

Exploratory Trading

Adam D. Clark-Joseph*

December 31, 2014

Abstract

To investigate how high-frequency traders (HFTs) can predict price-changes, I analyze novel, comprehensive, account-labeled message records from the E-mini S&P 500 futures market that allow me to identify and study HFTs' individual behaviors. I model how an HFT could actively learn about market conditions, by initiating small “exploratory” trades and observing other traders' responses. Empirical tests of the model's predictions provide evidence that HFTs in the E-mini use this technique to identify periods when prices are likely to change. These findings indicate that the HFTs' superior capacity to predict price-changes involves more than merely reacting to news faster than other traders. The empirical results also elucidate other connections between high-frequency trading and speed.

JEL classification: G14; G19

Keywords: High-frequency trading; Learning; Market microstructure

*Department of Finance, University of Illinois at Urbana-Champaign. E-mail: adcj@illinois.edu. I thank Jonathan Brogaard, Eric Budish, Terrence Hendershott, Alp Simsek, and Jeremy Stein for their detailed suggestions and discussion. I also thank seminar participants at Harvard, UBC Sauder, Indiana Kelley, Stanford GSB, Illinois, OSU Fisher, Dartmouth Tuck, Wharton, and NYU Stern for their useful feedback. I am particularly indebted to John Campbell and Andrei Shleifer for their invaluable advice and guidance. All remaining errors are my own. I gratefully acknowledge the support of an NSF Graduate Research Fellowship. The views expressed in this paper are my own and do not constitute an official position of the Commodity Futures Trading Commission, its Commissioners, or staff.

1 Introduction

High-frequency algorithmic traders are now responsible for almost half of the trading on financial exchanges. Analyzing modern markets requires not only characterizations of what high-frequency traders (HFTs) do, but also explanations of how and why they do those things. Many HFTs manage to predict price-movements with unprecedented success, and understanding how HFTs accomplish this feat is crucial for determining the economic mechanisms that underlie high-frequency trading and its effects. The standard assumption in the literature has been that HFTs can predict price-movements because—and only because—they digest and respond to new public information before slower traders can do so. Although using new information sooner than other traders unquestionably improves HFTs’ forecasting to some extent, the assumption that this sort of superior speed is the *only* mechanism driving HFTs’ exceptional predictions is both strong and restrictive, and this assumption’s validity has not been assessed empirically. Using novel data at the Commodity Futures Trading Commission (CFTC) that enable me to analyze individual HFTs’ actions, I study how HFTs in the E-mini S&P 500 futures market acquire superior information about imminent price-changes. Specifically, I investigate whether HFTs obtain any valuable private knowledge in a way other than just reacting to public information the fastest, and I find explicit evidence that in fact they do.

I model a simple framework in which an HFT places small aggressive (i.e., marketable) orders and actively learns about expected price-impact by observing the responses that these “exploratory” orders elicit. In the E-mini, as in many markets, aggressive order-flow exhibits strong predictability at short horizons. However, front-running predictable orders is profitable only when those orders have a sufficiently large price-impact, and price-impact is too small on average for indiscriminate front-running to be profitable. Through exploratory trading, the HFT gathers information that helps him to trade ahead of predictable orders at only those times when doing so will be profitable.¹ In general, the market activity following an arbitrary aggressive order often doesn’t convey information about expected price-impact, because the activity is just a response to the same stimulus that prompted the order rather than a true reaction caused by the order itself.

¹Time-variation in price-impact is a robust empirical fact in the data, and my results neither depend upon nor dictate the specific interpretation of that variation. Section 1.3 discusses some of the related empirical and theoretical literature.

The trader who placed the order can judge whether it causes the subsequent activity, but other traders cannot, and the possibility that they are observing the uninformative scenario diminishes how much they learn from the market activity. The HFT learns more from his exploratory order than other traders can learn from it, simply because he alone knows why his exploratory order was placed. By using exploratory trading, the HFT can leverage his seemingly inconsequential private knowledge about why he placed a particular order to obtain significant private information that helps him to predict price movements.

Using the unique CFTC data, I test the predictions of the exploratory trading model under the conservative assumption that HFTs' small aggressive orders all are exploratory in nature. There are many reasons other than exploration for which HFTs might place small aggressive orders, so treating all of those orders as exploratory dilutes the effects from any truly exploratory orders and therefore raises the bar for obtaining significant results. Nevertheless, consistent with the model's predictions, I find that a simple measure of the market response to the last small aggressive order by a given HFT helps to explain a significant component of that HFT's earnings on subsequent, larger aggressive orders, even after controlling for the market response to the last small aggressive order placed by anyone. Also as predicted, after controlling for the market response to the last small aggressive order by anyone, the market response to the HFT's last small aggressive order does not help to explain *other* traders' earnings on *their* subsequent, larger aggressive orders. As Section 1.2 explains in detail, although futures trades are typically thought of as just one segment of some cross-market play (hedging, arbitrage, etc.), for the purposes of this paper, empirical features of the E-mini and of the HFTs' behavior make it meaningful to analyze HFTs' earnings in the E-mini market by itself.

In principle, the two results above could be consistent with the alternative hypothesis that HFTs somehow possess long-lived private information about future prices and split up their orders over time (plausible in the context individual equities, if perhaps less so for index futures).² Unlike the exploratory trading theory, though, this informed-order-splitting story has the counter-factual implication that the market response to an HFT's small order will only help to explain the HFT's

²The recent theoretical models by Martinez and Rosu (2103), and Foucault *et al.* (2013), motivate this alternative hypothesis. These models include a long-lived "forecast error" component of information, in addition to the standard infinitesimally lived component of information (called "news" in the models).

earnings on a subsequent large order when the two orders have the same sign (i.e., buy or sell). Empirically, the market response to an HFT’s small order helps to explain the HFT’s earnings on the larger order regardless of whether both orders have the same sign.

Contrary to the prevailing assumption that HFTs’ superior information derives exclusively from superior reaction speed, the empirical evidence in this paper indicates that HFTs in the E-mini also obtain part of their informational edge through a separate channel: exploratory trading. At a minimum, this evidence about an important group of HFTs in an important market provides a significant counterexample to the standard “simply superior speed” assumption. As discussed in Section 1.3, though, my results are likely not unique to the E-mini, or even to futures markets, because the microstructure elements that enable exploratory trading are common features of numerous other markets and have well-established economic foundations.

Beyond illuminating several important issues specific to high-frequency trading and modern microstructure, my results also bear on an issue of broad significance to the field of finance, namely how the quantitative changes in the absolute speed at which markets operate have lead to qualitative changes in the economic mechanisms at work in those markets. Exploratory trading only yields substantive information when the time-delay between placing an order and observing the effects is sufficiently brief, and the temporal resolution of market data is sufficiently fine. Whereas superior reaction time is only a matter of relative speed—be it measured in months or microseconds—exploratory trading is tied to the absolute speed at which market activity occurs.

1.1 Related analyses of HFTs

HFTs are not all alike, but they share some distinctive features. They have the capacity to react to market events and news in milliseconds or less, they trade very frequently and unwind positions within minutes, and they usually end the trading day holding minimal net inventory. As Hagstromer and Norden (2013) document in the NASDAQ-OMX equities market, and Baron *et al.* (2013) document in the E-mini market, some HFTs basically behave like traditional market-makers and supply liquidity/immediacy, but many others do just the opposite, predominantly placing marketable orders, which consume liquidity/immediacy. Baron *et al.* directly compute individual HFTs’ trading profits in the E-mini and find that HFTs of both varieties tend earn large and stable

profits. Since HFTs’ trading profits are necessarily some combination of compensation for providing liquidity/immediacy, and gains from trading on information that other market participants do not have, the Baron *et al.* results suggest that HFTs of the second variety, at least, possess some kind of superior information. Reinforcing this conclusion, Brogaard *et al.* (2013) analyze aggregate HFT activity in a large sample of NASDAQ stocks and find that HFTs’ aggressive orders tend to go in the same direction as subsequent permanent price movements, and sufficiently so for the orders to be profitable on average.

The studies above indicate that HFTs enjoy some sort of valuable informational advantage, but the empirical literature offers less clarity about the nature and sources of the advantageous information. The standard assumption has been that HFTs’ ability to use new public information fastest is the sole source of their informational edge, and this premise underlies much of the theoretical work on high-frequency trading, including that of Biais *et al.* (2010), Jovanovic and Menkveld (2010), and Budish *et al.* (2013). This assumption tightly circumscribes the character, scope, and effects of high-frequency trading. Empirical evidence indicates that HFTs make use of publicly available information (*cf.* Brogaard *et al.*), and without question they can do so more quickly than other traders. However, there is no rigorous empirical basis for the strong assumption that this “superior speed” mechanism is the only source of HFTs’ exceptional information. The nature of HFTs’ superior information has far-reaching implications for HFTs’ effects on markets, for the structure and functioning of the high-frequency trading industry, and for optimal policy design, so understanding it is of first-order importance.

1.2 Studying HFTs’ superior information

When a trader initiates a transaction in the E-mini, i.e., places a so-called “aggressive” order, he mechanically pays his counterparty a fraction of the bid-ask spread, and so to profit on his aggressive order, the trader must correctly anticipate a price movement.³ I use profits on aggressive orders in the E-mini as a medium through which to study HFTs’ superior information. Much of my empirical analysis involves the profitability of individual aggressive orders. Because all E-

³If the trader “anticipates a price movement” because a mispricing in the ask (bid) presents an arbitrage opportunity, he pays part of the spread in only a nominal sense, because the spread-cost is offset by the simultaneous gain from buying (selling) the mispriced contracts. I thank an anonymous referee for this observation.

mini contracts of a given expiration date are identical, it is neither meaningful nor possible to distinguish among the individual contracts in a trader’s inventory, so there is generally no way to determine the exact prices at which a trader bought and sold a particular contract. As a result, it is typically impossible to measure directly the profits that a trader earns on an individual aggressive order. However, the cumulative price change following an aggressive order, normalized by the order’s direction (+1 for a buy, or -1 for a sell), can be used to construct a meaningful estimate of the order’s profitability, and this general approach is standard in the literature. I discuss implementation details in Section 4.1 and Internet Appendix B, but roughly speaking, the average expected trading profit from an aggressive order equals the expected permanent price movement in the order’s direction, minus trading/clearing fees and half the bid-ask spread.

Since HFTs exhibit great heterogeneity, aggregate HFT activity reveals little about what individual HFTs really do. Regulatory records that the Chicago Mercantile Exchange provides to the CFTC are currently some of the only data for U.S. markets disaggregated enough to be fully adequate for studying high-frequency trading at the level of individual HFTs. Kirilenko *et al.* (2010) pioneered the use of transaction data from these records to investigate high-frequency trading in the E-mini S&P 500 futures market during the “Flash Crash” of 2010. That paper introduces a data-driven scheme to classify trading accounts, and specifically to identify HFTs, using simple measures of overall trading activity, and of inter- and intra-day variation in net inventory position. In the present paper, I analyze a richer sample of E-mini data, and I build upon the techniques of Kirilenko *et al.* to identify the HFT accounts; I identify 30 HFTs in my sample.

As a group, the 30 HFTs earn roughly 40% of their trading profits in the E-mini from their aggressive orders. Examining these HFTs individually, however, reveals that although all of them make money in the E-mini, only eight of the 30 profit on average from their aggressive trading. For brevity, I refer to these eight HFTs as “A-HFTs,” and to the remaining 22 as “B-HFTs.” The B-HFTs may or may not possess unusually valuable information, but the A-HFTs definitely do. Therefore I focus on the A-HFTs and investigate the origins of their superior information, specifically, whether the A-HFTs obtain any of it through exploratory trading.

1.3 Microstructure foundations of exploratory trading

Exploratory trading is just a simple form of active learning in a financial market, an idea that dates back to the theoretical work of Leach and Madhavan in 1992 and 1993, if not further. At a mechanical level, exploratory trading involves nothing more than placing small aggressive orders, then learning about expected price-impact using the responses that the exploratory orders elicit from market-makers.⁴ From an economic perspective, exploratory trading is a device for obtaining knowledge from market-makers about the probability that orders in the near future will be followed by a permanent price change, i.e., the probability of what can loosely be termed “informed trading.”⁵ Prices in real markets are necessarily discrete, so quoted prices alone can never perfectly reveal or perfectly aggregate every individual market-maker’s private knowledge. This slight but inevitable heterogeneity in different market-makers’ beliefs makes it possible for an exploring trader to gather knowledge from some market-makers that other market-makers do not possess, and such knowledge enables the explorer to identify order flow that is more likely to be informed than some market-makers realize. Order flow exhibits strong short-run predictability in most markets, so the explorer can typically trade ahead of some of the identified informed orders, and thereby earn profits.

Starting with Hasbrouck’s work in 1991, short-run persistence in order-flow sign (buy vs. sell) has been a robust empirical finding across numerous markets. Equally robust and widespread is the finding that trades tend to cluster together in time. In their 2000 paper, Engle and Dufour document each of these features independently, and they further find that the autocorrelation in order-flow sign increases as orders arrive more closely together in time. In the E-mini data analyzed for the present paper, the signs of aggressive orders exhibit strong positive autocorrelation (the average probability that an aggressive order will have the same sign as the one before it is around 75%) and this autocorrelation becomes even stronger when the arrival rate of aggressive orders increases. These findings yield a picture of trading characterized by frequent, sporadic “bursts” of many orders with the same sign arriving in close succession.⁶

⁴I use “market-maker” as a heuristic short-hand for “trader with limit orders resting in the order book.” In the E-mini market, specifically, there are no official market-makers.

⁵This mechanical sense of “informed trading” is observationally equivalent to the traditional notion for equities of “trading based on private insider information.” However, the mechanical characterization is more suitable for the futures markets that I analyze, in which the appropriate analogue of “private insider information” may be unclear.

⁶See, for example, Ellul *et al.* (2007), Biais *et al.* (1995), and more generally, the literature review by Parlour

Such bursts of intense trading activity have two classic, conflicting explanations. Admati and Pfleiderer (1988) explain the clusters as coordinated liquidity trading among uninformed traders, while Easley and O'Hara (1992) posit that such bursts arise from information-based clustering of informed traders. Engle and Russell (1998) present evidence that some clusters of trades seem largely information-based, while other clusters of trades appear to be liquidity-based. If market-makers know the current probability of informed trading, standard theory suggests that the bid-ask spread should widen as that probability increases. Consistent with theory, Engle and Russell find that intense clusters of trades tend to be liquidity-based when the spread is narrow, and information-based when the spread is wide. Permanent price impact will tend to be large for an information-based burst of orders, but small for a liquidity-based burst.

Because order-flow sign is persistent, especially when orders are arriving rapidly, it is not difficult to aggressively buy (sell) ahead of future aggressive buy (sell) orders.⁷ However, indiscriminately trading ahead of the foreseeable remnants of a burst of orders tends to be unprofitable, since the subsequent price-change is generally smaller than the part of the spread you would pay to aggressively trade ahead of the foreseeable orders. To profit from order-flow predictability, a trader needs some private knowledge that helps him distinguish uninformed bursts of orders from informed ones. More specifically, the trader needs to be able to make this distinction more accurately than some market-makers.

If prices were continuous, competitive market-makers would, in equilibrium, post quotes at precisely the levels that would earn zero expected profit. In this case, through the spread, all the market-makers would reveal to one another, and to everyone else, all of their relevant knowledge about the expected adverse-selection risk. If the market-makers all know the same information, then knowledge obtained from the market-makers clearly could not give a trader an informational edge that would permit profit any market-maker's expense.

When prices are discrete, the situation is different. Competitive market-makers only have to post quotes within the same tick as the zero-profit price. Therefore the spread will not be perfectly revealing, different market-makers can hold slightly different beliefs, and knowledge obtained from

and Sepi (2008).

⁷I emphasize that exploiting this robust statistical property of order flow is *very* different from anticipating a particular trader's orders.

one (or more) market-maker could therefore grant an informational edge sufficient for a trader to profit at the expense of a different market-maker. Scope for exploratory trading arises precisely because of this possibility for heterogeneously informed market-makers. Moreover, the larger is the tick-size as a fraction of the spread, the greater is the potential heterogeneity in market-makers' information and beliefs. Exploratory trading closely ties to emerging research on the relationship between tick-size and high-frequency trading, e.g., Yao and Ye (2014).

The remainder of this paper is organized as follows: Section 2 introduces a simple model of exploratory trading along with the model's central predictions, and sets the empirical agenda. Section 3 describes the data set, presents some summary statistics, and precisely defines HFTs. Section 4 addresses the overall profitability of HFTs' aggressive orders and precisely characterizes the A-HFTs. Section 5 presents direct empirical tests of the exploratory trading model's key predictions, and section 6 examines the practical significance of exploratory information. Section 7 discusses extensions and implications of the empirical results, and Section 8 concludes.

2 Exploratory trading model

This section introduces a stylized model of exploratory trading that provides the framework for my empirical analysis. Because my empirical analysis centers on the A-HFTs' information, my model shares this tight focus, and it abstracts away from the detailed microstructure foundations discussed in Section 1.3.

2.1 Preliminaries

In an order-driven market, such as the E-mini, every regular transaction is initiated by one of the two executing transactors. The transactor who initiates is referred to as the "aggressor," while the opposite transactor is referred to as the "passor." The passor's order was resting in the order book, and the aggressor entered a new order that executed against the passor's preexisting resting order. If the best bid and best ask were held fixed, a trader who aggressively entered then aggressively exited a position would lose the bid-ask spread on each contract, whereas a trader who passively entered then passively exited a position would earn the bid-ask spread on each contract. Intuitively, aggressors pay for the privilege of trading precisely when they wish to do

so, and passors are compensated for the costs of supplying this “immediacy,” cf. Grossman and Miller (1988). These costs include fixed operational costs and costs arising from adverse selection, cf. Glosten and Milgrom (1985), Stoll (1989).

An aggressive order will execute against all passive orders at the best available price level before executing against any passive orders at the next price, so an aggressive order will only have a literal price-impact if it eats through all of the resting orders at the best price. In the E-mini market, it is rare for an aggressive order to have a literal price-impact, not only because there are typically enormous numbers of contracts at the best bid and best ask, but also because aggressive orders overwhelmingly take the form of limit orders priced at the opposite best (which cannot execute at the next price level).

Because the bid-ask spread in the E-mini is essentially constant, movements of the best bid, best ask, and mid-point prices are generally interchangeable, so unless otherwise noted, I restrict attention hereafter to price changes distinct from bid-ask bounce.

2.2 Model

Consider an order-driven market with discrete prices, and two periods $t = 1, 2$. Both the order book and order flow are observable. I refer to the aggregated quantities of the passive orders in the order book as “resting depth.”

2.2.1 The HFT

Consider a single trader, “the HFT,” who has the opportunity to submit an aggressive order at the start of each time-period. The HFT submits only aggressive orders, and these orders are limited in size to N contracts or fewer. Let q_t denote the signed quantity of the aggressive order that the HFT places in period t , where a negative quantity represents a sale, and a positive quantity represents a purchase. The HFT only trades contracts at the initial best bid or ask, so his orders affect resting depth in the order book but have no literal price-impact.

The HFT pays constant trading costs of $\alpha \in (0.5, 1)$ per contract. The lower bound of 0.5 on α corresponds to half of the minimum possible bid-ask spread, while the upper bound of 1 merely excludes trivial cases of the model in which aggressive orders will always be unprofitable for the

HFT.

2.2.2 Passive orders

There are two possible “liquidity states” (Λ) for the behavior of passive orders: accommodating ($\Lambda = A$) and unaccommodating ($\Lambda = U$). This is the same in both time-periods, $t = 1, 2$. With *ex-ante* probability u , $\Lambda = U$, and $\Lambda = A$ with *ex-ante* probability $1 - u$. Assume $0 < u < 1$, so that both liquidity states are possible.

Intuitively, aggressive orders have a small price-impact in the accommodating liquidity state, and a large price-impact in the unaccommodating liquidity state. The liquidity state characterizes the behavior of resting depth in the order book after an aggressive order executes—a generalization of price-impact appropriate for an order-driven market. When an aggressive buy (sell) order executes, it mechanically depletes resting depth on the sell (buy) side of the order book. Following this mechanical depletion, traders may enter, modify, and/or cancel passive orders, so resting depth at the best ask (bid) can either replenish, stay the same, or deplete further. The aggressive order’s impact is offset to some extent—or even reversed—if resting depth replenishes, whereas the aggressive order’s impact is amplified if resting depth depletes further. In the accommodating state ($\Lambda = A$) resting depth weakly replenishes, while in the unaccommodating state ($\Lambda = U$) resting depth further depletes.

Although the order book is observable, the static features of passive orders in the order book do not directly reveal the liquidity state Λ . Because the liquidity state relates to the dynamic behavior of resting depth after an aggressive order executes, Λ can only be deduced from the changes in the order book that follow the execution of an aggressive order.

As a baseline, assume that the HFT learns Λ prior to period 2 if and only if he places an aggressive order in period 1. This assumption is relaxed in Section 2.4.

2.2.3 Aggressive order-flow

At the end of period 2, traders other than the HFT place aggressive orders. Let the variable $\varphi \in \{-1, 0, +1\}$ characterize this aggressive order-flow. The realization of φ does not depend on the liquidity state, Λ , nor does it depend on the HFT’s actions; assume that $\varphi = +1$ and $\varphi = -1$

with equal probabilities $\mathbb{P}\{\varphi = +1\} = \mathbb{P}\{\varphi = -1\} = v/2$, and $\varphi = 0$ with probability $1 - v$. The variable φ is just a coarse summary of the order-flow—it does not represent the actual number of contracts. Intuitively, $\varphi = -1$ represents predictable aggressive selling, $\varphi = +1$ represents predictable aggressive buying, and $\varphi = 0$ represents the absence of predictable aggressive trading in either direction.

Note about exogeneity assumptions Aggressive order-flow is assumed to be independent of liquidity state for convenience; introducing dependence between aggressive order-flow and the liquidity state complicates the algebra but does not change the model in any interesting ways.

The assumption that neither aggressive order-flow nor the liquidity state depend on the HFT’s actions is more innocuous than it might initially seem, because the HFT will turn out not to do anything that would particularly stand out to other traders. In the first period, the HFT will either place a very small order, or no order. In the second period, if the HFT places an order, it may be large, but if it is, it will always be in the same direction as expected aggressive order-flow.

2.2.4 Prices and price-changes

Prices remain constant between periods 1 and 2, then at the end of period 2 the price changes by $y \in \{-1, 0, +1\}$. Together, the aggressive order-flow, φ , and the liquidity state, Λ , determine y as follows:

$$y = \begin{cases} \varphi & \text{if } \Lambda = U \\ 0 & \text{if } \Lambda = A \end{cases} \quad (1)$$

When the liquidity state is *unaccommodating* ($\Lambda = U$), the aggressive order-flow can affect the price, and $y = \varphi$. However, if the liquidity state is *accommodating* ($\Lambda = A$), aggressive order-flow does not affect the price, and $y = 0$ even when $\varphi \neq 0$. In the spirit of Easley and O’Hara/Engle and Russell, we would suppose that trading in the unaccommodating state is driven by information, so that the associated price movements would tend to be permanent.

2.2.5 Profits

The HFT's profit from the aggressive order he places in period t is given by

$$\pi_t = yq_t - \alpha |q_t| \quad (2)$$

$$= \varphi \mathbb{I}\{\Lambda = U\} q_t - \alpha |q_t| \quad (3)$$

Where $\mathbb{I}\{\Lambda = U\}$ is an indicator variable that equals 1 when $\Lambda = U$, and 0 when $\Lambda = A$. (See Section 4.1 for discussion of why this specification for profits is reasonable.)

Denote the HFT's total combined profits from periods 1 and 2 by

$$\pi_{total} := \pi_1 + \pi_2 \quad (4)$$

The HFT is risk-neutral and seeks to maximize the expectation of π_{total} .

Note that because $\mathbb{E}[y|\Lambda = U] = \mathbb{E}[\varphi] = 0$ and $\mathbb{E}[y|\Lambda = A] = 0$, the unconditional expectation of y is zero, as is the period-1 expectation of y . In expectation, the HFT will therefore lose money on any aggressive order he places in the first period.

2.2.6 Model time-line

Period 1 In period 1, the HFT has the opportunity to submit an aggressive order and then observe any subsequent change in resting depth. The HFT cannot observe the liquidity state directly, but he can infer the value of Λ from changes in resting depth if he places an aggressive order. Specifically, the HFT can conclude that $\Lambda = U$ if resting depth further depletes following his order, and that $\Lambda = A$ otherwise. If the HFT does not place an aggressive order in period 1, he does not learn Λ .

Period 2 At the start of period 2, the HFT observes the signal of future aggressive order-flow, φ . The HFT observes φ regardless of whether he placed an aggressive order in period 1 (this reflects the idea that aggressive order-flow is easy to predict on the basis of public market data). After the HFT observes φ , he once again has an opportunity to place an aggressive order. Finally, after the HFT has the chance to trade, aggressive order-flow characterized by φ arrives, then prices change

as determined by φ and Λ in equation (1).

2.3 Analysis of the model

By design, the model is not subtle, and determining the HFT's optimal strategy is straightforward (Appendix A contains full mathematical details). The HFT faces a trade-off between the direct trading costs of placing an exploratory order, and the informational gains from exploration.⁸ By placing a (costly) aggressive order in period 1, the HFT “buys” the perturbation needed to elicit a response in resting depth that reveals the liquidity state. Knowing the liquidity state enables the HFT, in period 2, to better determine whether he would profit by trading ahead of predictable aggressive order flow. Despite its simplicity, the model delivers testable implications of the hypothesis that a given trader engages in exploration.

2.3.1 Order-sizes and conditions for exploratory trading

For the HFT to weakly prefer to engage in period-1 exploratory trading with order-size $|q_1| \geq 1$, the expected gains from knowing the liquidity state in period two (which work out to equal $|q_2|vu(1 - \alpha)$) must weakly exceed the expected losses on the exploratory order itself (given by $-\alpha|q_1|$). In other words, a necessary condition for exploratory trading to occur in the model is

$$|q_2|vu(1 - \alpha) \geq \alpha|q_1| \tag{5}$$

$$\iff |q_2| \geq \frac{\alpha}{1 - \alpha} \left(\frac{1}{vu} \right) |q_1| \tag{6}$$

$$\Rightarrow |q_2| > |q_1| \tag{7}$$

where the final strict inequality follows from the assumptions that $\alpha \geq 0.5$, and $u < 1$. This “small exploratory order/large follow-up order” pattern arises because the exploratory orders are always costly in expectation, while the resulting exploratory information is only valuable when there is predictable aggressive order-flow in the next period (i.e., when $\varphi \neq 0$). The per-contract losses on exploratory orders will therefore be greater in magnitude than the per-contract profits

⁸Parameters of the model determine the relative costs and payoffs of exploration. I derive routine comparative statics in Appendix A, but because exogenous variation in N , v , α , or u is scarce, these comparative statics provide little in the way of testable implications.

on follow-up orders, so the total profits on follow-up orders will only exceed the total losses on exploratory orders if the follow-up orders are larger.

As a stand-alone result, this order-size pattern isn't particularly interesting or distinctive, but it plays an important supporting role in the empirics to follow.

2.3.2 Testable predictions

Assuming for the moment that candidate exploratory orders can be distinguished from the trader's other orders (I address this issue in Section 2.5), the model generates two key testable implications of the hypothesis that a given trader engages in exploration.

First, the model predicts that the market response following an exploratory order helps to forecast whether or not the explorer will place a follow-up order. The trader will not place a follow-up order if $\Lambda = A$, while he will place a follow-up order if $\Lambda = U$ and $\varphi \neq 0$. As noted in Section 2.3.1, the follow-up orders must tend to be larger than the exploratory orders, so the model implies that, holding fixed φ , the incidence of the trader's large aggressive orders will be higher when $\Lambda = U$ than when $\Lambda = A$, if the trader engages in exploration. In other words, *if a given trader engages in exploration, then the market response to his exploratory orders should help to explain the incidence of his larger aggressive orders (holding fixed the expectation of future aggressive order-flow analogous to φ).*

Next, the market response following an exploratory order also helps to forecast whether or not prices will change soon thereafter, according to equation (1). Because both the price-change and an exploring trader's decision to place a follow-up order (in the direction of an imminent price-change) will depend on Λ , the model implies that *if a given trader engages in exploration, then equation (1) will explain his earnings better than φ alone.* In other words, *the market response to his exploratory orders should help to explain his earnings on subsequent aggressive orders.*

2.4 Private knowledge from exploratory trading

Exploratory trading hinges on the fact that the HFT will learn more if he places an exploratory order than he would learn otherwise. In the model introduced earlier, the HFT could only observe a market response if he placed an exploratory order, since no one else submitted aggressive orders

in the first period. In reality though, other traders place aggressive orders all the time, so an HFT can observe a market response to an aggressive order (placed by someone else) even if he doesn't place an exploratory order himself. Because exploratory orders are costly in expectation, a necessary condition for exploratory trading—and a testable prediction of the model—is that *the HFT learns more from the market response to his own exploratory orders than he does from the market response to aggressive orders placed by other traders*. In an anonymous market (such as the E-mini), we obtain by symmetry the related prediction that *each other trader obtains no more useful information from the market response to the exploring trader's aggressive orders than they would from the market response to another, arbitrary trader's aggressive orders*.

The testable predictions above do not depend on the particular reason why the HFT learns more from the market response to his own orders than from the response to others' orders, but as mentioned in the introduction, there is a natural explanation for this. Sometimes, the changes in the order book following the arrival of an aggressive order are truly a response caused by the aggressive order, in which case the order book activity provides information about the liquidity state. Often, though, both the aggressive order and the subsequent order book activity are really just common responses to some third event, so there is no causal link between the aggressive order and the subsequent order book changes, and consequently the order book activity does not provide information about the liquidity state. If someone else placed the aggressive order, these two scenarios are indistinguishable to the HFT, so the possibility that he is observing the uninformative non-causal scenario attenuates the amount that he can learn from the market response to someone else's aggressive order. By contrast, if the HFT places an aggressive order himself, he can be entirely sure whether he did so for exogenous reasons, so the uninformative scenario need not be a concern. The HFT learns more about the liquidity state from his own aggressive orders than he does from those of traders because he can better infer causal effects from aggressive orders that he himself placed. (For completeness, in Appendix A, I formalize the arguments above using a variation of the baseline model.)

Although they would not be consistent with exploratory trading, there are possible scenarios in which an HFT might learn only as much, or perhaps even less, from the response to his own orders as he would from the response to others' orders. Whether an HFT truly learns more from

the market response to his own orders is an empirical question, and indeed, this is one of the questions that the empirical analysis in Section 5 and Section 6 helps to address.

2.5 Empirical agenda

Before attempting any empirical evaluation of the hypothesis that the A-HFTs engage in exploratory trading, suitable candidates for putative exploratory orders must be identified in some manner among the A-HFTs' aggressive orders. The results from Section 2.3.1 suggest that small, unprofitable aggressive orders are prime candidates. Empirical results presented in Section 4.3 indicate that the A-HFTs tend to lose money on their smallest aggressive orders, so I test the model's predictions under the assumption that all of the A-HFTs' small aggressive orders are exploratory. Given the myriad other reasons for which an A-HFT might place small aggressive orders, the assumption is conservative. The high probability that some of the orders are *not* exploratory only strengthens my results.

With that preliminary matter resolved, I turn to direct empirical tests of the model's key predictions. As a benchmark, I consider the market response following the last small aggressive order placed by anyone, which is public information. The empirical implications discussed earlier in this section can then be condensed into three central predictions, namely that relative to the public-information benchmark, information from the market response following an A-HFT's small aggressive orders:

PREDICTION.1 EXPLAINS A SIGNIFICANT ADDITIONAL COMPONENT OF THAT A-HFT'S EARNINGS ON
SUBSEQUENT AGGRESSIVE ORDERS, BUT

PREDICTION.2 DOES NOT EXPLAIN ANY ADDITIONAL COMPONENT OF OTHER TRADERS' EARNINGS ON
SUBSEQUENT AGGRESSIVE ORDERS, AND

PREDICTION.3 FURTHER EXPLAINS BY A SIGNIFICANT MARGIN THE INCIDENCE OF THAT A-HFT'S SUB-
SEQUENT LARGE AGGRESSIVE ORDERS

In Section 5, I introduce an explicit numeric measure of "market response," and in Section 5.3, I make precise the notion of "explaining earnings on subsequent aggressive orders," then I formally test the predictions above. The variables and functional forms used in these empirical tests follow closely from the structure of the baseline model and the predictions highlighted in Section 2.3.2.

3 High-frequency trading in the E-mini market

The E-mini S&P 500 futures contract is a cash-settled instrument with a notional value equal to \$50.00 times the S&P 500 index. Prices are quoted in terms of the S&P 500 index, at minimum increments, “ticks”, of 0.25 index points, equivalent to \$12.50 per contract. E-mini contracts are created directly by buyers and sellers, so the quantity of outstanding contracts is potentially unlimited.

All E-mini contracts trade exclusively on the CME Globex electronic trading platform, in an order-driven market. Transaction prices/quantities and changes in aggregate depth at individual price levels in the order book are observable through a public market-data feed, but the E-mini market provides full anonymity, so the identities of the traders responsible for these events are not released. Limit orders in the E-mini market are matched according to strict price and time priority; a buy (sell) limit order at a given price executes ahead of all buy (sell) limit orders at lower (higher) prices, and buy (sell) limit orders at the same price execute in the sequence that they arrived. Certain modifications to a limit order, most notably size increases, reset the time-stamp by which time-priority is determined.

E-mini contracts with expiration dates in the five nearest months of the March quarterly cycle (March, June, September, December) are listed for trading, but activity typically concentrates in the contract with the nearest expiration. Aside from brief maintenance periods, the E-mini market is open 24 hours a day, though most activity occurs during “regular trading hours,” namely, weekdays between 8:30 a.m. and 3:15 p.m. CT.

3.1 Description of the data

The data are account-labeled, millisecond-timestamped records at the CFTC of the so-called “business messages” entered into the Globex system between September 17, 2010 and November 1, 2010 for all E-mini S&P 500 futures contracts. These message records capture not only transactions, but also events that do not directly result in a trade, such as the entry, cancellation, or modification of a resting limit order. Essentially, business messages include any action by a market participant that could potentially result in or affect a transaction immediately, or at any point in

the future.⁹

I restrict attention to the December-expiring E-mini contract, ticker ESZ0. During my sample period, ESZ0 activity accounted for roughly 98% of the message volume across all E-mini contracts, and more than 99.9% of the trading volume. Trading volume in ESZ0 by the HFTs that I study is roughly 500 times greater than the total trading volume (by all traders) in all E-mini contracts other than ESZ0 combined, so cross-contract arbitrage is a negligible issue for my empirical analysis.

The price of an ESZ0 contract during the sample period was around \$55,000 to \$60,000, and (one-sided) trading volume averaged 1,991,252 contracts or approximately \$115 billion per day. Message volume averaged approximately 5 million business messages per day, and the number of aggressive orders executed per day averaged 132,127. The intensity of trading varies considerably throughout the day (aggressive orders typically arrive in tight clusters), so the median time interval between aggressive orders during regular trading hours is closer to 20 milliseconds than it is to the mean interval of roughly 200 milliseconds.

3.2 Defining “high-frequency trader”

Kirilenko *et al.* identify as HFTs those traders who exhibit minimal accumulation of directional positions, high inventory turnover, and high levels of trading activity. I, too, use these three characteristics to define and identify HFTs. To quantify an account’s accumulation of directional positions, I consider the magnitude of changes in end-of-day net position as a percentage of the account’s daily trading volume. Similarly, I use an account’s maximal intraday change in net position, relative to daily volume, to measure inventory turnover. Finally, I use an account’s total trading volume as a measure of trading activity.

I select each account whose end-of-day net position changes by less than 6% of its daily volume, and whose maximal intraday net position changes are less than 20% of its daily volume. I rank the selected accounts by total trading volume, and classify the top 30 accounts as HFTs. The original classifications of Kirilenko *et al.* and Baron *et al.* guided the rough threshold choices for inter-day

⁹Excluded from these data are purely administrative messages, such as log-on and log-out messages. The good-til-cancel orders in the order book at the start of September 2, and a small number of modification messages (around 2 – 4%) are also missing from these records. Because I restrict attention to aggressive orders, and I only look at *changes* in resting depth (rather than its actual level), my results are not sensitive to these omitted messages.

and intraday variation. Thereafter, since confidentiality protocols prohibit disclosing results for groups smaller than eight trading accounts, the precise cutoff values of 6%, 20%, and 30 accounts were chosen to ensure that all groups of interest would have at least eight members. My central results are not sensitive to values of these parameters.

Changing the 30-account cutoff to (e.g.) 15 accounts or 60 accounts does not substantially alter my results, because activity heavily concentrates among the largest HFTs. For example, the combined total trading volume of the 8 largest HFTs exceeds that of HFTs 9-30 by roughly three-quarters, and the combined aggressive volume of the 8 largest HFTs exceeds that of HFTs 9-30 by a factor of almost 2.5. The set of HFTs corresponds closely to the set of accounts with the greatest trading volume in my sample, so the set of HFTs is largely invariant both to the exact characterizations of inter-day and intraday variation in net position relative to volume, and to the exact cutoff values for these quantities. The 6% and 20% cutoffs are not remotely binding for the HFTs with the greatest trading volumes.

3.3 HFTs' prominence and profitability

Although HFTs constitute less than 0.1% of the 41,778 accounts that traded the ESZ0 contract between September 17, 2010 and November 1, 2010, they participate in 46.7% of the total trading volume during this period. In addition to trading volume, HFTs are responsible for a large fraction of message volume. During the sample period, HFTs account for 31.9% of all order entry, order modification and order cancellation messages. In aggregate, approximately 48.5% of HFTs' volume is aggressive, and this figure rises to 54.2% among the 12 largest HFTs. The HFTs also appear to earn large and stable profits. Gross of trading fees, the 30 HFTs earned a combined average of \$1.51 million per trading day during the sample period. Individual HFTs' annualized Sharpe ratios are in the neighborhood of 10 to 11.¹⁰

¹⁰These average trading profits reflect the total cumulative trading profits during my sample period, divided by the number of trading days. Total cumulative trading profits are computed using all transactions over the full course of my sample period, plus the marked-to-market value of each HFTs' net inventory position at the end of my sample period, minus the initial marked-to-market value of each HFTs' net inventory position at the start of my sample period. Positive (negative) initial inventory values are marked to market at the initial best ask (bid), while positive (negative) final inventory values are marked to market at the final best bid (ask); this yields more conservative estimates than marking to market at midpoint prices. Empirically, the initial/final net inventory values are tiny relative to the full cumulative profits from transactions.

Using the same methodology, I compute trading profits for each trading day in my sample period (for each individual HFT), and I calculate the standard deviation of those daily trading profits (for each individual HFT). I

The Chicago Mercantile Exchange reduces E-mini trading fees on a tiered basis for traders whose average monthly volume exceeds various thresholds. Trading and clearing fees were either \$0.095 per contract or \$0.12 per contract for the 20 largest HFTs, and were at most \$0.16 per contract for the remaining HFTs. Initial and maintenance margins were both \$4,500 per contract for all of the HFTs.

Hereafter, unless otherwise noted, I restrict attention to activity that occurred during regular trading hours. HFTs' aggressive trading occurs almost exclusively during regular trading hours (approximately 95.6%, by volume), and market conditions during these times differ substantially from those during the complementary off-hours.

4 HFTs' trading profits on aggressive orders

4.1 Measuring aggressive order profitability

The familiar “bookkeeping” approach used for computing trading profits in the preceding section is not suitable for measuring trading profits on individual aggressive orders because it inevitably commingles earnings from multiple orders. A more suitable general approach, fairly standard in the literature, involves examining the cumulative price change following an aggressive order, normalized by the order's direction (+1 for a buy, or -1 for a sell). Intuitively, the average expected trading profit from an aggressive order equals the expected favorable price movement, minus trading/clearing fees and half the bid-ask spread. This approach is particularly well-suited to estimating the profitability of individual HFTs' aggressive orders in the E-mini, both because the number of aggressive-order observations for each HFT is large, and because the bid-ask spread in the E-mini is essentially constant.

To obtain meaningful estimates, we must accumulate the price-changes following an HFT's aggressive order out to some time past the maximum horizon at which the HFT can predict price-movements. We can find a suitable accumulation period empirically by calculating cumulative direction-normalized price changes over longer and longer windows until their mean ceases to significantly change; beyond that point, the HFT displays no significant capacity to forecast addi-

divide average daily profits by this standard deviation, to get the Sharpe ratio for my sample period, then multiply that Sharpe ratio by $\frac{251}{32}$, i.e. $\frac{\text{total \# trading days in a year}}{\text{\# trading days in my sample}}$, to obtain the annualized Sharpe ratio.

tional price-changes. Using too long an accumulation period introduces extra noise, but it will not bias the estimates. I find that an accumulation period, measured in event-time, of 30 aggressive order arrivals is sufficient to obtain unbiased estimates, but for all of the empirical work in this paper I use an accumulation period of 50 aggressive order arrivals to allow a wide margin for error. Estimates do not significantly differ using an accumulation period of 200, or even 500 aggressive order arrivals instead of 50, so we can reasonably interpret the estimated price-movements to be permanent. See Internet Appendix B for further details.

As noted earlier, the bid-ask spread for the E-mini is almost constantly \$12.50 (one tick) during regular trading hours, and the HFTs in my sample face trading/clearing fees of \$0.095 to \$0.16 per contract, so the average favorable price movement necessary for an HFT's aggressive order to be profitable is between \$6.345 and \$6.41 per contract. Since trading/clearing fees vary across traders, I report aggressive order performance in terms of favorable price movement, that is, earnings gross of fees and the half-spread.

4.2 HFTs' overall profits from aggressive orders in the E-mini

To measure the overall mean profitability of a given account's aggressive trading, I compute the average cumulative price change following each aggressive order placed by that account, weighted by executed quantity and normalized by the direction of the aggressive order. As a group, the 30 HFTs in my sample achieve average aggressive order performance of \$7.01 per contract. On an individual basis, nine HFT accounts exceed the relevant $\$6.25 + \text{fees}$ profitability hurdle, and each of these nine accounts exceeds this hurdle by a margin that is statistically significant at the 0.05 level. One of these nine accounts is linked with another HFT account, and their combined average performance also significantly exceeds the profitability hurdle.

Overall, the HFTs vastly outperform non-HFTs, who earn a gross average of \$3.19 per aggressively-traded contract. However, these overall averages potentially confound effects of very coarse differences in the times at which traders place aggressive orders with effects of the finer differences more directly related to strategic choices. For example, if all aggressive orders were more profitable between 1 p.m. and 2 p.m. than at other times, and HFTs only placed aggressive orders during this window, the HFTs' out-performance would not depend on anything characteristically

high-frequency.

To control for potential low-frequency confounds, I divide each trading day in my sample into 90-second segments and regress the profitability of non-HFTs' aggressive orders during each segment on both a constant and the executed quantities of the aggressive orders. Using these local coefficients, I compute the profitability of each aggressive order by an HFT in excess of the expected profitability of a non-HFT aggressive order of the same size during the relevant 90-second segment. With these additional controls, only 27 HFT accounts continue to exhibit significant out-performance of non-HFTs, and only eight of the 27 accounts are among those whose absolute performance exceeded the profitability hurdle.

4.2.1 A-HFTs and B-HFTs

For expositional ease, I will refer to the eight HFT accounts that make money on their aggressive trades *and* outperform the time-varying non-HFT benchmark as "A-HFTs," and to the complementary set of HFTs as "B-HFTs." The eight A-HFTs have a combined average daily trading volume of 982,988 contracts, and on average, 59.2% of this volume is aggressive. The 22 B-HFTs have a combined average daily trading volume of 828,924 contracts, of which 35.9% is aggressive. Together, the eight A-HFTs place a daily average of 8,994 aggressive orders (during regular trading hours), with a mean size of 60.3 contracts and a median size of 10 contracts. The 22 B-HFTs together place an average of 31,113 aggressive orders per day (during regular trading hours), with a mean size of 8.3 contracts and a median size of 1 contract. Gross of fees, the A-HFTs earn a combined average of \$793,342 per day, or an individual average of \$99,168 per day, while the B-HFTs earn a combined average of \$715,167 per day, or an individual average of \$32,508 per day.¹¹ The highest profitability hurdle among the A-HFTs is \$6.37 per aggressively traded contract.

4.3 Identifying some potential exploratory orders

As noted in Section 2.5, to test the empirical predictions of the exploratory trading model, we must first specify some orders to treat as exploratory. Motivated by theory, I examine the A-HFTs' small aggressive orders. To make precise the meaning of "small" aggressive order, I specify a cutoff

¹¹The preceding descriptive statistics include the small amount of trading activity that occurred outside regular trading hours, except where noted otherwise.

Table 1: Summary Statistics for A-HFTs’ Small Aggressive Orders

Cutoff	Dollars Earned per Contract (95% CI)		AOs Below Cutoff % of All AOs	AOs Below Cutoff % of Aggr. Volume
	AOs Below Cutoff Size	AOs Above Cutoff Size		
1	(3.78, 3.89)	(7.59, 7.74)	24.31%	0.40%
5	(4.17, 4.29)	(7.62, 7.78)	43.74%	1.44%
10	(3.42, 3.55)	(7.71, 7.85)	54.64%	3.09%
15	(3.79, 3.92)	(7.71, 7.86)	56.75%	3.54%
20	(4.08, 4.20)	(7.75, 7.90)	60.82%	4.80%

Table 1 presents descriptive statistics about aggressive orders of varying sizes placed by the A-HFTs. The “Cutoff” column indicates the maximum size for orders included in the “below cutoff” statistics. All orders of sizes exceeding the specified cutoff are included in calculations of the “above cutoff” statistics. Columns two and three present 95% confidence intervals for the average gross earnings per contract (in dollars) among aggressive orders in the indicated size division. Column four reports the percentage of all aggressive orders placed by the A-HFTs with an order-size no greater than the indicated cutoff. Column five reports the A-HFTs’ aggressive volume from orders of size no greater than the indicated cutoff, as a percentage of the A-HFTs’ total aggressive volume.

for order-size and define an order to be “small” if and only if its size is no greater than the specified cutoff. Because there is no natural unique cutoff, I consider a range of different order-size cutoffs. The A-HFTs tend to profit on aggressive orders above size 20, so I consider various order-size cutoffs between 1 and 20.¹² Table 1 summarizes some important characteristics of the A-HFTs’ small aggressive orders.

A given A-HFT places an aggressive order of size 20 or less roughly once every 34 seconds on average, and this average interval drops to about 3 seconds during periods of intense market activity. Furthermore, A-HFTs’ aggressive orders of size 20 or less tend to lose money on average. Because the A-HFTs’ small aggressive orders tend to be unprofitable and to arrive rather frequently, they are at least plausible candidates to be exploratory orders. Clearly, these results don’t provide any compelling evidence of exploratory trading; orders with analogous features could also arise from

¹²Presenting statistics for every order-size cutoff between 1 and 20 could potentially reveal individually identifiable information, so I only present results for a regularly spaced subset of order-size cutoffs.

A-HFTs controlling risk, testing out new strategies, hedging, etc. However, identifying candidate exploratory orders provides the starting point for direct tests of the exploratory trading model’s sharper empirical predictions, which, it turns out, *do* provide compelling evidence of exploratory trading.

5 Testing the exploratory trading hypothesis

If the A-HFTs indeed engage in exploration using their small aggressive orders, the exploratory trading model generates the testable predictions presented in Section 2. The present section introduces empirical analogues of the two quantities from the exploratory trading model that appear in the model’s predictions: the market-response that reveals the liquidity state (Λ), and the signal of future aggressive order-flow, φ . I use these to directly test the predictions from Section 2.

The model’s first two predictions concern the explanatory power of market-response information for the earnings of subsequent aggressive orders, and I test these two predictions in the same empirical framework. The third prediction, concerning the incidence of large aggressive orders, requires a slightly different empirical approach, so I consider this prediction separately. I estimate results for the A-HFTs individually, but for compliance with confidentiality protocols, I present cross-sectional averages of these estimates. Empirically, these average results are representative of the results for individual A-HFTs.¹³

5.1 A simple measure of market response

The predictions in Section 2 involve the market response to a given A-HFT’s exploratory orders, which we have conjectured, for the purposes of testing, to be the A-HFT’s small aggressive orders. To make this precise, define an aggressive order to be “small” if that order’s submitted size is less than or equal to a specified size parameter, which I will denote by \bar{q} .

I characterize the market response to a small aggressive order using subsequent changes in

¹³Throughout the E-mini market, there exist assorted linkages between various trading accounts (as, for example, in the simple case where single firm trades with multiple accounts), so the trading-account divisions do not necessarily deliver appropriate atomic A-HFT units. Though the specifics are confidential, the appropriate partition of the A-HFTs is entirely obvious. For brevity, I use “individual A-HFT” as shorthand to “individual atomic A-HFT unit,” as applicable.

order book depth. I examine the interval starting immediately after the arrival and execution of a given small aggressive order and ending immediately before the arrival of the next aggressive order (which may or may not be small), and I sum the changes in depth at the best bid and best ask that occur during this interval.¹⁴ For symmetry, I adopt the convention that sell depth is negative and buy depth is positive. I also normalize these depth changes by the sign of the preceding small aggressive order to standardize across buy orders and sell orders.

To simplify the analysis and stack the odds against finding significant results, I initially focus only on the signs of the direction-normalized depth changes. These signs merely indicate whether or not resting depth moved further in the direction of the preceding small aggressive order—or in the language of the model, whether resting depth further depletes or weakly replenishes.

For a given value of \bar{q} , I construct the indicator variable Ω , with k th element Ω_k defined by

$$\Omega_k = \begin{cases} 1 & \text{if } DC(k; any, \bar{q}) > 0 \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

where $DC(k; any, \bar{q})$ denotes the direction-normalized depth change following the last small aggressive order (submitted by anyone) that arrived before the k th aggressive order. Similarly, I construct the indicator variable Ω^A , with k th element Ω_k^A defined by

$$\Omega_k^A = \begin{cases} 1 & \text{if } DC(k; AHFT, \bar{q}) > 0 \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

where $DC(k; AHFT, \bar{q})$ denotes the direction-normalized depth change following the last small aggressive order *submitted by a specified A-HFT* that arrived before the k th aggressive order.

Note the direct parallel between the omega variables and the binary liquidity states in the exploratory trading model.

¹⁴The best bid and best ask prices at which I measure changes in depth are the best bid and ask at start of the interval. The price levels at which changes in depth are recorded remain the same throughout an interval, even if the bid and/or ask prices move during the interval.

5.2 Order-flow signal

To test the exploratory trading theory, in addition to the measure of market response, we need something analogous to the signal of future aggressive order-flow, φ . Because we are ultimately interested in how future aggressive order-flow will affect prices, the task of finding an empirical analogue to φ simplifies to finding variables other than market-response measures that forecast price movements.

I select a handful of lagged market variables that forecast the cumulative price-change between the aggressive orders k and $k+50$, which I denote by y_k . These variables are: the signs of aggressive orders $k-1$ through $k-4$, the signed executed quantities of aggressive orders $k-1$ through $k-4$, and changes in resting depth between aggressive orders $k-1$ and k at each of the six price levels within two ticks of the best bid or best ask (with sell depth negative and buy depth positive, as before). To lighten notation, I concatenate these 14 variables in the row vector z_{k-1} . This vector, z_{k-1} , is the analogue of φ .

In the same way that price movements in the exploratory trading model can still be forecast to some extent by φ when the liquidity state is unknown, the variables in z_{k-1} should have some power to forecast y_k , even without the market-response omega variables. As a check on this and as a benchmark, I estimate the equation

$$y_k = z_{k-1}\Gamma + \epsilon_k \tag{10}$$

where Γ is a column vector of 14 coefficients. As desired, the estimated coefficients have the expected signs, and their joint significance is extremely high. I discuss the regression results directly, report coefficient estimates, and discuss the choice of explanatory variables in Internet Appendix C.

Naturally, the set of right-hand-side variables in equation (10) is not comprehensive, and many other variables can be added that could somewhat improve the price forecasts. However, the tests of the exploratory trading model's predictions do not rely on equation (10) as the means of controlling for public information, but rather rely on a different approach (described in the next section). The tests merely require that equation (10) have *some* forecasting power.

5.3 Testing predictions about explaining earnings

5.3.1 Explained earnings

The first testable prediction of the exploratory trading model is that the information from the market response following a given A-HFT's small aggressive orders will explain a significant additional component of that A-HFT's earnings on subsequent (large) aggressive orders, beyond what is explained by a public-information benchmark.

In this paper, the particular notion of “explaining earnings” that I employ involves computing what a trader's earnings are expected to be on the basis of some econometric forecast of price movements, and comparing that with the trader's actual earnings. For concreteness, consider a price forecast based on equation (10). Letting $\hat{\Gamma}$ denote the estimate of Γ , we have

$$\hat{y}_k = z_{k-1} \hat{\Gamma}$$

Given the sign of the k^{th} aggressive order, we can compute the forecast earnings on that order, conditional on the order's sign. Much as the direction-normalized cumulative price-change $sign_k * y_k$ provides an estimate of the true earnings on aggressive order k (see Section 4.1), the direction-normalized forecast cumulative price-change $sign_k * \hat{y}_k$ provides an estimate of the forecast earnings on aggressive order k .

Rather than working with the earnings on order k that are explained by a given econometric price forecast, it is convenient to work with the earnings on order k that are *not* explained by the specified forecast. I will refer to the earnings on order k that are *not* explained by the specified forecast as the “excess earnings on order k relative to [the specified forecast].” In the case above, the excess earnings on order k relative to the forecast from equation (10), denote it ξ_k , is given by

$$\begin{aligned} \xi_k &= sign_k * y_k - sign_k * \hat{y}_k \\ &= sign_k (y_k - \hat{y}_k) \end{aligned}$$

so ξ_k is simply the k^{th} regression residual multiplied by the sign of aggressive order k .¹⁵ The

¹⁵Because \hat{y}_k uses only information available prior to the arrival of the k^{th} aggressive order, there is no orthogonality constraint on the k^{th} regression residual and $sign_k$.

additional component of earnings on aggressive order k explained by some price forecast F , relative to some other price forecast G is given by $\xi_k^G - \xi_k^F$.

Finally, note that the all of the earnings discussed in this section are *per contract*.

5.3.2 Empirical strategy: overview

Though the implementation is slightly involved, my empirical strategy is straight-forward—it is basically just a diff-in-diffs approach. First, I augment the regression equation (10) from Section 5.2 using either:

1. Market response information from the last small aggressive order placed by anyone—i.e., Ω ,
or
2. Both market response information from the last small aggressive order placed by anyone,
and market response information from the last small aggressive order placed by a specified
A-HFT—i.e., both Ω *and* Ω^A

After estimating both of the regression specifications above, I find the additional component of earnings on larger aggressive orders explained by the second specification relative to the first one. The market response following an arbitrary small aggressive order is publicly observable. However, because the E-mini market operates anonymously, the distinction between a small aggressive order placed by a particular A-HFT and an arbitrary small aggressive order is private knowledge, available only to the A-HFT who placed the order. Because the market response information from the last small aggressive order placed by anyone is weakly more recent than the market response information from last small aggressive order placed by the A-HFT, comparing the second specification above to the first helps to isolate the effects attributable to private knowledge from effects attributable to public information.

Finally, I compare the additional explained earnings for the specified A-HFT to the additional explained earnings for all other traders. Intuitively, we want to verify that the A-HFT’s exploratory information provides extra explanatory power for the subsequent performance of the trader privy to that information (the A-HFT), but not for the performance of traders who aren’t privy to it (everyone else). Note that “everyone else” includes the A-HFTs other than the specified A-HFT.

Some A-HFT accounts and B-HFT/non-HFT accounts belong to the same firms, and various B-HFTs/non-HFTs may be either directly informed or able to make educated inferences about what one or more A-HFTs do. As a result, we should not necessarily expect exploratory information generated by an A-HFT's small orders to provide no explanatory power whatsoever for all other traders' performance. However, we should still expect the additional explanatory power for the A-HFT's performance to significantly exceed that for the other traders' performance.

Controlling for public information Comparing the second regression specification to the first one controls for the effects of most public information, but there could conceivably be some public information that is correlated with the market response to a specified A-HFT's small aggressive orders and yet uncorrelated with the market response to small aggressive orders placed by everyone else. One way to handle this concern is to compare the additional explained performance for the specified A-HFT to the additional explained performance for some other traders who use the same public information. Although trading objectives and sophistication vary widely across many participants in the E-mini market, all of the HFTs are sophisticated, profitable traders, with similar (very short) investment horizons, so it is extremely plausible that they all use very similar public data. Comparing the additional explained performance for the specified A-HFT to that for the other HFTs therefore serves as an added control for any lingering effects from public information.

5.3.3 Estimation procedure

In the model of exploratory trading presented in Section 2, exploratory information was valuable only in conjunction with information about future aggressive order flow. Following this notion, I incorporate market-response information by using the indicators Ω and Ω^A to partition the benchmark regression from Section 5.2.

Recall that Section 5.2 introduced the regression equation (10),

$$y_k = z_{k-1}\Gamma + \epsilon_k$$

where y_k denoted the cumulative price-change between the aggressive orders k and $k + 50$, and the vector z_{k-1} consisted of changes in resting depth between aggressive orders $k - 1$ and k , as

well as the signs and signed executed quantities of aggressive orders $k - 1$ through $k - 4$. Using the indicator Ω , I now partition the equation above into two pieces and estimate the equation

$$y_k = \Omega_k z_{k-1} \Gamma^a + (1 - \Omega_k) z_{k-1} \Gamma^b + \epsilon_k \quad (11)$$

Next, I use the indicator Ω^A to further partition (11), and I estimate the equation

$$\begin{aligned} y_k = & \Omega_k^A(k) \left(\Omega_k z_{k-1} \Gamma^c + (1 - \Omega_k) z_{k-1} \Gamma^d \right) + \\ & (1 - \Omega_k^A) \left(\Omega_k z_{k-1} \Gamma^e + (1 - \Omega_k) z_{k-1} \Gamma^f \right) + \epsilon_k \end{aligned} \quad (12)$$

The variables y_k and z_{k-1} denote the same quantities as before, and the Γ^j terms each represent vectors of 14 coefficients.

I estimate (11) and (12) for $\bar{q} = 1, 5, 10, 15, 20$, and for each specification I calculate the relative excess earnings of the specified A-HFT, and of all other trading accounts, on aggressive orders of size strictly greater than \bar{q} . As in Section 5.3.1, to compute the earnings of aggressive order k in excess of that explained by each regression specification, I normalize the k^{th} residual from the regression by the sign of the k^{th} aggressive order. I now also control for order-size effects directly by regressing the direction-normalized residuals (for the orders of size strictly greater than \bar{q}) on the (unsigned) executed quantities and a constant, then subtracting off the executed quantity multiplied by its estimated regression coefficient. Controlling for size effects in this manner makes results more comparable for different choices of \bar{q} . Size effects can be addressed by other means with negligible impact on the final results.

For each aggressive order larger than \bar{q} placed by the A-HFT under consideration, I compute the additional component of earnings explained by (12) relative to (11) by subtracting the order's excess earnings relative to (12) from its excess earnings relative to (11); I stack these additional explained components in a vector that I denote by Ξ_A . I repeat this procedure to obtain the analogous vector for everyone else, Ξ_{ee} .

Equation (12) has more free parameters than (11), so Ξ_A and Ξ_{ee} will both have positive means. However, additional explanatory power of (12) due exclusively to the extra degrees of freedom will, in expectation, manifest equally for all traders, so the extra degrees of freedom alone

Figure 1: Additional Earnings Explained by Exploratory Information




Figure 1 displays averages (with 95% confidence intervals) of the additional earnings per contract explained by (12) beyond what is explained by (11). This difference reflects added explanatory power that arises from including in (12) information about the market response to the last small aggressive order placed by a specified A-HFT (plus the slight mechanical increase that arises from introducing extra degrees of freedom). The horizontal axis specifies the cutoff, \bar{q} , for the maximum size of order defined to be “small.” The white circles mark this average computed among orders placed by an A-HFT, and the black squares mark this average computed among orders placed by everyone else. Estimates were run and means were computed for each individual A-HFT and the corresponding “everyone else”; the displayed numbers are cross-sectional averages of the individual estimates’ means.

should not cause Ξ_A and Ξ_{ee} to differ significantly.

5.3.4 Results on explaining earnings

I initially evaluate the first two empirical predictions of the exploratory trading model by comparing the additional explained component of earnings for each A-HFT (Ξ_A) to the additional explained component of earnings for all other traders (Ξ_{ee}). Figure 1 displays the cross-sectional means of Ξ_A and Ξ_{ee} for different values of \bar{q} . To formally compare the gain in explanatory power for the A-HFTs to the gain for everyone else, I construct 95% bootstrap confidence intervals for the difference of the pooled means $Mean(\Xi_A) - Mean(\Xi_{ee})$, displayed in Figure 2. Table 2 reports the numeric values from Figures 1 and 2.

Both of the tested predictions are borne out in these results. Information about the market activity immediately following an A-HFT’s smallest aggressive orders (in the form of Ω^A) improves our ability to explain that A-HFT’s earnings on larger subsequent aggressive orders by a highly

Figure 2: Difference in Additional Earnings Explained by Exploratory Information



Figure 2 depicts the difference for A-HFTs vs. for everyone else of the additional earnings per contract explained by exploratory information. Effectively, this is just the difference of the two series displayed in Figure 1. For a given A-HFT, and the corresponding “everyone else,” I find the additional earnings per contract explained by (12) beyond what is explained by (11), and I compute the difference between the average for the A-HFT, and the average for everyone else. Estimates were run and means were computed for each individual A-HFT, and the displayed points are cross-sectional averages of these individual estimates’ means (with 95% confidence intervals). The horizontal axis specifies the cutoff, \bar{q} , for the maximum size of order defined to be “small” when estimating (12) and (11).

significant margin, relative to using only information about the activity following any small aggressive order (in the form of Ω). Furthermore, the extra component of A-HFTs’ earnings on large aggressive orders explained by using Ω^A in addition to Ω is significantly greater than the extra component explained for other traders.

I refine my empirical evaluation of the first two predictions by comparing the additional explained component of earnings for each A-HFT to the additional explained component of earnings for the other HFTs. Consistent with the notion that certain HFTs may know something about what various A-HFTs are doing, the extra component of earnings explained by using Ω^A in addition to Ω is larger for the complementary set of HFTs than it is for the broader “everyone except the A-HFT of interest” group. Nevertheless, aside from the case of $\bar{q} = 1$, the average added explanatory power for each A-HFT is still significantly greater than is that for the complementary set of HFTs, as shown in Figure 3. See Table 2 for numeric values.

Table 2: Additional Earnings Explained by Exploratory Information

(Reported in hundredths of a cent per contract, with 95% confidence intervals)

	$\bar{q} = 1$	$\bar{q} = 5$	$\bar{q} = 10$	$\bar{q} = 15$	$\bar{q} = 20$
All Others	3.4 (1.8, 4.8)	8.2 (5.5, 10.9)	7.4 (4.1, 10.4)	10.1 (7.1, 13.3)	11.3 (7.7, 14.7)
A-HFTs	17.9 (6.9, 29.6)	65.9 (48.0, 85.0)	39.7 (22.1, 57.0)	53.3 (35.5, 70.5)	62.3 (45.3, 79.9)
A-HFT vs. All Others	14.5 (3.6, 25.7)	57.7 (39.3, 76.4)	32.3 (13.9, 50.0)	43.2 (25.4, 60.4)	51.0 (33.2, 68.6)
A-HFT vs. Other HFTs	3.8 (-7.6, 15.4)	45.8 (26.9, 65.1)	23.5 (5.3, 42.0)	32.3 (13.5, 50.3)	42.6 (23.8, 61.8)

Table 2 presents various cross-sectional averages of the estimated additional gross earnings per contract on aggressive orders of submitted size greater than \bar{q} explained by regression (12) in excess of that explained by regression (11). “A-HFT vs. All Others” denotes the additional explained earnings per contract for a given A-HFT minus those for all other traders, averaged across the A-HFTs; “A-HFT vs. Other HFTs” denotes an analogous quantity. Numbers reported for the A-HFTs are averages over the estimates for individual A-HFTs. The membership of “All Others” and “Other HFTs” depends upon the particular A-HFT being excluded, and the reported numbers are averages taken across the slightly different groups corresponding to each of the individual A-HFTs. Note that “other HFTs” includes the other A-HFTs.

The first two rows correspond to the results reported in Figure 1, the third row corresponds to Figure 2, and the fourth row corresponds to Figure 3.

Figure 3: Difference from Other HFTs in Additional Earnings Explained by Exploratory Information




Figure 3 depicts the average difference in additional earnings per contract explained by exploratory information for a given A-HFT vs for all other HFTs. For a given A-HFT, and the complementary set of all other HFTs, I find the additional earnings per contract explained by (12) beyond what is explained by (11), and I compute the difference between the average for the A-HFT, and the average for all the other HFTs. Note that “other HFTs” includes the other A-HFTs. Estimates were run and means were computed for each individual A-HFT, and the displayed points are cross-sectional averages of these individual estimates’ means (with 95% confidence intervals). The horizontal axis specifies the cutoff, \bar{q} , for the maximum size of order defined to be “small” when estimating (12) and (11).

5.4 Incidence of A-HFTs’ larger aggressive orders

In this subsection I test the exploratory trading model’s third prediction, namely that the market response to a given A-HFT’s small aggressive order provides significant explanatory power for the *incidence* of that A-HFT’s subsequent large aggressive orders, above and beyond that explained using the market response to the last small aggressive order placed by anyone. Since elements of the binary Ω -operators correspond almost directly to the binary liquidity-state Λ in the exploratory trading model, the incidence prediction can be made even more precise. In particular, all else being equal, the exploratory trading model predicts that an A-HFT will have a greater tendency to place large aggressive orders when $\Omega^A = 1$ than when $\Omega^A = 0$.

5.4.1 Empirical implementation

Much as the HFT in the model from Section 2 considered the signal of future aggressive order-flow as well as the liquidity state, A-HFTs consider public market data as well as exploratory information to decide when to place large aggressive orders. The size and direction of A-HFTs’ aggressive orders depend on the same variables that forecast price movements, or equivalently

on the forecasts of price movements themselves. On average, the signed quantity of an A-HFT's aggressive order should be an increasing function of the future price-change expected on the basis of public information. In this context, the exploratory trading model predicts that the expected future price-change will have a larger effect on the signed quantity of an A-HFT's aggressive orders when $\Omega^A = 1$ than it will when $\Omega^A = 0$.

To test the exploratory trading model's prediction about the incidence of A-HFTs' larger aggressive orders, I regress the signed quantities of a given A-HFT's aggressive orders on the associated fitted values of y from equation (11), partitioned by Ω^A . In other words, for a specified A-HFT and a given value of \bar{q} , I estimate the equation

$$q_k = \beta_0 (1 - \Omega_k^A) \hat{y}_k + \beta_1 \Omega_k^A \hat{y}_k + \epsilon_k \quad (13)$$

where q_k denotes the signed submitted quantity of the A-HFT's k th aggressive order, \hat{y}_k denotes the relevant fitted value of y_k from the public-information regression (11), and Ω^A is the usual indicator function. I restrict the β coefficients to be the same across all A-HFTs. Note that the fitted value \hat{y}_k includes the public market-response information through the inclusion of Ω_k in (11), so differences between β_0 and β_1 do not arise from any public information in Ω^A .

5.4.2 Results on explaining incidence

Table 3 displays the coefficient estimates from (13) for various values of \bar{q} . A Wald test rejects the null hypothesis $\beta_0 = \beta_1$ at the 10^{-15} level for all values of \bar{q} . As the exploratory trading model predicts, holding fixed the price-change expected on the basis of public information, the average A-HFT places significantly larger aggressive orders when $\Omega^A = 1$ than when $\Omega^A = 0$.

5.5 Possible alternative explanations of the results

Although the empirical results in this section confirm the predictions of the exploratory trading theory, that does not necessarily rule out alternative explanations for those results. To the extent that a movement of resting depth in the same direction as the last aggressive order is indicative of informed trading, the empirical results at first glance appear to be potentially consistent with the story that the A-HFTs (somehow) already possess private information and they split up their

Table 3: Differential Effects of Predicted Price-Changes on A-HFT Signed Order Size

	$\bar{q} = 1$	$\bar{q} = 5$	$\bar{q} = 10$	$\bar{q} = 15$	$\bar{q} = 20$
β_0 (β_0 Standard Error)	13.35 (0.094)	13.41 (0.093)	13.42 (0.095)	13.34 (0.095)	13.23 (0.094)
β_1 (β_1 Standard Error)	15.26 (0.162)	15.11 (0.169)	14.97 (0.160)	15.10 (0.159)	15.30 (0.160)

Table 3 reports coefficient estimates from the regression of the signed quantities of the A-HFTs' aggressive orders on the fitted values of cumulative future price changes from equation (11):

$$q_k = \beta_0 \left(1 - \Omega_k^A\right) \hat{y}_k + \beta_1 \Omega_k^A \hat{y}_k + \epsilon_k$$

The \bar{q} values specify the maximum size of order defined to be “small” when computing estimates.

The indicator Ω^A is the same market-response variable used elsewhere. The k th element of Ω^A equals unity if, following the most recent small aggressive order placed by a specified A-HFT, resting depth moved in the same direction as that order; otherwise, the k th element of Ω^A equals zero. The coefficients β_0 and β_1 respectively reflect the relative sizes of the average aggressive order that an A-HFT places when $\Omega^A = 0$, versus when $\Omega^A = 1$, holding \hat{y} fixed. Differences between β_0 and β_1 can be attributed to the private component of information in Ω^A because \hat{y} already incorporates the public-information analogue of Ω^A . The exploratory trading model predicts that $\beta_1 > \beta_0$.

orders as they trade on that information. While this alternative story would be more appropriate for individual stocks than for index futures such as the E-mini, it nevertheless merits consideration. However, closer examination indeed confirms that the order-splitting story is not a viable explanation of the empirical results.

There are three main types of order-splitting to consider, two of which can be readily dismissed. First, an A-HFT might split one large order into a near-instantaneous salvo of small orders—e.g., submit 500 1-contract orders in a millisecond. If all of these child-orders were large, or if they were all small (relative to \bar{q}), then they would not show up in my results, since I look at the market response to small aggressive orders, and the explained characteristics of larger orders. A second type of order-splitting that *would* appear in my results would be a salvo of alternating/mixed large and small orders. However, if an A-HFT actually did this kind of splitting and submitted all of the orders almost instantaneously, there would be essentially no chance that the resting depth would further deplete in the minuscule intervals between the arrival of these child orders. During my sample period, latency (the amount of time required for messages to be processed and pass

back and forth between a trader and the market) was several milliseconds in the E-mini market, so any further depletion in depth during an one- or two- millisecond salvo of orders almost certainly could not be a response to those orders, and hence almost certainly would not be an indication that those orders were informed.

The final type of order-splitting to consider cannot be so easily ruled out as an explanation. An A-HFT might split orders into small and large child orders, and submit them at each at least several milliseconds apart, such that depth depletion following the small ones would indeed appear in the data. Slow order-splitting of that type, though, necessarily implies that the A-HFTs possess private information which they do not trade on as quickly as they are able. This story leads to the same fundamental implication as the exploratory trading explanation, namely that the A-HFTs obtain some of their superior information through some channel other than merely reacting to public information faster than everyone else.

The slow-order-splitting story isn't quite observationally equivalent to the exploratory trading theory. In its simplest form, the order-splitting story would imply that the market response to an A-HFT's small aggressive order will only help to explain the A-HFT's earnings on a subsequent large aggressive order *if both the small order and the large order have the same sign*. The exploratory trading theory predicts that the market response to the small order will help to explain earnings on the large order, regardless of whether the two orders have the same sign. Consistent with the exploratory trading theory, and in contradiction to the slow-order-splitting story, the results in this section do not change qualitatively when we restrict attention to the market response to an A-HFT's last small aggressive order with the opposite sign from the present order. Furthermore, as the simulated trading strategy results in the next section will suggest, the market response to an A-HFT's small aggressive order can be used to better forecast the performance of a subsequent aggressive order in either direction.

6 Practical significance of exploratory information

The empirical evidence in Section 5 provides strong support for the hypothesis that the A-HFTs engage in exploratory trading as modeled in Section 2. However, while these results suggest that exploratory trading plays some part in how A-HFTs obtain the superior information that enables

them to profit from their aggressive orders, the results tell us little about how large that part is.

Estimates of the additional component of the A-HFTs’ aggressive-order earnings directly explained by the private information in Ω^A are likely to dramatically understate the true contribution of exploratory information, for two reasons. First, Ω^A is nearly the simplest possible characterization of exploratory information. Representations of exploratory information richer than Ω^A are extremely easy to construct. For example, an obvious extension would be to consider the not only the sign, but also the *magnitude* of the direction-normalized depth change following an exploratory order. Regardless of the particular representation of exploratory information used, though, the additional explained component of A-HFTs’ profits on the aggressive orders they place is likely to understate the true gains from exploration. As the simple model from Section 2 illustrates, exploratory information is valuable in large part because it enables a trader to avoid placing unprofitable aggressive orders. However, estimates of the additional explained component of profits on A-HFTs’ aggressive orders necessarily omit the effects of such avoided losses. While this bias, if anything, makes the preceding findings of statistical significance all the more compelling, it also complicates the task of properly determining the practical importance of exploratory information.

6.1 Simulated trading strategies

To investigate the gains from exploratory information, including the gains from avoiding unprofitable aggressive orders, I examine the effects of incorporating market-response information from small aggressive orders into simulated trading strategies. The key advantage of working with these simulated trading strategies is that avoided unprofitable aggressive orders can be observed directly.

The basic trading strategy that I consider is a simple adaptation of the benchmark regression from Section 5.2. I specify a threshold value, and the strategy entails nothing more than placing an aggressive order with the same sign as \hat{y}_k whenever $|\hat{y}_k|$ exceeds that threshold. To make this strategy feasible (in the sense of using only information available before time t to determine the time- t action) I compute the forecast of the future price movement, \hat{y}_k , using the regression coefficients estimated from the previous day’s data. I incorporate market-response information into this strategy by modifying the rule for placing aggressive orders to, “place an aggressive order (with the same sign as \hat{y}_k) if and only if all three of the following conditions hold:

- $|\hat{y}_k|$ exceeds its specified threshold,
- The direction-normalized depth-change following the last small aggressive order (placed by anyone) exceeds a specified threshold, and
- The direction-normalized depth-change following the last small aggressive order *placed by an A-HFT* exceeds a (possibly different) specified threshold.”

Choosing a threshold of $-\infty$ will effectively remove any of these conditions.

Each strategy yields a set of times to place aggressive orders, and the associated direction for each order. To measure the performance of a given strategy, I compute the average profitability of the indicated orders in the usual manner, with the assumption that these aggressive orders are all of a uniform size.

Relative to A-HFTs’ losses on small aggressive orders, the additional component of A-HFTs’ profits directly explained using Ω^A is smallest when $\bar{q} = 10$, and I present results for $\bar{q} = 10$ to highlight the impact of accounting for avoided losses on estimates of the gains from exploratory information. Results for other values of \bar{q} are similar.

6.2 Three specific strategies

All three threshold parameters affect strategy performance, so to emphasize the role of market-response information, I present results with the threshold for $|\hat{y}_k|$ held fixed. Varying the threshold for $|\hat{y}_k|$ does not alter the qualitative results. In particular, it is not possible by merely raising the threshold for $|\hat{y}_k|$ to achieve the same gains in performance that arise from incorporating exploratory information. The forecast \hat{y}_k uses coefficients estimated from the previous day’s data, and these forecasts exhibit increasing bias as the z_{k-1} observations assume more extreme values.

I consider a range of threshold values for the direction-normalized depth-change following the last small aggressive order placed by anyone, but, for expositional clarity, I present results for three illustrative threshold choices for the direction-normalized depth-change following the last small aggressive order placed by an A-HFT. Specifically, I consider thresholds of $-\infty$ (no A-HFT market-response information), 0 (the same information contained in Ω^A), and 417 (the 99th percentile value). Figure 4 displays the performance of these three strategies over a range of

Figure 4: Exploratory Information Improves Performance of Simulated Strategies




Figure 4 displays the estimated average gross earnings per aggressively traded contract for the three simulated trading strategies. The starting point for all three trading strategies is a simple linear forecast of the future price-change, call it \hat{y} , using the same lagged market variables as in the baseline regression, equation (10) (i.e., signs and signed quantities of the last four aggressive orders, and one lag of the changes in resting depth at prices within two ticks of the best bid and ask). The strategies also involve the change in resting depth following the last small aggressive order placed by anyone (normalized by the direction of that order), and similarly, the depth-change following the last small aggressive order placed by any A-HFT (again, normalized by that order's direction). The trading rule is to place an aggressive order with the same sign as \hat{y} whenever $|\hat{y}|$, and both depth-changes exceed their respective specified thresholds.

The three strategies differ in the threshold value for direction-normalized depth-change following the last small aggressive order placed by an A-HFT that must be satisfied in order for the strategy to enter a trade, as labeled in the figure. The threshold value for $|\hat{y}|$ is held fixed, and the horizontal axis is the threshold value (in percentiles) for the direction-normalized depth-change following arbitrary small aggressive orders.

threshold values for the market response following arbitrary small aggressive orders.

While the performance gains from incorporating A-HFT exploratory information are obvious, an equally important feature of the results above is more subtle. The A-HFTs' average gross earnings on aggressive orders over size 10 of \$7.78 per contract are well above the peak performance of the strategy that uses only public information, but substantially below the performance of the strategy that incorporates the A-HFTs' exploratory information with the higher threshold. This is exactly the pattern that we should expect, given that the former strategy excludes information that is available to the A-HFTs and the latter strategy includes information that is not available to any individual A-HFT, so these results help to confirm the relevance and validity of this simulation methodology.

6.2.1 Gains from exploration relative to losses on exploratory orders

Although the two strategies that incorporate exploratory information from the A-HFTs' small aggressive orders outperform the strategy that does not, the orders that generated the exploratory information were costly. To compare the gains from this exploratory information to the costs of acquiring it, I first multiply the increases in per-contract earnings for the two exploratory strategies (scaled by the respective number of orders relative to the public-information strategy) by the A-HFTs' combined aggressive volume on orders over size 10.¹⁶ I then divide these calibrated gains by the A-HFTs' actual losses on aggressive orders size 10 and under.

Figure 5 displays the calibrated ratio of additional gains to losses for each exploratory simulated strategy over a range of threshold values for the market response following arbitrary small aggressive orders. Using information from the A-HFTs' exploratory orders analogous to that in Ω^A , the additional gains are roughly 15% larger than the losses on exploratory orders. Whereas the extra component of the A-HFTs' performance directly explained using Ω^A represents less than 5% of A-HFTs' losses on exploratory orders, the analogous estimated performance increases more than offset the costs of exploration once we include the gains from avoiding unprofitable aggressive orders. In the case of the strategy that employs information from the A-HFTs' exploratory orders with the higher threshold, the estimated gains from exploration exceed the costs by more than one-third.

Even after netting out the calibrated losses on exploratory orders from the better-performing exploratory-information simulated strategy in 6.2, the simulated performance exceeds the maximum profitability hurdle among HFTs of \$6.41 per aggressively traded contract. An almost trivial trading strategy that incorporates exploratory trading appears to be profitable, suggesting very strongly that exploratory trading is at least sufficient to explain how a trader in the E-mini market could predict price-movements with accuracy adequate to consistently profit on average from her aggressive orders.

¹⁶The two strategies that incorporate exploratory information select subsets of the aggressive order placement times generated by the public-information-only strategy. Although the selected orders tend to be more profitable, they are also fewer in number.

Figure 5: Gains from A-HFT Exploratory Info Relative to Losses on Exploratory Orders




Figure 5 displays the calibrated ratio of additional gains to losses for the two exploratory simulated strategies across a range of threshold values (in percentiles) for the market response following arbitrary small aggressive orders.

7 Discussion

7.1 Broader opportunities for exploratory gains from aggressive orders

The empirical analysis in the preceding sections focused on the information generated by the A-HFTs' smallest aggressive orders. While these orders were the most natural starting point for an empirical study of exploratory trading, there is no theoretical reason why these small orders should be the sole source of exploratory information. In the baseline exploratory trading model, it was only to highlight the key aspects of the model that I assumed the HFT's period-1 order was expected to lose money and served no purpose other than exploration.

In principle, even aggressive orders that an A-HFT expects to be directly profitable could produce valuable, private, exploratory information. To investigate this possibility, I repeat the analysis of Section 5.3 setting $\bar{q} = 25, 30, 35, 40, 45, 50, 60, 75, 90$. The A-HFTs' incremental aggressive orders included with each increase of \bar{q} beyond $\bar{q} = 20$ are directly profitable on average, and yet the market response following these orders still provides significantly more additional explanatory power for the A-HFTs' performance on larger aggressive orders than it provides for that of other traders. See Figure 6, and see Table 4 for numeric results.

These results have the interesting implication that the A-HFTs enjoy natural and almost inevitable economies of scale—simply by being in the market and engaging in lots of aggressive

Figure 6: Difference in Additional Earnings Explained by Exploratory Information




Figure 6 is just an extended version of Figure 2, extended to larger values of \bar{q} . It depicts the difference for A-HFTs vs. for everyone else of the additional earnings per contract explained by exploratory information. For a given A-HFT, and the corresponding “everyone else,” I find the additional earnings per contract explained by (12) beyond what is explained by (11), and I compute the difference between the average for the A-HFT, and the average for everyone else. Estimates were run for each individual A-HFT, and the displayed solid line displays cross-sectional averages of these individual estimates (with 95% confidence bands in dotted lines). The horizontal axis specifies the cutoff, \bar{q} , for the maximum size of order defined to be “small” when estimating (12) and (11).

Table 4: Additional Earnings Explained by Exploratory Information

(Reported in hundredths of a cent per contract, with 95% confidence intervals)

	$\bar{q} = 25$	$\bar{q} = 30$	$\bar{q} = 35$	$\bar{q} = 40$	$\bar{q} = 45$	$\bar{q} = 50$	$\bar{q} = 60$	$\bar{q} = 75$	$\bar{q} = 90$
All Others	17.5 (13.5, 21.3)	14.7 (10.6, 18.2)	14.3 (10.6, 18.4)	16.2 (12.0, 20.5)	16.2 (11.9, 20.5)	19.9 (14.5, 25.3)	21.8 (16.2, 27.2)	17.1 (11.2, 23.0)	12.0 (5.3, 18.7)
A-HFTs	62.4 (43.7, 80.6)	69.7 (52.0, 89.4)	73.3 (55.1, 92.0)	85.0 (64.6, 103.7)	85.0 (65.9, 105.3)	100.3 (81.0, 121.9)	118.1 (98.5, 138.1)	107.3 (86.0, 129.3)	104.0 (82.1, 124.2)
A-HFT vs. All Others	45.0 (26.2, 63.9)	55.1 (37.0, 75.0)	59.0 (40.6, 78.1)	68.8 (48.3, 87.6)	68.8 (49.2, 89.6)	80.4 (59.9, 101.9)	96.4 (75.5, 117.2)	90.2 (68.7, 112.8)	92.0 (68.5, 115.1)
A-HFT vs. Other HFTs	36.0 (17.1, 55.4)	45.4 (27.2, 66.0)	48.9 (29.9, 68.4)	58.0 (36.2, 77.9)	56.0 (35.4, 77.6)	64.3 (43.5, 87.1)	78.6 (56.2, 99.2)	75.7 (53.5, 100.3)	74.6 (49.8, 99.1)

Table 4 extends Table 2 for more values of \bar{q} . The table presents various cross-sectional averages of the estimated additional gross earnings per contract on aggressive orders of submitted size greater than \bar{q} explained by regression (12) in excess of that explained by regression (11). “A-HFT vs. All Others” denotes the additional explained earnings per contract for a given A-HFT minus those for all other traders, averaged across the A-HFTs; “A-HFT vs. Other HFTs” denotes an analogous quantity. Numbers reported for the A-HFTs are averages over the estimates for individual A-HFTs. The membership of “All Others” and “Other HFTs” depends upon the particular A-HFT being excluded, and the reported numbers are averages taken across the slightly different groups corresponding to each of the individual A-HFTs. Note that “other HFTs” includes the other A-HFTs.

The extra explanatory power of (12) reflects the contribution from the private component of information (available to the A-HFT under consideration) manifested in Ω^A , as well as the effect of the extra degrees of freedom in (12) relative to (11). Since the effect from the extra degrees of freedom is the same (in expectation) for all traders, it has no impact on the *difference* in additional explained earnings per contract between an A-HFT and all other traders, or between an A-HFT and all other HFTs.

trading, they automatically generate lots of valuable, private information. Other, more obvious economies of scale and scope likely exist for high-frequency traders (e.g., tiered trading costs, applicability of similar algorithms across different markets), but the economies of scale arising from exploratory information appear to be new. The impressive performance of the extremely simple simulated strategies in Section 6 casts doubt on the standard fallback of “intellectual capital” as a barrier to entry. Although the A-HFTs earn positive profits on average, their marginal profits need not be strictly positive, so there may be no incentive for new A-HFTs to enter. However, should the structure of the A-HFT industry indicate the existence of some barriers to entry, the A-HFTs’ apparent economies of scale could potentially act as one such barrier. Industrial organization of high-frequency trading entities is an intriguing open area for future investigation, but detailed treatment lies beyond the scope of this paper.

7.2 Speed and exploratory trading

Evidence in this paper provides empirical justification for using the exploratory trading model to draw conclusions about real-world high-frequency trading. Further analysis of the exploratory trading model reveals natural connections between exploration and two important concepts of speed. These connections in turn help to illuminate the role that the two types of speed play in high-frequency trading. More importantly, although questions of speed arise most commonly in the context of high-frequency trading, the implications below are applicable much more broadly.

7.2.1 Low latency

One common measure of trading speed is *latency*—the amount of time required for messages to be processed and pass back and forth between a trader and the market. While low-latency operation and high-frequency trading are not equivalent, minimal latency is nonetheless a hallmark of high-frequency traders. HFTs can certainly react and communicate faster than some other market participants, but analogous differences in the *relative* reaction speed of various traders long predate high-frequency trading. For a trader who can identify profitable trading opportunities, there is obvious value to possessing latency low enough to take advantage of these opportunities before they disappear. The new insight from the exploratory trading model concerns the more subtle

matter of how low latency connects to the identification of such opportunities, that is, why it might matter for latency to be low in absolute terms.

In the model of exploratory trading developed in Section 2, the HFT’s inference about the liquidity state, Λ , on the basis of market activity following his aggressive order in period 1 implicitly depends on a notion related to latency. If we suppose that random noise perturbs the order book, say according to a Poisson arrival process, then the amount of noise present in the HFT’s observation of the market response in some interval following his aggressive order will depend on the duration of that interval. The duration of this interval will depend in large part upon the rate at which market data is collected and disseminated to the HFT, that is, the “temporal resolution” of the HFT’s data. Although this temporal resolution does not directly depend on the HFT’s latency, the temporal resolution of the HFT’s market information *does* implicitly constrain how quickly the HFT can learn about market events.

The finer temporal resolution required for low-latency operation enables low-latency traders to obtain meaningful—and empirically valuable—information about the market activity immediately following their aggressive orders, and this information degrades at coarser temporal resolutions. The empirical results from Section 5.3.4 provide a concrete illustration of this effect. The changes in resting depth immediately following an arbitrary aggressive order are less useful for forecasting price movements than are the analogous changes following an A-HFT’s aggressive order, but the two can only be distinguished (by the A-HFT) in data with a sufficient level of temporal disaggregation.

7.2.2 High frequency

Exploratory trading bears a natural relationship to the practice of placing large numbers of aggressive orders—what might be considered “high-frequency trading” in the most literal sense.

Any exploratory information generated by a given aggressive order is only valuable to the extent that it can be used to improve subsequent trading performance. Because exploratory information remains relevant for only some finite period, the value of exploratory information diminishes as the average interval between a trader’s orders lengthens. The exploratory trading model readily captures this effect if we relax the simplifying assumption that the liquidity state

Λ remains the same between periods 1 and 2. Suppose that Λ evolves according to a Markov process, such that with probability τ , a second Λ is drawn in period 2 (from the same distribution as in period 1), and with probability $1 - \tau$, the original value from period 1 persists in period 2. Intuitively, τ parametrizes the length of period 1, and this length increases from zero to infinity as τ increases from zero to unity. As τ tends towards unity—i.e., as the length of period 1 increases to infinity—the liquidity state in period 1 becomes progressively less informative about the liquidity state in period 2.

As discussed in Section 7.1, both theory and empirical evidence suggest that almost any aggressive order that a trader places generates some amount of exploratory information. Consequently, as a trader places aggressive orders in greater numbers, he will gain access to greater amounts of exploratory information. Furthermore, the average time interval between a trader’s aggressive orders necessarily shrinks as the number of those orders grows, so the exploratory information produced by each order tends to become more valuable to the trader. These synergistic effects dramatically magnify the potential gains from exploratory information for traders who place large numbers of aggressive orders.

8 Conclusion

This paper presents empirical evidence that HFTs use exploratory trading to obtain part of the superior information that enables them, among other things, to profitably predict price movements. In particular, these results demonstrate that HFTs do not obtain their informational advantage purely by reacting to public information milliseconds or microseconds sooner than other traders. Speed matters immensely, but by no means does it matter exclusively.

The theory of exploratory trading introduced in this paper sheds light on a number of issues related to HFTs, but it leaves many standing questions unresolved, and indeed, it raises several new questions. For example, exploratory trading could be considered a form of costly information acquisition (albeit an unusual one) which raises at least the possibility that HFTs uniquely contribute to the process of efficient price discovery. However, unlike traditional costly information acquisition, exploratory trading does not generate information that relates directly to the traded asset’s fundamental value, but that pertains rather to unobservable aspects of market conditions

that could eventually become public, ex-post, through ordinary market interactions. Furthermore, because exploratory trading operates through the market mechanism itself, exploration exerts direct effects on the market, distinct from the subsequent effects of the information that it generates. Comprehensive analysis of the myriad theoretical and empirical aspects of such issues lies beyond the scope of this paper, but the theory and evidence presented herein provide a starting point from which to more rigorously address the market-quality implications of high-frequency trading going forward.

A Exploratory trading model details

This appendix presents the full details of solving the models of Section 2.

A.1 Solving the baseline exploratory trading model (leading case)

A.1.1 Solving for the HFT's optimal trading strategy

When $\alpha > u$, the HFT will never place an order in period 2 if he doesn't know the liquidity state, and I focus on this case initially to provide a more intuitive exposition; results are qualitatively unchanged for $u \geq \alpha$, but for completeness, I analyze the general case in the next subsection.

I solve for the HFT's optimal trading strategy via backward induction.

Period 2 If the HFT learned the liquidity state during period 1, his optimal aggressive order in period 2 will depend on the values of both φ and Λ . The HFT's optimal strategy when he knows Λ is to set $q_2 = \varphi N$ if $\Lambda = U$, and to set $q_2 = 0$ if $\Lambda = A$. Taking expectations with respect to φ and then Λ , we find

$$\begin{aligned}\mathbb{E}[\pi_2 | \Lambda \text{ known}] &= Nv(1-\alpha)^*u + 0^*(1-u) \\ &= Nvu(1-\alpha)\end{aligned}\tag{14}$$

If the HFT did not learn the liquidity state during period 1, his (constrained) optimal aggressive order in period 2 will still depend on the value of φ , but it will only depend on the *distribution* of Λ , rather than the actual value of Λ . The HFT's optimal strategy when he does not know Λ is to set $q_2 = \varphi N$ when $u \geq \alpha$, and to set $q_2 = 0$ when $\alpha > u$. I assumed for simplicity that $\alpha > u$, so

$$\mathbb{E}[\pi_2 | \Lambda \text{ unknown}] = 0\tag{15}$$

Period 1 At the start of period 1, the HFT knows neither φ nor Λ , but he faces the same trading costs (α per contract) as in period 2. Consequently, the HFT's expected direct trading profits from a period-1 aggressive order are negative, given by

$$\mathbb{E}[\pi_1] = -\alpha |q_1|\tag{16}$$

Since there is no noise in this baseline model, and the HFT learns Λ perfectly from any aggressive order that he places in the first period, we can restrict attention to the cases of $q_1 = 0$ and $|q_1| = 1$.

We obtain the following expression for the difference in the HFT's total expected profits if he sets $|q_1| = 1$ instead of $q_1 = 0$:

$$\mathbb{E}[\pi_{total}|q_1| = 1] - \mathbb{E}[\pi_{total}|q_1 = 0] = Nvu(1 - \alpha) - \alpha \quad (17)$$

The HFT engages in exploratory trading if he sets $|q_1| = 1$, and he does not engage in exploratory trading if he sets $q_1 = 0$, so equation (17) represents the expected net gain from exploration. Exploratory trading is optimal for the HFT when this expected net gain is positive.

A.2 Solving the baseline exploratory trading model (general case)

Let s_t denote the sign of q_t .

A.2.1 Solving the model: period 2

If $\varphi = 0$, the HFT's optimal choice is to not submit an aggressive order in period 2, or equivalently, to set $|q_2| = 0$. If $\varphi \neq 0$, then it is optimal for the HFT to set $s_2 = \varphi$ (unless the optimal $|q_2|$ is zero), so we only need to determine the optimal magnitude, $|q_2|$. Because π_2 is linear in $|q_2|$ when s_2 is held fixed, we can restrict attention to corner solutions (0 or N) for the optimal choice of $|q_2|$ without loss of generality. Note that if $q_2 = 0$, then $\pi_2 = 0$, regardless of the values of φ and Λ .

Suppose that the HFT sets $|q_2| = N$. Without loss of generality, assume that $s_2 = \varphi \neq 0$. The HFT's period-2 profits are given by

$$\tilde{\pi}_2 = \begin{cases} N(1 - \alpha) & \text{if } \Lambda = U \\ -N\alpha & \text{if } \Lambda = A \end{cases} \quad (18)$$

where the tilde on $\tilde{\pi}_2$ denotes the fact that the HFT's choice of q_2 does not condition on the value of Λ .

HFT does not know Λ If the HFT does not know the value of Λ , then in the case where $\varphi \neq 0$, the HFT's expected period-2 profit if he sets $|q_2| = N$ is

$$\begin{aligned}\mathbb{E}[\tilde{\pi}_2 | \varphi \neq 0, |q_2| = N] &= uN(1 - \alpha) - (1 - u)N\alpha \\ &= (u - \alpha)N\end{aligned}\tag{19}$$

Taking expectations with respect to φ , we find that the *ex-ante* expectation of $\tilde{\pi}_2$ when the HFT sets $|q_2| = N$ (and $s_2 = \varphi$) is given by

$$\mathbb{E}[\tilde{\pi}_2 | |q_2| = N] = v(u - \alpha)N\tag{20}$$

When $u - \alpha < 0$, if the HFT did not know Λ , he would set $q_2 = 0$ rather than $|q_2| = N$. Hence the *ex-ante* expectation of $\tilde{\pi}_2$ is

$$\mathbb{E}[\tilde{\pi}_2] = \max\{v(u - \alpha)N, 0\}\tag{21}$$

HFT knows Λ Next, if the HFT *does* know the value of Λ , then he will set $|q_2| = N$ (and $s_2 = \varphi$) only when $\Lambda = U$ and $\varphi \neq 0$. Denoting the HFT's period-2 profits from this strategy by $\hat{\pi}_2$, we find

$$\mathbb{E}[\hat{\pi}_2 | \varphi \neq 0] = u(1 - \alpha)N\tag{22}$$

$$= (u - \alpha)N + \alpha(1 - u)N$$

$$\mathbb{E}[\hat{\pi}_2] = vu(1 - \alpha)N\tag{23}$$

$$= v(u - \alpha)N + v\alpha(1 - u)N$$

Note that

$$\mathbb{E}[\hat{\pi}_2] > \max\{v(u - \alpha)N, 0\}\tag{24}$$

so the HFT's expected period-2 profits are strictly greater when he knows Λ than when he doesn't know Λ .

A.2.2 Solving the model: period 1

At the start of period 1, the HFT knows neither φ nor Λ , but he faces the same trading costs, α , as he does in period 2. Consequently, the HFT's expected direct trading profits from a period-1 aggressive order are negative:

$$\begin{aligned}\mathbb{E}[\pi_1|q_1] &= \mathbb{E}[|q_1|(s_1 y - \alpha)|s_1, q_1] \\ &= |q_1| s_1 \mathbb{E}[y] - \alpha |q_1| \\ &= -\alpha |q_1|\end{aligned}\tag{25}$$

The second equality relies on the assumptions that φ and Λ (and hence y) are independent of s_1 and q_1 , while the final equality uses the fact that $\mathbb{E}[y] = 0$.

Since there is no noise in this baseline model, the HFT learns Λ perfectly from any aggressive order that he places in the first period with $|q_1| \geq 1$. An aggressive order of size greater than one yields no more information about Λ than a one-contract aggressive order in this setting, but the larger aggressive order incurs additional expected losses. Thus without loss of generality, we can restrict attention to the case of $q_1 = 0$ and the case of $|q_1| = 1$.

If the HFT sets $q_1 = 0$, he neither learns Λ nor incurs any direct losses in period 1, so his total expected profits are simply

$$\begin{aligned}\mathbb{E}[\pi_{total}|q_1 = 0] &= \mathbb{E}[\tilde{\pi}_2] \\ &= \max\{v(u - \alpha)N, 0\}\end{aligned}\tag{26}$$

Alternatively, if the HFT sets $|q_1| = 1$, his total expected profits are given by

$$\begin{aligned}\mathbb{E}[\pi_{total}| |q_1| = 1] &= -\alpha |q_1| + \mathbb{E}[\hat{\pi}_2] \\ &= vu(1 - \alpha)N - \alpha\end{aligned}\tag{27}$$

A.2.3 Comparative statics for model parameters

Recall that when the exogenous aggressive order-flow is described by $\varphi = 0$, the HFT does not have any profitable period-2 trading opportunities in either liquidity state. The probability that $\varphi \neq 0$, given by the parameter v , represents the extent to which the exogenous aggressive order-flow is predictable. To characterize how various parameters affect the viability of exploratory trading, I consider the minimal value of v for which the HFT finds it optimal to engage in period-1 (i.e., exploratory) trading. Denoting this minimal value by \underline{v} , we have

$$\underline{v} = \left(\frac{\alpha}{u}\right) \frac{1}{(1-\alpha)N} \quad (28)$$

The closer is \underline{v} to 0, the more conducive are conditions to exploratory trading, and by inspection, $\frac{\partial \underline{v}}{\partial \alpha} > 0$, $\frac{\partial \underline{v}}{\partial N} < 0$ and $\frac{\partial \underline{v}}{\partial u} < 0$.

The above results are intuitive. First, higher trading costs (α) tend to discourage exploratory trading. Second, when the HFT can use exploratory information to guide larger orders, the gains from exploration are magnified, so larger values of N tend to promote exploratory trading. Finally exploratory trading becomes less viable when u is smaller. The HFT will take the same action in period 2 when he knows that $\Lambda = A$ as when he doesn't know Λ , so when u is small, knowledge of the liquidity state is less valuable because it is less likely to change the HFT's period-2 actions.¹⁷

A.3 Solving the model of Section 2.4

A.3.1 Formalizing the intuitive argument

To make more rigorous the intuitive explanation of why the HFT could learn from the market response to his own orders than he could from the market response to an order placed by someone else, consider a variant of the baseline model from Section 2.2, in which now someone other than the HFT places an aggressive order at the beginning of period 1. With probability ρ , this aggressive order is the result of an unobservable informational shock, and resting depth further depletes following the order, regardless of the liquidity state Λ . Otherwise (with probability $1 - \rho$) resting

¹⁷When $u > \alpha$, the HFT will take the same action in period 2 when he knows that $\Lambda = U$ as when he doesn't know Λ , so knowledge of the liquidity state is less likely to change the HFT's period-2 actions when u is large. In the case of $u > \alpha$, equation (28) becomes $\underline{v} = \frac{1}{(1-u)N}$, and exploratory trading indeed becomes less viable as u approaches 1.

depth further depletes after the order if and only if the liquidity state is unaccommodating. Aside from this new aggressive order, all other aspects of the baseline model remain unchanged.

If the HFT places an aggressive order in period 1, his expected total profits are the same as they were in the baseline model, i.e.,

$$\mathbb{E}[\pi_{total} | |q_1| = 1] = Nvu(1 - \alpha) - \alpha \quad (29)$$

However, the HFT's expected profits if he does not place an order in period 1 are now higher than they were in the baseline model, because the HFT learns something from the depth changes following the other trader's aggressive order. If resting depth weakly replenishes after that order, the HFT learns with certainty that the liquidity state is accommodating (i.e., $\Lambda = A$), so the HFT will not submit an aggressive order in period 2, and his total profits will be zero. Alternatively, if resting depth further depletes following the aggressive order in period 1 (denote this event by g_1), we have

$$\begin{aligned} \mathbb{P}\{\Lambda = U | g_1\} &= \frac{\mathbb{P}\{\Lambda = U, \text{ and } g_1\}}{\mathbb{P}\{g_1\}} & (30) \\ &= \frac{\mathbb{P}\{g_1 | \Lambda = U\} \mathbb{P}\{\Lambda = U\}}{\mathbb{P}\{g_1 | \Lambda = U\} \mathbb{P}\{\Lambda = U\} + \mathbb{P}\{g_1 | \Lambda = A\} \mathbb{P}\{\Lambda = A\}} \\ &= \frac{1 * \mathbb{P}\{\Lambda = U\}}{1 * \mathbb{P}\{\Lambda = U\} + \rho * \mathbb{P}\{\Lambda = A\}} \\ &= \frac{u}{u + \rho(1 - u)} & (31) \end{aligned}$$

The HFT's optimal strategy when he does not know Λ is to set $q_2 = \varphi N$ when $\frac{u}{u + \rho(1 - u)} \geq \alpha$, and to set $q_2 = 0$ otherwise. Taking expectations with respect to Λ and φ , we find that the HFT's *ex-ante* expected total profits in this case are given by

$$\mathbb{E}[\pi_{total} | AO \text{ by someone else}] = \max \left\{ Nv \left(\frac{u}{u + \rho(1 - u)} - \alpha \right), 0 \right\} \quad (32)$$

The features of the baseline model discussed in Section 2.3 are qualitatively unchanged in the modified version, but now the “privacy” parameter ρ also exerts an influence. In the limiting

case where the depth change following an aggressive order placed by someone else is completely uninformative to the HFT (i.e., $\rho = 1$), equation (32) collapses down to equation (15) from the baseline model. At the opposite extreme, when the HFT learns the liquidity state perfectly from observing another trader's aggressive order (i.e., $\rho = 0$), the HFT's expected total profits are unambiguously lower if he places an aggressive order in period 1 himself. When the HFT can learn more about the liquidity state through mere observation, as he can when ρ is smaller, he has less incentive to incur the direct costs of exploratory trading.

References

- [1] Anat R. Admati and Paul Pfleiderer. A theory of intraday trading patterns: Volume and price variability. *Review of Financial Studies*, 1:3–40, 1988.
- [2] Matthew Baron, Jonathan Brogaard, and Andrei Kirilenko. The trading profits of high frequency traders. September 2013.
- [3] Bruno Biais, Thierry Foucault, and Sophie Moinas. Equilibrium algorithmic trading. Working Paper, October 2010.
- [4] Bruno Biais, Pierre Hillion, and Chester Spatt. An empirical analysis of the limit order book and the order flow in the paris bourse. *The Journal of Finance*, 50(5):1655–1689, 1995.
- [5] Jonathan Brogaard, Terrence Hendershott, and Ryan Riordan. High frequency trading and price discovery. April 2013.
- [6] Eric Budish, Peter Cramton, and John Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. Chicago Booth Working Paper, December 2013.
- [7] Alfonso Dufour and Robert F. Engle. Time and the price impact of a trade. *The Journal of Finance*, 55(6):2467–2498, December 2000.
- [8] David Easley and Maureen O’Hara. Time and the process of security price adjustment. *Journal of Finance*, 47:577–605, 1992.
- [9] Andrew Ellul, Craig W. Holden, Pankaj Jain, and Robert Jennings. Order dynamics: Recent evidence from the nyse. *Journal of Empirical Finance*, 14:636–661, 2007.
- [10] Robert F. Engle and Jeffrey R. Russell. Autoregressive conditional duration: A new model for irregularly spaced transaction data. *Econometrica*, 66(5):1127–1162, September 1998.
- [11] Thierry Foucault, Johan Hombert, and Ioanid Rosu. News trading and speed. Working Paper, May 2013.

- [12] Lawrence R. Glosten and Paul R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14:71–100, 1985.
- [13] Sanford J. Grossman and Merton H. Miller. Liquidity and market structure. *The Journal of Finance*, 43(3), July 1988.
- [14] Bjorn Hagstromer and Lars Norden. The diversity of high-frequency traders. *Journal of Financial Markets*, 16(4):741–770, 2013.
- [15] Joel Hasbrouck. Measuring the information content of stock trades. *The Journal of Finance*, XLVI(1), March 1991.
- [16] Boyan Jovanovic and Albert J. Menkveld. Middlemen in limit-order markets. Working Paper, November 2012.
- [17] Andrei Kirilenko, Mehrdad Samadi, Albert S. Kyle, and Tugkan Tuzun. The flash crash: The impact of high frequency trading on an electronic market. October 2010.
- [18] J.Chris Leach and Ananth N Madhavan. Intertemporal price discovery by market makers: Active versus passive learning. *Journal of Financial Intermediation*, 2(2):Pages 207–235, June 1992.
- [19] J.Chris Leach and Ananth N. Madhavan. Price experimentation and security market structure. *The Review of Financial Studies*, 6(2):pp. 375–404, 1993.
- [20] Victor H. Martinez and Ioanid Rosu. High frequency traders, news and volatility. Working Paper, March 2013.
- [21] Christine A. Parlour and Duane J. Seppi. *Handbook of Financial Intermediation and Banking*, chapter 3. Elsevier, 2008.
- [22] U.S. SEC. Findings regarding the market events of may 6, 2010, September 2010.
- [23] Hans R. Stoll. Inferring the components of the bid-ask spread: Theory and empirical tests. *The Journal of Finance*, 44(1):115–134, March 1989.

- [24] Chen Yao and Mao Ye. Tick size constraints, high frequency trading, and liquidity. Working Paper, July 2014.

Internet Appendices

B Measuring Aggressive Orders' Profitability

Calculating round-trip profits using a FIFO or LIFO approach is not a useful way to measure the profitability of individual aggressive orders. Even the most aggressive HFTs engage in some passive trading, so a FIFO/LIFO-round-trip measure would either confound aggressive trades with passive trades, or require some arbitrary assumption to distinguish between inventory acquired passively and inventory acquired aggressively (on top of the already-arbitrary assumption of FIFO or LIFO). A second, more general problem is that a measurement scheme based on inventory round-trips will always combine at least two orders (an entry and an exit), so such measurement schemes do not actually measure the profitability of *individual* aggressive orders.

In this appendix, I provide rigorous justification for the claim that the average expected profit from an aggressive order in the E-mini market equals the expected favorable price movement, minus trading/clearing fees and half the bid-ask spread. After presenting the formal proof, I discuss details of empirically estimating expected favorable price movement.

B.1 Preliminaries

Trading/clearing fees apply equally to both passively and aggressively traded E-mini contracts, so to simplify the exposition, I will initially ignore these fees. Similarly, I make the simplifying assumption that the bid-ask spread is constant, and identically equal to one tick; for the E-mini market, this assumption entails minimal loss of generality.

In the E-mini market, the profitability of individual aggressive orders can be considered in isolation from passive orders. Because E-mini contracts can be created directly by buyers and sellers, a trader's net inventory position does not constrain his ability to participate in a given trade¹⁸. As long as he can find a buyer, a trader who wishes to sell an E-mini contract can always do so, regardless of whether he has a preexisting long position. More generally, if a trader enters a position aggressively then exits it passively, he could have conducted the passive transaction even if he hadn't engaged in the preceding aggressive transaction. While a desire to dispose of passively-acquired inventory might *motivate* a trader to submit an aggressive order, the question

¹⁸The one exception would arise in the extremely rare event that a trader who did not qualify for a position-limit exemption held so many contracts (either long or short) that his inventory after the trade would exceed the position limit of 100,000 E-mini contracts. For HFTs, this minor exception can safely be ignored.

of underlying motivation is distinct from the question of whether the aggressive order was directly profitable.

B.2 Formal Argument

With these preliminaries established, I turn to the rigorous argument. Consider a trader who executes J aggressive sell orders of size one, and J aggressive buy orders of size one, for some large J . Following the remarks above, the trader's passive transactions can be ignored. Let the average direction-normalized price change after these aggressive orders be $\tilde{\vartheta} \equiv \vartheta \left(\frac{2J}{2J-1} \right)$ ticks for some ϑ that does not depend on J .

First, suppose that the trader always submits an aggressive sell after an aggressive buy, and always submits an aggressive buy after an aggressive sell. Without loss of generality, assume that the trader's first aggressive order is a buy. The trader's combined profit from all $2J$ aggressive orders is

$$\pi_{2J} = -a_1 + b_2 - a_3 + b_4 - \dots - a_{2J-1} + b_{2J} \quad (33)$$

$$= -a_1 + (a_2 - 1) - a_3 + (a_4 - 1) - \dots - a_{2J-1} + (a_{2J} - 1) \quad (34)$$

$$= -a_1 + a_2 - a_3 + a_4 - \dots - a_{2J-1} + a_{2J} - J \quad (35)$$

$$= -a_1 + (a_1 + \zeta_{b,1}) - (a_2 + \zeta_{s,2}) + (a_3 + \zeta_{b,2}) - \dots \quad (36)$$

$$\begin{aligned} & \dots - (a_{2J-2} + \zeta_{s,J}) + (a_{2J-1} + \zeta_{b,J}) - J \\ &= \sum_{i=1}^J (a_{2i-1} + \zeta_{b,i}) - \sum_{j=2}^J (a_{2j-2} + \zeta_{s,j}) - a_1 - J \end{aligned} \quad (37)$$

$$= \sum_{i=1}^J a_{2i-1} - \left(a_1 + \sum_{j=1}^{J-1} a_{2j} \right) + \sum_{i=1}^J \zeta_{b,i} - \sum_{j=2}^J \zeta_{s,j} - J \quad (38)$$

where a_k and b_k respectively denote the prevailing best ask and best bid at the time the k th aggressive order executes, $\zeta_{b,r}$ denotes the change in midpoint price following the r th aggressive buy order, and $\zeta_{s,r}$ denotes the change in midpoint price following the r th aggressive sell order. Note that $\vartheta \equiv \frac{1}{2J} \left(\sum_{r=1}^J \zeta_{b,r} + \sum_{r=1}^J (-\zeta_{s,r}) \right)$.

Next, taking expectations, we find

$$\mathbb{E}[\pi_{2J}] = \sum_{i=1}^J \mathbb{E}[a_{2i-1}] - \left(\mathbb{E}[a_1] + \sum_{j=1}^{J-1} \mathbb{E}[a_{2j}] \right) \quad (39)$$

$$+ \sum_{i=1}^J \mathbb{E}[\zeta_{b,i}] - \sum_{j=2}^J \mathbb{E}[\zeta_{s,j}] - J$$

$$= J\mathbb{E}[a_1] - \mathbb{E}[a_1] - (J-1)\mathbb{E}[a_1] + J\mathbb{E}[\tilde{\vartheta}] - (J-1)(-\mathbb{E}[\tilde{\vartheta}]) - J \quad (40)$$

$$= (2J-1)\mathbb{E}[\tilde{\vartheta}] - J \quad (41)$$

$$= (2J-1)\left(\mathbb{E}[\vartheta] \frac{2J}{2J-1}\right) - J \quad (42)$$

$$= J(2\mathbb{E}[\vartheta] - 1) \quad (43)$$

where the second equality uses the assumption that midpoint prices follow a martingale with respect to their natural filtration, together with the assumption of a constant bid-ask spread. From the final equality above, it follows immediately that the trader's average expected profit on an individual aggressive order is given by

$$\frac{1}{2J}\mathbb{E}[\pi_{2J}] = \mathbb{E}[\vartheta] - \frac{1}{2} \quad (44)$$

Finally, note that none of the calculations above relied on the assumption that the aggressive orders alternated between buys and sells (this only simplified the notation). It follows immediately from grouping together multiple aggressive orders of the same sign that the result would hold for orders of varying sizes, provided that the overall aggressive buy and aggressive sell volumes were equal.

Under the usual regularity conditions, as $J \rightarrow \infty$, $\tilde{\vartheta} \rightarrow_{A.S.} \lim_{J \rightarrow \infty} \mathbb{E}[\tilde{\vartheta}] = \mathbb{E}[\vartheta]$. \square

B.3 Obtaining Unbiased Estimates

Recall that the discussion in section 4.1 implied that we can estimate the profitability of an HFT's aggressive order using the (direction-normalized) accumulated price-changes following that aggressive order out to some time past the HFT's maximum forecasting horizon. If we choose too short an accumulation window, the resulting estimates of the long-run direction-normalized average price changes following an HFT's aggressive orders will be biased downward. This enables

us to empirically determine an adequate accumulation period by calculating cumulative direction-normalized price changes over longer and longer windows until their mean ceases to significantly increase

Market activity varies considerably in its intensity throughout a trading day, so event-time, which I measure in terms of aggressive order arrivals, provides a more uniform standard for temporal measurements than does clock-time. Empirically, an accumulation period of about 30 aggressive orders suffices to obtain unbiased estimates of the price movement following an HFT's aggressive order, but I consider results for an accumulation period of 50 aggressive orders to allow a wide margin for error. The mean direction-normalized price changes following individual HFTs' aggressive orders does not differ significantly for accumulation periods of 50, 200, or 500 aggressive orders, even if we distinguish aggressive orders by size. The same holds true for aggressive orders placed by non-HFTs. Using too long an accumulation period will not bias the estimates, but it will introduce unnecessary noise, so I opt for an accumulation period of 50 aggressive orders.

As I discuss at greater length in section C.1, future price movements are moderately predictable from past aggressive order flow and order book activity, but only at very short horizons. Of the variables that meaningfully forecast future price changes, the direction of aggressive order flow is by far the most persistent, but even its forecasting power diminishes to nonexistence for price movements more than either about 12 aggressive orders or 200 milliseconds in the future. The adequacy of a 30+ aggressive order accumulation period is entirely consistent with these results.

As a simple empirical check on the validity of direction-normalized cumulative price changes as a proxy for the profitability of aggressive orders, I use each HFT's explicit overall profits and passive trading volume, together with the profits on aggressive orders as measured by the proxy, to back out the HFT's implicit profit on each passively traded contract. The resulting estimates of HFTs' respective profits from passive transactions are all plausible from a theoretical perspective, and are comparable to non-HFTs' implicit performance on passive trades.

C Benchmark Regression Results

C.1 Variables that Forecast Price Movements

Bid-ask bounce notwithstanding, the price at which aggressive orders execute changes rather infrequently in the E-mini market. On average, only about 1 – 3% of aggressive buy (sell) orders execute at a final price different from the last price at which the previous aggressive buy (sell) order executed, and the price changes that do occur are almost completely unpredictable on the basis of past price changes. However, several other variables forecast price innovations surprisingly well.

In contrast to price innovations, the direction of aggressive order-flow in the E-mini market is extremely persistent and predictable. On average, the probability that the next aggressive order will be a buy (sell) given that the previous aggressive order was a buy (sell) is around 75%. In addition to forecasting the direction of future aggressive order-flow, the direction of past aggressive order-flow also forecasts future price innovations to statistically and economically significant extent, and forecasts based on past aggressive order signs alone are moderately improved by information about the (signed) quantities of past aggressive orders. Simple measures of recent changes in the order book offer further, yet modest, improvement in price forecasts.

The levels of resting depth in the order book, in addition to the changes in resting depth, also improve price forecasts slightly, but these stock variables cannot be reliably recovered in much of my data-set because a small number of modification messages (around 2 – 4%) are missing. These occasional missing modifications introduce only transient noise into flow variables such as changes in resting depth, but they have permanent effects on the corresponding stock variables. Fortunately, omitting resting-depth stock variables from the direct tests of the exploratory trading model’s predictions in section 5 is harmless. These tests use the explanatory variables in the benchmark regression (45) only as an empirical analogue of φ , the signal of future aggressive order-flow in the exploratory trading model. Thus the tests require only that the benchmark explanatory variables offer some predictive power, not that those variables control for all public information (I control for public information by other means, discussed in section 5).

C.2 Econometric Benchmark

For each trading day in my sample, I regress the cumulative price-change (in dollars) between the aggressive orders k and $k + 50$, denoted y_k , on lagged market variables suggested by the remarks in section C.1. Specifically, I regress y_k on the changes in resting depth between aggressive orders $k - 1$ and k at each of the six price levels within two ticks of the best bid or best ask, the signs of aggressive orders $k - 1$ through $k - 4$, and the signed executed quantities of aggressive orders $k - 1$ through $k - 4$. For symmetry, I adopt the convention that sell depth is negative and buy depth is positive, so that an increase in buy depth has the same sign as a decrease in sell depth. Denoting the row vector of the 14 regressors by z_{k-1} , and a column vector of 14 coefficients by Γ , I estimate the equation

$$y_k = z_{k-1}\Gamma + \epsilon_k \quad (45)$$

$$\begin{aligned} &:= \gamma_1 d_{k-1}^1 + \dots + \gamma_6 d_{k-1}^6 + \\ &\quad \gamma_7 \text{sign}_{k-1} + \dots + \gamma_{10} \text{sign}_{k-4} + \\ &\quad \gamma_{11} q_{k-1} + \dots + \gamma_{14} q_{k-4} + \epsilon_k \end{aligned} \quad (46)$$

Table 3 summarizes the estimates from the regression above, computed over my entire sample. All of the variables are antisymmetrical for buys and sells, and so have means extremely close to zero, but the mean magnitudes in the rightmost column of Table 3 provide some context for scale.

Table 5: Estimates from Benchmark Regression

	Coefficient ($\times 1000$)	Robust t-Statistic	Variable Avg. Magnitude
$d_{k-1}^{best\ bid-2}$	-0.90	-1.02	4.13
$d_{k-1}^{best\ bid-1}$	-2.08	-4.29	10.8
$d_{k-1}^{best\ bid}$	1.13	4.94	23.1
$d_{k-1}^{best\ ask}$	1.11	4.97	23.4
$d_{k-1}^{best\ ask+1}$	-2.03	-4.24	11.2
$d_{k-1}^{best\ ask+2}$	-1.60	-1.90	4.44
$sign_{k-1}$	1186	33.3	1
$sign_{k-2}$	753	20.2	1
$sign_{k-3}$	544	14.6	1
$sign_{k-4}$	472	13.4	1
q_{k-1}	4.09	9.29	12.6
q_{k-2}	2.66	6.59	12.6
q_{k-3}	1.85	4.66	12.6
q_{k-4}	1.16	2.98	12.6

Comparable results obtain using as few as two lags of aggressive order sign and signed quantity. Linear forecasts of y_k do not benefit appreciably from the inclusion of data on aggressive orders before $k-4$, or on changes in resting depth prior to aggressive order $k-1$. Because the price-change y_k is not normalized by the sign of the k th aggressive order, it has an expected value of zero, so I do not include a constant term in the regression. Including a constant term in the regression has negligible effect on the results.

Although the last several aggressive order signs do offer rather remarkable explanatory power, the respective distributions of resting depth changes and executed aggressive order quantities have much heavier tails than the distribution of order sign, so price forecasts are meaningfully improved by the inclusion of these variables.

The positive coefficients on the lagged aggressive order variables and on the depth changes at

the best bid and best ask are consistent with the general intuition that buy orders portend price increases, and sell orders portend price decreases. The negative coefficients on depth changes at the outside price levels require slightly more explanation.

Because the E-mini market operates according to strict price and time priority, a trader who seeks priority execution of his passive order will generally place that order at the best bid (or best ask); however, if the trader believes that an adverse price movement is imminent, he will place his order at the price level that he expects to be the best bid (ask) following the price change. It is relatively uncommon for prices to change immediately after an aggressive order in the E-mini market, but when prices do change, it is extremely rare during regular trading hours for the change to exceed one tick. As a result, the expected best bid (ask) following a price change is typically one tick away from the previous best, so it is not surprising that (e.g.) an increase in resting depth one tick below the best bid tends to precede a downward price change. These features of the E-mini market also shed some light on why changes in depth more than one tick away from the best (i.e., $d_{k-1}^{best\ bid-2}$ and $d_{k-1}^{best\ ask+2}$) are not significant predictors of future price movements.

D Supplemental Tables of Empirical Results

Table 6: Average Gross Earnings of Aggressive Orders, by Size and Trader Type (in Dollars per Contract)

\bar{q}	A-HFTs		B-HFTs		Non-HFTs	
	All AOs $\leq \bar{q}$	Incremental AOs	All AOs $\leq \bar{q}$	Incremental AOs	All AOs $\leq \bar{q}$	Incremental AOs
1	3.84	-	4.37	-	1.68	-
5	4.23	4.38	4.56	4.64	1.72	1.75
10	3.49	2.84	4.66	4.85	1.77	1.89
15	3.85	6.39	4.71	4.95	1.79	1.97
20	4.14	4.95	4.77	5.08	1.83	2.06
25	4.41	6.65	4.81	5.14	1.86	2.20
30	4.79	6.79	4.87	5.24	1.90	2.53
35	4.99	7.00	4.88	5.10	1.91	2.40
40	5.28	6.69	4.91	5.14	1.95	2.55
45	5.42	7.04	4.92	5.49	1.95	2.14
50	5.61	6.90	4.95	5.14	1.96	2.00
60	5.87	7.00	4.98	5.30	1.99	2.87
75	6.12	7.20	5.00	5.40	2.03	2.71
90	6.38	7.20	5.01	5.52	2.03	2.16
2000	7.65	7.99	5.67	7.63	3.19	4.20

Table 7: Fractions of Trader Types' Aggressive Orders Below Size Threshold \bar{q}

\bar{q}	A-HFTs		B-HFTs		Non-HFTs	
	% of All AOs	% of Aggr. Volume	% of All AOs	% of Aggr. Volume	% of All AOs	% of Aggr. Volume
1	24.31%	0.40%	39.48%	4.77%	53.79%	5.68%
5	43.74%	1.44%	76.09%	16.13%	83.26%	14.88%
10	54.64%	3.09%	84.10%	23.95%	89.87%	20.68%
15	56.75%	3.54%	88.01%	30.03%	91.57%	23.05%
20	60.82%	4.80%	90.70%	35.82%	93.27%	26.43%
25	62.38%	5.37%	92.36%	40.29%	94.41%	29.32%
30	64.62%	6.39%	94.14%	46.30%	94.98%	31.07%
35	65.82%	7.02%	94.97%	49.56%	95.23%	31.94%
40	68.27%	8.51%	96.03%	54.31%	95.66%	33.69%
45	69.29%	9.20%	96.32%	55.76%	95.80%	34.33%
50	71.07%	10.55%	97.66%	63.29%	97.09%	41.08%
60	73.81%	12.97%	98.54%	69.12%	97.36%	42.66%
75	76.65%	16.01%	99.02%	73.14%	97.64%	44.76%
90	80.68%	21.11%	99.20%	74.86%	97.83%	46.37%