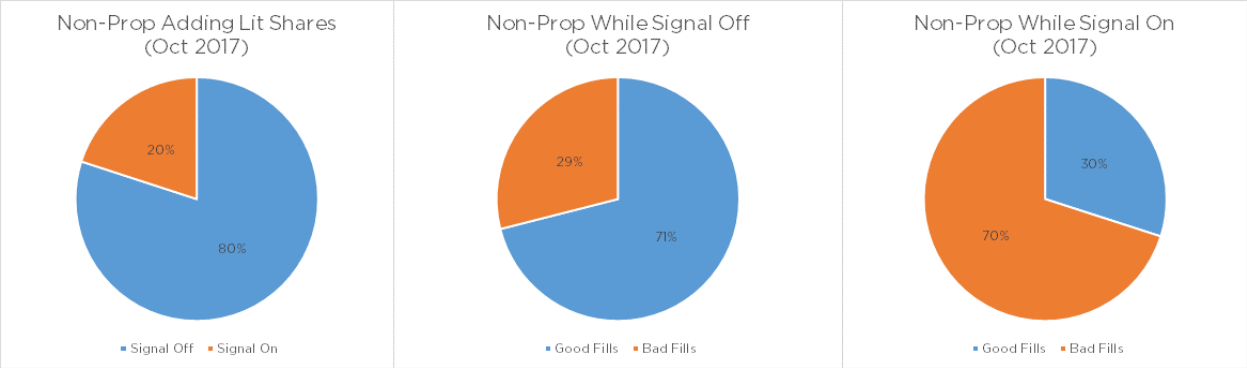

Non-proprietary trading firm experience when adding with displayed orders on IEX, good fills defined as not experiencing adverse selection on a 1-second basis; Source: IEX

The first chart shows that about a quarter of fills for lit adders occur after IEX has identified the market as being in transition. In other words, this quarter of instances represents trading as a result of aggressive orders from *reactive* traders. The two subsequent charts show the quality of execution when trading in normal circumstances versus when trading while the market is in transition. Note: in this case, we define good fills as those which do not experience adverse selection on a 1-second basis, and vice-versa. The results are stark: whereas the substantial majority of trades in a stable market are good fills, a whopping 78% of fills while the market is in transition experience adverse selection. This highlights the toxic nature of these reactive pick off traders.

[9] If the fee is successful at meaningfully reducing the prevalence of these aggressive strategies on IEX, both lit and hidden resting orders should see a tangible performance improvement. Even resting Discretionary Peg and Primary Peg orders, which do not exercise discretion during a crumbling quote, may see an improvement in performance as there are some instances where even those orders can get picked off if the signal experiences a "false negative" (i.e. fails to fire prior to a quote change).

To be clear, though, it is not only client-facing brokers who suffer as a result of aggressive pick off traders:

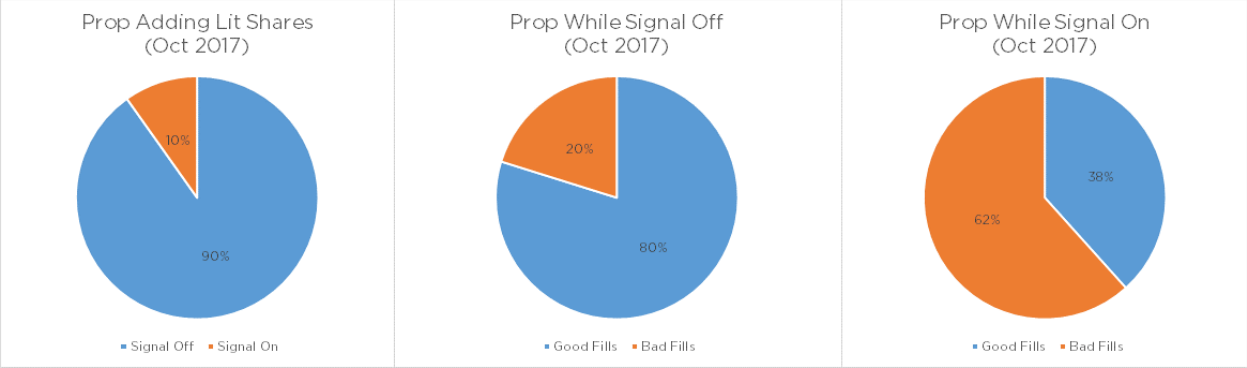

Proprietary trading firm experience when adding with displayed orders on IEX, good fills defined as not experiencing adverse selection on a 1-second basis; Source: IEX

While it is apparent from these charts that proprietary firms do a *better* job at protecting themselves, they are certainly still susceptible to getting picked off.

Now we can't expect 100% of reactive aggressive order flow to completely evaporate overnight, but we do anticipate the fee will be a powerful deterrent.

Criticism of the Fee

Our rule filing proposing this fee to the SEC was met with a small handful of harshly-worded comment letters, which we responded to point-by-point in a comment letter of our own. These letters contained multiple criticisms, but perhaps the most notable point was the notion that since the Signal does not always correctly predict an immediate price move, it is unfair to charge the fee in cases where the quote does not actually change.

Our stance is that this fee is designed to specifically and narrowly target aggressive order flow that arrives in the narrow window of time while the market is already in transition, whether or not the price actually winds up changing. We specifically target this narrow subset of trading activity as it is particularly detrimental to resting orders which contribute to price discovery and a stable market. Just as an aggressive trader seeking to profit during these moments would necessarily use a probabalistic model to determine when to trade, we use a probabalistic model to identify moments where the quote is in transition. While it is true that we do use the immediate future incidence of a price change as a significant factor in optimizing the design of the Signal using historical data, the bottom line is that this fee is specifically designed to apply to trades that occur while the crumbling quote indicator is on; it is not designed to apply to all trades that occur immediately prior to a price change. This distinction allows us to isolate such reactive strategies without impacting institutional routers that *initiate* a market sweep. Additionally, this difference makes it possible for us to provide the fee code in real-time on execution reports.

Another criticism we received homed in on the specific accuracy rate of the Signal, highlighting our own admission in a paper we published earlier this year that when tuning the Signal, we target a 50/50 success rate for predicting imminent ticks. These critics suggest that even if one grants that it is reasonable to use a probabilistic model to determine pricing, our model is simply too inaccurate and thus unfair.

To this point, we emphasize that the roughly 50/50 accuracy described in this paper refers specifically to the likelihood of a tick within the 2 milliseconds following a firing of the Signal. To be clear, when the Signal fires, it correctly predicts the direction of the next tick roughly 80% of the time. This means that a pick off strategy successfully executing off of this model would realize a profit within 2 milliseconds half of the time, and the other half of the time, it would have roughly a 60/40 win-loss ratio. In reality, speed-based trading is a highly competitive space with heavy barriers to entry and a continual technology arms race, so actually pulling off a strategy like this is no easy feat. However, since we are simply trying to identify moments in which such aggressive reactive strategies are likely to enter orders, and since we do not need to be the absolute fastest at identifying these moments (we only need to fire our Signal within 350us of the fastest trader), we believe this model serves our purposes effectively.

Expected Beneficiaries of the Fee

We expect the main beneficiary of this new fee will be market makers making lit markets. There are situations where a stock is in transition, and a market maker and a fast trader operating a pick off strategy get in a race: the market maker tries to cancel their order and the pick off strategy tries to clip them at a soon-to-be-stale price. This fee will disincentivize the predatory strategy and allow the market maker to adjust their quote.

Non-HFT Members posting lit orders will also benefit, although to a lesser extent. There are plenty of scenarios where an opportunistic reactive strategy swoops in to pick them off as well, and in these cases it is not even a race, as non-HFT members likely make little attempt to protect themselves. Still the fee will serve as a disincentive to the pick off strategy. As a result, in at least some cases, the stock will not wind up ticking as the lit order will simply remain in the market unexecuted. This will equate to a reduction in adverse selection for non-HFT lit adders in exchange for a modest reduction in fill rates, which we consider a positive trade-off for the adder.

The other beneficiary will be users of other resting midpoint peg and hidden orders who, unlike discretionary peg and primary peg users, do not benefit from crumbling quote protection. We expect that these orders will get picked off less often as well.

In addition to these primary effects, we expect there may be secondary benefits as well:

- Market makers can post thicker and tighter quotes as their adverse selection rates will be reduced
- Members with imperfect routers that suffer from information leakage and do not get ~100% fill rates will get electronically front-run less often on IEX

Ultimately, only time will tell the extent to which this initiative is successful, but we are very optimistic.

Conclusion

The U.S. equity market encompasses a wide range of trading participants and stakeholders, each employing tools, technology, and market structure know-how to accomplish its individual objectives. Institutional and retail investors seek to participate in the growth of corporations in exchange for capital allocation. Brokers enable investors by helping them navigate the markets as effectively as possible. Proprietary trading firms seek to connect disparate buyers and sellers across time, markets, geography, and securities, collecting a profit for this service in the process. These groups come together and interact on centralized exchanges and in dark pools. A healthy stock market is one that fulfills the needs of each of these participants by promoting competition, fairness, and transparency.

Whereas the HFT community has been vilified in the media over the past several years, we feel it is necessary to dive deeper into this umbrella term to understand the true picture. There are a multitude of proprietary trading strategies, many of which we consider to be productive and helpful to the market, and others that we consider to be purely predatory. At IEX, we have chosen to differentiate our exchange by

defining the specific strategies we consider to be harmful and then attempting to eliminate these practices on our exchange through market design or economic disincentives. At the same time, we are committed to maintaining the benefits brought on by advancements in technology, such as lower costs, faster speeds, and greater scale.

We have heard the question raised several times about whether this sort of intervention by a stock exchange is appropriate. After all, is it not the broker's role to protect its clients from predatory trading practices? Rather than picking winners and losers, shouldn't the exchange be a neutral utility designed to make trading as efficient and frictionless as possible? In general, we could not agree more. It is absolutely the broker's responsibility to get best execution for its clients. But unfortunately, there are certain black-and-white trading scenarios that exist in the very fabric of the trading system, in sub-millisecond windows where it is virtually impossible for a broker to swoop in and protect a client's order. After all, if a broker is not always the single fastest participant in the market, its customers are at risk of getting picked off in these moments. Even the most technologically sophisticated market makers—who are perhaps not constrained by the same level of regulatory burden as customer-facing brokers—cannot fully protect themselves in these moments. Further, it is unrealistic to expect that regulatory intervention would be appropriate or effective at thwarting these scenarios. As such, the exchange is the *only* party in a position to make a meaningful difference here.

At IEX, we see it as the exchange's duty to identify these clearly unproductive trading scenarios and eliminate them to the extent possible. The U.S. equity market has seen tremendous benefits from technological advances over the past twenty-five years, but it is still not perfect. We wish to keep these benefits while also making the market as fair and transparent as possible for investors.