JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 54, No. 2, Apr. 2019, pp. 729–758
COPYRIGHT 2018, MICHAEL G. FOSTER SCHOOL OF BUSINESS, UNIVERSITY OF WASHINGTON, SEATTLE, WA 98195
doi:10.1017/S0022109018000832

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

Anticipatory Traders and Trading Speed

Raymond P. H. Fishe, Richard Haynes, and Esen Onur*

Abstract

We examine whether speed is an important characteristic of traders who anticipate local price trends. These anticipatory participants correctly trade prior to the overall market and systematically act before other participants. They use manual and algorithmic order entry methods, but most are not fast enough to be high frequency traders (HFTs). Those anticipating price trends have impacts as if they are informed traders, while the case for anticipatory participants affecting the volume of other traders is rejected. A follow-up sample shows significant attrition in accounts and difficulty maintaining the anticipatory strategies. To identify anticipatory traders, we devise novel methods to isolate local price trends using order book data from the West Texas Intermediate (WTI) crude oil futures market.

I. Introduction

The growth of high frequency trading in financial markets has led to concerns that such traders may use processing speed to disadvantage other participants. Theoretical models suggest that high frequency traders (HFTs) may gain by entering new markets to compete with slower participants (see Hoffmann (2014), Biais, Foucault, and Moinas (2015)).[1] Consistent with theory, fast traders have an

*Fishe (corresponding author), pfishe@richmond.edu, University of Richmond Robins School of Business; Haynes, rhaynes@cftc.gov, U.S. Commodity Futures Trading Commission (CFTC); and Onur, eonur@cftc.gov, CFTC. We thank an anonymous referee, Jonathan Brogaard, Jennifer Conrad (the editor), Joel Hasbrouck, Scott Irwin, Steve Kane, Albert Menkveld, Scott Mixon, David Reiffen, Michel Robe, John Roberts, Sayee Srinivasan, and participants at the 2016 AFA meetings (High Frequency Trading session) for comments on earlier versions of this research. The coauthors of this paper prepared this research under the direction of the Office of the Chief Economist at the CFTC, where economists produce original research on a broad range of topics relevant to the CFTC's mandate to regulate commodity futures markets, commodity options markets, and the expanded mandate to regulate the swaps markets pursuant to the Dodd–Frank Wall Street Reform and Consumer Protection Act. These papers are often presented at conferences and many of these papers are later published by peer-review and other scholarly outlets. The analyses and conclusions expressed in this paper are those of the authors and do not reflect the views of other members of the Office of the Chief Economist, other CFTC staff, or the CFTC itself. All errors and omissions, if any, are the authors' own responsibility.

[1]Some researchers also suggest that anticipatory trading generates negative externalities, such as reducing liquidity provision, vitiating behavior as with quote stuffing, or inducing slower traders to depart (see, e.g., Ye, Yao, and Gai (2013), Biais et al. (2015), Foucault, Kozhan, and Tham (2014), Han, Khapko, and Kyle (2014), and Menkveld and Zoican (2017)).

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

advantage if they process news or intraday order information quickly to anticipate future order flows and prices. Liquidity costs may increase if these anticipatory traders discover smart routing of a parent order from child order flow (e.g., O'Hara (2015)). In this light, anticipatory trading may reduce market efficiency by temporarily raising (or lowering) price from what would occur without such participants (see Brunnermeier and Pedersen (2005)). However, if parent orders are informed, then anticipatory trading may increase efficiency by faster price discovery.

Brogaard, Hendershott, and Riordan (2014), Hirschey (2016), and Jiang, Lo, and Valente (2013) find support for anticipatory trading by selected HFTs. These analyses show that HFTs appear to identify and trade before short-horizon price movements in equity and fixed income markets. However, the correlation between non-HFT activity and subsequent price changes is also often positive in these results, so the relationship between the presumed comparative advantage of HFTs and the ability to anticipate price changes is unclear.

This study seeks to determine whether the ability to trade at high speeds is necessary for anticipatory trading. Our approach is different from the previous literature because we identify all sample traders, not just HFTs, who can repeatedly execute trades that anticipate subsequent intraday price changes ("anticipatory traders"). We do not condition the analysis on a pre-specified group of HFTs or use filters designed to isolate such traders from other participants.[2] Once anticipatory traders are identified, we examine if speed is a determining characteristic to confirm whether anticipatory strategies are a high frequency phenomenon.

To find anticipatory traders, we develop a novel method to identify temporary, directional changes in intraday prices, which we call local price paths. Conditional on these short-term price movements, we identify two types of anticipatory traders: those that act early during a directional price change ("Type E") and those that forecast an up-coming price reversal ("Type R"). Early refers to trades in the first 10% of path volume and to act before a price reversal means to trade correctly in the last 10% of the previous path's volume.

Our sample provides evidence that a relatively small fraction of traders, 4.1% (11.1%) out of 7,554 tested with 2.5% (5%) control rates, anticipate short-term directional price changes and/or price reversals. We find more Type R than Type E participants, with both groups containing algorithmic- and manual-entry traders. An investigation of Sharpe ratios for these traders confirms the relative success of these anticipatory strategies. However, most of these traders do not appear to be HFTs and, accordingly, our tests support the view that speed is not a significant determinant of anticipatory trading.[3]

[2]Biais and Foucault (2014) and U.S. Securities and Exchange Commission (2014) offer detailed discussions of the conditioning methods used in recent HFT research. Such methods may vitiate inferences about the overall population behavior.

[3]Note that our results do not imply that speed is unimportant in a general sense, as increases in participant speeds are connected to market quality improvements. For example, HFTs are associated with spread-decreasing limit orders (Carrion (2013), Hagströmer, Nordén, and Zhang (2014)), providing liquidity against temporary shocks (Brogaard, Hendershott, and Riodan (2014)), and spread decreases in general (Hendershott, Jones, and Menkveld (2011)) when they are the new entrants (Brogaard and Garriott (2019)). See Menkveld (2016) for a recent review of this literature.

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

The data examined are a subset of the regulatory information collected from exchanges by the U.S. Commodity Futures Trading Commission (CFTC). Our primary analysis uses anonymous account-level information from both order book and trade files for the West Texas Intermediate (WTI) crude oil futures contract traded on the Chicago Mercantile Exchange/New York Mercantile Exchange (CME/NYMEX). We use two samples for analysis. The initial sample is for the Dec. 2011 expiration and contains 48 trading days beginning on Sept. 12, 2011. The follow-up sample is for the Oct. 2014 expiration and covers 35 days beginning on Aug. 1, 2014. The WTI contract is one of the most active futures contracts with high volumes and daily participation by thousands of accounts.

Although most high frequency research uses equity market data, futures markets involve many HFTs and exhibit many features found in equity markets.[4] Both equity and futures exchanges offer subscribers real-time, low-latency access to order book information, trading volumes between venues are correlated through time, and market quality patterns are similar. Specifically, crude oil and equity price volatility are correlated, bid–ask spreads decrease as overall volumes increase, and execution speeds increase with the growth of HFTs.[5] One important difference is that equity markets offer multiple execution venues for the same security. Thus, a caveat to our findings is that while speed may be less important for anticipatory strategies in consolidated futures markets, the gains from quoting on multiple platforms may generate a comparative advantage for HFTs to use such strategies in equity markets.

To study the speed of Type E and Type R traders, we measure order book latencies, specifically the time between order entry and either cancellation or execution messages. We make speed inferences from all account orders as well as from a bootstrap simulation that gives equal weight to each participant account, thereby reducing the impact of skewness from high volume accounts (see Fishe and Smith (2018)). In these results, both Type E and Type R participants appear slower, on average, than other participants in both cancellation and execution speeds. Algorithmic traders are faster, but being algorithmic and either Type E or Type R does not further increase speed. Also, the average proprietary trader is faster than accounts trading on behalf of customers, and those who modify orders are, on average, much slower to both cancel and execute trades. Finally, Type E and Type R groups do not act faster when they are trading in the first or last 10% of path volume, indicating that they do not seem to be strategically using speed for their anticipatory strategy.

To find anticipatory traders, our methods examine thousands of trader histories and test statistics. We limit the possibility of false positives by controlling for the *false discovery rate* (FDR) as defined by Benjamini and Hochberg (1995) and

[4]See Fishe and Smith (2018) for a discussion of HFT research in commodity markets and the findings in Kirilenko, Kyle, Samadi, and Tuzun (2017) regarding HFT activity during the "Flash crash."

[5]Recent equity statistics are summarized in Angel, Harris, and Spratt (2015) and related references. The market quality and volatility characteristics of WTI crude oil futures are found in Brunetti, Büyükşahin, and Harris (2016), Raman, Robe, and Yadav (2016), Easley, Prado, and O'Hara (2016), and Haynes and Roberts (2017); the latter discusses speed characteristics for several commodities.

implemented by Storey (2002) and Fishe and Smith (2012). Our results use both 2.5% and 5% control rates for the FDR.

Local price paths are found by considering how a perfect foresight trader might act. This trader would initiate trades in the correct direction at or around price reversals and seek confirmation that price changes were directional and not merely random noise.[6] The sequential probability ratio test (SPRT) by Wald (1945) identifies such directional price paths. This approach provides a test for nonrandom, trending behavior and allows identification of turning points between up and down price paths.[7] The price path method uses information that traders do not have at the time of their trades, so it is a conservative approach and biased against finding such systematically successful participants.

Given that some participants are anticipatory traders, the question arises as to whether they appear informed. This question is hard to separate from "skilled" trading because both types of participants act before price changes. The data we examine identify who initiates a trade (the "aggressor" side) and whose order is executed from the book (the "passive" side). To help distinguish between informed and skill effects, we examine price impacts for aggressive trades and the ability to avoid adverse selection costs for passive trades. We find that aggressive trades by the Type E group have greater price impacts at 10- and 60-second intervals compared to other traders, suggesting an information effect. Moreover, they also seem to avoid higher adverse selection costs compared to other traders, making their passive trades well timed, suggesting a skill effect. The Type R group is not informed, as measured by the relative size of their price impacts, but may be weakly skilled at avoiding adverse selection costs.

We also examine whether Type E and Type R traders affect the trading volume of other participants. For example, they may incite other participants to trade, a so-called "momentum-ignition" effect, and thereby are not exactly anticipatory because they have caused others to trade.[8] Our results show no effects for Type E traders and possibly negative effects for Type R traders, which works against the momentum-ignition effect.

Finally, we examine whether these traders retain their skills using the 2014 sample and if speed is a factor in skill retention. We find substantial attrition among accounts, so that only 35% of the 2011 accounts are found trading 3 years later. Of those that remain, only 9% of the Type E group maintain an anticipatory strategy, while 21% of the Type R group continue with the reversal strategy. In total, the overlapping 2014 sample shows a reduction in the relative size of these Type E and Type R groups. We find that speed plays no role in whether participants continue with either the Type E or Type R strategies. Analogous to the difficulty of fund managers maintaining a "hot hand," these findings suggest that

[6]Our approach makes results conditional on local price paths in the same way that research on the Flash crash or other market events is conditional on trade paths around those events (e.g., Kirilenko et al. (2017), Menkveld and Yueshen (2019)).

[7]An alternative is to use a preset condition to define whether prices have changed sufficiently to identify turning points (Hautsch (2012)).

[8]Momentum-ignition and other follower-inducing strategies require that other participants believe there is a positive probability of very recent informed trades (Allen and Gale (1992)). These strategies are among those thought to be used by HFT firms (U.S. Securities and Exchange Commission (2014)).

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

anticipatory trading became more difficult to implement after 2011 and that the Type E strategy was more ephemeral (see Barras, Scaillet, and Wermers (2010)).

This paper proceeds as follows: Section II discusses the related literature. Section III describes the methods used to identify anticipatory traders, along with a general discussion of the price path approach. Section IV discusses the data and how we measure speed. Section V provides our analyses and results. Finally, Section VI offers several conclusions.

II. Related Literature

Several researchers offer evidence on anticipatory trading. Jiang et al. (2013) analyze how often transactions are in the right direction compared to subsequent price changes around macroeconomic announcements. Frino, Ibikunle, Mollica, and Steffen (2016) follow a similar approach for trading in Brent oil futures around the London fixing period. These studies report evidence of anticipatory trading but do not resolve the question of whether speed is a necessary factor. Importantly, Jiang et al. (2013) find that non-HFT limit orders are more informative for future prices than HFT limit orders, which suggests that anticipatory traders may be fast traders.[9]

Other more general studies by Brogaard et al. (2014) and Hirschey (2016) use a NASDAQ sample with HFT participation noted in the trade data. Brogaard et al. find that the correlation between net order flow for all sample HFTs and subsequent returns is positive, but it is short lived and quite low (less than 4% at 1 second and near 0 at 2 seconds). Interestingly, they find that non-HFTs demanding liquidity show higher, longer lived correlations with subsequent returns than HFTs demanding liquidity, implying that sub-groups excluding HFTs also appear informed. Hirschey (2016) also finds that liquidity-demanding trades by HFTs precede liquidity-demanding trades by non-HFTs. These studies leave unanswered the question of why some non-HFTs are predictive of returns if HFTs have anticipated their net orders. A possible explanation arising from our sample is that successful anticipatory strategies are found among both HFT and non-HFT groups, so the NASDAQ HFT filter may not sufficiently separate anticipatory from non-anticipatory behavior.

Anticipatory and momentum trading strategies may be similar. Bikhchandani, Hirshleifer, and Welch (1992) and Allen and Gale (1992) are momentum studies that may include anticipatory behavior.[10] Bikhchandani et al. develop a model in which informational cascades cause momentum reactions. Those who react first are analogous to anticipatory traders. Allen and Gale show that traders are affected by big traders because of the likelihood that they may be informed. Thus, smaller traders have an incentive to trade in the

[9]Clark-Joseph (2012) suggests that aggressive HFTs execute multiple, generally unprofitable smaller size orders for the purpose of obtaining order book information. When this information indicates a high propensity for a directional price movement, the firm then profits from well-timed larger orders.

[10]De Long, Shleifer, Summers, and Waldman (1990), Hirshleifer, Subrahmanyam, and Titman (1994), Chan, Jegadeesh, and Lakonishok (1996), and Dong, Polk, and Skouras (2019) offer examples of general momentum behavior.

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

same direction as a big trader, making the big trader appear anticipatory. These environments may create "following trades" by investors, implying that some traders may be misclassified as anticipatory. To sort out the degree to which such traders vitiate our results, we examine the overall volume behavior following trades by those identified as anticipatory.

Last, our study helps researchers understand how filters to identify HFTs may affect results. Others have observed that sample filters (inventory turnover, trading volume, cancellations, etc.) pose a problem for population inference. Biais and Foucault (2014) and the U.S. Securities and Exchange Commission (SEC) (2014)) discuss several of the filter methods used to classify data as algorithmic or HFT related. Biais and Foucault (p. 10) warn that, "One problem with this approach is that it may select HFTs with a specific trading style … while excluding others." Thus, rather than pre-condition the analysis on a subset of trader characteristics, we seek to infer such characteristics (specifically, speed of trading) by identifying all traders associated with a specific strategy.

III. Methods

Our analysis uses statistical methods to identify intraday price trends. The subsequent discussion explains how such paths are identified and how the FDR method finds participants who can systematically execute trades during selected segments on these paths.

A. Local Price Paths

Our approach to price paths locates sequences where price changes exhibit some degree of short-term predictability. We use the SPRT to define the local neighborhood size and to test for nonrandomness. In the resulting sequences, we search for local price extrema to identify turning points.

The process starts with the transaction price series from which we remove P_t if $P_t = P_{t-1}$, keeping the first price of each such sequence to preserve unique price levels. Then, we remove all sequences of contiguous bid–ask bounce. Specifically, all cases in which $\Delta P_t = -\Delta P_{t-1}$ are excluded. We retain the prices into and out of such sequences but remove the intermediate implied bid–ask trades. For example, consider the price sequence, $P_t \equiv \{3,4,5,5,4,4,5,4,5,6\}$. After the first filter, the sequence is $P_t^* \equiv \{3,4,5,4,5,4,5,6\}$, from which we retain $P_1^*, P_2^*, P_7^*,$ and P_8^* after the second filter. The purpose behind removing bid–ask bounce sequences is to exclude periods in which liquidity replenishment is sufficient to satisfy liquidity demand. These sequences may provide information but arguably provide little help identifying a price trend. Note that after we have identified the price trends, we restore *all* observations to the data set for identification of Type E and Type R participants.

The preceding procedure produces a sequence in which $\Delta P_t \neq 0 \, \forall_t$. We then define a set of candidate prices based on the SPRT test results. Within a group of K prices, the price P_{t^*} is a candidate for a local *minimum* price at trade t^* if

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

i) the count of *previous* price changes, $n^- \geq \delta K$, where $n^- = 1(\Delta P_t < 0)$ for $t \in \{t^*, t^*-1, t^*-2, \ldots, t^*-K\}$, and

ii) the count of *subsequent* price changes, $n^+ \geq \delta K$, where $n^+ = 1(\Delta P_t > 0)$ for $t \in \{t^*+1, t^*+2, \ldots, t^*+K\}$.

The parameter $\delta \in [0,1]$ acts to require price decreases before and price increases after a minimum price. This parameter is chosen based on the assigned power of the SPRT. The greater the power chosen for the test, the greater the number of decreases (or increases) required. The parameter K defines the local neighborhood of trades, which is how many trades before (or after) are required in a test. For example, $K = 10$ and $\delta = 1$ is an extreme case that derives from the near 100% power of the SPRT. In this case, a local minimum price candidate requires that all 10 prices before the candidate decrease stepwise toward the minimum and all 10 prices after the candidate increase stepwise away from the minimum. We use the same approach but reverse the inequalities in i) and ii) to define a candidate local *maximum* price.[11]

The properties of error rates guide the selection of the consistency parameter and the size of the local neighborhood. Assume the binary variable for the sign of any price change is binomial with the null hypothesis, $H_o : q = 0.5$, which normalizes the null distribution of sign changes to a random sequence. A participant who attempts to detect nonrandom price changes is focused on rejecting this null in the neighborhood of a candidate price extrema. Thus, it is useful to establish confidence that the null is rejected. The SPRT provides a means of identifying the appropriate size of the local neighborhood necessary to reject the null, as this test is uniformly most powerful against any other test in its expected stopping time (see Wald (1948)).

The SPRT computes the likelihood ratio for each successive observation in the trade sequence given a null and alternative hypothesis. It uses Type I (α) and Type II (β) error rates for these hypotheses to establish bounds for rejecting one hypothesis versus another. Our calculations set these error rates equal to 10%, which then feeds back to neighborhood size and the consistency parameter.

For neighborhood size, we simulate the number of trades necessary to reject the null ($H_o : q = 0.5$) against the alternative ($H_a : q = 0.8$). We use a strongly convincing alternative versus one closer to the null, as participants would not rely on a testing method if it required too many trades, perhaps more than might be observed in a local trend. Using small differences between the null and alternative hypotheses creates longer required sampling sequences. With 1,500 simulations, we find that if participants selected 17 observations, then in only 10% of the cases would they require more observations before rejecting the null. As this choice equals the required Type II error rate ($\beta = 10\%$), we use 17 observations on both sides of a candidate price to define the local neighborhood.

To define the consistency parameter, we use a choice that follows from the Type I error rate ($\alpha = 10\%$). Under the null hypothesis, this error rate is

[11] This method may produce cases in which multiple minimums or maximums are contiguous on a price path. We remove such cases by selecting a global maximum or minimum in such sequences. The final price paths alternate in the sign of ΔP_t, where this price change is from the beginning to the end of a path.

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

$\Pr[\sum x \geq x' | K = K^*] \leq 0.10$, where $x \in \{0, 1\}$ to indicate either a negative or positive price change. As the neighborhood size is set by the SPRT such that $K^* = 17$, the binomial distribution implies the cutoff for 10% arises when $10 < x' < 11$. We experimented on randomly chosen days with both choices and found that $x' = 10$ gave somewhat more paths, but the overlap was near 100% with $x' = 11$. As more paths make it more difficult to consistently trade in the correct direction, we use $x' = 10$ in our analysis. If at least 10 out of 17 price changes are observed with the appropriate sign, positive (negative) before a candidate price maximum (minimum) and negative (positive) after the candidate price, then that price is a valid "candidate" for the local extrema. This approach produces a consistency parameter δ approximately equal to 60%.[12] To choose among the set of valid candidate prices within the same neighborhood, we select based on the following conditions:

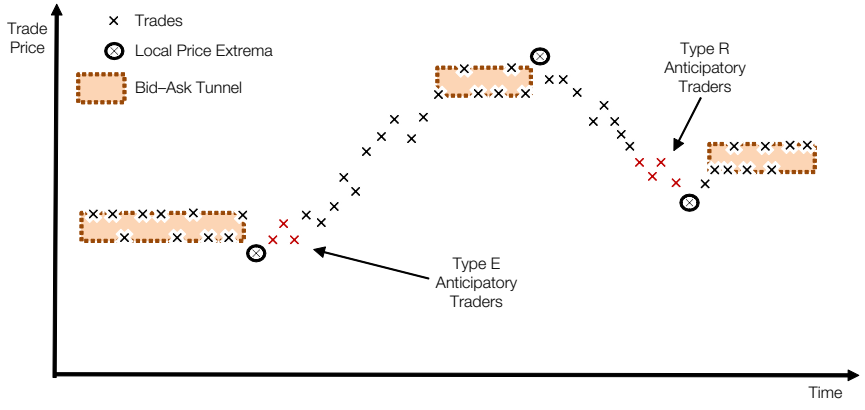
i) $P_{t^*}^{\min} = \min\{P_{t^*}, P_t\}\ \forall_{t \neq t^*}$, where $t \in \{t^* - K, t^* - K - 1, \ldots, t^* + K\}$ for a minimum, and

ii) $P_{t^*}^{\max} = \max\{P_{t^*}, P_t\}\ \forall_{t \neq t^*}$, where $t \in \{t^* - K, t^* - K - 1, \ldots, t^* + K\}$ for a maximum.

Figure 1 illustrates how to conceptualize the working of the price path algorithm and the relative location of anticipatory traders. The figure shows a sequence of trades (the "x" terms) for a portion of the day's trading. There are periods with

FIGURE 1
Price Paths and Anticipatory Traders

Figure 1 illustrates a sequence of trades (the "x" terms) from participant order flow. When new information about the value of the asset arrives or liquidity demand changes, the market price reacts until the information is impounded in the price structure or new liquidity arrives to resolve the imbalance. Non-price-moving liquidity-based trading is shown as trades along a bid–ask bounce sequence (shaded areas). Price reversals occur at the specified local price extrema (the circled trades). A participant that possesses the skills to process order flow and trade information to systematically forecast the short-term direction of trade prices is known as a Type E anticipatory trader. A participant who processes new information or market flow signals to systematically forecast the subsequent reversal in local trade price trend is known as a Type R anticipatory trader.



[12] We also simulated our results with Type I and Type II error rates equal to 5%. These simulations gave a neighborhood size of 24 observations to maintain the Type II error rate and a consistency parameter of approximately 63%.

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

market-making activity in which non-price-moving liquidity-based trading creates a bid–ask bounce sequence (shaded areas). When new information about the value of the asset arrives or liquidity demand changes, the market price reacts until the information is impounded in the price or new liquidity arrives to resolve the imbalance. Local price reversals occur at the specified price extrema (the circled trades).

In Figure 1, Type E traders possess skills to process order flow and information to systematically forecast the short-term direction of prices. These participants react quickly *after* a price reversal occurs. Type R traders may use strategies that analyze order book liquidity or new information to place limit order prices near upcoming price reversals. The trades of these participants occur *before* but close to the local price extrema.

B. Finding Anticipatory Traders

To identify the participants using of an anticipatory strategy, we implement the FDR method of Benjamini and Hochberg (1995) as applied by Storey (2002) and Fishe and Smith (2012), which adjusts for multiple testing problems.[13] The FDR controls for the expected proportion of false discoveries in our sample. By effectively adjusting test statistic critical values, the FDR method limits these expected mistakes to a pre-specified proportion of successful statistics. We use both 2.5% and 5% control rates, so we often report two sets of results. The FDR method gives greater confidence that the participants we identify as anticipatory traders are truly either Type E or Type R traders.

We use a volume metric to determine whether a trader may be classified as participating early or late in a price path. Let $x_{i,t}^{j} \in \boldsymbol{X}_j$ be the quantity traded by participant i at time t in the vector of all trades ($\boldsymbol{X}_j$) on price path j. If participant i is a buyer (seller) at time t in path j, then $x_{i,t}^{j} > 0 \left(x_{i,t}^{j} < 0\right)$. If participant i does not trade at time t in path j, then $x_{i,t}^{j} = 0$. The Heaviside function, $s_{i,t}^{j} = h\left(x_{i,t}^{j} \in \boldsymbol{X}_j; d\right)$, defines whether trade $x_{i,t}^{j}$ arises in the first dth percentile of path j's volume. The Heaviside function equals 1 if the trade is in the first dth percentile, and 0 otherwise. The price direction along path j is defined to be increasing (decreasing) if $\Delta P_j > 0$ ($\Delta P_j < 0$), where the first and last trade prices on the path are used to compute this difference.

Participants are identified as successful Type E traders on a given path if their trades occur in the first dth percentile of path volume and their trades are on the correct side for the path's price change. For a given value of d, we compute the sample frequency of successes for each participant:

[13] Recent applications of FDR include Barras et al. (2010) who seek to identify fund managers with positive alpha performance, Bajgrowicz and Scaillet (2012) who examine the success of technical trading rules, Fishe and Smith (2012) who identify the number of informed traders in several futures contracts, and Harvey, Liu, and Zhu (2014) who study threshold critical values necessary to claim a new risk factor after hundreds of asset pricing tests by previous researchers.

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

$$F_i^{\text{Type E}} = \frac{\sum_{j=1}^{J}\sum_{t=1}^{T_j}\mathbf{1}\left(x_{i,t}^j>0,\Delta P_j>0|s_{i,t}^j=1\right)+\mathbf{1}\left(x_{i,t}^j<0,\Delta P_j<0|s_{i,t}^j=1\right)}{\sum_{j=1}^{J}\sum_{t}^{T_j}\mathbf{1}\left(x_{i,t}^j\neq 0\right)}, \tag{1}$$

where $\mathbf{1}(\ldots)$ is the binary indicator function based on the given expression, T_j is the number of trades on path j, and J is the number of price paths in the sample.

The null hypothesis for our tests depends on several factors. For example, if a trader randomly submits buy and sell orders during the day in small sizes, then across all paths, we might expect to find about 10% of these trades in the first 10% of path volume ($d=10$). However, how many of these are expected to be successful, meaning that they are aligned with the price path direction? The answer is 5% if volume is equal between up and down paths, trade size is constant, and the participant submits an equal number of buy and sell orders. Unfortunately, these assumptions do not generally hold, so an exact null that defines randomness in order choice varies across participants.

Rather than seek a solution for such participant nuances, we choose a more restrictive null for our tests. The statistic in equation (1) indicates the proportion of participant i's trades that were executed in the first dth percentile of volume *and* were in the correct price direction. This proportion is conditional on a perfect foresight calculation of local price trends. If trades are randomly placed during the day, then a trader has a 10% chance of executing a buy or sell within the first 10% of path volume ($d=10$). Adjustments for order size or trade side will lower this fraction and, therefore, lower the null. Thus, to make it more difficult to find anticipatory traders, we use 10% as the null hypothesis. For each trader then, the Type E null hypothesis is $H_o: F_i^{\text{Type E}} = d/100$. This is a binomial test and will have statistical power in large samples.

For our empirical work, we set $d=10$ to identify Type E traders. To identify Type R traders, we consider the *last* 10th percentile of trading volume to be indicative of whether a trader uses information or foresight to anticipate the coming reversal of the price path. To measure success for Type R traders, we compute the proportion analogous to equation (1) using $d=90$ to define the Heaviside function:

$$F_i^{\text{Type R}} = \frac{\sum_{j=1}^{J}\sum_{t=1}^{T_j}\mathbf{1}\left(x_{i,t}^j>0,\Delta P_j<0|s_{i,t}^j=0\right)+\mathbf{1}\left(x_{i,t}^j<0,\Delta P_j>0|s_{i,t}^j=0\right)}{\sum_{j=1}^{J}\sum_{t}^{T_j}\mathbf{1}\left(x_{i,t}^j\neq 0\right)}. \tag{2}$$

The null hypothesis that we test to identify Type R traders is $H_o: F_i^{\text{Type R}} = (100-d)/100$. To ensure statistical power, we confine our investigation to participants with more than 30 trades.

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

IV. Data

The data are derived from audit trail files for the CME/NYMEX WTI light sweet crude oil futures contract. The WTI contract is traded worldwide on the Globex electronic platform. For the WTI contract, each 1 cent move in price represents a $10 change in contract value, which provides leveraged returns even for relatively small changes in price when weighed against the relatively low margin requirements.

Our initial sample covers the period from Sept. 12, 2011 to contract expiration (48 days). This period is selected based on the trading and open interest activity in the Dec. 2011 contract, which is the first or second most active month in the year. A follow-up sample was available from Aug. 1, 2014 to Sept. 19, 2014 (35 days) for the Oct. 2014 expiration, which was the nearby contract for most of this period. These data are used to examine whether anticipatory strategies survive through time. Both samples contain all trades and orders posted, modified, and/or canceled on the CME/NYMEX exchange, and identify who initiates a trade (the "aggressive" side) and the counterparty whose order is resting on the book (the passive side). We test for anticipatory trading in both samples, using participant accounts that have 30 or more trades. This cutoff is imposed to provide power to the FDR method.

To determine price paths using the SPRT described previously, we remove non-price forming trades from the sample, which are mainly transfers and offsets. We filter out spread trades where both sides are holding the spread, as these trades provide only relative price information. However, if the initial side of a spread trade is an outright, we keep that side's price if it is for the correct expiration.

Table 1 provides statistics on the trading volume, number of participants based on manual and algorithmic entry, and order book data calculated across days in both samples. The WTI contract is quite active in the two samples with an average of 178,447 (225,723) contracts traded each day in 2011 (2014). There is a daily average of 2,285 (4,173) accounts active in 2011 (2014). Most of these accounts use manual-entry methods to place orders.[14] Algorithmic accounts increase from a daily average of 421–704 between samples. Participants act to modify, on average, 62%–67% of new orders and eventually cancel an average of 82%–86% of those orders. Both samples show that the WTI crude oil contract is traded in a quote-driven market, with market orders, on average, less than 2.5% of daily orders. In this analysis, we do not examine stop orders, offsets, transfers, or other special-order types, such as trade at settlement trades.

V. Analysis

The first task is to estimate local price paths in the 2011 data. Then, we use the FDR method to assess whether any participants systematically trade in the

[14] A trader is manual or algorithmic based on a specific variable (Tag 1028) found in the exchange-created data set. See the CME Globex reference guide for additional details on this tag and reporting requirements (http://www.cmegroup.com/globex/files/GlobexRefGd.pdf).

TABLE 1
Daily Characteristics of Sample Data

Table 1 shows summary statistics for messages and orders in the WTI crude oil futures contract expiring in Dec. 2011 in Panel A and expiring in Oct. 2014 in Panel B. These data are for outright or outright-to-spread orders, as these orders are price informing for each expiration. The 2011 sample covers all trading days from Sept. 12 to Nov. 18, 2011 and the 2014 sample covers Aug. 1 to Sept. 19, 2014. The table provides average, median, standard deviation, minimum, and maximum statistics across sample days for volume, number of unique traders (algorithmic or manual order entry methods), message types (entry, modification, or cancellation), and order types (market or limit). Note that message and order counts may not add up exactly, as some order types arise from different message origins, such as a market order being generated from a modified message that was originally a limit order.

Characteristics	Average	Median	Std. Dev.	Minimum	Maximum
Panel A. 2011 Sample: Dec. 2011 Expiration					
Trading volume	178,447	53,456	181,564	8,243	565,140
Count of					
Algorithmic traders	421	280	318	47	890
Manual traders	1,864	1,088	1,474	242	4,388
Order entry messages	819,799	506,624	562,948	7,649	2,095,500
Modification messages	552,411	390,795	400,678	5,475	1,579,307
Cancellation messages	705,919	480,767	451,195	3,837	1,764,984
Market orders	11,075	3,711	11,429	431	34,926
Limit orders	816,439	505,980	559,644	7,408	2,087,561
Panel B. 2014 Sample: Oct. 2014 Expiration					
Trading volume	225,723	270,396	156,284	8,401	486,328
Count of					
Algorithmic traders	704	934	397	29	1,262
Manual traders	3,469	4,692	2,287	243	6,533
Order entry messages	781,966	762,280	301,255	216,777	1,415,200
Modification messages	484,818	500,726	137,082	220,791	829,322
Cancellation messages	640,131	663,227	226,690	192,038	1,134,177
Market orders	17,673	20,012	12,099	1,068	41,268
Limit orders	776,262	760,971	297,646	215,233	1,402,853

correct direction on a path. After identifying such traders, we examine whether anticipatory traders are rewarded by computing Sharpe ratios by accounts. We then examine their characteristics relative to other traders, specifically cancellation and execution speeds. Next, the question of whether such traders are informed is addressed as well as whether the volume of other traders is affected by anticipatory traders. Finally, we compare the 2011 and 2014 samples to investigate whether the Type E and Type R groups retain their skills 3 years later.

A. Local Price Paths

We apply the SPRT method to find local price paths for each day in the sample. Table 2 reports summary statistics derived from all price paths as well as paths that contain anticipatory traders. Path information is summarized by price direction. The data show the average and median path returns in percentages, average path duration in seconds, average path volume, average number of trades, and average number of participants in a price path. Average and median returns are nearly equal between up and down paths, with up paths about 12.5% longer in duration. Up paths also average greater volume and trades, so the near equal returns imply that prices trend up slower than they trend down. There are somewhat more than 80 participants, on average, on these paths, suggesting a competitive environment for price discovery.[15]

[15] Additional summary data on the sample and price paths are available from the authors or at https://www.cftc.gov/idc/groups/public/%40economicanalysis/documents/file/oce_anticipatory_supplemental.pdf.

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

TABLE 2
Local Price Path Characteristics

Table 2 shows summary statistics for all local price paths in the Dec. 2011 WTI crude oil futures contract. The paths are computed from intraday prices on outright trades or spread trades with one side being outright. The sample covers all trading days from Sept. 12 to Nov. 18, 2011. Statistics are computed by month and path direction for all price paths, price paths with Type E participants, and price paths with Type R participants. The statistics are average and median path returns in percentages, average path duration in seconds, average path volume in contracts, and the average number of trades. The average number of unique participants is shown for all paths. Type E and Type R participants are selected using the 5% FDR control rate.

					Average			
Description	Path Direction	Path Count	Avg. Path Return (%)	Median Path Return (%)	Path Duration (seconds)	Path Volume	No. of Trades	No. of Unique Participants
All price paths	Down	7,903	−0.148	−0.119	115.6	569.5	447.6	88.8
	Up	7,916	0.149	0.118	102.8	513.7	402.3	82.0
Price paths with Type E participants	Down	2,455	−0.193	−0.156	190.7	892.8	692.8	121.6
	Up	2,290	0.204	0.166	175.6	876.2	678.1	120.5
Price paths with Type R participants	Down	7,140	−0.148	−0.120	114.4	612.8	482.3	95.2
	Up	6,996	0.152	0.122	102.5	562.5	441.1	89.1

B. Identifying Anticipatory Participants

The FDR method uses the binomial statistics in equations (1) and (2) to identify Type E and Type R participants. This method produces a two-tailed test in which our focus is on the upper tail results. We examine results for 2.5% and 5% FDR control rates. Note that the 5% control rate includes the 2.5% control rate accounts, as it is less restrictive. The 2.5% (5%) tests found 112 (301) Type E and 196 (542) Type R participant accounts. We also identified 4 (25) overlaps between Type E and Type R groups for the 2.5% (5%) control rate. These accounts appear to have exceptional timing skills. Within these two groups, we find that 150 are algorithmic traders: 52 in the Type E group and 98 in the Type R group in the 5% control rate tests. The remaining traders use manual order entry methods.

Table 2 also provides path characteristics when Type E and Type R accounts are active anywhere on a price path. For these entries, the average number of unique participants refers to anticipatory traders, not to all traders. There are, on average, 1–2 Type E and 7 Type R traders active on these paths. Notably, when Type E accounts are active, average and median path returns are greater in absolute value, the volume and trade count increases, and the path duration is greater. When Type R accounts are active, average and median path returns are about the same as the all path statistics. Thus, Type E accounts may be somewhat more successful than Type R accounts in identifying paths with larger price changes. Subsequently, we explore whether this suggests informed trading by either anticipatory type.

C. Aggregate Sharpe Ratios

Sharpe ratios examine if the traders identified as anticipatory use a strategy that offers benefits relative to the risk. Otherwise, such strategies will become ephemeral as other, more dominant strategies are recognized or such traders lose their skills. Menkveld (2013) computes Sharpe ratios using high frequency trading data. Our approach follows his method, which incorporates trading fees and margins. The primary difference here is that we do not know a participant's inventory at the start of the sample because our sample begins months after trading

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

starts in the Dec. 2011 expiration. Thus, we focus on the gross cash flows from trades to compute Sharpe ratios.

Following Menkveld (2013), a participant's average gross cash flow from trading on a given day is given by

$$\bar{\pi}_i = \frac{1}{T}\sum_{t=1}^{T} -\Delta\left(n_t^i P_t\right) - \left|\Delta n_t^i\right|\tau, \tag{3}$$

where T is the total number of trades for participant i during the day, n_t^i measures the cumulative net position after the tth trade, P_t is the price of the futures contract, and τ captures the exchange and clearing fees paid by the participant.[16] The capital employed (m_i) is estimated from the NYMEX daily margin requirements for the highest absolute position held during the day.[17] Daily returns are estimated by $r_i = \bar{\pi}_i - r_f m_i$, where the risk-free rate (r_f) is proxied using daily 1-year constant maturity Treasury bill rates. These returns are averaged across trading days to compute the mean return (μ_i) and standard deviation (σ_i) for the Sharpe ratio, which is given by $S_i = a(r_i)\mu_i/\sigma_i$, where $a(r_i)$ is an annualization factor that may be used to adjust for autocorrelation in returns (Lo (2002), Table 2).

Table 3 shows the estimated Sharpe ratios for Type E and Type R anticipatory traders identified using the 2.5% and 5% FDR control rates. Sharpe ratios for the "All Other" group and the sub-group of participants who are both Type E and Type R are also included. Sharpe ratio results for an unadjusted annualization factor ($\sqrt{254}$) and one adjusted for first-order autocorrelation (coefficient ρ) are shown in the table.

Sharpe ratios reveal that the Type E group performs better than the Type R group, with both groups exceeding the average performance of the All Other group. Notably, the overlapping participants within the Type E and Type R groups show exceptional performance. This latter group is both trading early and late on a price path, which would be an optimal reversal strategy. For comparison, Nagel (2012) finds annualized Sharpe ratios between 4.91 and 9.58 using a daily reversal strategy on individual equities between 1998 and 2010. This range is like the 4.78 to 5.48 range reported for the overlapping Type E and Type R participants. The daily return ratios reported by Menkveld ((2013), Table 3) for a single HFT are also close to these values.

Nagel (2012) reports a realized Sharpe ratio of 0.26 for the CRSP value-weighted index, like those for the All Other participants group. As these participants are close to representative of the overall futures market, they may perform like an overall market index. In addition, because the contract here is a derivative, there is a zero-sum effect that suggests the average gross cash flows aggregated over all trades during the life of a contract approach zero (without fees). Although the sample is only over the ending months of the contract, this

[16] The cumulative net position is set to zero at the start of the sample. As the position accumulates, the first trade on a new day is netted against the last cumulative position of that participant, generally the previous day's last net position. The end-of-day mark-to-market approach is not appropriate here because the initial inventory position of each participant is unknown.

[17] Exchange and clearing fees may vary across participants because of exchange incentive programs. We could classify many participants into specific programs, but not all. Thus, we applied the worst-case fees (and margins) to all participants, thereby reducing the estimated Sharpe ratios.

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

TABLE 3
Sharpe Ratios for Anticipatory Traders

Table 3 shows annualized Sharpe ratios for selected groups of participants. The Sharpe ratios are estimated for each participant using the method described by Menkveld (2013). Trade-to-trade cash flows are averaged daily to compute mean returns, which net out capital employed based on daily margin requirements and an estimate of the risk-free rate obtained from daily 1-year Treasury bill yields. Daily returns are averaged across trading days to compute the mean and standard deviation of daily returns used in the Sharpe ratio. Both annualized and adjusted ratios are computed. The adjusted ratios follow the methods in Lo (2002) for a first-order autocorrelation model of returns. The reported ratios are averaged over each participant within a group. Groups are those participants identified as Type E or Type R using 2.5% and 5% control rates (CR) for the FDR tests and the collection of "All Other" participants who are those found not significant in these FDR tests. The overlap group includes participants who are in both Type E and Type R groups. Heteroscedastic-consistent *p*-values are shown for the estimated first-order autocorrelation coefficient (ρ). "NA" indicates that the number of observations is too small for meaningful inferences.

	Type E Participants		Type R Participants		Overlaps for Type E and Type R		All Other Participants	
Characteristic	2.5% CR	5% CR	2.5% CR	5% CR	2.5% CR	5% CR	2.5% CR	5% CR
Sharpe ratio	1.046	0.830	0.523	0.531	NA	4.784	0.371	0.376
Adj. Sharpe ratio	1.126	1.046	0.531	0.535	NA	5.858	0.353	0.362
Rho	−0.082	−0.139	−0.023	−0.015	NA	−0.109	0.043	0.034
p-value of ρ	0.304	0.179	0.569	0.651	NA	0.434	0.151	0.292
Count of participants	112	301	196	542	NA	25	7,250	6,736

effect may also help explain the relatively low Sharpe ratios found for the All Other group.

D. Speed Characteristics

A central question is whether anticipatory traders are HFTs in the usual sense of the term. The SEC (2014) offers five criteria to identify HFTs: i) high speed in routing and executing orders, ii) use of colocation services, iii) short time between establishing and liquidating positions, iv) use of a submit and cancel approach to orders, and v) ending of trading sessions with near zero inventories. We have information on items i), iii), and iv) from the order book data, which is the source for participant speeds.

A participant's speed is measured by the duration between initial order submission time and the exit time of an order.[18] Specifically, we calculate the time between order entry and either cancellation or execution.[19] For cancels, the duration is the difference between the time the CME/NYMEX receives the cancel message and the initial order confirmation time. If an order executes, duration is the difference between the CME/NYMEX confirmation time and the initial order confirmation time. This execution duration depends on a host of factors that affect the order book, such as liquidity flows and new information about price, as well as the initial and subsequent decisions of the participant placing the order. A comparison of these speeds across all orders for Type E and Type R participants is shown in Table 4.

In Table 4, Panels A and B report average (minutes) and median (seconds) speeds, respectively, and Panel C reports the distribution of messages by Type E, Type R, and All Other participants. There are clear differences between average and median speeds across participant types. Average speeds are slow, measured in minutes, suggesting that account-level behavior is very different than market-level

[18] We do not use modification messages to measure speed because they do not remove an order from the book, so they are not part of the exit accounting.

[19] Orders may also exit by administrative action, but we exclude these from the sample.

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

TABLE 4
Speed Characteristics from Order Book Messages

Table 4 shows the average (minutes) and median (seconds) durations between four message types: order entry to either order cancellation or execution and order entry to cancellation or execution with no intermediate modification messages (No Mods). These calculations are grouped by Type E, Type R, and "All Other" participants and are partitioned into manual-entry and algorithmic-entry methods. The Type E and Type R participants are those found using the 5% control rate in FDR tests. Averages and medians shown are across messages for each group and entry type. Panel A shows summary statistics for averages across messages, Panel B shows median statistics, and Panel C shows the distribution and total number of messages by message type: cancellation, modification, or execution. * and ** indicates that tests of whether Type E or Type R is different from the corresponding average for All Other participants are significant at the 5% and 1% levels, respectively. Wilcoxon and Median tests are used to test whether the medians are also different between Type E or Type R and the All Other group durations. The least significant result is reported for the later tests.

	Type E Participants		Type R Participants		All Other Participants	
Ending Message	Man-Entry	Algo-Entry	Man-Entry	Algo-Entry	Man-Entry	Algo-Entry
Panel A. Average Duration between Messages (minutes)						
Cancellation	25.04**	7.59**	5.82**	0.68**	9.44	1.29
Cancellation (No mods)	24.76**	3.98**	5.11**	0.51**	8.67	1.11
Execution	16.19**	7.03	8.07**	2.76	11.55	3.03
Execution (No mods)	4.73**	2.03*	3.33**	0.58	2.24	0.59
Panel B. Median Duration between Messages (seconds)						
Cancellation	54.56**	0.83*	60.07**	7.93**	23.18	0.73
Cancellation (No mods)	35.49**	0.27	59.92**	7.42**	20.11	0.40
Execution	4.32**	0.01**	17.43**	1.83**	8.18	0.10
Execution (No mods)	0.07**	0.00**	9.21**	0.18**	0.87	0.04
Panel C. Message Distribution						
Cancellation	8.5%	2.8%	32.7%	13.4%	15.8%	57.8%
Modification	4.3%	59.5%	22.7%	84.0%	12.8%	33.1%
Execution	87.3%	37.7%	44.6%	2.6%	71.4%	9.1%
Total messages	68,128	51,454	294,194	5,409,423	3,995,440	54,335,721

speeds.[20] These median and average speed differences show that skewness affects the statistics for all groups, including the All Other category. Thus, we focus on the median data for inferences from Table 4.

Considering the SEC criteria, we see that the cancel speed of Type E and Type R accounts is slower than the median cancel speed of All Other participants, with or without modification messages. Type R algorithmic accounts are more than 10 times slower than the All Other algorithmic accounts. All groups use cancellations, but Type E (particularly the algorithmic entry) is below the frequencies shown in Panel C for All Other accounts, while Type R is above for manual entry and below for algorithmic entry. For algorithmic comparisons, both Type E and Type R fall short of using a "submit and cancel" approach compared to other participants. The fact that there are many manual-entry participants in both groups argues against the SEC criteria matching well with these groups.

Execution speeds for Type R accounts are also slower than the All Other group. However, execution speeds for Type E accounts are faster, which is likely due to the greater use of market orders and marketable limit orders by Type E participants. This is consistent with the high execution percentage reported in Panel C of Table 4, suggesting that Type E accounts place orders with the intention of execution, with few cancellations. In contrast, Type R participants have much

[20] Market-level speed computes the difference between two successive cancellations or executions without regard to who is the participant. For example, Hasbrouck and Saar (2013) find 2–3 millisecond peaks in hazard rates that measure market responses after a quote improvement, although there are long tails in these rates.

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

higher modification rates compared to All Other and lower execution rates. Type R accounts, thus, may be modifying to search for the end of a price path. In general, these statistics do not suggest that HFT-like speeds are necessary to implement an anticipatory strategy.

To validate our interpretation of the summary statistics in Table 4, we specify a regression model with characteristics partitioned to identify anticipatory groups. In these regressions, the effects of the anticipatory participants are measured relative to other omitted groups. Table 5 shows regression results using order cancellation speed and order execution speed as dependent variables. The transformation, natural logarithm of 1 plus the cancellation (execution) speed (in seconds with fractional milliseconds), is used.[21] The variables FIRST_10%_PATH_VOL, LAST_10%_PATH_VOL, and BETWEEN_10%–50%_PATH_VOL indicate whether the message occurs during the first 10% of path volume, last 10%, or during the first half of volume that excludes the first 10%, respectively. The TYPE_E_TRADER and TYPE_R_TRADER variables indicate whether the message is by a Type E or Type R participant, respectively. Type E and Type R participants are those identified using the 5% control rate for the FDR. There are interaction terms in selected models that interact as described: Type E trader in the first 10% volume bin (TYPE_E_AND_FIRST_10%_VOL), Type R trader in the last 10% volume bin (TYPE_R_AND_LAST_10%_VOL), Type E trader that is also algorithmic (TYPE_E_AND_ATS), and Type R trader that is also algorithmic (TYPE_R_AND_ATS). Additional binary variables are MODIFIED_ORDER if there are any modification messages to the original order; PROPRIETARY if the trade is made from a proprietary account; the ALGORITHM variable, which is 1 if this message was submitted by a computer algorithm; and AGGRESSIVE_SIDE in execution regressions to indicate if this message initiated the trade. The aggressive side of a trade is assigned by the exchange so no mistakes are made by using other approaches (e.g., Lee–Ready algorithm). As all variables are binary, the intercept captures the omitted categories.

The parentheses below each coefficient in Table 5 show p-values, computed using heteroscedastic-consistent standard errors. The adjusted R^2 for each model is shown at the bottom of the table. The sample size is 30,870,518 observations for cancel durations and 6,738,138 observations for execution durations. The intercept term captures the average speed of the omitted group, which in the most general case (Models II and IV) is a non-anticipatory, nonproprietary, manual participant, during the 50%–90% of path volume with no order modifications, and not aggressive if this is an execution message.

Focusing on the cancel regressions, the overall average cancellation speed of the omitted group is 27.361 seconds, which decreases to an average of 3.094 seconds if these are algorithmic participants. The estimates in Models I and II show that Type E and Type R participants are slower than the omitted group until we control for other order characteristics. Specifically, Model II shows that the average cancellation speed of manual Type E (Type R) traders is 38.375 (68.622)

[21] The dependent variables in Table 5 are ln(1 + duration) because the order confirmation time of a market order equals its execution time, so duration will be 0 for these orders.

TABLE 5
Participant Cancellation and Execution Speeds

Table 5 shows regression results for the speed of cancellation or execution controlling for characteristics of the participants, position within a price path, and other features. The speed of cancellations equals the difference between the time the NYMEX receives a cancellation (or confirms an execution) message and the confirmation time of the original order. Both measures are in seconds with fractional milliseconds. The dependent variable is the natural logarithm of 1 plus the cancellation (execution) speed. The variables FIRST_10%_PATH_VOL, LAST_10%_PATH_VOL, and BETWEEN_10%_50%_PATH_VOL are binary variables indicating whether the message occurs during the first 10% of path volume, last 10%, or during the first half of volume that excludes the first 10%, respectively. TYPE_E and TYPE_R are binary variables indicating whether the message is made by a Type E or Type R trader, respectively. Additional binary variables are MODIFIED_ORDER if there are any modification messages to the original order; PROPRIETARY if the message is made by a proprietary account; AGGRESSIVE_SIDE for the side submitting an execution message; and ALGORITHM (ATS), which is 1 if this side of the message was submitted by a computer. There are interaction terms in selected models that interact as described: TYPE_E_AND_FIRST_10%_VOL, TYPE_R_AND_LAST_10%_VOL, TYPE_E_AND_ATS, and TYPE_R_AND_ATS. The intercept captures the omitted categories. The *p*-values are computed using heteroscedastic-consistent standard errors and are shown in parentheses. The "Bootstrap" columns show coefficient results from 1,500 random samples in which one message for each account is chosen only once per sample. The 95% confidence intervals are shown in parentheses below the average coefficient value for these bootstrap columns. * indicates that the bootstrap coefficient is significant at the 5% level. Type E and Type R traders are defined using the 5% FDR control rate. The cancel regression sample size is 30,870,518 observations and the execution regression sample size is 6,738,138 observations.

	Dependent Variable					
	ln(1 + Cancel Speed)			ln(1 + Execution Speed)		
Independent Variable	I	II	Bootstrap	III	IV	Bootstrap
INTERCEPT	1.072 (<0.001)	3.345 (<0.001)	4.301* (4.24; 4.36)	0.884 (<0.001)	1.951 (<0.001)	2.876* (2.82; 2.93)
FIRST_10%_PATH_VOL	−0.035 (<0.001)	−0.027 (<0.001)	0.038 (−0.09; 0.16)	−0.022 (<0.001)	−0.041 (<0.001)	−0.092* (−0.16; −0.03)
LAST_10%_PATH_VOL	−0.046 (<0.001)	−0.036 (<0.001)	−0.002 (−0.12; 0.11)	−0.058 (<0.001)	−0.031 (<0.001)	0.079* (0.01; 0.15)
BETWEEN_10%_50%_PATH_VOL	−0.014 (<0.001)	−0.013 (<0.001)	0.029 (−0.05; 0.11)	−0.052 (<0.001)	−0.048 (<0.001)	−0.109* (−0.15; −0.07)
TYPE_E_TRADER	1.930 (<0.001)	0.328 (<0.001)	0.455* (0.18; 0.74)	0.570 (<0.001)	0.504 (<0.001)	0.126 (−0.07; 0.32)
TYPE_R_TRADER	1.430 (<0.001)	0.898 (<0.001)	0.423* (0.23; 0.60)	1.388 (<0.001)	1.205 (<0.001)	0.409* (0.25; 0.56)
MODIFIED_ORDER		1.094 (<0.001)	1.080* (0.96; 1.19)		0.875 (<0.001)	1.582* (1.48; 1.67)
PROPRIETARY		−0.516 (<0.001)	−1.975* (−2.05; −1.91)		−0.049 (<0.001)	−0.432* (−0.50; −0.35)
ALGORITHM (ATS)		−1.935 (<0.001)	−0.553* (−0.63; −0.48)		−0.474 (<0.001)	−0.063* (−0.10; −0.02)
AGGRESSIVE_SIDE					−1.562 (<0.001)	−2.696* (−2.73; −2.66)
Interaction Terms						
TYPE_E_AND_FIRST_10%_VOL		−0.024 (0.772)	−0.176 (−1.03; 0.72)		−0.171 (<0.001)	−0.109 (−0.46; 0.26)
TYPE_R_AND_LAST_10%_VOL		0.074 (<0.001)	−0.285 (−0.81; 0.26)		0.107 (0.004)	0.080 (−0.29; 0.46)
TYPE_E_AND_ATS		0.582 (0.029)	−0.098 (−0.72; 0.57)		−0.060 (0.01)	−0.312 (−0.65; 0.06)
TYPE_R_AND_ATS		0.113 (<0.001)	−0.043 (−0.40; 0.30)		−0.618 (<0.001)	0.269 (−0.07; 0.59)
Adj. R^2	2.8%	14.5%	12.1%	2.7%	28.8%	39.7%
Regression (*p*-value)	(<0.001)	(<0.001)	(11.4%–12.8%)	(<0.001)	(<0.001)	(38.9%–40.5%)

seconds. Type E (Type R) traders who are algorithmic are slower than the omitted algorithmic group, with an average cancellation speed of 9.169 (10.252) seconds. From these estimates, the algorithmic and modified order coefficients have the largest impact on average cancellation speeds.

Table 5 also shows that cancels during the first and last 10% of path volume are marginally faster for all traders. However, the interaction terms for Type E

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

and Type R participants show no significant effect for Type E participants and a slowdown in speed for Type R participants during the last 10% of path volume. Combined with the preceding results, these findings suggest that fast cancellation speeds are not a significant characteristic of anticipatory trading.

Table 5 includes a cancel regression under the "Bootstrap" column heading. These are averages of coefficient estimates from 1,500 random samples (with replacement) in which each trader account is chosen *once* per sample. This gives equal weight to each account so that the numerous orders from HFT participants do not skew the resulting coefficients (Fishe and Smith (2018)). The 95% confidence intervals from these simulations are shown below the average coefficient value, with an asterisk indicating significance at the 5% level.

The bootstrap results show that there are fewer significant variables in the cancel regressions. Of those that are significant, the Type E and Type R dummies continue to indicate slower cancel speeds relative to the omitted group. The skewed effects of HFT volumes are clear in the new algorithmic trader coefficient, which decreases by 71% in absolute value from Model II. Interestingly, the proprietary trader coefficient nearly quadruples, suggesting that those participants are much faster than the omitted group. While inferences about the duration of orders ending in a cancel message are correctly derived from Models I and II, inferences about participants from such message durations are better reflected in the bootstrap results. These bootstrap results continue to show that being fast as measured by cancellation speeds is not a meaningful determinant of who is an anticipatory trader.

Table 5 also presents speed results using trade executions. In addition to the variables in Table 5, we include a dummy variable for whether a participant was on the aggressive side of the trade (AGGRESSIVE_SIDE). Model IV coefficient estimates evaluated at the mean show a sample execution speed of 0.818 seconds. This is somewhat faster than equity execution speeds for NASDAQ- and NYSE-listed securities in 2011 and 2012 (Angel et al. (2015)). Part of the reason for this low average is the effect of aggressive orders, which lower the base speed of the omitted group from 6.034 to 2.233 seconds. In isolation, a manual Type E (Type R) participant is 0.655 (2.337) seconds slower with executions than the base speed of the omitted group. However, being an algorithmic trader increases these speeds, with the base group executing in an average of 3.38 seconds and the Type E (Type R) group being slower with an average speed of 5.832 (6.878) seconds. Even so, these anticipatory traders are not faster than non-anticipatory traders as measured by execution speeds.

Some of these observations no longer hold in the bootstrap results for the execution sample. These coefficients were created in the same manner as described for the cancellation sample. The interaction between Type E and Type R groups with the first and last 10% of volume is not significant, so there is no acceleration of speed in those periods. Also, the interaction terms between Type E or Type R groups and the algorithmic variable are not significant. Thus, in an equal-weighted sample, Type E and Type R algorithmic participants are no faster than All Other algorithmic participants.

E. Measuring Information and Volume Effects

Two important issues in this analysis are whether anticipatory traders have an information advantage and whether the volume of other traders is affected by these participants. We first estimate price impacts to examine whether information is provided or identified by anticipatory traders. Price impacts are computed for both aggressive and passive transactions, with the latter measuring adverse selection costs. Following the methods in Bessembinder and Venkataraman (2010), aggressive price impacts are measured relative to a true underlying value, which is estimated from midpoint price changes:

$$\text{(4)} \qquad \text{Price impact of trade} \equiv \text{PI}_{it} = 100 q_{it} \left(m_{it}^{n} - m_{it}^{0} \right) / m_{it}^{0},$$

where m_{it}^{0} is the midpoint of the bid–ask spread at the time of trade t for participant i, m_{it}^{n} is the midpoint n seconds after the tth trade, and q_{it} is a buy-sell indicator that equals –1 (1) if the trade is a sell (buy). The change in midpoints is the market's assessment of the private information in the trade. These impacts are computed on a percentage basis. The sign of q_{it} is reversed for passive-side impacts, which then are interpreted as adverse selection costs to liquidity providers (see Subrahmanyam and Zheng (2016)). To make inferences about participants, we average PI_{it} across the trades for each participant and report averages across all participants by trader type: Type E, Type R, or All Other.[22] These data are shown in Table 6 along with effective half spreads for comparison.

TABLE 6
Price Impacts by Anticipatory Participants

Table 6 shows price impacts to examine whether Type E or Type R anticipatory traders affect market price after a trade. The price impact equals the difference in spread midpoint immediately before and at either 10 or 60 seconds after a trade, scaled by the before midpoint and signed by the trade direction: +1/–1 for buy or sell aggressive side and the reverse on the passive side. Price impacts and adverse selection costs are reported in basis points. Panel A reports the aggressive-side and Panel B reports the passive-side results, where price impact captures adverse selection costs on the passive side. Effective half spreads are shown for comparison. The table is divided by the groups found using 2.5% and 5% FDR control rates. Means and medians are shown for each performance measure by anticipatory trader type and for All Other traders. Price impacts and adverse selection costs are computed across all trades for each participant and then averaged across participants. Participants with less than 10 trades with complete order book data (both begin and end records) are excluded from the sample. Participant counts are shown at the bottom of each panel. Results for pairwise *t*-tests (Satterwaithe when variances are unequal) and Wilcoxon and Median tests (based on least significant *p*-value) are reported by 5% (*) and 1% (**) levels. These tests compare the mean and median values for each anticipatory trader group to the corresponding "All Other" trader group summary measures. All values are reported in basis points.

	FDR Control Rate of 2.5%				FDR Control Rate of 5%					
	Type E		Type R		Type E		Type R		All Other	
Market Performance Measure	Mean	Median	Mean	Median	Mean	Median	Mean	Median	Mean	Median
Panel A. Participant Trades on Aggressive Side										
Price impact (10 sec)	1.565**	1.253**	0.686	0.553	1.199**	0.831**	0.453	0.387*	0.591	0.524
Price impact (60 sec)	2.224**	1.373**	0.796	0.756	1.820**	1.503**	0.691	0.672	0.534	0.542
Effective half spread	1.439**	1.131*	1.046*	0.859*	1.285*	1.069*	1.006**	0.827**	1.117	0.968
Participant counts	105		186		273		519		10,858	
Panel B. Participant Trades on Passive Side										
Impact/Adverse selection cost (10 sec)	0.801*	0.523**	1.397	1.287	1.010**	0.918**	1.335	1.215	1.487	1.213
Impact/Adverse selection cost (60 sec)	0.232*	0.057**	1.018	1.093	0.462**	0.308**	0.969*	1.029	1.458	1.171
Effective half spread	1.113*	0.926*	1.520**	1.258**	1.222	1.036	1.416**	1.216**	1.278	1.023
Participant counts	104		186		291		517		10,704	

[22] Weighting by trade size to calculate each participant's average generated similar results.

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

The results in Table 6 are for anticipatory groups selected using both the 2.5% and 5% FDR control rates. Impacts are computed for midpoints found 10 and 60 seconds after a trade. Panel A shows aggressive-side and Panel B shows passive-side results. Participants with less than 10 trades with complete order book data (both beginning and ending records) are excluded from the sample. Both averages and medians across participants are shown in basis points, with levels of significance indicated by asterisks for tests comparing anticipatory trader impacts to those for the All Other group.

Table 6 shows that price impacts average between 48% and 53% of effective half spreads for All Other participants (Panel A), which is comparable to equity results that use millisecond data and accurate quote updates.[23] For the aggressive Type E group, however, they are higher and generally exceed the effective half spread. Thus, aggressive trades by Type E participants represent a loss to liquidity providers, suggesting an information effect. The loss is greater at the 60-second interval, which is not unexpected because the Type E group trades early in a local price path. When Type E participants provide liquidity (Panel B), they face lower adverse selection costs compared to other traders, suggesting a skill for selective timing on the order book. In contrast, when the Type R group is aggressive, price impacts are no different than the All Other group, but there is weak evidence from the mean that Type R participants face lower adverse selection costs than the All Other group. These findings suggest that Type R participants do not bring price-moving information to the market, but Type E participants are likely price informed and are also skilled at finding counter-parties who are non-informed traders.

We also examine how participant speed affects price impacts and adverse selection costs. Table 7 presents these regression results. The Base model controls for the average trade size (AVG_TRADE_SIZE) executed by each participant to capture the liquidity impact from larger size trades. This model also includes a measure of average participant speed (AVG_SPEED), which is the time between order entry and execution, and an algorithmic-entry dummy variable to control for electronic trading. The Full model includes Type E and Type R dummies, interaction terms between algorithmic- and anticipatory-type traders and between participant speed and anticipatory type. Coefficient estimates appear above p-values (shown in parentheses), which are computed using heteroscedastic-consistent standard errors.

Table 7 results show that Type E anticipatory traders affect price impacts and adverse selection costs, supporting the findings in Table 6.[24] Trades from the Type E group have higher price impacts at both 10- and 60-second intervals and face lower adverse selection costs. The Type E group also is likely informed and skilled in these results, as their aggressive trades impact market price and their

[23] Holden and Jacobsen (2014) examine the differences between two versions of publicly available trade and quote data for equities, one with second and the other with millisecond time stamps. Although they use a longer impact window (5 minutes), their "first best" millisecond summaries imply that price impacts average 59%–62% of effective spreads.

[24] For interpretation purposes, the intercept of 0.0023 in the 10-second aggressive-side Base model represents 0.23 basis points.

TABLE 7

Speed Effects on Price Impacts and Adverse Selection Costs

Table 7 shows regression estimates of price impacts for aggressive trades and adverse selection costs for passive trades. Impacts and selection costs are computed from equation (4) with the future midpoint price measured at either 10 or 60 seconds after the trade and with the trade sign reversed for passive-side adverse selection costs. Independent variables are the average participant speed in seconds (AVG_SPEED), computed from the time of execution minus the time of order submission for each trade; average trade size by participant (AVG_TRADE_SIZE), equal to volume executed divided by total executed order count; dummy variables for whether the participant is Type E (TYPE_E_TRADER) or Type R (TYPE_R_TRADER), based on the 5% FDR control rate, and if the participant is an algorithmic account (ALGORITHM (ATS)); and interaction terms between Type E (TYPE_E_AND_ATS) and Type R (TYPE_R_AND_ATS) dummy variables and the algorithmic indicator, as well as the measure of average participant speed (TYPE_E_AND_AVG_SPEED, TYPE_R_AND_AVG_SPEED). In the Base model, the intercept captures the impact of all manual-entry participants, while in the Full model, it captures the impact of All Other participants who use manual-entry methods, as the effects of Type E and Type R participants are measured directly. The participant sample size is 11,650 for aggressive-side and 11,512 for passive-side regressions. The *p*-values are computed using heteroscedastic-consistent standard errors and are shown in parentheses below each coefficient. [a]indicates the coefficient is scaled up by a factor of 1,000 to ease reporting.

	Aggressive Side/Price Impact				Passive Side/Adverse Selection Cost			
	10-Second Impact		60-Second Impact		10-Second Impact		60-Second Impact	
Independent Variable	Base	Full	Base	Full	Base	Full	Base	Full
INTERCEPT	0.0023	0.0023	0.0037	0.0036	0.0110	0.0110	0.0106	0.0108
	(0.019)	(0.015)	(0.013)	(0.019)	(<0.001)	(<0.001)	(<0.001)	(<0.001)
AVG_SPEED[a]	0.0010	0.0011	0.0003	0.0004	0.0010	0.0010	0.0008	0.0008
	(0.005)	(0.005)	(0.531)	(0.486)	(<0.001)	(<0.001)	(0.096)	(0.095)
AVG_TRADE_SIZE	0.0024	0.0023	0.0012	0.0009	0.0010	0.0011	0.0016	0.0019
	(0.004)	(0.005)	(0.349)	(0.477)	(0.024)	(0.016)	(0.022)	(0.011)
ALGORITHM (ATS)	0.0015	0.0017	0.0017	0.0021	0.0015	0.0012	0.0008	0.0007
	(0.003)	(0.002)	(0.073)	(0.036)	(0.002)	(0.015)	(0.416)	(0.487)
TYPE_E_TRADER		0.0068		0.0139		−0.0045		−0.0109
		(<0.001)		(<0.001)		(0.018)		(0.012)
TYPE_R_TRADER		−0.0001		0.0046		−0.0009		−0.0040
		(0.946)		(0.080)		(0.312)		(0.045)
Interaction Terms								
TYPE_E_AND_ATS		−0.0020		−0.0019		0.0015		−0.0036
		(0.946)		(0.762)		(0.656)		(0.596)
TYPE_R_AND_ATS		−0.0020		−0.0084		0.0052		0.0050
		(0.485)		(0.056)		(0.011)		(0.122)
TYPE_E_AND_AVG_SPEED[a]		−0.0025		−0.0031		0.0009		0.0045
		(0.556)		(0.668)		(0.828)		(0.517)
TYPE_R_AND_AVG_SPEED[a]		−0.0017		−0.0035		−0.0015		−0.0051
		(0.568)		(0.583)		(0.487)		(0.310)
Adj. R^2	0.8%	1.0%	0.1%	0.3%	0.7%	0.8%	0.2%	0.4%
Regression (*p*-value)	(<0.001)	(<0.001)	(0.111)	(<0.001)	(<0.001)	(<0.001)	(<0.001)	(<0.001)

passive trades execute when subsequent prices move less against their side of the book than found with passive orders by All Other participants.

The Type R results show significantly lower selection costs for passive trades and positive, marginally significant (10% level) price impacts for aggressive trades, both at 60-second intervals. Those who are algorithmic have higher selection costs at 10-second intervals. This evidence supports the view that while the Type R group is successful at identifying price reversals, the faster members may act more than 10 seconds before a reversal, too soon to catch the reversal, but not too soon in the 60-second interval. These results support those in Table 6 and make it difficult to claim that the Type R group brings price information to the market. Because they are anticipatory, they are skilled at predicting reversals, which requires knowledge of price pressure and when such pressure may abate (see Hendershott and Menkveld (2014)); however, compared to Type E

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

participants, they are less skilled at determining whether incoming orders are from noise or liquidity traders.

Estimates for both Base and Full models also show that participant speed has a positive (but small) effect on all participants' price impacts and adverse selection costs at 10-second intervals. In effect, aggressive trades from slower participants have greater price impacts than those from relatively faster participants, which implies that price-moving orders are not confined to faster traders. At 60 seconds, the effects are not significant. Also, there is no significant interaction effect for speed and anticipatory trader types, even at the 10-second interval. Thus, the speed of anticipatory participants does not influence the magnitude of their price impacts or adverse selection costs.

Algorithmic participants are faster than manual-entry traders. Table 7 shows that they generate greater price impacts and adverse selection costs at 10-second intervals, but only significant price impact effects at the 60-second interval. These results do not vary when algorithmic anticipatory traders are isolated with interaction terms, except that Type R participants affect adverse selection costs at the 10-second interval. In addition, the faster traders in the passive results face lower adverse selection costs (faster implies the participant speed variable is smaller) but there is no differential effect for Type E or Type R participants. This confirms that speed matters, on average, but it is not a consistent factor in the price impact or adverse selection costs of Type E or Type R participants.

Of the remaining variables, trade size effects are significant and positive for 10-second intervals but less meaningful at 60 seconds, as other volume has accumulated. This result is expected, as larger trade sizes are found to have larger impacts on liquidity (see, e.g., Bessembinder and Kaufman (1997)), which would be more important over a shorter interval.

Our second results address whether anticipatory traders affect the volume of trading by other market participants. If such traders increase volume by other participants, then price moves may be "caused" by their anticipatory trading, a momentum-ignition effect. In this event, they are not "anticipatory" because, instead of indirectly anticipating other's actions, they have caused them. Like price impacts, we examine how volume changes around participant trades. Table 8 shows regression results for volume impacts, which are measured by volume over 10 seconds (60 seconds) after a trade minus volume 10 seconds (60 seconds) before the trade, divided by volume 10 seconds (60 seconds) before the trade. This relative change in volume is averaged for each participant across aggressive and passive trades, separately, similar to the price impact calculations. Independent variables are those used in Table 7.

The regression p-values in Table 8 show that these independent variables are not significant when volume changes are compared at 10-second intervals. For 60-second intervals, we find no effects for Type E participants but that aggressive Type R participants are associated with lower subsequent volumes. This is the opposite of a momentum-ignition effect. Both trade size and algorithmic participants tend to increase volumes on the aggressive side but have no effect on the passive side. The effects of faster order entry by algorithms may be less important while waiting to execute in the limit order queue but quite important for aggressive traders trying to claim the best quote. Algorithms often react simultaneously,

TABLE 8
Participant Volume Impacts and Speed

Table 8 computes volume impacts using volume over 10 seconds (60 seconds) after a trade minus volume 10 seconds (60 seconds) before the trade, divided by volume 10 seconds (60 seconds) before the trade. The relative change in volume is averaged for each participant across aggressive and passive trades, separately. Independent variables are the average participant speed in seconds (AVG_SPEED), computed from the time of execution minus the time of order submission; average trade size by participant (AVG_TRADE_SIZE), equal to volume executed divided by total executed order count; dummy variables for whether the participant is Type E (TYPE_E_TRADER) or Type R (TYPE_R_TRADER), based on the 5% FDR control rate, and if the participant is an algorithmic account (ALGORITHM (ATS)); and interaction terms between Type E (TYPE_E_AND_ATS) and Type R (TYPE_R_AND_ATS) dummy variables and the algorithmic indicator, as well as the measure of average participant speed (TYPE_E_AND_AVG_SPEED, TYPE_R_AND_AVG_SPEED). In the Base model, the intercept captures the average relative volume impact for all manual-entry participants, while in the Full model, it captures the relative volume impact of All Other participants who use manual-entry methods. The sample size is 11,458 for the aggressive side and 11,317 for the passive side. The *p*-values are computed using heteroscedastic-consistent standard errors and are shown in parentheses below each coefficient.

	Aggressive Side/Relative Volume				Passive Side/Relative Volume			
	10 Second		60 Second		10 Second		60 Second	
Independent Variable	Base	Full	Base	Full	Base	Full	Base	Full
INTERCEPT	0.4837 (<0.001)	0.4899 (<0.001)	0.1697 (<0.001)	0.1741 (<0.001)	0.7899 (<0.001)	0.7946 (<0.001)	0.2973 (<0.001)	0.2980 (<0.001)
AVG_SPEED	0.0178 (0.284)	0.0187 (0.282)	0.0319 (0.001)	0.0308 (0.002)	0.0331 (0.041)	0.0350 (0.048)	0.0504 (<0.001)	0.0503 (<0.001)
AVG_TRADE_SIZE	0.0417 (0.104)	0.0407 (0.112)	0.0373 (0.014)	0.0376 (0.013)	0.0197 (0.376)	0.0202 (0.371)	−0.0026 (0.721)	−0.0026 (0.725)
ALGORITHM (ATS)	0.0587 (0.039)	0.0646 (0.032)	0.0657 (<0.001)	0.0629 (<0.001)	−0.0098 (0.808)	−0.0128 (0.766)	0.0163 (0.279)	0.0166 (0.297)
TYPE_E_TRADER		0.0426 (0.503)		−0.0419 (0.183)		−0.1067 (0.227)		−0.0610 (0.164)
TYPE_R_TRADER		−0.1408 (0.003)		−0.0913 (<0.001)		−0.0703 (0.171)		0.0131 (0.561)
Interaction Terms								
TYPE_E_AND_ATS		0.0205 (0.893)		0.0249 (0.729)		0.1949 (0.414)		0.0319 (0.694)
TYPE_R_AND_ATS		−0.1188 (0.150)		0.0492 (0.392)		−0.0208 (0.817)		−0.0275 (0.534)
TYPE_E_AND_AVG_SPEED		0.0722 (0.360)		0.0335 (0.346)		0.0288 (0.673)		0.1235 (0.154)
TYPE_R_AND_AVG_SPEED		−0.0411 (0.455)		0.0157 (0.781)		−0.0243 (0.496)		−0.0368 (0.142)
Adj. R^2	0.1%	0.2%	0.6%	0.7%	0.0%	0.0%	0.5%	0.6%
Regression (*p*-value)	(0.084)	(0.078)	(<0.001)	(<0.001)	(0.187)	(0.635)	(<0.001)	(<0.001)

so clustering by ATS participants may increase volume. Interestingly, the average participant speed matters on both aggressive and passive sides, but the effect is that slower participants are associated with higher subsequent volumes at 60-second intervals. This result suggests that slower participants (possibly liquidity or noise traders) may attract other participants to trade.

F. Skill Retention

Last, we examine whether these traders retain their skills in a subsequent sample and whether speed plays a role in skill retention. A drawback here is that the WTI crude oil futures data from 2014 show substantial attrition among the 2011 accounts, with only 35% of the 2011 accounts included. This attrition rate is larger than that found by Boyd and Kurov (2012) in energy futures after side-by-side electronic trading began in 2006. Their data predict approximately 50% attrition after 2.5 years. Several factors contribute to attrition, including oil price volatility, rule changes under Dodd–Frank raising capital requirements and costs,

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

and flows into and out of energy markets after the 2008–2009 financial crisis.[25] Notably, the bankruptcy of MF Global on Oct. 31, 2011 forced many customers to find alternative clearing arrangements, which led to new account designations.[26] Thus, a caveat to these out-of-sample results is that they may understate the success of the two anticipatory groups identified in 2011. That is, some of these anticipatory participants may be successful under a new executing broker account number in 2014.

Table 9 summarizes our findings for the 2014 sample. Panel A shows the size of the Type E and Type R groups in both 2011 and 2014, Panel B shows the counts for those who are found in both samples, and Panel C identifies the number of overlaps who retain the anticipatory strategy. The percentage of sample figures in Panel A indicates that the anticipatory strategy was slightly more difficult to execute in 2014: There are 24% (11%) fewer anticipatory participants at the 2.5% (5%) FDR control rate, mainly due to fewer Type R participants. These

TABLE 9

Type E and Type R Participants in 2011 and 2014

Table 9 shows counts of Type E and Type R participants from FDR tests using crude oil futures sample data in 2011 and 2014. The 2011 sample is for the December expiration and includes all trading days between Sept. 12 and Nov. 18. The 2014 sample is for the October expiration and includes trading days from Aug. 1 to Sept. 18. Significant accounts are those identified as Type E or Type R using 2.5% and 5% control rates (CR) in the FDR tests. Panel A identifies Type E and Type R participants found in the two samples and the sample size. Panel B identifies the number of Type E and Type R participants who traded in both samples. There were 2,669 accounts identified as trading in both samples. Panel C identifies whether the Type E and Type R participants found in 2011 retained those characteristics in the 2014 sample. The last two columns in Panel C report the retention counts by control rate and the percentage of that count relative to the "type" totals observed in 2014 for the overlapping accounts.

	2011 Sample		2014 Sample	
Descriptor	2.5% CR	5% CR	2.5% CR	5% CR
Panel A. Type E and Type R Participants in 2011 and 2014 Samples				
Type E participants	112	301	101	322
Percentage of sample	1.5%	4.0%	1.2%	4.0%
Type R participants	196	542	133	425
Percentage of sample	2.6%	7.2%	1.6%	5.3%
Sample size	7,554		8,087	
Panel B. Only Accounts Included in Both 2011 and 2014 Samples				
Type E participants	35	81	38	95
Percentage of sample	1.3%	3.0%	1.4%	3.6%
Type R participants	65	138	50	131
Percentage of sample	2.4%	5.2%	1.9%	4.9%
Sample size	2,669		2,669	
Panel C. Skill Retention in 2014 for Accounts in Both Samples				
Type E in 2011			1 (2.6%)	7 (7.4%)
Type R in 2011			5 (10.0%)	27 (20.6%)

[25] These factors also contributed to a decline in Futures Commission Merchants (FCM) and rotation of customers to new executing brokers. See McCrank (2015) or https://www.brookings.edu/research/dwindling-numbers-in-the-financial-industry/.

[26] At the time, MF Global was the eighth largest bankruptcy in U.S. history. It managed over 36,000 customer accounts, many of which were involved with futures and options trading (Gibbons (2012)). Approximately 800 of our 2011 sample participants used MF Global as the executing broker, and we suspect that many of these entities are still active but trading with another commodities broker.

decreases are not as large for the 2011 participants that overlap the 2014 data, as shown in Panel B. In effect, those who do continue to trade 3 years later are more experienced and appear better able to maintain a strategy compared to newer participants.

Only a small number of the Type E overlaps (7 participants at the 5% control rate (CR), which includes 1 at the 2.5% CR) continue to implement that strategy in 2014. Skill retention for Type R participants is higher (27 participants or 21%), indicating that this group may have adopted a more durable strategy. As noted, these counts are likely lower bounds because many participants had to find new executing brokers during the intervening years. Thus, some unidentified accounts in 2011 may be among those chosen as successful in the 2014 sample. Even so, those that are identified as overlapping both samples indicate that the Type E, or early trading, strategy was the relatively more difficult strategy to maintain.

Finally, we examine the speed of the 2,669 participants who overlapped between the 2011 and 2014 samples using data for both samples. Table 10 reports the results following the regression methods and including the control variables discussed in Table 5. Results are shown for both cancellation and execution speeds using a combined 2011 and 2014 sample. We include binary variables for those who are Type E or Type R participants in the 2011 or 2014 samples and an overlap variable for those who are anticipatory in both samples. We also include interaction terms to isolate the anticipatory participants who are algorithmic traders. Both trade-based and bootstrap-by-participant coefficient estimates are presented. Our focus is on the bootstrap results, as those are for an average participant.

Table 10 shows that manual-entry cancel and execution speeds decrease between samples, the latter results offering support to Hoffmann's (2014) theory that slower traders will place orders with lower execution probabilities when competing with faster traders. Algorithmic-entry speeds are faster overall, and these speeds increased for cancellations and executions in 2014. Aggressive traders have significantly faster execution speeds but modifying an order slows execution speeds as well as cancel speeds. In contrast, being a proprietary trader implies faster cancellations. Cancel speeds for Type E participants are slower in 2011 but are no different than other participants in 2014. In addition, execution speeds for Type E participants show no differences compared to other participants in both samples. The interaction terms between Type E and algorithmic entries also show no differences compared to other participants. These results support the view that speed does not play a role in whether a participant continues with the Type E strategy in 2014.

For manual-entry Type R participants, cancel and execution speeds are slower in 2011. Those who are Type R in 2014 are only slower with cancellations compared to All Other participants. No other variables for the Type R groups are found to be significant, including the interaction effect with algorithms. These findings support the results for Type E traders; that is, speed does not play a role in whether Type R traders are more likely to retain such strategies in 2014.

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

TABLE 10

Speed Characteristics of Participants in Both 2011 and 2014 Samples

Table 10 shows regression estimates of cancel and execution speeds after combining the 2011 and 2014 samples. The dependent variable is the natural logarithm of 1 plus the cancellation (execution) speed. Binary variables identify observations in the 2014 sample (SAMPLE_2014), if the trader is an algorithm (ALGORITHM (ATS)), if the order was modified before execution or cancellation (MODIFIED_ORDER), whether the trader uses a proprietary account (PROPRIETARY), if the execution was the aggressive side (AGGRESSIVE_SIDE), if the trader is Type E or Type R in 2011 or 2014 samples (TYPE_E_2011, TYPE_E_2014, TYPE_R_2011, TYPE_R_2014), and if the trader is Type E or Type R in both 2011 and 2014 samples. There are interaction terms between algorithmic traders and i) the 2014 sample, ii) Type E participants in the 2011 or 2014 samples, iii) Type R participants in the 2011 or 2014 samples, and iv) Type E or Type R traders that are these types in both samples. The sample size is 2,669 accounts, all of which traded in both 2011 and 2014. The intercept captures the omitted categories, which are manual-entry participants. For these regressions, Type E and Type R traders are defined using the 5% FDR control rate. The *p*-values are computed using heteroscedastic-consistent standard errors and are shown in parentheses below each coefficient. The "Bootstrap" column shows coefficient results from 1,500 random samples in which one trade for each account is chosen only once per sample. The 95% confidence intervals are shown below the average bootstrap coefficient value. * indicates that the bootstrap coefficient is significant at the 5% level. The sample size of the cancel (execution) speed regression is 17,732,173 (4,445,802).

	Dependent Variable					
	ln(1 + Cancel Speed)			ln(1 + Execution Speed)		
Independent Variable	I	II	Bootstrap	III	IV	Bootstrap
INTERCEPT	3.707 (<0.001)	3.209 (<0.001)	4.346* (4.24; 4.43)	2.417 (<0.001)	2.277 (<0.001)	2.776* (2.66; 2.89)
SAMPLE_2014	−0.262 (<0.001)	0.122 (<0.001)	0.276* (0.16; 0.40)	0.984 (<0.001)	1.007 (<0.001)	0.797* (0.68; 0.91)
ALGORITHM (ATS)	−2.839 (<0.001)	−2.355 (<0.001)	−1.221* (−1.38; −1.06)	−0.813 (<0.001)	−0.688 (<0.001)	−0.342* (−0.49; −0.18)
SAMPLE_2014_AND_ATS	0.963 (<0.001)	0.571 (<0.001)	−1.000* (−1.23; −0.76)	−0.784 (<0.001)	−0.826 (<0.001)	−0.977* (−1.17; −0.78)
MODIFIED_ORDER	1.680 (<0.001)	2.049 (<0.001)	0.959* (0.00; 4.89)	4.316 (<0.001)	4.368 (<0.001)	4.043* (3.66; 4.41)
PROPRIETARY	−0.141 (<0.001)	−0.190 (<0.001)	−1.763* (−2.61; −0.83)	1.015 (<0.001)	1.029 (<0.001)	−0.738 (−1.39; 0.02)
AGGRESSIVE_SIDE				−1.668 (<0.001)	−1.643 (<0.001)	−1.871* (−1.97; −1.77)
TYPE_E_2011		1.376 (<0.001)	0.690* (0.19; 1.18)		0.292 (<0.001)	−0.169 (−0.51; 0.18)
TYPE_E_2014		0.875 (<0.001)	−0.076 (−0.45; 0.31)		0.332 (<0.001)	−0.155 (−0.48; 0.18)
TYPE_E_2011_2014		−1.330 (<0.001)	−0.858 (−2.09; 0.46)		−0.306 (<0.001)	−0.034 (−1.11; 1.18)
TYPE_R_2011		0.998 (<0.001)	0.304* (0.02; 0.59)		1.396 (<0.001)	0.362* (0.08; 0.67)
TYPE_R_2014		0.630 (<0.001)	0.512* (0.19; 0.83)		0.196 (<0.001)	−0.043 (−0.34; 0.26)
TYPE_R_2011_2014		−0.346 (<0.001)	−0.028 (−0.67; 0.61)		−0.197 (<0.001)	0.584 (−0.15; 1.25)
ATS Interaction with						
TYPE_E_2011		−0.402 (<0.001)	−0.624 (−1.13; 0.96)		0.100 (<0.001)	0.109 (−0.46; 0.66)
TYPE_E_2014		−0.311 (<0.001)	−0.700 (−1.45; 0.08)		−0.815 (<0.001)	−0.224 (−0.70; 0.29)
TYPE_E_2011_2014		−1.266 (<0.001)	−0.961 (−2.88; 1.22)		0.280 (<0.001)	−0.003 (−1.43; 1.38)
TYPE_R_2011		−0.073 (0.002)	−0.024 (−0.58; 0.54)		−1.517 (<0.001)	0.039 (−0.20; 0.99)
TYPE_R_2014		0.915 (<0.001)	0.644 (−0.19; 1.55)		0.467 (<0.001)	0.058 (−0.49; 0.62)
TYPE_R_2011_2014		−0.186 (<0.001)	2.340 (−0.13; 4.56)		−0.115 (<0.001)	−1.066 (−2.32; 0.27)
Adj. R^2	8.5%	9.4%	11.2%	28.8%	29.6%	19.5%
Regression (*p*-value)	(<0.001)	(<0.001)	(10.0% −12.4%)	(<0.001)	(<0.001)	(17.8% −21.2%)

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

VI. Conclusions

We identify and examine a select group of participants who can anticipate short-term, intraday price movements. The purpose is to understand whether such strategies require high-speed actions or, rather, are populated with HFTs. We identify two groups of anticipatory traders: Type E are prescient early in a price path and Type R trade immediately before a reversal in the price path. Sharpe ratios show that their risk-adjusted performance exceeds that of other participants. Importantly, there are both algorithmic- and manual-entry traders within these groups, so speed, per se, is neither necessary nor sufficient to conduct anticipatory trading.

Regression results show that several factors affect message speed. Algorithmic and proprietary traders are faster than manual-entry, customer-directed participants. Those who use modifications are much slower than other participants are, and aggressive traders have demonstrably faster execution speeds, likely because they are entering market orders. Type E and Type R groups are found to be slower than other traders, confirming the main result that speed is not an essential characteristic of anticipatory trading.

We also investigate whether anticipatory traders are informed and whether they cause momentum trading by others. There is price impact evidence that Type E traders may be informed, and strong evidence that they are skilled at timing passive trades to avoid adverse selection costs. Type R traders show little evidence of an information effect and are, at best, weakly skilled at avoiding adverse selection costs. Volume momentum-ignition effects are not found for these anticipatory traders.

Last, we examine whether anticipatory traders retain their skills 3 years later in a similar sample. Attrition from the original sample qualifies these claims. Even so, the evidence suggests that the skills held by Type E participants are relatively ephemeral but that Type R participants show more durability with their strategies, with speed playing no role in strategy survival for either group of participants.

References

Allen, F., and D. Gale. "Stock-Price Manipulation." *Review of Financial Studies*, 5 (1992), 503–529.

Angel, J. J.; L. E. Harris; and C. S. Spatt. "Equity Trading in the 21st Century: An Update." *Quarterly Journal of Finance*, 5 (2015), 1–39.

Bajgrowicz, P., and O. Scaillet. "Technical Trading Revisited: False Discoveries, Persistence Tests, and Transaction Costs." *Journal of Financial Economics*, 106 (2012), 447–684.

Barras, L.; O. Scaillet; and R. Wermers. "False Discoveries in Mutual Fund Performance: Measuring Luck in Estimated Alphas." *Journal of Finance*, 65 (2010), 179–216.

Benjamini, Y., and Y. Hochberg. "Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing." *Journal of the Royal Statistical Society, Series B*, 57 (1995), 289–300.

Bessembinder, H., and H. Kaufman. "A Comparison of Trade Execution Costs for NYSE and NASDAQ-Listed Stocks." *Journal of Financial and Quantitative Analysis*, 32 (1997), 287–310.

Bessembinder, H., and K. Venkataraman. "Bid–Ask Spreads: Measuring Trade Execution Costs in Financial Markets." In *Encyclopedia of Quantitative Finance*, R. Cont, ed. New York, NY: John Wiley & Sons (2010), 184–189.

Biais, B., and T. Foucault. "HFT and Market Quality." *Bankers, Markets & Investors*, 128 (2014), 5–19.

Biais, B.; T. Foucault; and S. Moinas. "Equilibrium Fast Trading." *Journal of Financial Economics*, 116 (2015), 292–313.

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

Bikhchandani, S.; D. Hirshleifer; and I. Welch. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades." *Journal of Political Economy*, 100 (1992), 992–1026.
Boyd, N. E., and A. Kurov. "Trader Survival: Evidence from the Energy Futures Market." *Journal of Futures Markets*, 32 (2012), 809–836.
Brogaard, J., and C. Garriott. "High Frequency Trading Competition." *Journal of Financial and Quantitative Analysis*, forthcoming (2019).
Brogaard, J.; T. Hendershott; and R. Riordan. "High Frequency Trading and Price Discovery." *Review of Financial Studies*, 27 (2014), 2267–2306.
Brunetti, C.; B. Büyükşahin; and J. H. Harris. "Speculators, Prices, and Volatility." *Journal of Financial and Quantitative Analysis*, 51 (2016), 1545–1574.
Brunnermeier, M., and L. Pedersen. "Predatory Trading." *Journal of Finance*, 60 (2005), 1825–1863.
Carrion, A. "Very Fast Money: High Frequency Trading on the NASDAQ." *Journal of Financial Markets*, 16 (2013), 680–711.
Chan, L. K. C.; N. Jegadeesh; and J. Lakonishok. "Momentum Strategies." *Journal of Finance*, 51 (1996), 1681–1713.
Clark-Joseph, A. "Exploratory Trading." Working Paper, University of Illinois at Urbana–Champaign (2012).
De Long, J. B.; A. Shleifer; L. Summers; and R. Waldman. "Positive Feedback Investment Strategies and Destabilizing Rational Speculation." *Journal of Finance*, 14 (1990), 379–395.
Dong, L.; C. Polk; and S. Skouras. "A Tug of War: Overnight versus Intraday Expected Returns." *Journal of Financial Economics,* forthcoming (2019).
Easley, D.; M. L. de Prado; and M. O'Hara. "Discerning Information from Trade Data." *Journal of Financial Economics*, 120 (2016), 269–286.
Fishe, R. P. H., and A. D. Smith. "Identifying Informed Traders in Futures Markets." *Journal of Financial Markets*, 15 (2012), 329–359.
Fishe, R. P. H., and A. D. Smith. "High Frequency Trading of Commodities." In *Commodities: Markets, Performance, and Strategies*, H. K. Baker, G. Filbeck, and J. Harris, eds. New York, NY: Oxford University Press (2018), 447–468.
Foucault, T.; R. Kozhan; and W. W. Tham. "Toxic Arbitrage." *Review of Financial Studies*, 30 (2014), 1053–1094.
Frino, A.; G. Ibikunle; V. Mollica; and T. Steffen. "Anticipatory Trading in Brent Futures: Evidence from the Unregulated Dated Brent Benchmark." Working Paper, University of Edinburgh (2016).
Gibbons, J. W. "Report of the Trustee's Investigation and Recommendation: MF Global." United States Bankruptcy Court, Southern District of New York, Case No. 11-2790 (MG) SIPA (June 4, 2012).
Hagströmer, B.; L. Nordén; and D. Zhang. "How Aggressive Are High-Frequency Traders?" *Financial Review*, 49 (2014), 395–419.
Han, J.; M. Khapko; and A. Kyle. "Liquidity with High-Frequency Market Making." Working Paper, Stockholm School of Economics (2014).
Harvey, C. R.; Y. Liu; and H. Zhu. "… and the Cross-Section of Expected Returns." *Review of Financial Studies*, 29 (2014), 5–68.
Hasbrouck, J., and G. Saar. "Low-Latency Trading." *Journal of Financial Markets*, 16 (2013), 646–679.
Hautsch, N. *Econometrics of Financial High-Frequency Data*. New York, NY: Springer (2012).
Haynes, R., and J. S. Roberts. "Automated Trading in Futures Markets." White Paper, Office of the Chief Economist, U.S. Commodity Futures Trading Commission (2017).
Hendershott, T.; C. Jones; and A. Menkveld. "Does Algorithmic Trading Improve Liquidity?" *Journal of Finance*, 66 (2011), 1–33.
Hendershott, T., and A. Menkveld. "Price Pressures." *Journal of Financial Economics*, 114 (2014), 405–423.
Hirschey, N. H. "Do High-Frequency Traders Anticipate Buying and Selling Pressure?" IFA Working Paper, London Business School (2016).
Hirshleifer, D.; A. Subrahmanyam; and S. Titman. "Security Analysis and Trading Patterns When Some Investors Receive Information before Others." *Journal of Finance*, 49 (1994), 1665–1698.
Hoffmann, P. "A Dynamic Limit Order Market with Fast and Slow Traders." *Journal of Financial Economics*, 113 (2014), 156–169.
Holden, C. W., and S. Jacobsen. "Liquidity Measurement Problems in Fast, Competitive Markets: Expensive and Cheap Solutions." *Journal of Finance*, 69 (2014), 1747–1785.
Jiang, J. G.; I. Lo; and G. Valente. "High-Frequency Trading around Macroeconomic News Announcements: Evidence from the U.S. Treasury Market." Working Paper, Bank of Canada (2013).
Kirilenko, A.; A. Kyle; M. Samadi; and T. Tuzun. "The Flash Crash: High Frequency Trading in an Electronic Market." *Journal of Finance*, 72 (2017), 967–998.

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press

Lo, A. W. "The Statistics of Sharpe Ratios." *Financial Analysts Journal*, 58 (2002), 36–52.
McCrank, J. "Ranks of Commodities Brokers Dwindle as U.S. Futures Industry Evolves." *Reuters News Service* (July 2, 2015).
Menkveld, A. J. "High Frequency Trading and the *New Market* Makers." *Journal of Financial Markets*, 16 (2013), 712–740.
Menkveld, A. J. "The Economics of High-Frequency Trading: Taking Stock." *Annual Review of Financial Economics*, 8 (2016), 1–24.
Menkveld, A. J., and B. Z. Yueshen. "The Flash Crash: A Cautionary Tale About Highly Fragmented Markets." *Management Science*, forthcoming (2019).
Menkveld, A. J., and M. Zoican. "Need for Speed? Exchange Latency and Liquidity." *Review of Financial Studies*, 30 (2017), 1188–1228.
Nagel, S. "Evaporating Liquidity." *Review of Financial Studies*, 25 (2012), 2005–2039.
O'Hara, M. "High Frequency Market Microstructure." *Journal of Financial Economics*, 116 (2015), 257–270.
Raman, V.; M. A. Robe; and P. K. Yadav. "Liquidity Provision under Stress: The Fast, the Slow, and the Dead." Working Paper, American University (2016).
Storey, J. D. "A Direct Approach to False Discovery Rates." *Journal of the Royal Statistical Society*, 64 (2002), 479–498.
Subrahmanyam, A., and H. Zheng. "Limit Order Placement by High-Frequency Traders." *Borsa Istanbul Review*, 16 (2016), 185–209.
U.S. Securities and Exchange Commission. "Equity Market Structure Literature Review, Part II: High Frequency Trading." Staff of the Division of Trading and Markets (Mar. 18, 2014).
Wald, A. "Sequential Tests of Statistical Hypotheses." *Annals of Mathematical Statistics*, 16 (1945), 117–186.
Wald, A. "Optimum Character of the Sequential Probability Ratio Test." *Annals of Mathematical Statistics*, 19 (1948), 326–339.
Ye, M.; C. Yao; and J. Gai. "The Externalities of High Frequency Trading." Working Paper, University of Illinois at Urbana–Champaign (2013).

https://doi.org/10.1017/S0022109018000832 Published online by Cambridge University Press