The Good, the Bad, and the Ugly: How algorithmic traders impact institutional trading costs ☆

Tālis J. Putniņš [a,b,*] and Joseph Barbara [c]

[a] *University of Technology Sydney*
[b] *Stockholm School of Economics in Riga*
[c] *Australian Securities and Investments Commission*

September, 2020

Abstract

We show that behind the aggregate effects of algorithmic and high-frequency trading (AT/HFT) lies rich heterogeneity in how individual algorithms impact institutional trading costs. Using unique order-level trader-identified regulatory data, we find that the most "harmful" algorithmic traders double institutional trading costs. "Beneficial" algorithmic traders offset much of this increase. We find no evidence that speed (e.g., being an HFT) is associated with an increased likelihood of being a harmful trader. We identify other characteristics that distinguish harmful and beneficial algorithmic traders. The heterogeneity explains why AT/HFT appear detrimental to some investors despite being beneficial or benign in aggregate.

JEL classification: G14

Keywords: algorithmic trading, high-frequency trading, liquidity, trading costs, implementation shortfall, predatory trading

☆ Putniņš gratefully acknowledge funding from the Centre for International Finance and Regulation (CIFR) under grant number F005. This study expands on analysis undertaken by the Australian Securities and Investments Commission (ASIC) market surveillance to identify potential systemic sources of poor quality execution outcomes for investors. The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of ASIC, CIFR or CIFR's consortium members. This paper has benefitted from comments and suggestions from David Easley, Jonathan Brogaard, David Gallagher, Philip Gharghori, Adrian Looi, Maureen O'Hara, Vincent van Kervel, Ioanid Roşu, and seminar participants at the Q-Group Colloquium, the UBS Conference on Investment Management and Capital Markets, Chi-X Sell-side Roundtable, Macquarie Securities Buy-side Roundtable, Quorum 15: International Equity Quorum, the J.P. Morgan Quantference, the UK Financial Conduct Authority, Macquarie Global Quantitative Research Conference, Hong Kong Equity Trader Roundtable, Paul Woolley Centre for Capital Market Dysfunctionality Annual Conference, 6th Behavioral Finance and Capital Markets Conference, Goethe University Frankfurt, Research Centre for Sustainable Architecture for Finance in Europe, the Baltic International Centre for Economic Policy Studies, the Stockholm School of Economics in Riga, BI Business School Oslo, the University of South Australia, the Australian Securities and Investments Commission, and the University of Technology Sydney.

* Corresponding author: Tel.: +61 2 9514 3088. Fax: +61 2 9514 7722. Email: talis.putnins@uts.edu.au. Post: UTS Business School, PO Box 123 Broadway, NSW 2007, Australia.

1. Introduction

Technology has fundamentally transformed how trading occurs on financial markets, but not everyone agrees that it is for the better. Few changes in securities markets have ever generated as much debate and disagreement as algorithmic and high-frequency trading (AT and HFT). On one hand, many academic and regulatory studies find that AT/HFT *in aggregate* is beneficial (e.g., lowering spreads and improving price discovery) or at worst benign. Yet, at odds with this view, many institutional investors claim that finding liquidity for large orders has become more difficult in modern markets, they face greater price slippage in trading large orders, and their trading costs have suffered as a result.[1] Investors often blame predatory trading, order anticipation strategies, or "back running" by algorithmic traders, in particular HFTs.[2] In response, new trading mechanisms are emerging to cater for institutional investors and their concerns about increasing market impact costs.[3]

In this paper, we seek to reconcile these conflicting views by examining the impact of AT/HFT on institutional trading costs, rather than simple liquidity measures, and explicitly accounting for the rich heterogeneity in the cross-section of AT/HFT. We show that some AT/HFT are extremely beneficial to markets and to institutional investors, while others consistently increase institutional trading costs. We propose that this heterogeneity helps understand the conflicting views about the effects of AT/HFT. While *in aggregate*, AT/HFT are beneficial or benign, an investor that disproportionately interacts with harmful AT/HFT will face higher trading costs.[4] Such investors may naturally view contemporary markets and the growth in AT/HFT with skepticism, with their concerns reflected in financial media, lobbying effort, and the emergence of alternative trading mechanisms. We show that the harmful AT/HFTs are balanced by a segment of highly beneficial AT/HFT that considerably decrease market impact costs of institutional investors. We go further and analyze what characteristics distinguish the beneficial AT/HFT from the harmful ones.

---

[1] For example, "as big institutional buyers and sellers, if we can't find blocks we have to trade in smaller sizes, across multiple venues using algos ... which leaves us open to being taken advantage of by HFT and other participants" (Richard Nelson of T. Rowe Price, quoted in Global Trading, Nov 21, 2015).

[2] Brunnermeier and Pedersen (2005) and Yang and Zhu (2020) provide theoretical evidence of such strategies, and Hirschey (2020) and van Kervel and Menkveld (2019) provide empirical evidence suggesting that some HFTs employ such strategies.

[3] Large institutions increasingly rely on block crossing networks, dark pools, and closing auctions where liquidity is consolidated. NYSE, LSE, and Chi-X Europe have announced the introduction of additional batch auction mechanisms for liquidity consolidation. New trading venues such as Plato in Europe and Luminex in the US specifically claim to shield users from HFTs to lower trading costs.

[4] Our reference to a group of "harmful" AT/HFTs is from the perspective of an institutional investor primarily concerned about the effects on their trading costs. This group of AT/HFTs may bring some benefits to the market such as faster price discovery.

Our ability to analyze the cross-sectional heterogeneity in the universe of AT/HFT stems from our unique regulatory audit trail data. The data contain all trades and quotes for all participants in the Australian equities markets, identified at the most granular possible level of individual traders.[5] Using these data, we are able to aggregate "child orders" into the "parent orders" of institutional investors. With the parent orders in hand, we accurately measure institutional trading costs using implementation shortfall, or total cost of trading the parent order, accounting for the order's price impact through the course of trading.[6] To the best of our knowledge no other study has examined such a comprehensive institutional trading cost measure across *all* institutional investors in a market. We are also able to observe the AT/HFT traders in the data as the most active (highest volume) non-directional traders: traders that turn over positions within relatively short time frames, in contrast to institutional investors who tend to "buy and hold". These 187 traders account for approximately half of the trading volume and contain a mix of fast traders (HFT) and non-HFT algorithmic traders (AT).

Our first finding is that the individual traders (AT/HFT accounts) tend to neatly cluster into two distinct groups: (i) those that systematically increase institutional trading costs and thus appear "toxic" to institutional investors, although they may bring other benefits to the market such as improved price discovery, and (ii) those that systematically decrease institutional costs and thus appear "beneficial" to institutional investors.[7] Our basic approach involves regressing institutional trading costs (implementation shortfall of large parent orders) on the activity of each of the AT/HFTs. The coefficients provide "toxicity" scores for each of the traders, resulting in a cross-sectional distribution of toxicity. We use instrumental variables to pin down causality.

The high level of granularity in our data poses additional challenges. Working at the trading account level, our regressions effectively have 187 right-hand side variables (one for each of the individual AT/HFT accounts). Even if none of these traders have any relation with institutional trading costs, some will appear toxic (others beneficial) purely by statistical chance. This is analogous to the problem of disentangling skill from luck in the cross-section of fund managers. Consequently, we borrow from the funds management literature and use bootstrap simulations to adjust our toxicity estimates to account for the variation that is expected from

[5] The Australian equities market is similar to the US and other major equity markets with respect to the types of trading platforms, the level of HFT trading activity, the major trading firms that participate in the market and the trading technology that they use, and the level of institutional holdings. We elaborate on this point in the data / institutional details section of the paper.

[6] Implementation shortfall compares the average price at which the entire parent order is executed to the price at the start of the parent order.

[7] We use the term "toxic" to refer to traders that increase institutional trading costs. The term "toxic" is also used in the market microstructure literature from the perspective of a market maker to refer to informed or unbalanced order flow, which can cause losses to the market maker.

statistical chance. The bootstrap simulations provide strong evidence that the toxicity and beneficial trading observed in the data are not just artefacts of statistical variation.

The magnitudes of the effects of toxic and beneficial traders are economically meaningful. The toxic traders increase the cost of executing large parent orders by more than ten basis points, roughly doubling the cost. The additional trading costs equate to around $437 million per annum for large institutional orders in the top 200 stocks. The effects of the toxic traders are offset by beneficial traders that reduce institutional trading costs. Consequently, AT/HFT *as a group* have little or no net effect on institutional trading costs. The aggregate effects mask considerable heterogeneity and large effects within segments of the AT/HFT population.

We show that the heterogeneity in the cross-section of AT/HFT has important implications. One implication is that how institutions execute large orders is likely to have a considerable impact on trading costs. This is because institutions that disproportionately trade with toxic traders will experience higher trading costs. The magnitudes suggest that carelessly managed execution can have a material effect on a fund's performance. Minimizing trading costs can involve strategies that change the probability of interacting with toxic traders.[8] Institutional investors are becoming increasingly important market participants with their holdings of US large-cap equities having increased to approximately 80%.[9]

Another implication of the heterogeneity is that market structure changes that increase the amount of AT/HFT could have positive or negative effects depending on whether the changes disproportionately encourage toxic or beneficial traders. This fact may explain why studies of AT/HFT that use different exogenous events as instruments often arrive at different conclusions about the effects of AT/HFT. For example, the finding of Brogaard et al. (2015) that the introduction of co-location at Nasdaq OMX Stockholm improves market quality, implies that this event disproportionately encouraged the activity of the beneficial AT/HFT.

Who are the toxic traders? Toxic traders are likely to trade *with* institutional order flow rather than *against* it. While some toxic traders might intentionally exploit institutional order flow (e.g., predatory, order anticipation, or back-running strategies), others unintentionally amplify institutional trading costs by trading on common entry/exit signals, or through active

[8] For example, investing in the design of execution algorithms to better conceal trading intentions (better mixing of order sizes, times, types, venues), dynamically adjusting the parameters of an execution algorithm (e.g., aggressiveness or trading pace) in response to the level of toxicity, searching for off-market block crossings when toxicity is high, or pausing a parent order in such conditions, and so on.

[9] Based on 13F institutional holding filings, institutions hold more than 80% of the free float of large-cap US equities in 2015, compared to just 50% in the year 2000 (NBIM "Asset Manager Perspectives Report", 02/2015).

participation in price discovery in the presence of order flow imbalances. Therefore not all toxic traders are necessarily harmful to the market overall—some might increase institutional trading costs by improving price discovery around institutional parent orders. In contrast, the beneficial traders are likely to be liquidity-providing intermediaries that "lean against the wind" and thereby attenuate the price pressure that arises from large institutional orders.

We also shed light on the characteristics of toxic traders. Across a number of characteristics, we consistently find that speed is not associated with toxicity—there is no evidence that HFTs are any more toxic than non-HFTs. For example, a trader's share of dollar volume is unrelated to their toxicity. Prior studies show that the fastest traders have the highest volume shares and high volume is used as the defining characteristic that separates HFTs from other non-directional traders (e.g., Kirilenko, Kyle, Samadi, and Tuzun, 2017). Other characteristics corroborate this finding: toxicity is unrelated to the speed of order amendments, the frequency of fast orders, the consistency with which a trader extracts intraday trading profit, and order-to-trade ratios. On the basis of our evidence, concerns voiced by institutional investors about HFTs in particular being the culprits responsible for increased trading costs are misdirected.

We find that the typical toxic trader tends to concentrate their activity in a subset of stocks and is fairly consistently active in that "preferred habitat". We find that the preferred habitat tends to be smaller stocks. The typical toxic trader also tends to close a larger proportion of their positions before the end of the day and subsequently holds less inventory overnight than the typical beneficial trader.

Our study also has a number of regulatory applications and policy implications. First, our approach to measuring the toxicity of individual traders can be used as a market surveillance/monitoring tool.[10] Some manipulative trading strategies that exploit other market participants are likely to produce relatively high toxicity scores and can therefore be detected with our approach. Second, the finding that speed is unrelated to toxicity suggests that regulation or policies that disproportionately restrict fast traders or HFTs are unlikely to result in substantial improvements to institutional trading costs.

[10] A similar approach is being used by the Australian Securities and Investments Commission.

Related literature

On a broad level, this paper is related to the literature on the effects of AT and HFT on market quality. For good surveys of this literature see Jones (2013) and Menkveld (2016). Most of these studies find that AT/HFT in aggregate are beneficial, lowering bid-ask spreads and improving price discovery, or at worst benign (e.g., Hendershott, Jones, and Menkveld, 2011; Hasbrouck and Saar, 2013; Menkveld, 2013; Brogaard, Hendershott, and Riordan, 2014; Boehmer, Fong, and Wu, 2020; Brogaard, Hagströmer, Nordèn, and Riordan, 2015; and van Kervel, 2015).

Our study differs from this literature in three main ways. First, we focus on the heterogeneity across AT/HFT at the granular level of individual traders, rather than their aggregate effects. We show that there is a lot going on beneath the surface of the aggregate effects and that the heterogeneity has important implications. Two notable exceptions to the tendency to analyze aggregate effects is Hagströmer and Nordèn (2013) who separate HFTs into market making and "opportunistic", and Boehmer, Li, and Saar (2018) who use principal component analysis to identify three underlying strategies common to most HFT firms. Both of these papers highlight the heterogeneity among HFT and importantly they show the heterogeneity is important for understanding HFT behavior. However, they do not examine the effects of these groups of HFT on institutional trading costs, nor do they partition HFTs further than the two or three categories.

Other studies recognize that the effects of HFTs depend on whether they use their speed to provide liquidity more competitively or use speed to pick off stale quotes and thereby impose adverse selection costs on market makers (e.g., Biais, Foucault, and Moinas, 2015; Foucault, Kozhan, Tham, 2017; Brogaard, Hendershott, and Riordan, 2017). These studies imply that AT/HFT are not monolithic—different trading strategies should have different effects—but they do not analyze the heterogeneity among individual traders and they do not examine how AT/HFT impact institutional trading costs.

The second difference between our study and most of the AT/HFT literature is that by using regulatory data to reconstruct institutional parent orders, we examine comprehensive measures of institutional trading costs rather than simple liquidity measures. Accurate measurement of institutional trading costs is important to our study to be able to reconcile the conflict between institutional investor concerns and existing evidence on the effects of AT/HFT. Simple liquidity measures such as bid-ask spreads and depth are not necessarily a good reflection of institutional trading costs because the cost of executing a large parent order depends to a large extent on the price impact ("slippage") that occurs through course of executing the series of child

orders. Using trading costs for a self-reported sample of institutional investors, Brogaard, Hendershott, Hunt, and Ysusi (2014) find that HFTs as a whole have a negligible effect on institutional trading costs. This result is consistent with our findings that the effects of toxic traders are offset by a group of beneficial active traders and fast traders are no more toxic than slow traders.[11]

The third difference is that while most of the recent studies in this literature focus on HFTs, we cast our net wider and analyze a broad cross-section of traders that includes HFT and non-HFT. We do so because predatory traders are not necessarily HFTs. Speed is not essential to exploit institutional investors that manage orders in the market over horizons of hours or days. Starting from a broader set of traders has the advantage that it allows us to ask whether speed is a characteristic of traders that increase costs; the evidence indicates that it is not.

Our paper is also related to literature on trading strategies that intentionally exploit institutional investor order flow and thereby increase their trading costs. In a theoretical model, Brunnermeier and Pedersen (2005) show that strategic traders often provide liquidity to institutional investors but, at times, predate on distressed institutions that are forced to liquidate a position. When strategic traders become predatory, they initially trade in the same direction as the institutional investor, amplifying price impact (increasing implementation shortfall), before closing the position. Yang and Zhu (2020) show in a two-period Kyle-type model, that strategic traders engage in "back-running" informed institutional orders, i.e., identifying (with noise) the presence of a large informed institution and then trading in the same direction.

Three recent studies provide empirical evidence consistent with HFTs engaging in order anticipation, predatory, and back-running strategies at different frequencies. Hirschey (2020) finds that on Nasdaq, HFTs' aggressive trades predict non-HFT order flow over the next 30 seconds. Korajczyk and Murphy (2019) examine HFTs in Canada that use predominantly (>80%) passive orders and find that they reduce liquidity provision to (and switch to competing with) large "stressful" institutional trades. Van Kervel and Menkveld (2019) examine HFT trading around the orders of four large institutions in Stockholm and find that HFTs tend to "lean against the wind" initially (provide liquidity to the institutional orders), but then after a few hours, trade "with the wind", amplifying institutional trading costs. These studies are complimentary to ours as they illustrate trading patterns that are likely to be responsible for some of the toxicity that we quantify. Our study differs in several ways. First, we examine a broader cross-section of traders and show that it is not only HFTs that impose toxicity on institutional orders, and in fact

[11] A similar study by Tong (2015) reaches different conclusions to Brogaard et al. (2015) regarding the effects of HFT as a group on institutional trading costs.

there is no evidence that HFTs are any more toxic than non-HFTs. Second, we do not limit the analysis to specific strategies such as order anticipation/predatory/back-running strategies and instead identify all traders that systematically increase institutional trading costs irrespective of how they do so. Third, we explicitly allow for heterogeneity across active traders and analyze the characteristics that distinguish toxic traders from the rest.

Finally, this paper is also related to the smaller literature on institutional trading costs more generally. Anand, Irvine, Puckett, and Venkataraman (2012) characterize the heterogeneity across *institutional investors* and brokers in their trade execution abilities. In contrast, we characterize heterogeneity across *active traders* and intermediaries that either harm or benefit institutional trading costs. Anand et al. (2012) find considerable dispersion in trading-desk and broker skill, and show that the trade implementation process is economically important and can contribute to relative portfolio performance. Our findings support the conclusion that the trade implementation process can have first-order effects on performance because the increases in costs from disproportionately interacting with (or leaving orders exposed to) toxic traders are considerable (same order of magnitude as unconditional mean implementation shortfall).

2. Data, trader types, and trading costs

2.1. Sample and institutional details of the market

Our sample covers trading in the largest 200 Australian equities (ASX 200 Index constituents) during the 13 month period September 1, 2014 to September 31, 2015 (273 trading days).[12] At the stock-day level, this gives us a panel with 52,873 observations. We use unique trader-identified regulatory audit trail data to construct comprehensive estimates of institutional trading costs and identify the most active non-directional traders (AT/HFT) and measure their trading activity.[13] The identification of individual traders occurs through the "origin of order" identifiers that the main market regulator, the Australian Securities and Exchange Commission (ASIC), collects within the regulatory data feed obtained under the Market Integrity Rules.[14]

[12] Because stocks enter and leave the ASX 200 during the sample period, in total our sample includes 225 stocks, although not all stocks are in the sample at all times.

[13] ASIC process the audit data and provide only aggregated data that is purged of confidential information. We exclude block trades and after-hours trades, and include all trades executed through any of the lit or dark trading venues in Australia.

[14] Under the Market Integrity Rules, all market participants must provide to ASIC (via market operators) information about each order submitted to and trade executed on a market. Among the fields that must be submitted is an identifier for the "the person on whose instructions the Order is submitted or Transaction was executed", which allows ASIC to identify all the order and trades originating from individual traders or entities within each broker.

The Australian equities market is in the top-ten largest in the world by market capitalization, with total market capitalization of around $1.6 trillion.[15] It has two lit trading venues, the Australian Securities Exchange (ASX) and Chi-X Australia, and a number of dark trading venues (an exchange-operated dark pool, Centre Point, and 17 broker-operated dark pools during our sample period). Additionally there is an "upstairs" market for block trades, which are allowed to be negotiated off-market at any price if they exceed size thresholds ($1 million, $0.5 million or $0.2 million, depending on the liquidity category of the stock). The two lit venues operate centralized electronic limit order books, which trade approximately 52% and 8% of total dollar volume, respectively. Their technical protocol is effectively the same as that used on Nasdaq (it is owned by Nasdaq OMX Group). The ASX opening and closing auctions account for a further 15% of volume. Block trades account for 15% of volume and below-block size dark trading accounts for the final 10% of volume.[16] Therefore, there is considerable fragmentation of trading in the Australian equities market, comparable to Canada and several European countries (somewhat less extreme than in the US) and the types of trading venues and platforms are similar to those in the US.

Many of the participants in the Australian equities market are major international banks and electronic trading firms, using similar trading technology as they use elsewhere. These include Goldman Sachs, Merrill Lynch, UBS, Bank of America, Citigroup, Deutsche Bank, J.P. Morgan, GETCO, Citadel, and others.

Volume in Australian equities is around $5.5 billion per day during Q1 2015, or around $25 million per stock per day for the 200 stocks in our sample. Effective bid-ask spreads during the same time period in our sample have a value-weighted average of 11 basis points (bps).[17]

High-frequency trading accounts for around 28% of volume in our sample of stocks during the first quarter of 2015 (ASIC, 2015), which is comparable to estimates of HFT trading activity in UK equities (27%; Aquilina and Ysusi, 2016), US E-mini S&P 500 futures market before the May 2010 flash crash (34%; Kirilenko et al., 2017), Canadian equities (33%, averaging the 20% estimated by Brogaard, Hendershott, and Riordan (2019) and the 46% estimated by Boehmer, Li, and Saar (2018)), slightly lower than US large-caps (42%; Brogaard et al., 2014) and slightly higher than US small-caps (18%; Brogaard et al., 2014). Slightly over 80% of Australian equity market capitalization is held by institutions (Bradrania et al., 2017), which is comparable to the estimated 80% in US large-caps.

[15] Given that trading occurs in Australian dollars (AUD), throughout this paper we use "$" to refer to AUD unless stated otherwise. At the start of our sample, AUD 1.00 is equal to approximately USD 0.93.
[16] See ASIC "Equity market data for quarter ending March 2015" available at http://asic.gov.au.
[17] See ASIC "Equity market data for quarter ending March 2015" available at http://asic.gov.au.

In summary, the Australian equities market is similar in most respects (market structure, trading platforms, fragmentation, level of HFT trading activity, institutional holdings, and market participants) to other major, developed equities markets.

2.2. Classification of trader types

Our study classifies market participants along two dimensions as illustrated in Figure 1. The first, directionality, is the tendency for a market participant to either buy or sell a given security in a given interval of time, but not both buy and sell. Directional market participant (the left two quadrants) are fundamental buyers and fundamental sellers—investors moving into or out of a position that is not quickly reversed. We refer to them as "investors". In contrast, non-directional participants (those that both buy and sell a given security within a relatively short period of time—the right two quadrants) comprise intermediaries such as market makers and traders. We refer to these as "traders".

The second dimension along which we partition market participants is dollar volume of their trading. Large directional market participants (top left quadrant) are the "institutional investors" that are used in our trading cost measurement. Small directional market participants (bottom left quadrant) include retail investors and small institutions. For this group, simple liquidity measures, such as the bid-ask spread, are likely to closely approximate trading costs.

Non-directional traders that account for substantial proportions of volume (top right quadrant) almost certainly are algorithmic traders, with the highest volume traders likely to be HFTs. We label this group of high-volume non-directional traders "active traders"; they are the focus of this study as the traders that could systematically increase or decrease institutional trading costs.

Our focus on high-volume traders is primarily because for a trader to have a meaningful impact on the trading costs of large institutional investors they are likely to trade in considerable volume. Our classification of "active traders" as the largest non-directional traders has similarities with data-driven definitions of HFT (e.g., Kirilenko et al., 2017; ASIC, 2015; Brogaard et al., 2019), but is not as narrow. An advantage of casting the net wider and including a mix of HFT and non-HFT AT in the set of active traders is that it allows us to ask whether HFTs are any more or less toxic than non-HFTs. In doing so, we do not formally partition our set of active traders into HFT and non-HFT subsets, but rather we examine whether trader-level characteristics typically associated with HFT (including speed, order-to-trade ratio, near-zero inventory, and sophistication) differ between toxic and non-toxic traders.

Non-directional traders that account for small shares of volume (bottom right quadrant) are likely to use strategies similar to those of active traders but on a smaller scale or with less sophistication/automation. This category is likely to involve some retail trading and some opportunistic strategies that only occasionally trade. Traders in this category are therefore unlikely to have a material effect on institutional trading costs.

< Figure 1 here >

For each market participant (trading account), we measure non-directional dollar volume as the dollar volume of buying that is accompanied by corresponding selling of a given security within a period of one week. This definition ensures that fundamental buying and fundamental selling, where positions are held for at least one week, does not contribute to non-directional volume.[18] After summing non-directional dollar volume for each trading account, we classify the highest non-directional volume traders as "active traders" (those that trade on average at least $8 million of non-directional volume per day).[19] We impose two additional (not particularly restrictive) requirements that ensure the active traders have sufficient breadth and continuity: (i) they are active in at least ten of the weeks in our 13 month sample, and (ii) they trade an average of at least 20 securities per day. This procedure results in 187 "active traders". Once an account is classified as an active trader that classification stays with the account throughout the sample.

For each stock-day it, we measure the fraction of that day's double-counted dollar volume (the dollar volume of the buy side of all trades plus that of the sell side) executed by active trader k as $Activity_{itk}$. We also define the binary variable $Presence_{itk} = 1$ if $Activity_{itk} > 0$ (i.e., trader k is present in stock-day it) and 0 otherwise.

Figure 2 shows the 187 active traders' percentage of dollar volume in stock quartiles through time (summing $Activity_{itk}$ across all 187 active traders and expressing it as a percentage). The participation rates of active traders do not vary substantially across the stock quartiles, ranging from an average of around 46% in the lowest dollar volume quartile to around 52% in the second highest quartile (these estimates are reported in Table 1). The pattern across quartiles is not monotonic; the highest proportional activity of active traders is in the second highest dollar volume quartile. There is a slight upward trend through time; active traders

[18] We confirm that our definition of active traders sufficiently separates them from fundamental institutional investors by checking that the trading of active traders does not result in unidirectional parent orders that are used in estimation of institutional trading costs.

[19] This cutoff (and the subsequent filters) is based on regulatory judgement about the types of traders that could have a material effect (good or bad) on institutional trading costs.

account for around 43% of dollar volume at the start of the sample and around 53% at the end, roughly one year later.

Table 1 indicates that the 187 active traders (AT/HFT) collectively account for around 48% of trading (dollar volume), which is close to twice the share of "pure" HFTs during the sample.[20] On an average stock-day, 64 of the active traders are present and trading the given stock.

< Figure 2 here >

< Table 1 here >

2.3. Institutional trading costs

Measuring institutional trading costs from the regulatory trader-level data involves three steps. First, we reconstruct institutional unidirectional "parent orders" as follows:

a) aggregate all trades for each account within a given stock-day to obtain a parent order;
b) classify parent orders as *unidirectional* if all trades are in one direction (all buying or all selling); and
c) classify unidirectional parent orders as *institutional* if their size (dollar volume) exceeds the median size of all unidirectional parent orders traded that stock-day and the parent order is "worked" in the market for at least two hours (there are at least two hours between the first and last child order in the parent order).

Next we measure the execution costs of the institutional unidirectional parent orders using implementation shortfall (Perold, 1988; Anand et al., 2012). Put simply, the implementation shortfall (or total cost) of executing a parent order is the average price at which the order is executed compared to the market price at the time the parent order execution starts (before the order impacts prices). For parent order j in stock i on day t, implementation shortfall is calculated as:

$$IShorfall_{itj} = \left[\frac{VWAP_{itj} - P0_{itj}}{P0_{itj}}\right] D_{itj} \quad (1)$$

where $VWAP_{itj}$ is the volume-weighted average execution price for the parent order, $P0_{itj}$ is the price at the time of the first trade in the parent order, and D_{itj} is the direction of the parent order (+1 for buys and –1 for sells).

[20] ASIC (2015) estimate that HFTs (using a data-driven identification procedure involving total dollar volume, inventory, order-to-trade ratio, number of fast messages, holding time, and sophistication) account for around 28% of dollar volume in our sample of stocks during Q1 2015.

Finally, we calculate the volume-weighted average implementation shortfall for all unidirectional institutional parent orders in each stock-day. The resulting measure, $IShorfall_{it}$, is measured in bps.

Figure 3 Panel A shows the large institutional parent orders as a percentage of dollar volume through time and Table 1 provides descriptive statistics on the institutional orders. Overall, the large institutional parent orders account for approximately 19.3% of traded dollar volume. Their share of volume is not monotonic across stock quartiles. Figure 3 Panel B plots the simple average of $IShorfall_{it}$ in stock dollar volume quartiles through time. $IShorfall_{it}$ decreases monotonically with the dollar volume of the stocks (lower trading costs in larger or more traded stocks), averaging around ten bps in the highest volume quartile and around 24 bps in the lowest volume quartile. Table 1 indicates that the pooled mean of $IShorfall_{it}$ is around 16 bps, which is very similar to the 17 bps average implementation shortfall for institutional orders in the US, as reported by Anand et al. (2012) for 2007 (the most recent non-crisis year of their sample).

< Figure 3 here >

3. Analysis of toxicity

3.1. The basic approach to measuring toxicity

At the core of our approach to measuring toxicity is the following regression (which we refine in various ways later including using instrumental variables and bootstrap simulations):

$$IShorfall_{it} = \alpha + \sum_{k=1}^{187} \gamma_k Activity_{itk} + \varepsilon_{it} \quad (2)$$

Recall $IShorfall_{it}$ is the implementation shortfall (total execution cost) of large institutional parent orders and $Activity_{itk}$ is the activity (fraction of that stock-day's dollar volume) of active trader k. The estimated coefficients $\widehat{\gamma_k}$ measure the "toxicity" of each active trader k. Positive (negative) values of γ_k indicate that the trader's activity is associated with an increase (decrease) in institutional trading costs. Traders that increase (decrease) institutional trading costs significantly are termed "toxic" ("beneficial"). We double cluster standard errors by stock and by date in all regressions to account for dependencies within the panel data.

To translate the toxicity measures into basis point impacts on institutional trading costs, we define gross toxicity of an active trader as their toxicity per unit of activity ($\widehat{\gamma_k}$) multiplied by their average share of dollar volume ($\overline{Activity_k}$), i.e., $GrossToxicity_k = \widehat{\gamma_k}\overline{Activity_k}$. A trader's gross toxicity is their average basis point impact on institutional trading costs.

Visual representation of the toxicity estimates is telling. Figure 4 plots each of the 187 traders in a two dimensional space. The vertical axis measures the trader's toxicity estimate, $\widehat{\gamma_k}$,

(on a log scale) with positive (negative) numbers indicating the trader is associated with increased (decreased) institutional trading costs.[21] The horizontal axis measures the consistency of the trader's toxicity (log of the standard error of $\widehat{\gamma_k}$) with lower values indicating greater consistency in how the trader impacts institutional trading costs. The size of the circles indicates the statistical significance of the toxicity estimate, with the smallest circles being toxicity estimates that are not statistically significant at the 10% level, through to the largest circles indicating statistical significance at the 1% level.

The figure shows quite strikingly that the 187 active traders (the AT/HFT accounts) cluster fairly neatly into two distinct groups: (i) those traders whose trading activity is associated with systematically increased institutional trading costs ("toxic" traders) and (ii) those associated with systematically decreased institutional costs ("beneficial" traders). Very few traders have toxicity estimates around zero. Many traders' toxicity coefficients are statistically different from zero. Some degree of statistical significance is expected by chance when testing so many coefficients. We address this issue later with bootstrap simulations and show the number of statistically significant coefficients is considerably greater than what would be expected by chance alone.

Institutional investors are likely to be concerned not only about an active trader's average toxicity level, but also the consistency with which the trader imposes toxicity, which we measure using the standard error of the toxicity estimate. Therefore, the collection of active traders that have maximum toxicity for a given level of variation in their impact on institutional trading costs or a minimum level of variation in their impact for a given level of toxicity form, what we term, the "toxicity frontier". Figure 4 illustrates the approximate toxicity frontier for our sample.

< Figure 4 here >

We test variations to the functional form of the regression above before proceeding to instrumental variables models and bootstrap simulations. First, we add stock fixed effects, then date fixed effects, then both sets of fixed effects. The latter of these models is:

$$IShorfall_{it} = \alpha_i + \mu_t + \sum_{k=1}^{187} \gamma_k Activity_{itk} + \varepsilon_{it} \quad (3)$$

The fixed effects eliminate many potential confounding effects and subsume most potential control variables. The fixed effects subsume stock characteristics that have mainly cross-sectional variation (e.g., market capitalization, institutional holdings, spread/depth, whether the

[21] Because the toxicity estimates can take positive and negative values, the conversion to a log scale is done in a way that preserves the sign, i.e., $sign(\widehat{\gamma_k})log(|\widehat{\gamma_k}| + 1)$.

stock's spread is constrained by its tick size causing limit order queuing, percentage of volume traded in the dark, and so on). They also subsume time-series variables including market conditions (e.g., VIX, market-wide realized volatility/volume/returns, and so on). Fixed effects in our setting also come at a cost—they absorb much of the time-series and cross-sectional variation that is useful in identifying traders' toxicity.

Figure 5 illustrates the results from the toxicity regressions with fixed effects. The distributions of toxicity are qualitatively similar with and without fixed effects, pointing away from the presence of significant confounding factors or omitted variables bias. The levels of statistical significance in the toxicity estimates are also similar with fixed effects.

Although the distributions and significance of toxicity estimates are similar with and without fixed effects, it is possible that the ordering of traders with respect to their toxicity is different, e.g., a trader estimated to be toxic could swap places with a trader estimated to be beneficial in two different models without changing the distributions of toxicity. To investigate this possibility we compute Pearson parametric correlations and Spearman non-parametric rank correlations for the trader toxicity estimates across the different models (with/without fixed effects). The correlations are all very high (above 0.8) indicating that even the relative ranking of traders on the toxicity scale is not overly affected by the inclusion of fixed effects. Given these results, we proceed with the simpler model with no fixed effects as it places less burden on parameter estimation (important in the bootstrap simulations) and allows us to exploit more dimensions of variance in further tests and instrumental variables models.

< Figure 5 here >

We report two further variations on the basic toxicity regression (2) in the Appendix. The first is a model that considers whether an active trader's presence in the market (trading non-zero volume, $Presence_{itk}$) affects institutional trading costs, or whether their level of activity (share of dollar volume, $Activity_{itk}$) affects institutional trading costs, or both have independent effects. We simultaneously include $Activity_{itk}$ and $Presence_{itk}$ in a regression similar to (2) and find that while the mere presence of some active traders has an effect on institutional trading costs independent of how active they are, presence does not subsume activity in explaining the effects of active traders on trading costs. The statistical significance of the toxicity estimates using $Activity_{itk}$ remain approximately the same when $Presence_{itk}$ is included in the toxicity regressions. The relative ranking of traders on a toxicity scale is also not overly affected, with correlations around 0.8 between the rankings generated by the two approaches.

Second, we transform the bounded activity measures, $Activity_{itk}$, into unbounded variables via a logit transformation and re-estimate the toxicity regression. We find that the transformation compresses the difference between toxic and beneficial traders, but it does not overly affect inference about toxicity levels, statistical significance of toxicity, or the relative ranking of traders on a toxicity scale (correlations around 0.9 between the rankings).

3.2. Instrumental variables models

To deal with the potential endogeneity of active traders' activity and pin down their causal effects on institutional trading costs, we turn to instrumental variables. Given the heterogeneity in our sample of active traders, which is a deliberate result of casting a wider net across traders than most studies (e.g., HFT studies) it is unlikely that exogenous market-wide changes will be useful in identifying the effects of our cross-section of active traders. This is because any market structure change, such as a platform upgrades for latency reduction or co-location, will affect some traders but not others (e.g., latency reductions might increase HFT activity) or affect different traders in different ways. Instead, we use exogenous information on the individual traders as instruments, more precisely, their lagged trading activity in a given stock. This approach is similar in spirit to Sarkar and Schwartz (2009) who also use lags of endogenous variables as instruments in a microstructure setting.

We estimate the following two-stage least squares instrumental variables (2SLS IV) model:

$$Activity_{itk} = \mu + \sum_{\tau=1}^{5} \beta_\tau Activity_{i,t-\tau,k} + \epsilon_{itk} \quad (4)$$

$$IShorfall_{it} = \alpha + \sum_{k=1}^{187} \gamma_k \widehat{Activity}_{itk} + \varepsilon_{it} \quad (5)$$

In the first stage (4), the trading activity of each active trader k in stock-day it is regressed on that active trader's activity in the same stock in each of the past five trading days. In the second stage (5), we estimate a similar toxicity regression as previously, but replacing the activity variables with their fitted values from the first stage, $\widehat{Activity}_{itk}$.

The first requirement of instrumental variables—that they are correlated with the endogenous explanatory variables—is clearly satisfied. The correlations between lagged and current activity of a given trader in a given stock are between 0.54 and 0.62 (depending on the lag) and F-tests of whether the instruments in the first stage are statistically significant produce p-values well below 1%. The second requirement—that the instrumental variables are exogenous with respect to the dependent variable (not correlated with the error term)—is satisfied as a result of the temporal difference between the instrumental variable measurement and the dependent variable measurement. Past activity cannot respond to current market conditions or current levels

of institutional trading costs. The only possible contamination is through persistence in the conditions that affect institutional trading costs. We rule out this possibility with three additional tests: (i) adding lagged $IShorfall_{it}$ as a control variable in the first and second stages to absorb persistence in institutional trading costs, (ii) adding time fixed effects to absorb time-series trends or persistent changes in market conditions, and (iii) omitting the first and second lags in the first-stage regression and using only the third, fourth, and fifth lags as instruments. Our results are largely unchanged in these three additional tests.

Figure 6 illustrates the results from the 2SLS IV toxicity regressions (5). The distributions of toxicity, as previously, show strong clustering of traders into a group of traders that systematically tend to increase institutional trading costs and a second group that tend to decrease these costs. In general, fewer traders have a statistically significant relation with institutional trading costs than in the OLS models, which could be the result of eliminating endogeneity in the estimates, but it could also be due to the decrease in statistical power. The Pearson parametric correlation (Spearman non-parametric rank correlation) of trader gross toxicity estimated via OLS and the 2SLS IV models is 0.70 (0.48) indicating that using instrumental variables somewhat changes inference about individual traders and their relative ranking on a toxicity scale. Given these results, we proceed with the 2SLS IV models in the further analysis of toxicity.

< Figure 6 here >

3.3. Quantifying toxicity beyond statistical chance

One of the challenges in analyzing toxicity at the level of individual traders is that if sufficiently many traders are analyzed some will, purely by statistical chance, appear toxic even if they have no true relation with institutional trading costs. Put differently, with a non-zero Type 1 error rate, the number of Type 1 errors increases with the number of tests performed (number of traders tested for toxicity).

The challenge of disentangling true and spurious toxicity at the individual trader level is similar to the problem of disentangling skill from luck in the cross-section of fund managers. Given a sufficiently large number of fund managers, some will beat their benchmark over many consecutive periods purely by chance and thus appear skilled. In neither context (measuring skill or measuring toxicity) is it sufficient to adjust critical values because the distributions of alpha and toxicity are complex and not necessarily Normal. We therefore borrow from the fund

management literature and use bootstrap simulations to quantify excess toxicity i.e., the amount that exceeds what would be expected by statistical chance.

For the bootstraps, we follow Kosowski, Timmermann, Wermers, and White (2006). The details are in Appendix B. The essence of the procedure is as follows. Estimate the 2SLS IV model second stage and save the residuals (the procedure works the same with the OLS model, but the 2SLS IV models have the advantage of addressing endogeneity). Simulate data on implementation shortfall by sampling from the residuals (with replacement) for each stock. In the simulated data, by construction, none of the active traders have an underlying relation with institutional trading costs (the "zero-toxicity null") other than what might emerge spuriously. Estimate the 2SLS IV model on the simulated data saving the toxicity estimates. Repeat the simulate-then-estimate steps 1,000 times to build distributions of the levels of toxicity and their statistical significance that would be expected purely by chance. Finally, compare the actual estimated toxicity levels and significance to those that would be expected by chance. Throughout the bootstrap, we use double clustered standard errors (clustered by stock and by date).

< Table 2 here >

Table 2 reports the bootstrap results. Panel A reports estimates from the 2SLS IV model using the *actual* data, i.e., the model estimated previously. The columns *min* through to *max* (with P5, P25, P75, P95 being the 5^{th}, 25^{th}, 75^{th}, 95^{th} percentiles) describe the distribution of the 187 toxicity estimate t-statistics (one estimate for each of the 187 active traders). The median t-statistic is close to zero (–0.08). The 5% most toxic (beneficial) traders have toxicity t-statistics above 2.12 (below –2.37). The minimum and maximum toxicity t-statistics across the 187 active traders are –3.61 and +4.00, respectively. The last four columns of the table indicate how many of the 187 active traders have toxicity t-statistics beyond the threshold given in the column heading. In the actual data (Panel A), four traders are toxic with t-statistics above +3 and 12 traders with t-statistics above +2. Similarly, four traders are "beneficial" with t-statistics less than –3 and 15 traders with t-statistics less than –2.

Panel B of Table 2 reports a sample of the results using simulated data: simulation iterations 1–6 and 1,000, purely as illustrations of the simulation process. Recall that each iteration of the bootstrap procedure creates a panel dataset similar to the actual dataset (but with zero toxicity by design) and on that dataset estimates the 2SLS IV toxicity model, saving the 187 toxicity estimates. The results from an individual iteration (an individual row of Panel B) by themselves are not particularly useful (we report them simply to illustrate the bootstrap

procedure); what is useful is the distribution created by the 1,000 iterations. Each iteration forms a point in the "bootstrap distribution" of each statistic (column). The last row of Panel B reports means (across the 1,000 iterations) of the number of active traders with toxicity t-statistics beyond a certain threshold.

Having built a bootstrap distribution for the toxicity t-statistics under the "zero-toxicity null" (Panel B), we can test whether the toxicity t-statistics estimated on the actual data (Panel A) deviate from what is expected by chance. Panel C expresses the toxicity t-statistic estimates using actual data (Panel A) in terms of percentiles in the bootstrap distribution. For example, that value of 52 in the column *min* in Panel C indicates that the minimum toxicity t-statistic of in the actual data (–3.61) falls in the 52nd percentile of the minimum t-statistics in the 1,000 simulated datasets. Thus, the most negative t-statistic (most statistically beneficial trader) in the actual data is no more extreme than would be expected by chance. Put differently, in 52% of the simulated datasets (in which there is zero true toxicity by design), the most negative t-statistic across the 187 traders is more negative (larger absolute value) than the most negative t-statistic in the actual data. Similarly, the largest positive t-statistic in the actual data is no more extreme than would be expected by chance (it is in the 50th percentile of the bootstrap distribution). This suggests the actual data do not contain extreme individual outliers. Also, the median active trader in the actual data is no more or less toxic than what would be expected under the zero-toxicity null.

The 5th and 25th percentiles of the toxicity t-statistics in the actual data (the 5% and 25% most beneficial traders), however, *are* more beneficial than would be expected under the zero-toxicity null. The toxicity t-statistics for those groups of traders are in the 0th percentile (the 0–1% segment of the bootstrap distribution). Similarly, the 5% most toxic traders in the actual data are also more toxic than would be expected by chance; their toxicity t-statistics are in the 99th percentile (the 99–100% segment of the bootstrap distribution). These results indicate that it is very unlikely (less than 1% probability) that one would find the levels of toxicity estimated for the 5% most beneficial and 5% most toxic traders by chance. Thus the bootstrap results reject the zero-toxicity null hypothesis at a 99% confidence level.

The last four columns of Table 2 provide an additional way of quantifying toxicity beyond statistical chance. The *Mean* row in Panel B indicates that under the zero-toxicity null, we would expect by chance 5.55 of the 187 active traders to appear to be statistically toxic at 95% confidence (toxicity t-statistic > +2) and 1.20 to be statistically toxic at 99.5% confidence (toxicity t-statistic > +3). In the actual data (Panel A) we observe that in fact 12 (rather than 5.55) of the active traders are statistically toxic at 95% confidence, and four (rather than 1.20) are statistically toxic at 99.5% confidence, i.e., considerably more than would be expected under the

zero-toxicity null. In fact, Panel C indicates that the probability of observing that many statistically toxic traders at the 95% and 99.5% confidence levels is less than 1% and less than 2%, respectively. Similarly for the beneficial traders, Panel B indicates that we would expect by chance 6.87 of the 187 active traders to appear to be statistically beneficial at 95% confidence (t-statistic < –2) and 1.24 to be statistically beneficial at 99.5% confidence (t-statistic < –3). In the actual data (Panel A) we observe that in fact 15 of the active traders are statistically beneficial at 95% confidence, and four are statistically beneficial at 99.5% confidence, with the probabilities of observing that many beneficial traders by chance being less than 1% and less than 2%, respectively (Panel C).

Repeating the bootstrap analysis for $GrossToxicity_k$ rather than the t-statistics of the toxicity estimates leads to similar conclusions, namely, that the level of toxicity in the 5% most toxic and most beneficial traders is beyond what would be expected by chance.

3.4. Net toxicity

The bootstrap analysis, building on the 2SLS IV model, indicates that there are groups of traders in the data that have a causal effect on institutional trading costs (both positive and negative) beyond what would be expect by chance. We now turn to quantifying their impacts on institutional trading costs.

In terms of economic significance, what matters most is the ultimate impact on institutional trading costs. This is best measured by $GrossToxicity_k$, which is trader k's basis point impact on the implementation shortfall of large institutional orders. Accounting for what would be expected by chance is done by subtracting, for each trader, the expected $GrossToxicity_k$ for that trader, estimated from the bootstrap distribution under the zero-toxicity null ($\mathrm{E}[Toxicity_k]$). This process effectively takes the integral of the difference between actual and expected $GrossToxicity_k$ across segments of the $GrossToxicity_k$ distribution (or across the whole distribution). Continuing the analogy of skill in the cross-section of fund managers, calculating the excess gross toxicity for a group of traders is like quantifying the net alpha generated by a group of fund managers in excess of the alpha that is the result of chance (luck).

Aggregating excess $GrossToxicity_k$ across groups of traders involves summing their $GrossToxicity_k$ estimates. For the active traders that have statistically significant toxicity estimates ($\widehat{\gamma_k}$), positive or negative, their impact on institutional trading costs (accounting for what is expected by change) is:

$$ExcessGrossToxicity_{ToxicTraders} = \sum_k (GrossToxicity_k - \mathrm{E}[Toxicity_k])\mathbf{1}_{\{t_k>2\}} \quad (6)$$

$$ExcessGrossToxicity_{BenficialTraders} = \sum_k (GrossToxicity_k - \mathrm{E}[Toxicity_k])\mathbf{1}_{\{t_k<-2\}} \quad (7)$$

where $\mathbf{1}_{\{t_k>2\}}$ and $\mathbf{1}_{\{t_k<-2\}}$ are indicator functions for whether trader k is significantly toxic ($\widehat{\gamma_k}$ t-statistic > +2) or significantly beneficial ($\widehat{\gamma_k}$ t-statistic < –2), respectively.

We find that the 12 significantly toxic traders increase the average implementation shortfall for large institutional orders by 10.3 bps beyond what is expected by chance ($ExcessGrossToxicity_{ToxicTraders}$ = 10.3). That is an economically meaningful magnitude given the pooled sample unconditional mean implementation shortfall of 16.4 bps (Table 1), and the value-weighted average effective bid-ask spread is around 11 bps. Another way of thinking about this impact is that without the toxic traders, institutional trading costs would be approximately 10.3 bps lower, or around 6.1 bps. Therefore toxic traders more than double the costs of executing large parent orders. Note that these cost impacts are per parent order. To obtain the impact on a fund's returns (which could be larger) one has to scale up by the fund's annual turnover.

In dollar terms, large institutional orders account for around 19% of dollar volume, and traded dollar volume is approximately $25 million per stock per day during our sample.[22] All up, that implies that a 10.3 bps increase in $IShorfall_{it}$ as a result of the toxic traders equates to increased trading costs of around $437 million across all large institutional orders in the top 200 stocks during a one-year period (assuming 220 trading days).

At the same time, the 15 significantly beneficial active traders decrease the average implementation shortfall for large institutional orders by 8.9 bps beyond what is expected by chance ($ExcessGrossToxicity_{BeneficialTraders}$ = –8.9). This effect is also economically meaningful compared to average implementation shortfall and is similar in magnitude to the increase in implementation shortfall caused by the toxic traders. In the absence of the 15 significantly beneficial active traders, institutional trading costs would be around 25.3 bps (16.4 + 8.9). Therefore the beneficial traders reduce institutional trading costs by slightly more than one-third. In dollar terms, this translates to a reduction in institutional trading costs of around $375 million across all large institutional orders in the top 200 stocks during a one-year period.

These dollar estimates for the impact of toxic and beneficial traders represent a lower bound for two reasons. First, they only capture the largest institutional trades, not medium trades, which are also likely to be affected to some extent. And second, they are computed for the top 200 stocks, ignoring other stocks.

Netting the effects of the significantly toxic and significantly beneficial active traders, implies a small net effect: around +1.4 bps or an increase in annual costs of around $62 million

[22] With large institutional orders accounting for 19% of dollar volume, their dollar volume is $2 \times 19\% \times \$25\text{mil}$ to account for both the buying and selling sides of trades.

across all large institutional orders. Thus, while toxic and beneficial traders individually have large effects on institutional trading costs, their net effect is close to zero.

We also measure the net effects across all active traders, irrespective of whether they have statistically significant toxicity estimates or not, again accounting for what is expected by chance. The absence of statistical significance does not rule the possibility that a trader has a true impact on institutional trading costs, so it is worth assessing the impact of all active traders. Incorrectly including traders in the aggregation that have no effect on institutional trading costs will not bias the net toxicity estimate, it merely adds noise. Net excess toxicity across all active traders is estimated as:

$$NetExcessToxicity = \sum_k (GrossToxicity_k - \mathrm{E}[Toxicity_k]) \quad (8)$$

We find that net excess toxicity across all active traders is near zero: –0.79 bps. The sign of the point estimate implies a net benefit from the active traders as a whole, reducing costs of executing large institutional orders by around $33 million per annum. However, this estimate is not statistically distinguishable from zero at the 5% level using parametric or non-parametric tests.

The results above help understand the conflicting views between buy-side institutions and existing studies about the effects of AT/HFT, in particular, why there are significant concerns about AT/HFT increasing institutional trading costs, despite evidence that AT/HFT as a whole are beneficial or benign. Toxic traders have an economically meaningful detrimental effect on institutional trading costs. An institutional investor that disproportionately interacts with toxic traders will face higher trading costs. An institutional investor might disproportionately interact with toxic traders as a result of their investment style (their entry/exit signals correlating with toxic trader activity), as a result of their size, or as a result of carelessly managed execution that allows predatory algorithms to detect the institutional investor's trading intentions and anticipate the remaining child orders. The toxicity magnitudes that we report above imply that carelessly managed execution can have a material effect on a fund's performance. Importantly, the effects of toxic traders are balanced by a segment of highly beneficial AT/HFT that considerably decrease market impact costs. These findings explain why from the perspective of some buy-side institutions (those that disproportionately interact with toxic traders), AT/HFT appear to do more harm than good, despite the evidence that *in aggregate* AT/HFT are benign or beneficial.

Our findings about the heterogeneity among AT/HFT also have implications for how institutional investors can decrease trading costs. The results suggest that considerable trading cost savings could be achieved by avoiding toxic traders or at least minimizing interactions with them. Doing so might involve monitoring toxicity levels and then adjusting the execution strategy for large parent orders in response to toxicity levels. For example, when toxicity is high

an institutional investor might spend more time/effort searching for an off-market block counterparty, might trade the parent order at a slower rate to leave less of a signal to predatory algorithms, increase the proportion of child orders routed to dark pools, and possibly even pause the parent order execution and restart it at a later time. All these actions involve a cost and are therefore only desirable if the toxicity level is sufficiently high. Furthermore, the trading cost savings need to be balanced against the costs of more sophisticated execution strategies and technology, as well as opportunity costs if execution is delayed.

The heterogeneity in the cross-section of AT/HFT also provides a potential reason why studies of AT/HFT that use different exogenous events as instruments arrive at different conclusions about the net effects of this group of traders. An event that disproportionately encourages the trading activity of the beneficial traders or gives them an advantage will tend to improve trading cost dimensions of market quality. The opposite is true for events that benefit or encourage the group of toxic traders. Brogaard et al. (2015) for example, find that the introduction of co-location at Nasdaq OMX Stockholm improves market quality, implying that this event disproportionately encouraged the activity of beneficial AT/HFT.

3.5. Characteristics of toxic traders

Having estimated the toxicity of each active trader, we now ask what characteristics distinguish toxic traders from beneficial ones? Toxic traders are likely to use different trading strategies compared to beneficial traders. Given that one of the main determinants of implementation shortfall for large orders is whether others are trading in the same direction (e.g., ASIC, 2015) toxic traders are likely to trade *with* institutional order flow (in the same direction) rather than *against* it. Trading with institutional order flow could result from a trader intentionally exploiting institutional order flow, for example, predatory trading strategies (e.g., Brunnermeier and Pedersen, 2005), order anticipation algorithms (e.g., Hirschey, 2020), and strategies that seek to identify and "back-run" large informed orders (e.g., Yang and Zhu, 2020). However, trading with institutional order flow could also result from a trader sharing common entry/exit signals or through active participation in the price discovery process during periods of imbalance. Therefore not all toxic traders necessarily intentionally exploit institutional investors. Furthermore, not all toxic traders are necessarily harmful to the market overall—some might contribute to price discovery around institutional parent orders. In contrast, the beneficial traders are likely to be liquidity-providing intermediaries (informal market makers) that "lean against the wind" and thereby attenuate price pressure from large institutional orders.

The differences in trading strategies employed by toxic and beneficial traders imply that they should differ in some measureable trading characteristics and patterns in activity. We analyze the distinguishing features of toxic traders by estimating regressions of trader-level toxicity on trader characteristics:

$$\widehat{Toxicity}_k = \alpha + \sum_c \beta_c Characteristic_{c,k} + \epsilon_k \quad (9)$$

where the trader-level toxicity measures ($\widehat{Toxicity}_k$) are either toxicity per unit activity ($\widehat{\gamma_k}$), gross toxicity ($\widehat{GrossToxicity}_k$), excess gross toxicity ($GrossToxicity_k - \mathrm{E}[Toxicity_k]$), or the statistical significance of the trader's toxicity estimate (t-statistic of $\widehat{\gamma_k}$), all obtained from the 2SLS IV toxicity models. We first explore trader characteristics such as level of activity, speed, sophistication, and order placement, followed by patterns in trading activity (when and where toxic traders trade).

< Table 3 here >

Table 3 reports the results of the regression in (9) using excess gross toxicity as the dependent variable (the results are similar using the other toxicity estimates). The first characteristic we test is the trader's activity, measured as the trader's average share of dollar volume. Higher volume traders are likely to be more sophisticated, but importantly also faster—several papers show, both theoretically and empirically, that relative speed has a large impact on a trader's share of volume (e.g., Roşu, 2016). Model 1 in Table 3 shows that a trader's share of volume is not significantly related to their toxicity. To the extent that high volumes are related to speed, this finding suggests that speed is not a distinguishing characteristic of toxic traders.

We further explore the relation between speed and toxicity using other measures of speed, sophistication, and order placement characteristics, many of which have been used as defining features of HFT. The first of these is the average holding time of long or short positions that are closed within a day.[23] Model 2 in Table 3 shows that *Holding Time* is also not significantly related to toxicity and if anything, the point estimate suggests toxic traders have *longer* holding times. To the extent that HFTs have short holding times, this evidence also suggests that HFTs are no more toxic than non-HFTs.

We define fast orders as order amendments sent within 500 milliseconds of the order placement and use fast orders to construct two related measures: the number of fast orders (*Number Fast*), and the average speed of fast orders (*Order Amend Time*). *Number Fast* indicates

[23] For a given trader in a given stock-day we construct first-in-first-out (FIFO) "pipes" for long and short positions and measure the average time a position spends in the pipe. Positions held overnight do not contribute to this measure.

how active a trader is in "managing" submitted orders. *Order Amend Time* is a proxy for the speed of the trader's technology. Model 3 shows that neither of these characteristics is significantly related to toxicity, consistent with the conclusion that fast traders are not more toxic on average than slower traders. The point estimate on the number of fast orders is positive in Model 3 but becomes negative when we control for the total number of orders (Model 8).

We measure each trader's efficiency in generating intraday trading profits as the Sharpe ratio of daily trading profit (counting only those positions that are closed within the day). The Sharpe ratio of trading profit is a measure of sophistication (sophisticated traders consistently make money from their trading) and is another of the characteristics that distinguishes HFT. We find that toxic traders are not more efficient in generating intraday trading profits than others (statistically insignificant coefficient in Model 4), in fact, the insignificant point estimate goes the other way.

Models 5 and 6 show that a trader's order-to-trade ratio, which measures the degree of strategic order submission and active order management (and is also a characteristic of HFT), is not significantly related to toxicity. Model 5 tests the relation between a trader's toxicity and their average number of submitted orders, holding their dollar volume fixed, and Model 6 tests the order-to-trade ratio directly. Although not statistically significant, the point estimates are consistent with the notion that toxicity is positively related to excessive quoting. The average order-to-trade ratio across all active traders is around 22. Consistent with our previous results, Model 5 shows that a trader's dollar volume is not significantly related to their toxicity.

The broad collection of characteristics examined thus far consistently point to the conclusion that the HFTs in our sample are no more likely to be toxic than non-HFTs. On one hand, this conclusion seems reasonable in light of the evidence on the timing of predatory and back-running strategies. For example, van Kervel and Menkveld (2019) show that HFTs initially trade against institutional parent orders and only seem to be able to detect (and then exploit) them several hours after the start of the parent order. Speed is not necessary for predatory trading or back-running orders, but speed is important in market making and arbitrage. Consistent with this finding, Hagströmer and Nordèn (2013) find that market making HFTs have lower latency and higher order-to-trade ratios than other HFTs, who in turn have latency and order-to-trade ratios on par with non-HFTs.

At the same time, the evidence suggesting that HFTs are no more toxic than non-HFTs suggests that concerns voiced by institutional investors about HFTs being the culprits responsible for increased trading costs might be misdirected. While some HFTs may be predatory traders, it seems HFTs are no more predatory than other slower traders.

Model 7 in Table 3 shows that toxic traders on average close a larger percentage of their positions by the end of the day, i.e., they hold less inventory overnight. Across all 187 active traders, around 22% of their dollar volume on average is in the form of intraday round-trip trades, with a standard deviation of around 27%. The effect size (coefficient of *Percentage Traded*) is economically meaningful and larger in magnitude once we control for other characteristics (Model 8).[24] If electronic market makers tend to "go home flat" (some of the earlier literature on electronic market making suggests they do) then it might seem surprising that *Percentage Traded* is positively related to toxicity. However, several recent studies that track inventory of electronic market makers (e.g., Malinova and Park, 2016) show that they do not go home flat in individual securities and in fact hold large inventory positions overnight. Rather than going home flat in individual securities, electronic market makers often try to go home flat in terms of risk (either netting beta-adjusted positions across stocks, beta hedging their net inventory, or multi-factor hedging their inventory). In contrast, order anticipation strategies might hold little to no inventory overnight and therefore have a high percentage of their positions closed by the end of each day.

Model 8 includes all the speed, sophistication, and order placement characteristics together. Our main conclusions hold when the characteristics are tested together.

< Table 4 here >

In Table 4 we examine differences in the trading patterns of toxic and beneficial traders. We start with the consistency of a trader's activity. For each trader, we separately compute the cross-sectional standard deviation of their activity (their activity is measured by their share of total traded dollar volume) and the time-series standard deviation of their activity. Model 1 in Table 4 shows that toxic traders tend to have higher cross-sectional standard deviation and lower time-series standard deviation than other traders (although the latter is only marginally statistically significant). The high cross-sectional standard deviation of active traders' activity suggests that, on any given day, toxic traders tend to concentrate their activity in a subset of stocks rather than trading equally across the market portfolio. The low time-series standard deviation suggests toxic traders tend to be more consistently trading their preferred stocks day-to-day. Thus, toxic traders appear to have a preferred habitat of stocks.

[24] The coefficient estimate in Model 4 implies that a one-standard deviation (27%) increase in the percentage of volume made up by intraday round-trip trades for the 12 significantly toxic traders would increase their detrimental effect on institutional implementation shortfall by around one bp in total.

What is the preferred habitat of toxic traders? We measure each trader's share of dollar volume in large (top quartile) medium (second quartile) and small (bottom two quartiles) stocks, as well as expressing their activity in small stocks relative to their overall activity. Models 2 and 3 in Table 4 show that toxic traders are more active in small stocks and less active in large stocks.

We also examine whether toxic traders are more active in times of market stress by measuring each trader's relative activity during a period in which the market fell sharply and exhibited high volatility (August 5–24, 2015). Models 4 and 5 in Table 4 indicate that there is no statistical difference in the activity of toxic traders (relative to other traders) during the period of market stress compared to other times. This is consistent with their relatively low time-series standard deviation. It is worth noting that institutional orders as a share of volume also slightly decline during this period (Figure 2).

Model 6 of Table 4 combines all of the activity characteristics in one regression. The characteristics that are individually significant remain significant after controlling for other characteristics. The R-squared reaches a maximum of 7% in Model 6, suggesting that there are many factors (which could be predominantly unobservable) beyond those included in our regression models that explain variation in toxicity. Additionally, measurement error in the toxicity estimates is also likely to contribute to the low R-squared.

In summary, the characteristics associated with HFT are unrelated to toxicity, including volume, speed of order amendments, frequency of fast orders, consistency with which they extract intraday trading profit, and order-to-trade ratios. The evidence suggests that HFTs are not more toxic than non-HFTs. Toxic traders do differ from non-toxic traders in a few regards—they tend to concentrate their activity in a subset of stocks, which they trade fairly consistently, they are more active in smaller stocks, and are more likely to close positions before the end of the day and hold less inventory overnight.

3.6. Robustness tests

We find qualitatively similar results across a number of robustness tests, including: (i) subperiod tests, omitting a period of one month in which the market fell sharply (August 2015); (ii) stock and/or time fixed effects in the toxicity regressions; (iii) logit transformations of the toxic trader activity measures to give unbounded variables; (iii) controlling for the presence of active traders in the toxicity regressions (in addition to their activity); (iv) using different toxicity estimates (toxicity per unit activity, gross toxicity of each trader, and the t-statistic of the trader's toxicity estimate) in the bootstrap analysis and analysis of trader characteristics. The results of

some of these robustness tests are mentioned above in the corresponding section of the paper, others are omitted for conciseness.

4. Conclusions

Behind the veil of their aggregate effects lies a rich cross-sectional heterogeneity in the effects of individual of algorithmic and high-frequency traders (AT/HFT). We find strong evidence that some of these traders systematically increase institutional trading costs while others decrease these costs, i.e., some of these traders appear "toxic" to an institutional investor while others appear "beneficial". The effects of these two groups of traders on institutional trading costs are economically meaningful. Toxic traders increase institutional trading costs by more than ten bps on large parent orders, roughly doubling these costs. In dollar terms this translates to additional trading costs of around $437 million per annum for large institutional orders in the top 200 stocks. The effects of the toxic traders are offset by a group of beneficial traders that significantly decrease those costs. Consequently, the small aggregate effects of AT/HFT on institutional trading costs mask considerable heterogeneity and large effects within subsets of the AT/HFT population.

An implication of our findings, in particular the pockets of toxicity within the AT/HFT population, is that the trading technology used by institutional investors is likely to have material impact on their trading costs. Institutions that disproportionately trade against toxic traders will experience higher trading costs. These findings help understand the concerns raised by institutional investors about AT/HFT increasing their trading costs. At the same time, our results reconcile these concerns with the evidence that AT/HFT as a group seem to be benign or beneficial—the negative effects of toxic traders are offset by significantly beneficial traders. The findings also suggest that by monitoring toxicity levels and adjusting trading strategies in response to the toxicity level, institutions that are able to limit their exposure to toxic traders could potential obtain substantial reductions in trading costs.

Who are the toxic traders? Toxic traders are likely to trade *with* institutional order flow rather than *against* it. While some toxic traders might intentionally exploit institutional order flow (e.g., predatory, order anticipation, or back-running strategies), others unintentionally amplify institutional trading costs by trading on common entry/exit signals, or through active participation in price discovery in the presence of flow imbalances. In contrast, the beneficial traders are likely to be liquidity-providing intermediaries that "lean against the wind" and thereby attenuate the price pressure that arises from large institutional orders.

Examining the characteristics of toxic traders, we consistently find that speed is not associated with toxicity. For example, a trader's share of dollar volume is not related to their toxicity (fast traders have high volume shares). Similarly, toxicity is not related to the speed of order amendments, the frequency of fast orders, the consistency with which a trader extracts intraday trading profit, and order-to-trade ratios. These results all suggest that HFTs are not more toxic to institutional investors than non-HFTs. On the basis of our evidence, concerns voiced by institutional investors about HFTs in particular being the culprits responsible for increased trading costs are misdirected. From a policy perspective, our results suggest that regulation or policies that disproportionately restrict fast traders or HFTs are unlikely to result in substantial improvements to institutional trading costs.

We find that the typical toxic trader tends to concentrate their activity in a subset of stocks and is fairly consistently active in that "preferred habitat". We find that the preferred habitat tends to be smaller stocks. The typical toxic trader also tends to close a larger proportion of their positions before the end of the day and subsequently holds less inventory overnight than the typical beneficial trader.

The focus of this paper has been on the diversity among AT/HFT. Future work might explore the diversity of institutions in contemporary markets. Our finding that the effects of toxic and beneficial traders are substantial in magnitude suggests that the extent to which an institution interacts with these two types of traders or the sophistication with which they execute large orders can have a considerable impact on their trading costs. This raises questions such as do some institutions systematically get exploited? If so which ones and why? To what extent are naïve or unsophisticated execution algorithms to blame? Even in markets before the proliferation of AT and HFT, an unsophisticated execution strategy would result in higher trading costs. Is the institutional investor concern about AT/HFT anything more than this age old issue re-cast in the new high-tech environment?

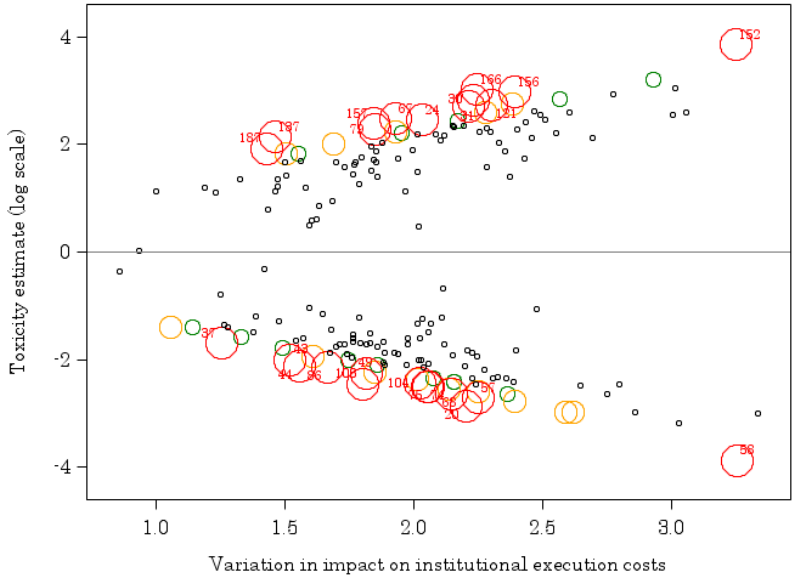
Appendix A: Further tests of the toxicity functional form

We test the independent effects of presence and activity by estimating the following regression:

$$IShorfall_{it} = \alpha + \sum_{k=1}^{187} \beta_k Presence_{itk} + \sum_{k=1}^{187} \gamma_k Activity_{itk} + \varepsilon_{it} \quad \text{(A.1)}$$

giving two measures of toxicity for each active trader, $\widehat{\beta_k}$ and $\widehat{\gamma_k}$. The results are shown in Figure A.1 below:

Panel A: Toxicity of active trader activity, measured by $\widehat{\gamma_k}$



Panel B: Toxicity of active trader presence, measured by $\widehat{\beta_k}$

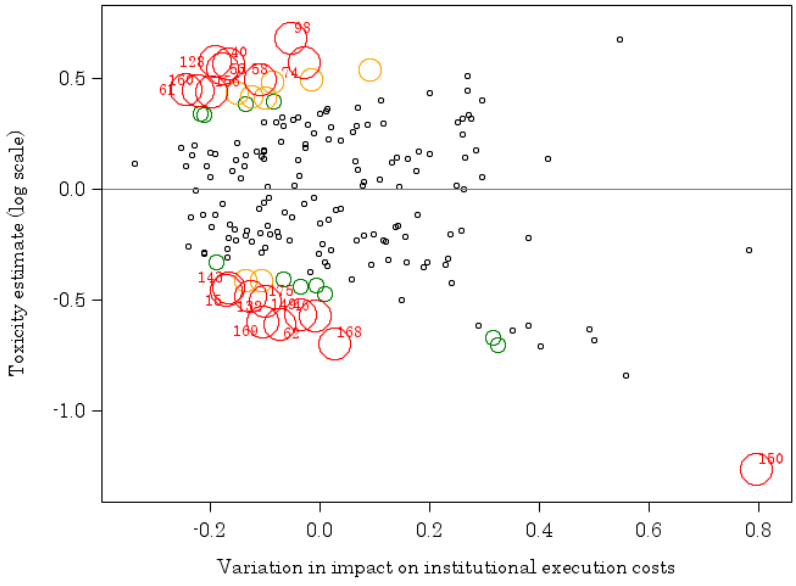


Fig. A.1. Independent effects of active trader activity and presence.

We test the sensitivity of the toxicity estimation procedure to transformation of the activity measures from bounded variables to continuous ones via logit transformations:

$$LogitActivity_{itk} = \ln\left(\frac{Activity_{itk}+0.01}{1-Activity_{itk}+0.01}\right) \quad \text{(A.2)}$$

and estimating the following regression:

$$IShorfall_{it} = \alpha + \sum_{k=1}^{187} \gamma_k LogitActivity_{itk} + \varepsilon_{it} \quad \text{(A.3)}$$

The results are shown below in Figure A.2.

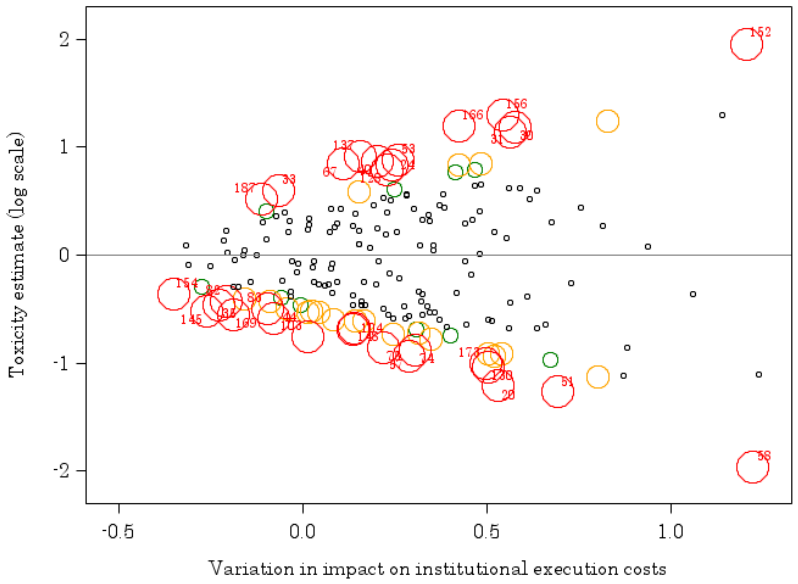


Fig. A.2. Toxicity, estimated using a continuous measure of active trader activity.

Appendix B: Bootstrap procedure

The bootstrap procedure is as follows:

1. Estimate the 2SLS IV models saving the residuals and 187 t-statistics.
2. For each stock i, draw with replacement from its residuals to create a pseudo-time-series of resampled residuals in such a way that re-orders the original time-series. Retain the original values of $Activity_{itk}$ in their original chronological order (relaxed in robustness tests).
3. Construct the pseudo-values of $IShorfall_{it}$ using the resampled residuals, imposing the null hypothesis of zero toxicity, i.e., $\gamma_k = 0 \ \forall k$.
4. Estimate the second stage of the IV models saving the 187 t-statistics.
5. Repeat steps 2–4 many (1,000) times to build a bootstrap distribution of the t-statistics for the toxicity estimates γ_k.
6. Compare the mean, median, quartiles, min, and max of the cross-sectional 187 t-statistics from step 1 against the bootstrap distributions for each of these measures. For example, the bootstrap distribution for the max t-statistic across traders is constructed as the distribution of the maximum t-statistic generated in each of the 1,000 iterations of the bootstrap.

References

Anand, A., P. Irvine, A. Puckett, and K. Venkataraman, 2012, Performance of institutional trading desks: An analysis of persistence in trading costs, Review of Financial Studies 25, 557–598.

Aquilina, M., and C. Ysusi, 2016, Are high-frequency traders anticipating the order flow? Cross-venue evidence from the UK Market, Financial Conduct Authority Occasional Paper 16.

ASIC, 2015, Australian Securities and Investments Commission Report 452, Review of high-frequency trading and dark liquidity.

Biais, B., T. Foucault, and S. Moinas, 2015, Equilibrium fast trading, Journal of Financial Economics 116, 292–313.

Foucault, T., R. Kozhan, and W.W. Tham, 2017, Toxic arbitrage, Review of Financial Studies 30, 1053–1094.

Boehmer, E., K. Fong, and J. Wu, 2020, Algorithmic trading and market quality: International evidence, Journal of Financial and Quantitative Analysis (forthcoming).

Boehmer, E., D. Li, G. Saar, 2018, The competitive landscape of high-frequency trading firms, Review of Financial Studies 31, 2227–2276.

Bradrania, R., A. Grant, P.J. Westerholm, and W. Wu, 2017, Fool's mate: What does CHESS tell us about individual investor trading performance? Accounting and Finance 57, 981–1017.

Brogaard, J., B. Hagströmer, L. Nordèn, and R. Riordan, 2015, Trading fast and slow: Colocation and market quality, Review of Financial Studies 28, 3407–3443.

Brogaard, J., T. Hendershott, S. Hunt, and C. Ysusi, 2014, High-frequency trading and the execution costs of institutional investors, Financial Review 49, 345–369.

Brogaard, J., T. Hendershott, and R. Riordan, 2014, High-frequency trading and price discovery, Review of Financial Studies 27, 2267–2306.

Brogaard, J., T. Hendershott, and R. Riordan, 2017, High frequency trading and the 2008 short-sale ban, Journal of Financial Economics 124, 22–42.

Brogaard, J., T. Hendershott, and R. Riordan, 2019, Price discovery without trading: Evidence from limit orders, Journal of Finance 74, 1621–1658.

Brunnermeier, M.K., and L.H. Pedersen, 2005, Predatory trading, Journal of Finance 60, 1825–1863.

Hagströmer, B., and L. Nordèn, 2013, The diversity of high-frequency traders, Journal of Financial Markets 16, 741–770.

Hasbrouck, J., and G. Saar, 2013, Low-latency trading, Journal of Financial Markets 16, 646–679.

Hendershott, T., C.M. Jones, and A.J. Menkveld, 2011, Does algorithmic trading improve liquidity? Journal of Finance 66, 1–33.

Hirschey, N.H., 2020, Do high-frequency traders anticipate buying and selling pressure?, Management Science (forthcoming).

Jones, C.M., 2013, What do we know about high-frequency trading?. Working paper.

Kirilenko, A.A., A.S. Kyle, M. Samadi, and T. Tuzun, 2017, The Flash Crash: High frequency trading in an electronic market, Journal of Finance 72, 967–998.

Korajczyk, R., and D. Murphy, 2019, High frequency market making to large institutional trades, Review of Financial Studies 32, 1034–1067.

Kosowski, R., A. Timmermann, R. Wermers, and H. White, 2006, Can mutual fund "stars" really pick stocks? New evidence from a bootstrap analysis, Journal of Finance 61, 2551–2595.

Malinova, K., and A. Park, 2016, Liquidity provision and market making by HFTs, Report prepared for the Investment Industry Regulatory Organization of Canada.

Menkveld, A.J., 2013, High frequency trading and the New Market makers, Journal of Financial Markets 16, 712–740.

Menkveld, A.J., 2016, The economics of high-frequency trading: Taking stock, Annual Review of Financial Economics 8, 1–24.

Perold, A., 1988, The implementation shortfall: Paper versus reality, Journal of Portfolio Management 14, 4–9.

Roşu, I., 2016, Fast and slow informed trading, Journal of Financial Markets 43, 1–30.

Sarkar, A., Schwartz, R.A., 2009, Market sidedness: Insights into motives for trade initiation, Journal of Finance 64, 375–423.

Tong, L., 2015, A blessing or a curse? The impact of high frequency trading on institutional investors, Working Paper.

Van Kervel, V., 2015, Competition for order flow with fast and slow traders, Review of Financial Studies 28, 2094–2127.

Van Kervel, V., and A.J. Menkveld, 2019, High-frequency trading around large institutional orders, Journal of Finance 74, 1091–1137.

Yang, L., and H. Zhu, 2020, Back-running: Seeking and hiding fundamental information in order flows, Review of Financial Studies 33, 1484–1533.

Table 1
Descriptive statistics

This table reports descriptive statistics for a number of variables calculated at the stock-day level (stock i on day t). $\$Volume_{it}$ is the dollar volume of trades per stock-day, $ATshare_{it}$ is the "active trader" share of dollar volume ("active traders" are the 187 traders (AT/HFT) with the highest non-direction dollar volume), $ATcount_{it}$ is the number of active traders trading a given stock on a given day, $InstoShare_{it}$ is the large institutional order share of dollar volume (large institutional orders are unidirectional parent orders that are worked in the market for at least two hours and exceed the stock-day's median unidirectional parent order value), and $IShorfall_{it}$ is implementation shortfall for large institutional orders. Quartiles are by dollar volume, with 1 (4) being the highest (lowest) volume stocks.

Quartile	Statistic	$\$Volume_{it}$ ($mil)	$ATshare_{it}$ (%)	$ATcount_{it}$ (#)	$InstoShare_{it}$ (%)	$IShorfall_{it}$ (bps)
Panel A: Pooled sample						
	Mean	22.3	48.4%	63.9	19.3%	16.4
	Std.dev.	43.2	13.1%	16.6	10.2%	57.8
	25^{th} percentile	3.1	40.3%	52.0	11.8%	–7.0
	median	7.9	49.3%	63.0	18.3%	8.7
	75^{th} percentile	22.1	57.5%	76.0	25.5%	32.4
	Observations	52,873	52,873	52,873	52,873	52,873
Panel B: By quartile						
1	Mean	67.7	48.1%	82.0	17.9%	10.4
2	Mean	15.4	51.5%	70.4	20.0%	14.2
3	Mean	6.3	48.7%	58.9	20.9%	16.4
4	Mean	2.6	45.6%	46.7	18.4%	23.6

Table 2
Bootstrap results

This table reports results from bootstrap simulations that quantify toxicity beyond that which is expected by statistical chance (the bootstrap procedure is described in Appendix B). Panel A reports results from the 2SLS IV model estimated using the actual data. The columns *min* through to *max* (where *P* = "percentile") describe the distribution of the 187 toxicity estimate t-statistics (one t-statistic for each of the 187 active traders (AT/HFTs) in the sample). The last four columns indicate how many of the 187 active traders have toxicity t-statistics beyond the threshold given in the column heading (e.g., the column with heading $t < -3$ displays the number of active traders with toxicity t-statistics less than –3). Panel B reports examples of the results using simulated data (simulation iterations 1–6 and 1,000 for illustration). Each iteration of the simulation creates a panel dataset similar to the actual dataset (with active traders and institutional trading costs each stock-day, using residuals from the actual dataset) but with no relation between the active traders and institutional trading costs (other than spurious relations). For each simulated dataset, we estimate the same 2SLS IV model and report the results in the same format as for the actual data. Additionally, the last row of Panel B reports means (across the 1,000 iterations) of the number of active traders with toxicity t-statistics beyond a certain threshold. Panel C reports where each of the results using actual data sit in the distribution generated by the 1,000 simulated datasets (in terms of a percentile). Percentiles are recorded as the starting point of the range, i.e., 0 is the 0–1% bucket, 99 is the 99–100% bucket. For example, the value 52 in the column *min* indicates that the minimum t-statistic of –3.61 in the actual data corresponds to the 52nd percentile of the minimum t-statistics in the 1,000 simulated datasets. All t-statistics are based on double clustered standard errors (by stocks and by date).

	Distribution of t-values for the 187 active trader toxicity estimates							Number of active traders with toxicity t-values beyond various significance thresholds			
Iteration	min	P5	P25	median	P75	P95	max	$t < -3$	$t < -2$	$t > +2$	$t > +3$
Panel A: Actual data											
	–3.61	–2.37	–1.10	–0.08	0.63	2.12	4.00	4	15	12	4
Panel B: Simulated data											
1	–2.58	–1.75	–0.78	–0.04	0.59	1.67	2.39	0	4	6	0
2	–3.09	–1.87	–0.62	0.13	0.67	1.43	3.13	1	7	2	1
3	–2.69	–1.73	–0.83	–0.01	0.69	1.71	4.32	0	7	4	1
4	–3.79	–1.75	–0.76	0.06	0.85	1.58	2.81	1	6	4	0
5	–3.04	–1.91	–0.77	–0.12	0.68	1.83	2.71	2	9	5	0
6	–3.00	–1.88	–0.84	0.00	0.67	1.59	3.28	0	4	5	1
…											
1000	–3.28	–1.85	–0.69	0.06	0.80	1.90	3.23	2	6	7	2
Mean								1.24	6.87	5.55	1.20
Panel C: Actual data in terms of percentiles of the bootstrap distributions											
	52	0	0	27	31	99	50	98	99	99	98

Table 3
Characteristics of toxic traders

This table reports results from regressions of trader-level toxicity on trader-level characteristics:

$$\widehat{Toxicity}_k = \alpha + \sum_c \beta_c Characteristic_{c,k} + \epsilon_k$$

where the dependent variable is estimated gross toxicity in excess of expected toxicity. The c trader-level characteristics are as follows. *Average Activity* is the trader's average share of dollar volume. *Holding Time* is the average holding time (in '000 seconds) of long or short positions that are closed within a day. *Number Fast* is the number of order amendments (in millions) sent within 500ms of the order submission. *Order Amend Time* is the average time (in seconds) between an order submission and its amendment for amendments sent within 500ms of the order submission. *Sophistication* is a measure of each trader's efficiency in generating intraday margin from traded stock (Sharpe ratio of daily trading profit). *$Volume* is the natural log of the trader's average daily traded dollar volume. *Number of Orders* is the natural log of the trader's average daily number of submitted orders. *Order-to-trade Ratio* is the average of the number submitted orders divided by the number of trades executed by the trader each stock-day. *Percentage Traded* is the average percentage of the trader's dollar volume that is in the form of intraday round-trip trades (i.e., both buying and selling a stock within a day). For a given trader, this percentage is calculated each stock-day, then averaged across stocks (with dollar volume weighting) and across days (equal weighting). T-statistics are reported in parentheses. ***, **, and * indicate statistical significance at 1%, 5%, and 10% levels, respectively.

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
Intercept	–0.01	–0.03	–0.01	0.01	0.35	0.20	–0.06	0.99*
	(–0.17)	(–0.55)	(0.12)	(0.24)	(0.78)	(0.46)	(–1.31)	(1.95)
Average Activity	1.24							
	(0.11)							
Holding Time		0.01						0.008
		(0.71)						(0.48)
Number Fast			2.66					–0.827
			(0.99)					(–0.27)
Order Amend Time			–0.10					–0.13
			(–0.30)					(–0.37)
Sophistication				–0.05				–0.09
				(–0.80)				(–1.22)
$Volume					–0.04	–0.01		–0.08**
					(–1.20)	(–0.49)		(–2.24)
Number of Orders					0.03			0.03
					(1.58)			(1.36)
Order-to-trade						0.0007		
Ratio						(1.21)		
Percentage Traded							0.25**	0.40**
							(2.07)	(2.53)
R-Squared	0.02%	0.28%	0.67%	0.36%	0.39%	0.93%	2.34%	6.24%

Table 4
Characteristics of toxic trader activity

This table reports results from regressions of trader-level toxicity on characteristics of trader-level activity:

$$\widehat{Toxicity}_k = \alpha + \sum_c \beta_c ActivityCharacteristic_{c,k} + \epsilon_k$$

where the dependent variable is estimated gross toxicity in excess of expected toxicity. The c trader-level activity characteristics are as follows. *Cross-sectional StdDev* for a trader is calculated by taking their daily cross-sectional standard deviations of their activity (share of dollar volume) and then averaging those daily cross-sectional standard deviations. *Time-series StdDev* for a trader is calculated by taking, for each stock, the time-series standard deviation of their activity and then averaging those time-series standard deviations. *Average Activity* is the trader's average share of dollar volume. *Activity in Medium* and *Activity in Small* are the trader's average share of dollar volume in medium stocks (second quartile) and in small stocks (bottom two quartiles), with *% Activity In Small* being a measure of activity in small stocks relative to the trader's overall activity. *Activity In Falling* is similarly defined as the trader's average share of dollar volume during a period when the market fell sharply (August 5–24, 2015), with *% Activity In Falling* being the trader's activity during that period relative to their average activity overall. T-statistics are reported in parentheses. ***, **, and * indicate statistical significance at 1%, 5%, and 10% levels, respectively.

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Intercept	0.00	–0.02	–0.27	–0.01	–0.03	–0.27
	(0.04)	(–0.45)	(–4.32)***	(–0.18)	(–0.65)	(–3.86)***
Cross–sectional StdDev	32.27					44.03
	(2.00)**					(2.27)**
Time-series StdDev	–32.23					–39.33
	(–1.68)*					(–1.97)*
Average Activity		–304.28		–22.08		–8.24
		(–2.04)**		(–0.77)		(–0.62)
Activity In Medium		87.76				
		(1.15)				
Activity In Small		222.6				
		(2.89)***				
% Activity In Small			0.83			0.90
			(3.81)***			(3.93)***
Activity In Falling				21.43		
				(0.95)		
% Activity In Falling					0.06	–0.02
					(0.65)	(–0.19)
R-Squared	1.87%	5.56%	4.19%	0.51%	0.08%	7.00%

	Directional trading	Non-directional trading
High volume	"*Institutional investors*" Large fundamental buyers or sellers	"*Active traders*" (AT/HFT) Algorithmic market making and short holding horizon strategies including various arbitrage algorithms and predatory trading
Low volume	"*Retail and small institutional investors*" Small fundamental buyers or sellers	"*Other*" Non-algorithmic intermediation, small short holding horizon traders, opportunistic traders

Fig. 1. Classification of traders in two dimensions. This figure shows qualitatively how the two main groups of traders that we analyze are classified ("*Institutional investors*", top left quadrant, and "*Active traders*" (AT/HFT), top right quadrant).

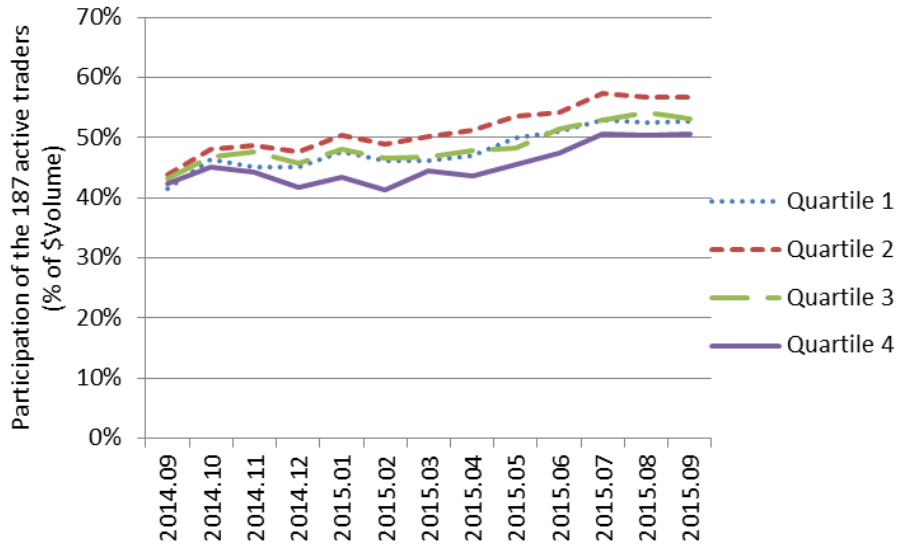
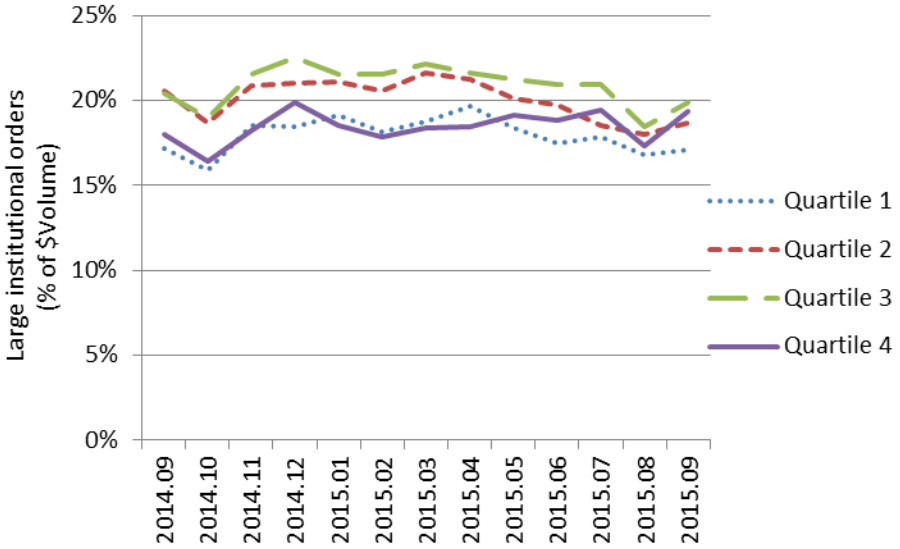


Fig. 2. Active trader share of dollar volume in stock quartiles through time. The "active traders" are the 187 traders (AT/HFT) with the highest non-directional dollar volume (buys that are accompanied by sells in the same security within a week and vice versa) throughout the sample period. Their share of volume is measured each stock-day as the dollar volume of their buys and their sells normalized by the total dollar volume of all buys and all sells. We then compute equal-weighted averages of their share of dollar volume in quartiles of stocks each month. The quartiles are by dollar volume with Quartile 1 being the highest volume stocks.

Panel A: Large institutional orders as a percentage of dollar volume



Panel B: Implementation shortfall for the large institutional orders

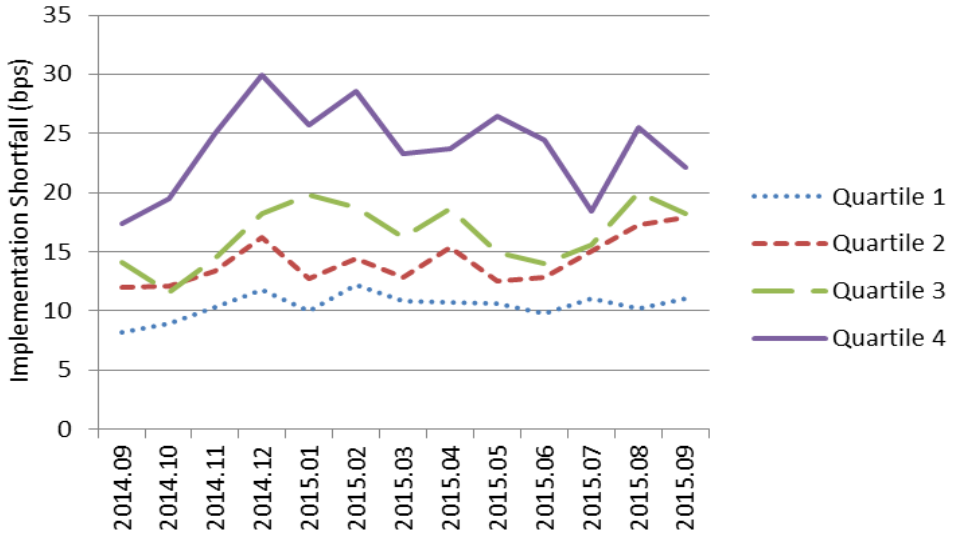


Fig. 3. Large institutional orders and their trading costs for stock quartiles through time. This figure shows large institutional orders as a percentage of dollar volume (Panel A) and their average implementation shortfall in bps (Panel B). Large institutional orders are unidirectional parent orders that are worked in the market for at least two hours and exceed the median size of unidirectional parent orders that stock-day. We calculate value-weighted average implementation shortfall for each stock-day and then take the equal-weighted average across stocks in a given quartile each month. The quartiles are by dollar volume with Quartile 1 being the highest volume stocks.

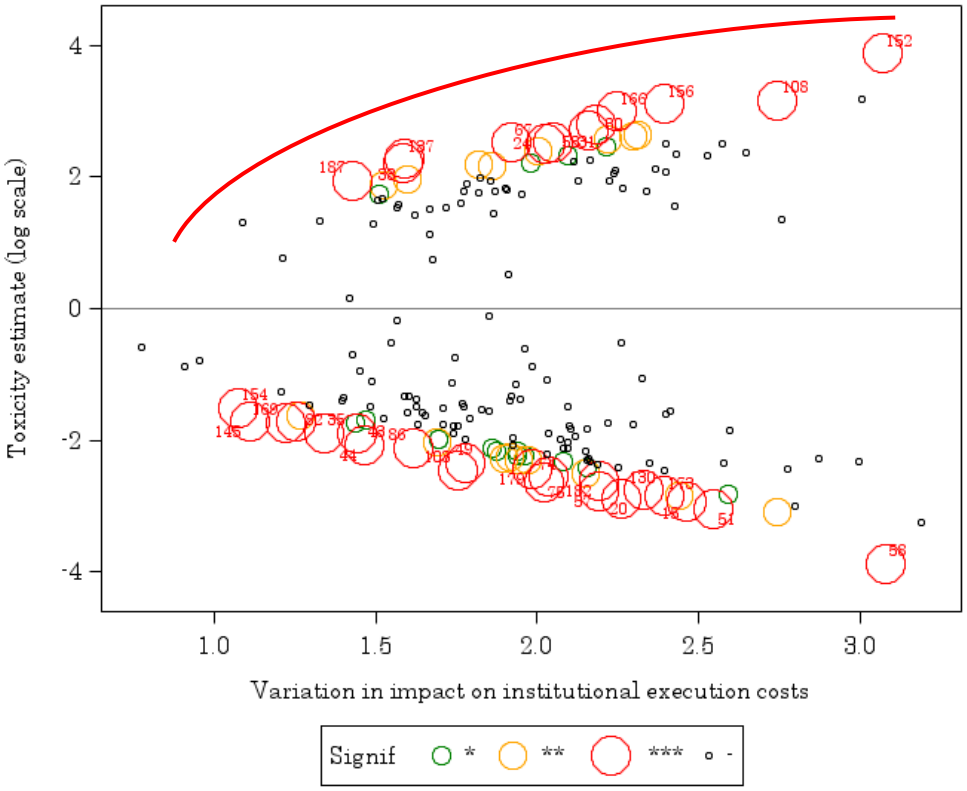


Fig. 4. Toxicity of active traders using the baseline OLS approach. Each circle on this figure represents one of the 187 active traders (AT/HFT accounts). The vertical axis measures the trader's toxicity (expressed in log terms) with positive (negative) numbers indicating the trader is associated with increased (decreased) institutional trading costs. Toxicity estimates are obtained from a regression of institutional trading costs each stock-day on each of the trader's shares of dollar volume that stock-day. The horizontal axis measures the variation in the trader's impact on institutional trading costs (log of the standard error of the toxicity estimate) with lower values indicating greater consistency in impact. The superimposed curve is the approximate "toxicity frontier", i.e., the collection of active traders that have maximum toxicity for a given level of variation in their impact on institutional trading costs or a minimum level of variation in their impact for a given level of toxicity. The size of the circles indicates the statistical significance of the toxicity estimate, with the smallest circles being toxicity estimates that are not significant at the 10% level, followed by significant at the 10%, 5%, and 1% levels (largest circles). Statistical significance is based on double clustered standard errors (by stocks and by date). Numbers next to the largest circles are masked (anonymized) trader identifiers.

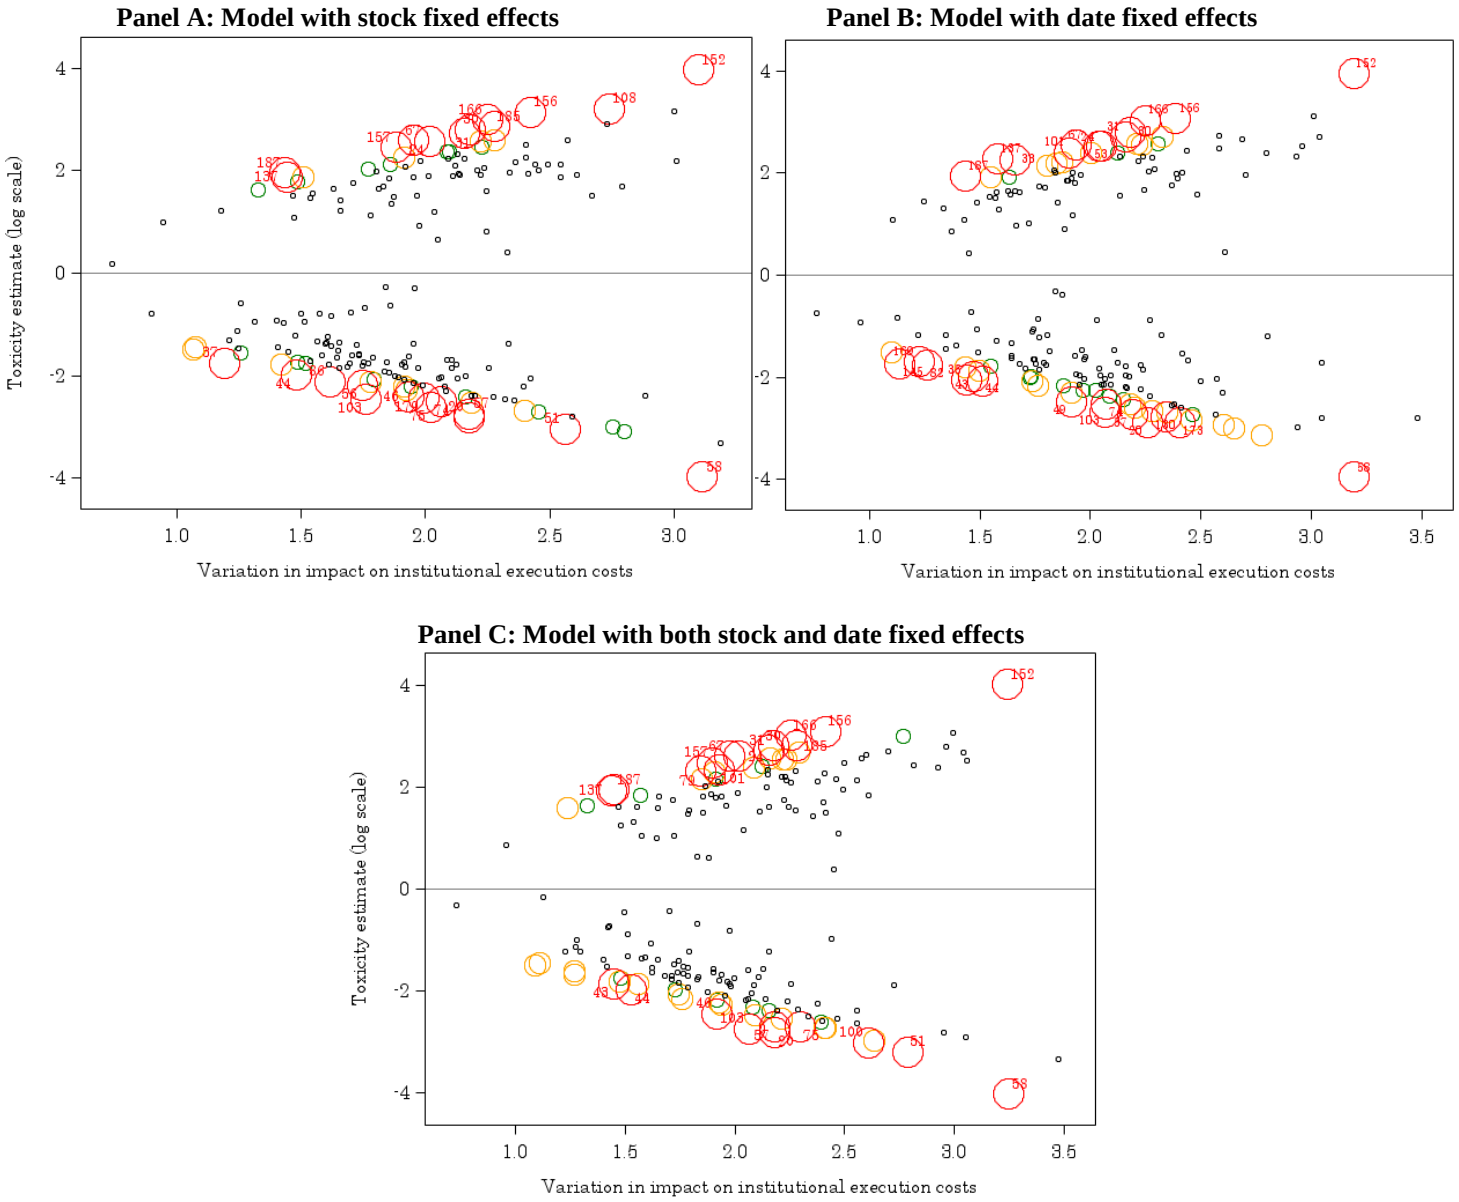


Fig. 5. Toxicity of active traders using OLS approach and fixed effects. Each circle on this figure represents one of the 187 active traders (AT/HFT accounts). The vertical axis measures the trader's toxicity (expressed in log terms) with positive (negative) numbers indicating the trader is associated with increased (decreased) institutional trading costs. The horizontal axis measures the variation in the trader's impact on institutional trading costs (log of the standard error of the toxicity estimate) with lower values indicating greater consistency in impact. Toxicity estimates are derived from an OLS regression of institution trading costs each stock-day on each of the trader's share of dollar volume that same stock-day, similar to Figure 4, but with stock fixed effects (Panel A), date fixed effects (Panel B), and both stock and date fixed effects (Panel C). The size of the circles indicates the statistical significance of the toxicity estimate, with the smallest circles being toxicity estimates that are not significant at the 10% level, followed by significant at the 10%, 5%, and 1% levels (largest circles). Statistical significance is based on double clustered standard errors (by stocks and by date).

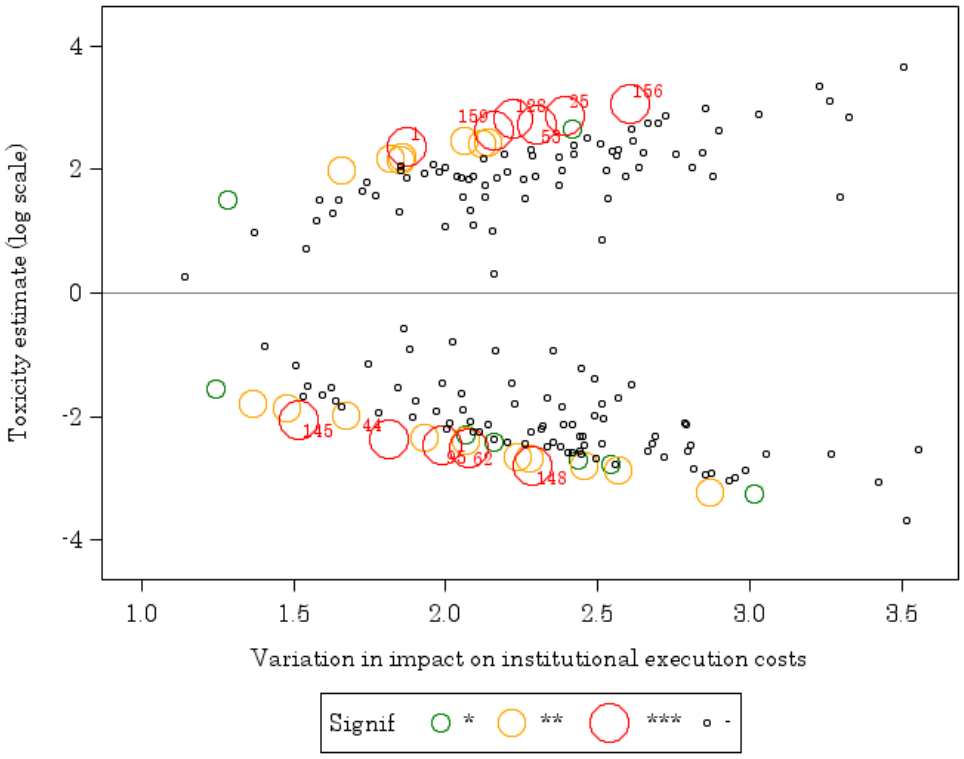


Fig. 6. Toxicity of active traders using 2SLS IV regressions. Each circle on this figure represents one of the 187 active traders (AT/HFT accounts). The vertical axis measures the trader's toxicity (expressed in log terms) with positive (negative) numbers indicating the trader is associated with increased (decreased) institutional trading costs. Toxicity estimates are derived from two-stage least squares instrumental variables regressions in which active trader activity is instrumented with lags of their activity in the same stock. The horizontal axis measures the variation in the trader's impact on institutional trading costs (log of the standard error of the toxicity estimate) with lower values indicating greater consistency in impact. The size of the circles indicates the statistical significance of the toxicity estimate, with the smallest circles being toxicity estimates that are not significant at the 10% level, followed by significant at the 10%, 5%, and 1% levels (largest circles). Statistical significance is based on double clustered standard errors (by stocks and by date). Numbers next to the largest circles are masked (anonymized) trader identifiers.