

# governmentattic.org

"Rummaging in the government's attic"

Description of document:

Four (4) studies performed under contract to the Commodity Futures Trading Commission (CFTC):

- Identifying Informed Traders in Futures Markets, 2011
- <u>Algorithmic Finance: Discovering the ecosystem of an</u> <u>electronic financial market with a dynamic machine-</u> <u>learning method</u>, 2013
- <u>Gaussian Process-Based Algorithmic Trading Strategy</u> <u>Identification</u>, 2015
- <u>The Flash Crash: High-Frequency Trading in an</u> <u>Electronic Market</u>, 2017

Requested date:

Release date: 03-June-2019

Posted date: 09-September-2019

Source of document: Freedom of Information Act Request FOIA Compliance Office Commodity Futures Trading Commission Three Lafayette Centre 1155 21st Street, NW Washington, DC 20581 CFTC FOIA Online Request Form

17-December-2018

The governmentattic.org web site ("the site") is a First Amendment free speech web site, and is noncommercial and free to the public. The site and materials made available on the site, such as this file, are for reference only. The governmentattic.org web site and its principals have made every effort to make this information as complete and as accurate as possible, however, there may be mistakes and omissions, both typographical and in content. The governmentattic.org web site and its principals shall have neither liability nor responsibility to any person or entity with respect to any loss or damage caused, or alleged to have been caused, directly or indirectly, by the information provided on the governmentattic.org web site or in this file. The public records published on the site were obtained from government agencies using proper legal channels. Each document is identified as to the source. Any concerns about the contents of the site should be directed to the agency originating the document in question. GovernmentAttic.org is not responsible for the contents of documents published on the website.

-- Web site design Copyright 2007 governmentattic.org --

### **U.S. COMMODITY FUTURES TRADING COMMISSION**



Three Lafayette Centre 1155 21<sup>st</sup> Street, NW, Washington, DC 20581 www.cftc.gov

June 3, 2019

RE: 19-00031-FOIA Digital/electronic copy of the final report and final presentation for studies performed under contract to CFTC

This is in response to your request dated December 17, 2018 under the Freedom of Information Act seeking access to copies of the final report and final presentation for various studies performed under contract to the CFTC. In accordance with the FOIA and agency policy, we have searched our records, as of December 21, 2018, the date we received your request in our FOIA office.

We have located 163 pages of responsive records. You are granted full access to the responsive records, which are enclosed.

If you have any questions about the way we handled your request, or about our FOIA regulations or procedures, please contact Rosemary Bajorek at 202-418-5912, or Jonathan Van Doren, our FOIA Public Liaison, at 202-418-5505.

If you are not satisfied with this response to your request, you may appeal by writing to Freedom of Information Act Appeal, Office of the General Counsel, Commodity Futures Trading Commission, Three Lafayette Centre, 8<sup>th</sup> Floor, 1155 21<sup>st</sup> Street, N.W., Washington, D.C. 20581, within 90 days of the date of this letter. Please enclose a copy of your original request and a copy of this response.

Additionally, you may contact the Office of Government Information Services (OGIS) at the National Archives and Records Administration to inquire about the FOIA mediation services they offer. The contact information for OGIS is as follows: Office of Government Information Services, National Archives and Records Administration, Room 2510, 8601 Adelphi Road, College Park, Maryland 20740-6001, email at ogis@nara.gov; telephone at 202-741-5770; toll free at 1-877-684-6448; or facsimile at 202-741-5769.

Sincerely,

Rosemary Bajorek Attorney-Advisor

## **Identifying Informed Traders in Futures Markets**

**Raymond P. H. Fishe** 

Aaron D. Smith

First Draft: April 2010 This Version: March 2011

#### ABSTRACT

We use daily positions of futures market participants to identify informed traders. These data cover the period from 2000 to mid-2009 and contain 8,921 unique traders. We identify between 94 and 230 traders as overnight informed and 91 as intraday informed with little overlap between these two groups. Floor brokers/traders are over-represented in the overnight informed group, suggesting that ability to process order flow information creates success at this horizon. The intraday informed group is dominated by managed money traders/hedge funds and swap dealers, with commercial hedgers significantly under-represented in this group. Also, we find that trader characteristics such as experience, average position size, amount of trading activity, and type of positions held offer significant predictive power for who is informed. An analysis of daily trader profits confirms that our methods select highly profitable traders.

*Keywords:* Commodities, False Discovery Rate, Forecasting Ability, Informed Traders

JEL classification: G10, G13

**Fishe**: CFTC and Department of Finance, Robins School of Business, University of Richmond, Richmond, VA 23173. Tel: (+1) 804-287-1269. Email: pfishe@richmond.edu. **Smith**: CFTC and Department of Agricultural and Resource Economics, University of California - Davis, Davis, CA 95616, Tel: (+1) 530-752-2138. Email: adsmith@ucdavis.edu. We thank seminar participants at the University of Illinois – Urbana-Champaign, University of California – Davis, NBER-NSF Time Series Meeting, Commodity Futures Trading Commission, and especially Brad Barber, Roger Edelen, Steve Kane and David Reiffen for detailed thoughts. All errors are our own responsibility. This paper reflects the opinion of its authors only, and not that of the CFTC, the Commissioners, or any other staff at the CFTC.

## **Identifying Informed Traders in Futures Markets**

#### I. INTRODUCTION

Informed traders are an essential feature of market microstructure models, but there is little research that establishes who is an informed trader.<sup>1</sup> Such traders are thought to be anonymous to market makers and other participants, but eventually the information they possess is impounded in market prices. Most researchers detect the presence of informed traders from price responses to order flow. Because permanent price responses signal informed trades, consistent profits gained from positions or trading activity provide an indicator of who is informed. However, data limitations make trader identities unavailable to most previous studies so the characteristics and profits of the informed are generally unknown.<sup>2</sup>

In this paper, we identify informed traders via profits using a unique dataset that provides trader positions and selected trader characteristics in futures markets. We find traders whose actions show they hold valuable short-term price information and develop methods to separate these cases from thousands of other participants whose profits (if any) arose due to luck in the sampling process. From the subset we identify as informed, we use inverse regression techniques

<sup>&</sup>lt;sup>1</sup> Research that identifies specific informed traders is usually based on insider trading cases (e.g., Meulbroek (1992), Cornell and Sirri (1992) and Fishe and Robe (2004)), but none of these cases address informed trading in futures markets. Other researchers have examined whether groups are differentially informed, such as institutional traders (Chakravarty (2001)) or floor brokers (Anand and Subrahmanyam (2008)), but have not isolated which participants in these groups cause their results. Exceptions include Keim and Madhavan (1995), who analyze informed and other motives for trading by 21 institutions in which they have order flow data, and Boehmer, Jones and Zhang (2008), who find that non-program short sales by institutions and proprietary traders yield significant alphas at up to 60-day holding periods.

<sup>&</sup>lt;sup>2</sup> The extensive research that examines market behavior and documents the presence of informed traders include trade indicator models (e.g., Glosten and Harris (1988) and Huang and Stoll (1997)), variance decomposition models (e.g., George, Kaul, and Nimalendran (1991)), vector autoregression models (e.g., Hasbrouck (1988) and Menkhoff and Schmeling (2010)), structural likelihood models (e.g., Easley, Kiefer, O'Hara, and Paperman (1996)), as well as general aggregate models and special cases (e.g., Evans and Lyons (2002) and Ito, Lyons, and Melvin (1998)). Also, Hasbrouck (1995) develops models to identify the market that is the source of price discovery, which by implication implies that that market (or market structure) attracts informed traders.

to analyze their characteristics and to examine how their net trading and end-of-day positions set them apart from other participants. This approach complements that of Menkhoff and Schmeling (2010), who use the model of Hasbrouck (1991) to relate price impacts from trades in the dollar/rouble FX market to trader and market characteristics such as order size, timing, total volume, and origin.<sup>3</sup>

We use data on trader positions from 2000 to mid-2009 for twelve futures markets. We find that traders who hold information about intraday price changes are not the same as those who hold information about the next day's price—the overnight informed. Depending on the test and reference price, we identify from 94 to 230 traders as overnight informed and 91 traders as intraday informed out of 8,921 unique traders. Thus, our methods parse informed traders into two groups for short horizons, one day or less.

These two types of informed traders are analogous to the *ex ante* and *ex post* notion of informed trading.<sup>4</sup> The ex ante informed are thought to be those who *possess* a fairly precise signal about future returns, such as insiders trading in advance of an earnings announcement. The ex post informed are thought to be those who *process* information about order flows to develop accurate predictions of future returns (e.g., Evans and Lyons (2002, 2008) in foreign exchange markets, Brandt and Kavajecz (2004) for U.S. Treasury securities, and Deuskar and Johnson (2010) for the S&P 500 index). Our results suggest that the overnight informed are

<sup>&</sup>lt;sup>3</sup> Our results also complement recent studies of equity markets that focus on investor characteristics, particularly those that analyze retail brokerage records and mutual fund returns. Odean (1999) and Barber and Odean (2000) find that retail investors in equity markets tend to lose money trading, which would suggest that they lack information as a group. Grinblatt, Keloharju and Linnainmaa (2009) find that high-IQ investors in Finland outperform low-IQ investors in stock-picking ability. Nicolosi, Peng and Zhu (2009) find that individual investors learn from past trading experience to become better investors.

<sup>&</sup>lt;sup>4</sup> Examples of research on *ex ante* informed trading include Kim and Verrecchia (1991), Easley, Kiefer, O'Hara and Paperman (1996), Fishe and Robe (2004), Pasquariello and Vega (2007), Boehmer, Jones and Zhang (2008), and Christophe, Ferri, and Hsieh (2010). Examples of research based on *ex post* processing of information include Harris and Raviv (1993), Kandel and Pearson (1995), Kim and Verrecchia (1997), Green (2004) and Love and Payne (2008).

efficient processors of information like the ex post informed; these results are also consistent with a noisy rational expectations model of trading (e.g., Grundy and Kim (2002)). The intraday informed are like the ex ante informed as they appear to possess the best signals about very short horizon price changes and trade to capitalize on this knowledge.

We use end-of-day positions and the daily change in these positions to identify informed traders. As such, our methods identify a *subset* of the informed population: those whose information advantage can be extracted from overnight holdings and net daily trading. Informed traders whose information advantage is realized by specific intraday trades, such as measured by price impact models (e.g., Hasbrouck (1991)) may not be selected by our methods, unless their net daily trading reflects these informed trades. Importantly, our statistical procedure avoids confounding the lucky with the informed by imposing a conservative standard for identifying the informed. Thus, in order to be certain about who is selected, we may miss some informed traders who had only moderate trading success or who were successful only for a short time.

The small numbers of informed traders we identify represent a larger fraction of open interest than their numbers suggest. For example, overnight informed traders comprise between 1.2 and 3.3 percent of open interest depending on the test versus the 1.1 and 2.8 percent suggested by their numbers. Thus, they generally hold larger positions than the average trader. Correspondingly, we find that the average daily profits per trader are significantly greater for the informed than those of the not informed group. For the overnight informed in crude oil, we find the largest differences in average daily profits: \$45,237 for the informed versus losses \$3,401 for the uninformed. For intraday informed, the largest difference in average daily profits is found in natural gas: \$95,885 for the informed versus losses of \$2,693 for the not informed.

Our results also show that floor brokers/traders (FBT) are over-represented and commercial firms—those with an underlying reason to hedge—tend to be under-represented among overnight informed traders. These findings support those of Kurov and Lasser (2004) for exchange locals and Anand and Subrahmanyam (2008) for floor traders and specialists in equity markets.<sup>5</sup> The overnight informational advantage of FBTs is likely to stem from their access to order flow as they are not likely to have better access to market fundamentals than commercial firms, nor are they likely to conduct superior analysis compared to hedge funds. Consistent with the latter observation, we also find that the intraday informed are dominated by money managers/hedge funds (MMT) and swap dealers, with commercial firms again significantly under-represented in these results.

Although line-of-business variables are a significant predictor of representation in our informed counts, variables that measure trader characteristics have even stronger predictive power, consistent with Menkhoff and Schmeling's (2010) results. Specifically, a trader's experience (time in the market), frequency of trading activity, and position size have significant effects on representativeness in an informed group, as do the types of positions taken. Using inverse regression methods (e.g., Li (1991)), we estimate that intraday informed traders have 15% more experience, 39% more activity and hold 58% larger positions than the average trader. Informed traders, particularly overnight informed, are generally more likely to trade on both sides of the market (i.e., both long and short).

<sup>&</sup>lt;sup>5</sup> Our methods are different from Kurov and Lasser (2004) and Anand and Subrahmanyam (2008) as they examine information shares for traders in a group, not the sequence of positions or trades for specific participants. Effectively, they classify traders and then ask if the group is informed. In contrast, we ask who is informed based on evaluating the forecasting ability of each trader and then determine the characteristics of those we have identified relative to the population. Not everyone is informed in a group, so we avoid group biases by first starting with individual data.

We also find that simultaneously holding positions in more contract expirations affects representation in the informed group, but this effect differs between the overnight and intraday informed. The overnight informed tend to hold positions in more expirations, again consistent with FBTs processing information from order flows, while the intraday informed hold positions in fewer expirations, consistent with selective trading by MMTs and swap dealers based on precise signals.

To complement our informed tests, we develop methods to identify which participants demand and supply liquidity based on intraday trading activity, and which participants consistently behave as contrarian or momentum traders. We identify 90 traders whose net position changes are sufficiently large to act like liquidity suppliers and 115 who appear to be liquidity demanders. Our analysis also finds that contrarian traders (n=912) outnumber momentum traders (n=444) by more than two-to-one.

In addition, the composition of those who are liquidity demanders and suppliers suggest that commercial firms are unlikely pay a liquidity premium to hedge and may instead tend to receive a liquidity premium. Our results show that commercial firms (hedgers) are over-represented in the group of liquidity suppliers and that MMTs are over-represented in the group of liquidity demanders. This result differs from those using earlier aggregate data, which showed that commercial hedgers brought price pressure to futures markets when they adjusted their positions (deRoon, Nijman and Veld (2000)). As our sample period includes a substantial increase in index fund participation, our evidence suggests that the normal role of commercial hedgers has changed from that of demanders to suppliers of liquidity (Harris and Buyuksahin (2009), Tang and Xiong (2009)).

This research builds on a small, but insightful body of previous work in commodity futures markets. Early studies of informed trading focused on the forecasting ability of futures traders used data on small traders at a single brokerage firm (Hieronymus (1977) and Teweles, Harlow and Stone (1977)) or placed traders into aggregate groups from monthly or semi-monthly observations (Houthakker (1957), Rockwell (1964) and Chang (1985)). Hartzmark (1987), who obtained a proprietary sample of end-of-day positions from the Commodity Futures Trading Commission (CFTC) covering nine commodities from mid-1977 to 1981, offered a more comprehensive view. He found that commercial traders as a group earn significant daily profits compared to non-commercial traders, suggesting that commercial traders are informed.

A follow-on study by Hartzmark (1991), however, finds that the daily forecasting ability of large, more frequent traders in seven futures markets is not statistically different from that generated by luck. Moreover, his findings for commercial traders show that superior forecasters are limited to the pork bellies market, a finding supported by Leuthold, Garcia and Lu (1994). Similarly, Phillips and Weiner (1994) found no evidence of daily profits, but did find some intraday profits for large integrated oil companies in the crude oil market, suggesting that commercial firms had better information. Recent research by Dewally, Ederington and Fernando (2010) finds that profits of individual traders in energy futures are largely determined by whether they hold net positions opposite those of commercial firms, which suggests that commercial firms are likely not informed traders, consistent with our results.

These previous studies differ in their conclusions on whom, if anyone is informed in futures markets. One reason for the various results is that past studies use different statistical methods. Importantly, these methods make no allowance for the multiple-testing problem found in these

studies.<sup>6</sup> The multiple-testing problem arises because there are thousands of traders in futures markets, which gives rise to thousands of test statistics, a fraction of which are expected to be significant because of chance. To control for these chance effects, we implement the false discovery rate (*FDR*) method of Benjamini and Hochberg (1995) and Storey (2002) to isolate who is informed in our tests. This method makes a meaningful difference as the median p-value across all our tests is 0.00018, which suggests that multiple-testing bias would be pronounced if we used the classical 5-percent critical values.

Our results also complement the long literature on mutual fund performance, which finds that most mutual fund managers do not offer positive alpha to investors (e.g., Wermers (2000) and Fama and French (2009)). However, recent work shows that some subgroups of funds exhibit significant information processing or forecasting ability. For example, Chen, Jegadeesh, and Wermers (2000) and Kosowski, Timmermann, Wermers, and White (2006) find that the best-performing managers of growth-oriented funds make significant excess returns. Barras, Scaillet and Wermers (2010) estimate the proportion of non-zero alpha fund managers using multiple testing methods similar to those employed here. For actively managed, open-end funds, they find superior skills existed in between 10% and 20% of funds in the early to mid-1990s, but trended down since then. In contrast, our informed fraction varies from 1.1% to 2.8% out of the entire sample and shows significant variation across commodities.

This paper is structured as follows. The next section develops our methodology. We explain how we isolate informed from other traders and how the *FDR* method is applied to our sample. Section III describes the CFTC data and provides summary statistics for the explanatory

<sup>&</sup>lt;sup>6</sup> The statistical techniques we employ are similar to those used by Barras, Scaillet and Wermers (2010) to find positive alpha mutual fund managers. Their primary focus is on estimating the population probability of a positive alpha fund, whereas we seek to identify informed and liquidity traders and examine their characteristics. An important difference in methods is that we develop an alternative approach to identifying the probability that a trader is null in the population.

variables used in our empirical analysis, which is described in Section IV. We use an inverse regression technique to make inferences about the representation of trader characteristics and business lines in the informed population. Our conclusions are offered in Section V. The Appendix explains how we implement the *FDR* method for our hypothesis tests.

#### **II. METHODOLOGY**

To identify informed traders, we must define and measure trading success. Within each trading day, a participant may make many trades in multiple contracts on a particular commodity. These contracts differ by time to expiration. Let  $\{OI_{0,t}^k, OI_{1,t}^k, OI_{2,t}^k, \dots, OI_{j_k,t}^k\}$  denote the sequence of positions held by a trader in contract *k* on day *t*, where a positive value indicates a long position and a negative value indicates a short position. This trader begins the day with position  $OI_{0,t}^k$  and makes  $J_k$  trades to end the day with position,  $OI_{j_k,t}^k$ . Our data report the daily open interest positions of reporting traders in each contract in each market, which means that we observe  $OI_{0,t}^k$  and  $OI_{j_k,t}^k$ , but we do not observe intra-day position changes.

Aggregating over all trades (j=1, 2, ..., J) and contracts (k=1, 2, ..., K) on day *t*, we write profit as

$$\pi_t^* = \sum_{k=1}^K \sum_{j=0}^{J_k} OI_{j,t}^k \left( P_{j+1,t}^k - P_{j,t}^k \right), \tag{1}$$

where  $P_{j,t}^k$  denotes the price of contract k at the time of trade j. The initial day t price,  $P_{0,t}^k$ , equals the previous day's closing price and the final price,  $P_{j_k+1,t}^k$ , is the day t closing price. We rewrite equation (1) as

$$\pi_t^* = \sum_{k=1}^K OI_{0,t}^k \left( P_{J_k+1,t}^k - P_{0,t}^k \right) + \left( OI_{J_k,t}^k - OI_{0,t}^k \right) \left( P_{J_k+1,t}^k - P_{*,t}^k \right), \tag{2}$$

where

$$P_{*,t}^{k} = \frac{\sum_{j=1}^{J_{k}} (OI_{j,t}^{k} - OI_{j-1,t}^{k}) P_{j,t}^{k}}{\sum_{j=1}^{J_{k}} (OI_{j,t}^{k} - OI_{j-1,t}^{k})}$$

represents a *reference* price. It is a size weighted average of the trade prices for a given trader and is defined only if the closing open interest differs from day *t*-1 to *t*.

Equation (2) shows how we can represent daily profits as a function of observables and a single unobserved variable,  $P_{*,t}^k$ . Because the *j* subscript is redundant, we consolidate notation and rewrite (2) as

$$\pi_t^* = \sum_{k=1}^K OI_{t-1}^k \Delta P_t^k + \Delta OI_t^k (P_t^k - P_{*,t}^k),$$
(3)

where  $\Delta P_t^k = P_t^k - P_{t-1}^k$  denotes the change in closing price between day *t*-1 and *t* for the *k*<sup>th</sup> expiration and  $\Delta OI_t^k = OI_t^k - OI_{t-1}^k$  denotes the change in closing open interest position between day *t*-1 and *t*.

We refer to the first term in equation (3) as *position profit* and the second term as *trading profit*.<sup>7</sup> Position profit measures the potential profit if a trader holds a closing position throughout the following day. Trading profit measures the incremental profit from the net position change. Trading profits are unobserved in our data because we do not know the reference price,  $P_{*,t}^k$ . Any proxy for  $P_{*,t}^k$  that is measured before the closing price,  $P_t^k$ , such as the opening price or the daily midpoint, confounds actual trading profits with foregone trading profits. For example, if we were to use the opening price to proxy for  $P_{*,t}^k$ , then a momentum trader who responds late in the day to intra-day price changes would appear to have positive trading profits even though she traded after the price change rather than before it.

<sup>&</sup>lt;sup>7</sup> Equation (3) represents trading profits from net daily position changes. Of course, traders may have no net position changes and many offsetting trades in a day. We do not capture these actions in our sample because we do not observe individual trades.

The profit expression in (3) shows how previous research is formulated. Hartzmark (1987, 1991) and Leuthold, et al. (1994) approximate daily profit by assuming that  $P_{*,t}^k = P_t^k$ , which implies they use only the first term in (3) and thereby calculate position profits. Dewally, Ederington and Fernando (2010) use the average of the closing prices between day *t*-1 and *t* to estimate  $P_{*,t}^k$ . Thus, they estimate trading profits in addition to position profits. This approach is useful for their goal, which is to estimate total daily profits. However, because we aim to identify informed traders, it will not work for our goal.

We develop our approach in the following subsections, beginning with the position profit method of Hartzmark and Leuthold, et al. and concluding with how the *FDR* method is used to control the multiple-testing problem.

#### A. Measuring Forecasting Ability Using Position Profits

To assess the forecasting ability of traders across all contract expirations, we define two *profit rules* using the binary variable theta ( $\theta$ ) to measure success:

(i) Position Profits: 
$$\theta_t^p = 1 \quad iff \quad \sum_{k=1}^K OI_{t-1}^k \Delta P_t^k > 0$$
  
= 0 otherwise.

(ii) Lagged Trading Profits: 
$$\theta_t^c = \text{undefined iff } \sum_{k=1}^K |\Delta OI_{t-1}^k| = 0$$
  
= 1 iff  $\sum_{k=1}^K \Delta OI_{t-1}^k \Delta P_t^k > 0$   
= 0 otherwise.

The lagged trading profits rule refines the position profits rule by including only days in which a trader's net position changes and then by measuring the forecasting success implied by the position change on day *t*-1. Specifically, it recognizes that  $OI_{t-1}^k = OI_{t-2}^k + \Delta OI_{t-1}^k$ , which

implies that position profits can be decomposed as  $\sum_{k=1}^{K} OI_{t-2}^{k} \Delta P_{t}^{k} + \Delta OI_{t-1}^{k} \Delta P_{t}^{k}$ . Lagged trading profits therefore measure the incremental effect on position profits of net trading decisions on day *t*-1.<sup>8</sup> We use the qualifier "lagged" to distinguish this measure from the within-day trading profits in equation (3).

These profit rules aggregate positions and trades over all expirations within a day and reflect the futures portfolio held by a trader in a commodity because they weight by position size. A larger position taken in one contract could override a smaller position taken in another contract. By aggregating across positions, the profit rules can also capture informed traders with information about the spread between contract prices but no information about price levels.

We use a binary measure of success ( $\theta$ ) rather than profit levels for several reasons. First, the binary measure maps directly to success in a continuous measure; that is, are profits positive or not. We want to investigate whether a trader is consistently successful, not whether she is a large or small success. Second, we have a long sample period over which trading volume and open interest increased substantially, so the size of profits in later periods will not be comparable to earlier periods due to inflation. We do not want our choice of inflation adjustment to affect these tests. Third, for a given expiration, profits sum to zero across traders, which is a further constraint on tests using profit levels. The binary measure is not similarly constrained because it is not dependent on the amount of profits. Finally, outliers in the level of trader profits could easily dominate profit level statistics, which would demand strong assumptions and some type of statistical adjustment.

<sup>&</sup>lt;sup>8</sup> Leuthold, Garcia and Lu (1994) implement the position profits rule differently by examining only days for which there was a change in open interest in the trader's account. They argue that these days capture a demonstrative action—trading activity—and thus are likely to represent the immediate opinion of a trader. The profit rules, (i) and (ii), allow us to evaluate separately profits from positions held, whether those positions are managed actively or passively, and incremental profits from trading activity.

In spite of its benefits, we recognize that a binary measure can cause us to miss some types of informed traders. Specifically, we will likely miss those who profit from skewness in returns; that is, traders who make money by taking frequent small losses and occasional large gains. We also may miss traders with fleeting information—those who are informed only infrequently and hold positions for only a few trading days—or those who have a successful trading system that stopped working after other traders discovered it. Thus, as noted earlier, our results represent a subset of the informed trader population.

#### B. Identifying Overnight Informed Traders

Consistent with previous research, we compute the binary success statistics using a close-toclose price window to identify the overnight informed traders. Thus, we examine whether a trader's positions correctly forecast the change from yesterday's closing price to today's closing price ( $P_t^k$ ).

For each trader, we test the null hypothesis that the trader is successful half of the time; that is, we test  $E(\theta_t^p) = 0.5$  for position profits and  $E(\theta_t^c) = 0.5$  for lagged trading profits. We call this an *unconditional test* for forecasting ability. Thus, traders who randomly take positions in such markets expect positive profits on 50 percent of days, assuming a symmetric distribution for price changes. This null is reasonable as the long literature on price behavior in single commodity futures shows little evidence of systematic daily bias in one direction or another (e.g., Erb and Harvey 2006). While there may be no systematic price bias, some traders may benefit from trends, such as those described by Moskowitz, Ooi, and Pedersen (2010). To adjust for possible trends, the analyses of Chang (1985), Hartzmark (1991), and Leuthold, et. al. (1994) use a test suggested by Henriksson and Merton (HM, 1981), called the *HM test*. This is a conditional test with the null of no forecasting ability defined by:

$$\Pr(i \in Long \mid \pi_i^L > 0) + \Pr(i \in Short \mid \pi_i^S > 0) = 1, \tag{4}$$

where *i* is a trader index,  $\pi_i^L$  are the profits in cases where trader *i* holds a net long position, and  $\pi_i^S$  are the profits in cases where trader *i* holds a net short position. A trader whose long or short positions are informative will have on average more than 50% of her long positions turn out profitable, more than 50% or her short positions turn out profitable, or both, such that the combined probabilities in equation (4) exceed unity. This test is identical to Fisher's (1922) exact test for independence in a 2x2 contingency table. The logic of the HM test is that forecasts are not valuable if they cannot discriminate sufficiently in both up and down markets. Thus, the HM and unconditional tests have power in different directions: the HM test is most sensitive to the success rate. Thus, we conduct both tests to identify overnight informed traders.

#### C. Identifying Intraday Informed Traders

The analysis in the previous section focuses on finding overnight informed traders using endof-day trades and positions prior to next day's prices. However, informed traders may have incentive to adjust positions during the trading day particularly if their information is time sensitive, such as the case with an earnings announcement. These trades may reveal information to the market, and thereby remove such traders from the set we identify as overnight informed.

Our goal in this section is to show how end-of-day position data may be used to identify intraday informed traders separate from other trader types, specifically liquidity demanders and suppliers, momentum traders, and contrarians. This problem reduces to unbundling the strategies of these trader types. Table 1 shows how we characterize trading strategies under the assumption that there are six trader types.<sup>9</sup> These types are labeled informed, uninformed, large liquidity demanders, large liquidity suppliers, momentum, and contrarian traders. The informed were discussed above; they trade correctly in advance of price changes. The uninformed trade incorrectly as judged by subsequent prices. Momentum and contrarian traders are effectively the causal inverse of informed and uninformed traders because they adjust open interest *in response to price changes*, not before them (Conrad and Kaul (1998)). The large liquidity demander and supplier are symbiotic traders. With finite depth at a given price, a large liquidity demander is expected to pay a premium to trade (Grossman and Miller (1988)). Correspondingly, the large liquidity supplier is expected to receive a premium for facilitating such trades.

The classification in Table 1 shows that the informed, momentum and large liquidity demanders are observationally equivalent for a given change in daily prices. That is, each trader type changes open interest in the same direction as the open-to-close price change. In the context of equation (3), these trader types differ by the value of their unobserved reference price  $P_{*,t}^k$ ; for informed traders, this reference price occurs before a price change, for momentum traders it occurs after a price change and for large liquidity demanders it occurs contemporaneously with price changes. Similarly, the uninformed, large liquidity suppliers and contrarian traders are also observationally equivalent as they all change open interest inversely with daily price changes. The problem then is to separate these combined types, so that we can isolate the intraday informed from the liquidity demanders and suppliers.

To start, define the following success rule to determine whether a trader's open interest changes are consistent with intraday price changes:

<sup>&</sup>lt;sup>9</sup> A noise or random trader represents a seventh trader type. This trader's open interest changes randomly relative to price changes in the day.

(iii) Intraday Trading Profits:  $\theta_t^d$  = undefined *iff*  $\sum_{k=1}^{K} |\Delta OI_t^k| = 0$ = 1 *iff*  $\sum_{k=1}^{K} \Delta OI_t^k (P_t^k - P_{o,t}^k) > 0$ = 0 otherwise.

Using the open-to-close price change, the intraday trading profits rule in (iii) identifies two groups of traders: (1) those whose change in open interest moves directly with intraday price changes ( $E(\theta_t^d) > 0.5$ ) and (2) those whose change in open interest moves against intraday price changes ( $E(\theta_t^d) < 0.5$ ). Each of these two groups is expected to contain the three types of traders as shown in the columns in Table 1. Thus, the identification problem requires two additional pieces of information to partition trader types. We develop two joint tests to solve this problem.

#### C.1. Overnight Close-to-Open Price Effects

Consider the group composed of the informed, liquidity demanders and momentum traders. We can separate out momentum traders by their expected response to the overnight close-to-open price change,  $P_{o,t}^k - P_{t-1}^k$ . After the market opens, momentum traders will use this information in their strategy. To the extent that  $P_{*,t}^k - P_{o,t}^k$  is small relative to  $P_{o,t}^k - P_{t-1}^k$ , the close-to-open price change will be relatively more important in momentum trades. In contrast, informed traders are not expected to be affected by the close-to-open price change as by definition their information is forward looking. Also, liquidity demanders will view such information as noise. Thus, we expect to find a positive correlation between  $\Delta OI_t^k$  and  $(P_t^k - P_{0,t}^k)$ .

We therefore define the following momentum-trading rule to measure the propensity for a trader's position to respond to the previous overnight price change:

(iv) Momentum Trading: 
$$\theta_t^m$$
 = undefined  $iff \sum_{k=1}^{K} |\Delta OI_t^k| = 0$   
= 1  $iff \sum_{k=1}^{K} \Delta OI_t^k (P_{o,t}^k - P_{t-1}^k) > 0$   
= 0 otherwise.

Momentum traders exhibit the quality that  $E(\theta_t^m) > 0.5$ , whereas for informed and liquidity traders, we expect to find  $E(\theta_t^m) = 0.5$ . We can reverse the preceding arguments to identify contrarian traders as those with  $E(\theta_t^d) < 0.5$  and  $E(\theta_t^m) < 0.5$ , as distinct from the uninformed and liquidity suppliers, who exhibit  $E(\theta_t^d) < 0.5$  and  $E(\theta_t^m) = 0.5$ .

#### C.2. Joint Tests for Intraday Informed Traders

We can separate the informed and liquidity demand traders from momentum traders if we identify those who *do not* respond directly to close-to-open price changes. In effect, we combine two independent events—the close-to-open and the open-to-close price changes—each of which occur with probability 0.5 under the null of no relationship between open interest changes and prices. The joint probability that a trader is null for both price changes is 0.25, which gives rise to the first joint test, which we call the "pairs" test:

Pairs Test 
$$H_0: E[(1 - \theta_t^m)\theta_t^d] = 0.25.$$

In the pairs test, we expect informed and liquidity demand traders to do no worse than luck predicts when judged by close-to-open price changes, but to exceed the 50% chance of success when judged by open-to-close price changes. When the two price change events are combined,

informed and liquidity demand traders' success is expected to exceed 0.25. However, momentum traders' success rate is expected to exceed 0.5 for both price changes, so reversing the definition for close-to-open success (i.e., not predictive of this change) reduces the probability of the combined events to less than 0.25 for momentum traders.<sup>10</sup> Corollary arguments may be made for contrarian and uninformed traders. Thus, significant upper-tail (lower-tail) traders in the pairs test will identify the informed and liquidity demanders (uniformed and liquidity supply).

The pairs test finds liquidity demanders and informed traders jointly. To extract liquidity demanders, we extend the horizon from the intraday close,  $P_t^k$ , to the opening price on the next day,  $P_{o,t+1}^k$ . Large liquidity demanders are not expected to have positions that predict subsequent price changes. Generally, if a large liquidity demander moves prices on a particular day, then we expect a price reversal the following day as illiquidity dissipates. Such a reversal generates negative position profits for large liquidity demanders, which will improve the precision of this next-day price filter. In contrast, informed traders expect their trades and resulting end-of-day open interest positions to yield profits; otherwise, they could have continued to trade on day *t* until the market had fully priced their information. Thus, we test whether those who are intraday successful based on changes in their open interest are also informed about next day's close-to-open price change,  $P_{o,t+1}^k$ , based on their day *t* net position.<sup>11</sup>

Effectively, our approach extends the pairs test to include the following day's price change. We call this the *triple test* as it combines all of the results noted above about informed,

<sup>&</sup>lt;sup>10</sup> This claim always holds when the chance a momentum trader's open interest responds directly to close-to-open prices is the same as the chance of responding to open-to-close prices. To the extent that the chance is less for close-to-open prices, this test may contain some leakage of momentum traders into the resulting significant group.

<sup>&</sup>lt;sup>11</sup> If informed traders make profitable trades within day t and do not close their positions by the end of the day even when they expect no subsequent gains, then their end-of-day t positions may not be predictive. The test is still valid; however, as we would expect informed traders not to consistently lose money against next day's price changes versus the likelihood that a liquidity demander faces a price reversal.

momentum and large liquidity demanders as well as those parallel results for uninformed, contrarian, and large liquidity suppliers. The null probability for the triple test is 0.125 with the null hypothesis given by:

Triple Test 
$$H_0: E[(1 - \theta_t^m)\theta_t^d \theta_{t+1}^p] = 0.125.$$

#### C.3. Refining the Triple Test

The triple test produces a single set of traders that may be considered an upper bound on the number of significant intraday informed traders. The reason for the upper bound is because the three components of the test allow flexibility in how to reject the null. For example, we may reject the null if a trader's day t position changes are negatively correlated with previous overnight close-to-open price changes but positively correlated with intraday price changes and not generally predictive of tomorrow's close-to-open price change. In effect, the triple test is flexible relative to the null, and thus may not narrow the groups in Table 1 down to a well defined set of informed traders.

To overcome the "flexibility" limitation, we impose the original premise established in Table 1 about informed traders. Specifically, the informed are those whose intraday position change is significantly (and positively) correlated with the intraday open-to-close price change. Thus, using the significant results in the upper tail of the triple test, we classify a trader as informed if that trader is also significant in a single hypothesis test comparing her position change to the intraday open-to-close price change.

Now to identify large liquidity demanders, we first find the overlap between traders who were not significant in the triple test (i.e., not in the upper tail) and those traders who were significant in the upper tail of the pairs test. Then, we designate as liquidity demanders those traders who are also significant when their position change is compared to the concurrent intraday price change. Finally, to identify the large liquidity suppliers, we apply the same steps used to identify large liquidity demanders, but these steps apply to the lower tail of our tests.

#### D. Multiple Testing for Predictive Ability

Because we have thousands of traders, the usual significance levels give rise to a multiple testing problem (Miller, 1981). The classical approach to multiple-testing limits the family-wise error rate, which is the probability of rejecting at least one true hypothesis. This approach sets a high bar for rejection because it controls the probability of making a *single* type I error. In our application, we are not averse to falsely rejecting the null hypothesis for a few traders if doing so enables us to discover numerous informed traders. Our objective is not to avoid rejecting any true hypothesis, but rather to identify a set of significant traders with a known error rate. Thus, we apply the framework developed by Benjamini and Hochberg (1995) to control the false discovery rate.<sup>12</sup> The *FDR* equals the proportion of rejected hypotheses that are in fact true.

To explain this method, suppose there are three types of traders, null traders (no predictive ability), uninformed traders (negative predictive ability) and informed traders (positive predictive ability). We denote the population proportion of null traders by  $\pi_0$ . For each trader (j = 1, 2, ..., n), we calculate a statistic,  $z_j$ . This statistic has an asymptotic standard normal distribution under the null hypothesis and is centered away from zero under the two alternative hypotheses. Suppose that, for each trader, we use a critical value, c, to test the one-sided hypothesis that the trader is

<sup>&</sup>lt;sup>12</sup> Another alternative is that proposed by Kosowski, Timmermann, Wemers and White (2006) in their assessment of mutual fund performance. They convert the multiple tests into a single test using the sequence of *z*-statistics. For some *k*, they test the null hypothesis that the  $k^{th}$  most successful manager makes statistically significant excess returns, using a bootstrap procedure to estimate the empirical distribution of this order statistic. Their test addresses a different hypothesis than considered here. Specifically, they test the null hypothesis that no manager is skilled against the alternative hypothesis that at least the top *k* fund managers are skilled, while we test separately whether each individual trader is skilled controlling the composition of the group for whom we reject this null hypothesis.

informed. Now suppose that we randomly pick the  $j^{th}$  trader from among those for whom the null hypothesis is rejected. The *FDR* is the probability that this trader is null, defined by:

$$FDR(c) = \Pr(j \in \{\text{null}\} \mid z_j > c)$$

$$= \frac{\Pr(z_j > c \mid j \in \{\text{null}\}) \Pr(j \in \{\text{null}\})}{\Pr(z_j > c)}$$

$$= \frac{\Pr(z_j > c \mid j \in \{\text{null}\}) \pi_0}{\Pr(z_j > c)}$$
(5)

Using *FDR* to control the size of the false discovery group, we choose the minimum c such that  $FDR(c) \le 0.05$ , or rather we seek rejection results that are expected to contain no more than 5 percent null cases. With this control, we reject the null hypothesis for each trader with a *z*-statistic that exceeds c.

The *FDR* approach has three useful features for our application. First, the critical value c is independent of the number of traders under consideration because it controls the proportion of null traders for whom the null hypothesis is rejected. In contrast, controlling family-wise error rate requires the critical value to increase as the number of hypothesis tests increases. Our primary objective is to separate informed from lucky traders. If a particular trader is successful, we want to find that trader, so we do not want a rule for separating luck from skill that depends on the number of other traders considered.

The second useful feature is that the *FDR* approach automatically adjusts the critical value depending on the proportion of traders in the population who are null ( $\pi_0$ ). As this proportion increases, the chance increases that a particular successful trader is merely lucky. Consequently, the method chooses a larger critical value when the proportion of null traders is greater.

Finally, the *FDR* method adjusts the critical value depending on the location of the informed traders. If the null and alternative hypotheses are close together, then there is a greater chance of

confounding the null and alternative hypotheses, which leads to a more conservative (greater) critical value. On the other hand, if the null and alternative hypotheses are far apart, then the critical value can be chosen more aggressively. The Appendix explains how we implement the *FDR* method using our data.

#### III. DATA

The data for this study are derived from the Large Trader Reporting System (LTRS) maintained by the CFTC. The LTRS provides end-of-day positions for all traders who exceed its mandatory reporting thresholds, which vary across commodities. According to the CFTC, the aggregate of all traders' positions reported in the LTRS usually represents 75 to 95 percent of the total open interest in any given market.

The sample studied here includes twelve commodities, with three each from grains (corn, soybeans and wheat), metals (copper, gold and silver), and energy (WTI crude oil, heating oil and natural gas). In addition, we include cotton, soybean oil, and sugar futures contracts. These data cover the period from January 2000 to May 2009 and include 8,921 unique traders.

The LTRS also reports a trader's business line activity, which is self-reported on Form 40 to the CFTC. This information is used by the CFTC to classify traders as "commercial" interests, which generally implies that they have an operating interest or holding in the underlying commodity. The business lines that make up commercial interests in our sample include: dealers or merchants (AD) who are usually wholesalers, manufacturers (AM) who are generally fabricators or refiners, producers (AP), agricultural/natural resources (AO) and other companies that are end users, and commodity swaps/derivative dealers (AS), which aggregates both swap dealers, arbitrageurs and broker dealers. In our results below, we aggregate AD, AM, AP and AO into one commercial group, but keep swap dealers separate as these tend to be financial firms and in later years this group includes commodity index traders.

Traders that have not identified an underlying hedging purpose are labeled "noncommercials" and are not eligible for hedging exemptions to CFTC position limits. The noncommercial groups include floor brokers and traders (FBT), hedge funds/managed money (MMT) and non-registered participants (NRP). The hedge funds/managed money group includes commodity pool operators, commodity trading advisors, associated persons and any others who manage money for clients. The non-registered participants are traders with positions large enough to meet the CFTC reporting requirements, but do not have to register under the rules of the Commodity Exchange Act. These non-registered participants are generally smaller financial firms. There are a small number of traders with specialty designations, such as non-U.S. commercial bank, insurance company, corporate treasurer, etc., and these traders and NRPs are combined into an "Other" category.

Table 2 provides summary statistics for our sample data by trader-specific characteristics in Panel (A); Panel (B) shows the distribution of traders across business lines; and Panel (C) shows the representativeness of the reported trader data collected by the CFTC. To generate trader characteristics we average positions for each trader over time as described in Table 2. The statistics in Panel (A) equal the average of the trader characteristics across traders by commodity.

We characterize traders by their experience, which is measured by the number of days they hold an open position. On average traders have 1.4 years of trading experience with substantial variation across commodities. Traders are also fairly large as judged by average position size and fairly active in most commodities. The average participant trades on more than half of the days in which they hold open positions. Participants also typically hold positions in multiple expirations, with traders in the energy commodities holding on average the most contract expirations. About 25% of all traders use options-on-futures positions (not shown) with crude oil traders showing the greatest use of options (33.6%) and copper traders the least use (8.4%).

We characterize the one-sided choices of trades with average net long and net short futures position sizes, both relative to absolute futures position size. To generate these statistics, we divide traders into those who hold more long positions than short positions on average and those who hold a greater number of short than long positions. For all net long (or short) traders, we then calculate the proportion of all positions that are long (or short). This proportion would equal one for a net long trader who never holds a short positions. For example, Table 2 shows that across all crude oil traders who are net long on average, the ratio of average long to average total position is about 48%. Across all commodities, this average is about 60% for both net long and net short traders. For comparison, in the entire sample about half of all traders are net long on average, 20% of traders are always long, and 14% are always short.

Panel (B) shows that almost half of traders are non-registered participants or specialized traders, who average nearly 46% in the sample. MMTs are the second most populous group at 23%, with all commercials at about 18% of participants. The ADs are the largest group among commercials with about 10% of the sample participants. The FBT also constitute about 10% of the sample.

Panel (C) confirms the claim that the data reported by the CFTC represents between 75 and 95 percent of the total open interest of all traders. This panel also shows the distribution of trader counts by commodity. The grains, corn, soybeans and wheat attract the greatest number of participants, with the energy complex and gold futures containing about 60-70% of these

23

numbers. The other commodities have fewer traders, but in each case our sample contains hundreds of traders. In this light, each commodity presents a multiple testing problem as we seek to identify who is informed.

#### **IV. ANALYSIS**

#### A. Identifying Overnight Informed Traders

We implement the *FDR* approach to detect forecasting ability using one-sided tests on each tail of the sample distribution. We use a 5% level of significance to determine the critical value for the false discovery rate. Table 3 reports the overnight informed counts using the profit rules described above. These counts provide a lower bound on the actual number of informed traders. Specifically, we may omit some informed traders who were only mildly successful because we could not select them without also selecting other traders who were merely lucky.<sup>13</sup> Our *FDR* criterion implies that at least 95% of traders in the overnight informed group are informed.

The unconditional and HM tests shown in the Table 3 summarize the daily success of each trader's lagged trading or position profits as measured against close-to-close prices. This table reports the count and proportion of successful traders for each test, and average percent of open interest represented by informed traders for the position profits rule. For both tests, the *FDR* approach evaluates the significance of the resulting p-values across all traders to identify those who are informed.

Let us first focus on the lagged trading profits, where we judge success by whether net trades made on day t-1 provide positive profits when evaluated at day t prices. This test identifies

<sup>&</sup>lt;sup>13</sup> We also implemented two-sided tests and investigated close-to-open prices with very similar results. In addition, we ran tests based on combining lagged trade and position profit rules; that is, use the lagged trade profit rule if there is a trade, otherwise use the position profit rule. These results are consistent with our findings here and are available on request.

almost no informed traders. Out of 8,921 unique traders across commodities, we find a total of five informed traders. Thus, very few traders change their positions systematically one day in advance of price changes the next day. This result suggests that the net effect of trading reveals information that is assimilated into prices within the trading day.

Table 3 also shows results for the position profits rule. Compared to the lagged trading profits rule, this rule detects many more overnight informed traders across all commodities. Silver exhibits the largest number of informed in the unconditional test, with 93 selected as informed (12% of traders; 14.3% of OI). In the HM test, copper produces the highest count at 51 informed traders (8.1% of traders; 7.9% of OI). Other commodities also show meaningful counts for informed traders, particularly corn, soybean oil and soybeans. In total, we identify 246 informed traders (230 unique) in the unconditional tests and 96 traders (94 unique) in the HM tests. The small number of overlaps shows that we find only a few traders who are informed in multiple commodities.

The HM test shows that about half of informed traders appear to benefit from participating during trending episodes. Specifically, the overnight informed counts for copper, corn, silver, and soybean oil are significantly reduced for the HM test. In general, the position profits rule identifies between 1.1% and 2.6% of our sample as overnight informed traders, with the higher percentage skewed by copper, corn, silver, and soybean oil counts.

Table 3 also reports the range of critical values arising from implementing the *FDR* method using a five-percent level of significance for each commodity. The table contains 48 counts across commodities of which 25 are zero. A zero count implies that the *FDR* exceeds 0.05 for all observed z statistics. Of the remaining 23 entries, four critical values lie between two and three, 16 between three and four, and the remaining three between four and five. In all tests, the

commodity with the highest informed count exhibits the smallest selected critical value, which underlies the point that the counts in Table 3 signify the number of informed traders we can identify rather than the number of informed *per sé*. If the informed traders do not have sufficiently large z statistics, then controlling the false discovery rate requires a larger critical value and a smaller count irrespective of the total number of informed. The median critical value across all tests is 3.57, which corresponds to a one-sided p-value of 0.00018. This suggests that multiple-testing bias would be severe if we had applied classical critical values in these tests.

Overall, the informed counts in Table 3 show that relatively few, but nonetheless a meaningful number of traders consistently hold end-of-day *positions* that reveal superior information. Importantly, we find almost no traders whose lagged *trading* decisions consistently predict subsequent profits.

#### B. Characteristics of Overnight Informed Traders

We are interested in both identifying informed traders and describing their characteristics; specifically, whether their business activities suggest that they are commercial traders, hedge funds, floor brokers, or swap dealers, and how their trading activity compares to the average trader. Thus, we condition on the informed already found to identify their characteristics. This approach is called an inverse regression (e.g., Li (1991)). The task is essentially the opposite of discriminant analysis, which uses observed characteristics to classify among observations. To accomplish our goal, we run a linear regression of the binary variable indicating membership in the informed group (y) on a matrix of characteristics X, specified as follows:

$$y = X\beta + \varepsilon = \begin{bmatrix} x_j & X_{-j} \end{bmatrix} \begin{bmatrix} \beta_j \\ \beta_{-j} \end{bmatrix} + \varepsilon,$$
(6)

where  $E(X'\varepsilon) = 0$ ,  $x_j$  denotes a particular variable of interest, and  $X_{j}$  denotes the other variables in *X*. The coefficient  $\beta_j$  can be written as

$$\beta_{j} = \frac{E(y_{i}x_{ji}) - E(y_{i}\hat{x}_{ji})}{E(x_{ji} - \hat{x}_{ji})^{2}}$$

$$= \frac{E(x_{ji}|y_{j} = 1) \operatorname{Pr}(y_{j} = 1) - E(\hat{x}_{ji}|y_{j} = 1) \operatorname{Pr}(y_{j} = 1)}{E(x_{ji} - \hat{x}_{ji})^{2}}$$
(7)

where  $\hat{x}_j = X_{-j} (X'_{-j} X_{-j})^{-1} X'_{-j} x_j$  is the linear projection of  $x_j$  onto the column space of  $X_{-j}$ . To re-interpret  $\beta_j$  based on characteristics, we re-write this expression to obtain

$$\tilde{\beta}_{j} \equiv \beta_{j} \frac{E(x_{ji} - \hat{x}_{ji})^{2}}{\Pr(y_{j} = 1)} = E(x_{ji} | y_{j} = 1) - E(\hat{x}_{ji} | y_{j} = 1)$$
(8)

Thus, by scaling the regression coefficient appropriately, we obtain a measure of the difference between the expected value of  $x_j$  for all the observations in the informed group (i.e.,  $y_j = 1$ ) and the expected value of  $x_j$  from a prediction based on the other variables in X. In effect, we estimate a forward regression of y on X to derive an inverse regression interpretation.

For example, suppose  $x_j$  denotes the log of the average size of positions held by trader *j*, and  $X_{-j}$  contains only a constant. Then  $\hat{x}_{ji} = n^{-1} \sum_{j=1}^{n} x_{ji}$  is the sample mean of  $x_{ji}$ . It follows that

$$\tilde{\beta}_j = E(x_{ji} \mid y_j = 1) - E(x_{ji}), \qquad (9)$$

which is the difference between the average log-position-size of informed traders and the corresponding average *across all traders*. Similarly, if  $x_{ji}$  is a dummy variable signifying whether trader *i* is a managed-money or hedge fund trader (denoted MMT) and  $X_{ij}$  contains a set

of dummy variables signifying whether trader *i* is of another type, then  $\hat{x}_{ji} = 0$  and the adjusted coefficient becomes

$$\tilde{\beta}_j = Pr(x_{ji} = 1 | y_i = 1). \tag{10}$$

In this case,  $\tilde{\beta}_j$  measures the proportion of traders in the informed group who are classified as MMT. If  $X_{-j}$  includes a trader characteristic such as log-position-size, then  $\tilde{\beta}_j$  would measure the representation of MMT traders  $x_j$  in the informed group holding log-position-size constant.

Table 4 shows the inverse regression results for traders selected by the unconditional and HM tests based on the position profits rule. These regressions pool observations across commodities and include fixed effects by commodity. We pool across commodities because not every commodity revealed informed trading. However, the sign and significance of the parameters is quite similar in separate regressions for commodities with more than fifteen informed traders. When estimating standard errors, we correct for heteroscedasticity using White's (1980) estimator and cluster by trader. The coefficient estimates in Table 4 are those produced by the projection in equation (8) with the significance of the underlying coefficient shown next to the estimate of  $\tilde{\beta}_i$ .

We show three models in Table 4 for each test. We show results including business line dummy variables alone, trader characteristic variables alone, and both combined. In terms of  $R^2$ , the trader characteristics better predict membership in the informed group than the business line variables. However, the business line variables provide insights into the characteristics of the informed. The relative change in the business line coefficients across these models helps filter out the effects of business type from those due to the characteristics of traders. For identification purposes, we exclude the constant term and show the sample proportion of the business line variables for comparison. Thus, a group is over-represented among the informed if the coefficient on that group's dummy variable exceeds the sample proportion for that group.

To understand how to interpret the coefficients in Table 4, consider how the projected effect for the Commercial group changes from model (1) to model (3). The significant coefficient of 0.06 in model (1) implies that the proportion of overnight informed traders who are commercials is 6 percent, which is lower than the overall sample representation of 21 percent. Because commercial traders are generally hedgers, this result suggests that hedgers are underrepresented in the overnight informed group based on close the unconditional test, which also hold for the HM test. In model (3), this coefficient is 0.02 or 2 percent and insignificant, which suggests that controlling for other trader characteristics removes the representation of the "commercial" distinction from the informed group.

Comparing the results from the unconditional test to the HM test, we tend to see relatively more FBTs selected as over-representative by the HM test. Correspondingly, the HM test selects relatively fewer traders from the other business lines. In the HM test, FBTs comprise 40% of the informed, which significantly exceeds the 14% representation that FBTs have in the full HM test sample. This estimate drops to 27% when we add trader characteristic variables.

The over-representation of FBTs, who tend to be market makers, among the informed suggests that the ability to understand and process order flow information determines much of what it means to be overnight informed. These traders are less likely to possess private information about flows of the physical commodity than commercial traders, and they are less likely to trade based on sophisticated technical models than managed money traders. However, their role as liquidity providers places them in a position to track and predict order flow. This result is consistent with Evans and Lyons (2008), who show that most of the effect of

macroeconomic news on the deutschmark/dollar exchange rate is transmitted through order flow, as well as the findings of Anand and Subrahmanyam (2008) for intermediaries trading equities on the TSX. Thus, these FBTs appear to behave as the superior information processors described in microstructure models, such as Kim and Verrecchia (1997).

Among the trader characteristics, the net long and net short variables have larger negative coefficients in the HM results than in the unconditional results. Model (6) shows that the average net position of informed traders is 19% less long and 22% less short than the average trader in that commodity. This result shows that traders whose net position is predominately on one side of the market are much less likely to be selected as informed than traders such as FBTs whose net position alternates on both sides of the market. Moreover, the similar magnitude of these coefficients shows that it is being on one side of the market that is under-represented in the informed group, rather than being long or short *per se*. In the unconditional test, however, net long traders are not significantly less likely to be in the informed group. This result shows that the HM test differs from the unconditional test by filtering out those traders who consistently made good forecasts by being predominately net long during a period of increasing prices.

Among business categories, the commercial firms are under-represented as informed traders. The MMTs show a slight over-representation in the unconditional test, but not in the HM test. This result suggests that some MMTs may benefit from taking long positions during a period of rising prices. The swap dealers and index traders (AS) are not distinctive and generally insignificant in these results. The "other" category, which includes non-registered participants, shows no real indication of being under- or over-represented after controlling for trading characteristics. In the full models, the experience variable stands out as always significant and positive. As it is defined in logs, the coefficient implies that informed traders are between 10% and 25% more experienced than the average trader, *ceteris paribus*. This result may reflect selection bias to some degree as less successful traders may exit the market after repeated failures. Informed traders in the HM test tend to hold larger size positions and be more active than the average trader, which is consistent with the results of Menkhoff and Schmeling (2010) in a foreign exchange market. The size coefficient increases in the full model estimates, which reveals that informed traders within a particular business line are more likely to be larger than the average trader in that business line.

We also estimated the models in Table 4 using only a commercial/non-commercial dummy variable, rather than the five business line dummy variables. The commercials were defined as the AD, AM, AP, AO, and AS groups for these regressions. This split matches the standard definition of hedgers and speculators used by the CFTC and in numerous academic papers. The results showed that non-commercials were over-represented among the informed, which is consistent with the disaggregated results discussed above.

Although FBTs feature prominently in the informed groups, most FBTs are not informed. In fact, for the HM test sample, the average success rate of FBTs is the lowest of any business-line group with the average success rate highest for AS firms. If we had aggregated all FBTs into a single category rather than studying individual accounts, then we would have missed the fact that some FBTs have predictive ability. Thus, our results differ from what would be found if first we had classified firms by business line and then conducted our tests in a manner similar to past studies.

#### C. Identifying Trader Types from Intraday Activity

The lagged trading profits tests in Table 3 selected almost no informed traders, which may indicate that information held by informed traders makes its way into prices during the trading day. Thus, we now analyze intraday price changes using the methods introduced above. Table 5 reports results using the pair and triple tests and the next day open-price filter to identify the intraday informed and liquidity traders.

Table 5 shows trader counts for various trader types by commodity. We use a test of the null hypothesis  $E(\Theta_i^m) = 0.5$  to identify momentum and contrarian traders and the triple test to isolate informed and liquidity traders. As described in Section II(c), we only designate informed and liquidity traders who are also significant when their position change is compared to the concurrent intraday price change. We find relatively low counts for liquidity suppliers and demanders, with about 1% of all traders in each of these categories. These low counts may be due to the fact that only large traders are likely to move prices when they demand liquidity. The liquidity supplier counts show only three duplicates across commodities, but there is greater cross-commodity overlap in liquidity demanders. There are 115 unique traders in the total count of 152, of which 92 appear in a single commodity and 14 in two commodities. No participant is an identified liquidity demander in more than six commodities.

With the triple test, we find a total count of 158 informed traders, which represents 91 unique traders. Of these 65 show up as informed in a single commodity, 12 in two commodities and the remaining 14 in three or more commodities. We explore cross-commodity patterns below. We find only eight uninformed traders, i.e., those who are significant traders in the left tail of the triple test (not shown in Table 5).
Consistent with the overnight informed results, the triple test identifies relatively few intraday informed traders. As a percentage of all traders, these counts suggest that soybean oil, natural gas, soybeans and wheat include a larger fraction of informed traders. In contrast, the results in Table 3 suggest that informed traders are more concentrated in silver and copper, with nearly 8% of traders as overnight informed in copper for both the HM and unconditional tests. These differences suggest that information manifests itself differently between commodities, with the result that information held by traders may have varying horizons.

The counts for momentum and contrarian traders are sizable across nearly all commodities. On average across commodities, 8% of futures traders follow momentum strategies and 11% follow contrarian strategies. There is also significant overlap across commodities, so the total number of unique traders is less than the totals of 834 and 1,170 shown in Table 5. Specifically, we find a total of 444 unique momentum traders, which is five percent of all traders. Of these, 298 are identified in a single commodity, 64 in two commodities, 36 in three commodities, and the remaining in four or more commodities. A total of 10 firms are identified as momentum traders in at least 10 of the 12 commodities. We find a total of 912 unique contrarian traders, including 742 who are contrarian in a single commodity, 117 in two commodities, 34 in three commodities and 19 in four or more commodities. We identify no firms as contrarian in more than eight commodities. Momentum and contrarian traders are more concentrated in the metals, although the corn market also exhibits a substantial fraction of identified contrarian traders.

#### D. Characteristics of Traders Identified from Intraday Activity

Table 6 reports inverse regression results for liquidity, informed, momentum, and contrarian traders identified by the intraday tests. The liquidity demanders and suppliers selected by the triple test show stark differences in composition. Commercial traders comprise 42% of liquidity

suppliers and just 2% of liquidity demanders, whereas managed money traders comprise only 7% of liquidity suppliers but 52% of liquidity demanders. The FBTs are over-represented among liquidity suppliers, consistent with their traditional role as market makers, but they are also somewhat over-represented among liquidity demanders. Swaps dealers are also over-represented in both groups, and the Other category is under-represented in both groups, which is consistent with this group acting as noise traders.

The trader characteristics variables show that liquidity demanders hold positions 57% larger in size, but they are no more experienced than the average trader. In contrast, liquidity suppliers hold positions that are insignificantly different in size, and they are 41% more experienced than the average trader. Liquidity demanders hold concurrent positions in about 41% fewer contract expirations and are less likely to be on one side of the market than the average trader. Both groups tend to be more active traders than average. The full models in (3) and (6) show that trader characteristics explain much of the difference in representation of the various business lines in the liquidity trader groups. In particular, the representation of managed money in the liquidity demander group drops from 0.52 to 0.37 when we control for trader characteristics. Similarly, the representation of commercial firms among liquidity suppliers is close to its representation in the sample (25%) once we control for experience, trading activity and a tendency to be on the short side of the market.

These results suggest that there are essentially no commercial firms who pay a liquidity premium to hedge. Rather, a subset these firms, especially dealers and merchants, tend to *supply* liquidity services. In other words, they trade in the opposite direction to concurrent price changes, but do not possess information to predict the following day's price change. It appears instead as if selected large managed money traders and hedge funds may pay for immediacy and

are met by liquidity from experienced commercial firms and FBTs. This result is the opposite of that obtained by deRoon, Nijman and Veld (2000) who found that, in aggregate, commercial hedgers tend to apply price pressure when they adjust their positions.

Both the liquidity demanders and the informed in Table 6 change open interest in the same direction as prices during the day but do not follow overnight price changes. The difference between these two groups is that the informed tend to hold overnight positions that better predict the next day's opening price. Table 6 shows that the informed tend to be even larger, more active, and present in fewer contract expirations than the liquidity demanders. In addition, MMTs and swaps dealers are substantially over-represented and commercial firms are under-represented in the informed group. The FBTs comprise 13% of the informed, which is the same as their representation in the whole sample and much less than their representation in the overnight informed. This result suggests that FBTs do not provide information during the trading day, but rather MMTs and swap dealers bring information to the markets.

When comparing momentum to contrarian traders, the largest difference is the representation of MMT and commercial traders. Managed money traders are strongly over-represented among the momentum group; they comprise 56% of momentum traders and just 13% of contrarians. Conversely, commercial firms show over-representation among contrarians; they comprise 46% of these traders and only 8% of momentum traders. Each of these two business lines constitutes about a quarter of all traders.

The trading characteristic effects show that momentum traders hold positions in about 34% fewer contract expirations than the average trader. Contrarian traders hold 8% more contract expirations than average. Thus, momentum traders are more concentrated on the term structure than the average trader. In contrast, contrarians tend to hold a more diverse set of contracts, as

35

indicated by their presence in a greater number of expirations. These results paint a picture of a subset MMTs following momentum strategies and being matched with hedgers. Both contrarians and momentum traders tend to be about 20% more experienced than the average trader.

In sum, our intraday analysis shows that commercial traders are under-represented among liquidity demanders and over-represented among liquidity suppliers. The MMTs are strongly over-represented in the liquidity demander and with swap dealers are over-represented among informed traders. The FBTs are over-represented among liquidity suppliers. Liquidity demanders and the informed tend to be large, active traders who trade in few contract expirations at a time.

#### E. Profits of Informed Traders

For robustness, we examine whether the FDR approach has identified traders who are indeed profitable as judged against other traders in the sample not selected as informed. To do this, we compute the daily profits for each type of trader: the informed and the collection of those who are not right-tail informed. These daily profits are computed using equation (3) with the reference price equal to the midpoint between the open and closing price. For the intraday informed, we only evaluate the profits generated for the second term in equation (3) (i.e., trading profit) as that sample conditions on those traders who have a position change. We use as our measure the average daily profit per trader because total profits and profits per position would not be comparable; the former being affected by the number of traders and the latter affected by the number of expirations.

Table 7 reports the results of our tests comparing profits in panels (A) and (B). We apply two tests to the daily profit per trader data: the Wilcoxon rank-sum non-parametric test because the data may not be normal and the t-test for the difference between means in the event that our statistics are approximately normal. The table shows results for the informed counts of both

position tests in Table 3 and for the intraday informed in Table 5. The one-tail p-value for the Wilcoxon rank-sum is shown along with the mean rank-sums for both informed and not informed participants by commodity. The t-test is reported below the Wilcoxon results. The one-tail p-value for the t-test is computed assuming unequal variances.

Panel (A) reports the overnight informed results. For the informed found in the unconditional tests, the Wilcoxon mean ranks and p-values show significant differences in all commodities. This also generally holds for the corresponding t-tests that the average daily profits per traders, except for natural gas and wheat. For the HM tests, the Wilcoxon results are similarly strong, except for gold and wheat. The t-tests also show that natural gas, soybean, and soybean oil are not significant. The t-tests fail here because there is significant variation in these data partly due to the small count of informed in these commodities, which makes these data fat tailed and generally non-normal.

Panel (B) shows the profit results for intraday informed traders versus the non-informed. Here the not informed include all the other trader types: momentum, contrarian, liquidity demanders and suppliers. With the exception of heating oil—only three informed traders—the Wilcoxon tests show that the informed earn significantly higher daily profits per trader than the uninformed. The t-tests also support this conclusion, but for soybeans and heating oil the results are not significant.

In general, these results indicate that the FDR method implemented above has identified traders who have an exceptional information advantage compared to other traders.

#### F. Information Across Tests and Commodities

We have provided several counts of informed traders depending on which test and what prices were used for the test. To see whether the overnight and intraday tests select the same

37

informed traders, we investigate the amount of overlap between traders found to be informed in our tests—that is, the unconditional and HM tests—and the triple test. Table 8 provides these results by commodity.

Table 8 compares counts found for the various tests by reporting the unique count of informed traders by commodity, removing any double counting of traders who are significant in both tests and by showing the percent of traders who overlapped in these paired comparisons. The fraction of the overlap varies by commodity, but is quite low overall with no comparison exceeding a 10% overlap, which is the percentage reached by crude oil and nearly by corn for the unconditional and triple tests. We conclude that the overnight informed found in the lagged profit position tests are different than the informed found in the intraday tests. This result is consistent with the model of Kim and Verrecchia (1997), who describe two types as informed: those possessing private information signals and those who profit by processing information efficiently. We view the intraday informed as possessing private signals and the overnight informed superior information processors.

We also compared whether traders overlap across commodities in a given test. We know from the total and unique counts at the bottom of Table 3 that there is very little overlap for overnight informed traders. Our investigation of intraday informed in Table 5 showed a similar result. In general, there appears a great deal of information specialization across these twelve commodities.

#### **V.** CONCLUSIONS

In this article, we characterize individual traders in futures markets by the positions they hold and their trading performance. We find that few if any traders systematically trade one day in a way that increases profits the following day. However, we identify between 1.1% and 2.5% of the traders who hold positions at the end of a day that systematically predict the next day's price change. A high proportion of such traders are floor brokers, although it is their trading behavior, particularly their experience, presence on both sides of the market, and their larger than average size that characterizes them more than their business classification as floor brokers, *per sé*. Neither commercial firms, nor hedge funds/managed money show a propensity for holding positions that are profitable the following day.

We also find numerous traders who move either with (or against) prices during the day; their positions change systematically in the same (or opposite) direction as prices. From these traders, we identify two groups whose actions suggest that they tend to either supply or demand liquidity. This screen allows us to isolate the group who may be intraday informed. Across these groups we find that commercial firms tend to provide liquidity to the hedge funds/managed money traders who demand it. We also find that hedge funds/managed money and swap dealers are over-represented among the intraday informed, with commercial hedgers significantly under-represented.

In addition, we find that the intraday informed have only a small overlap with the informed group identified as overnight informed. The characteristics of these two informed groups suggest that information that leads to superior forecasting abilities varies across such traders. Possibly, end-of-day informed traders may gain their skills from the ability to process and forecast order flows and intraday informed traders from possessing signals that require trade actions, such as those from macro-type information or commodity specific announcements. An analysis of daily trader profits confirms that our methods have selected highly profitable traders.

#### REFERENCES

- Anand, Amber, and Avanidhar Subrahmanyam, 2008, Information and the Intermediary: Are Market Intermediaries Informed Traders in Electronic Markets?, *Journal of Financial and Quantitative Analysis* 43, 1-28.
- Barber, Brad M. and Terrance Odean, 2000. Trading is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors, *Journal of Finance* 55 (2), 773–806.
- Barras, Laurent, Scaillet, Oliver, and Russ Wermers, 2010, False Discoveries in Mutual Fund Performance: Measuring Luck in Estimating Alphas, *Journal of Finance* 65(1), 179-216.
- Benjamini Y. and Y. Hochberg, 1995, Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing, *Journal of the Royal Statistical Society*, Series B, 57, 289-300.
- Biais, Bruno, Glosten, Larry, and Chester Spatt, 2005, Market Microstructure: A Survey of Microfoundations, Empirical Results, and Policy Implications, *Journal of Financial Markets*, 8 (2), 217-264.
- Boehmer, E., Jones, C., and X. Zhang, 2008, Which shorts are informed?, *Journal of Finance* 63, 491–527.
- Brandt, M. W., and K. A. Kavajecz, 2004, Price Discovery in the U.S. Treasury Market: The Impact of Order Flow and Liquidity on the Yield Curve, *Journal of Finance* 59, 2623-2654.
- Chakravarty, Sugato, 2001, Stealth Trading: Which Traders' Trades Move Stock Prices?, Journal of Financial Economics 61, 289-307.

- Chang, Eric C., 1985, Returns to Speculators and the Theory of Normal Backwardation, *Journal of Finance* 40(1), 193-208.
- Chen, Hsiu-Lang, Jegadeesh, N., and R. Wermers, 2000, The Value of Active Mutual Fund Management: An Examination of the Stockholdings and Trades of Fund Managers, *Journal of Financial and Quantitative Analysis*, 343-368.
- Christophe, Stephen E., Ferri, Michael G., and Jim Hsieh, 2010, Informed Trading Before Analyst Downgrades: Evidence from Short Sellers, *Journal of Financial Economics* 95, 85-106.
- Conrad, Jennifer S., and Gautam Kaul, 1998, An Anatomy of Trading Strategies, *Review of Financial Studies* 11, 489-519.
- Cornell, B., and E. R. Sirri, 1992, The Reaction of Investors and Stock Prices to Insider Trading. Journal of Finance 47(3), 1031–1059.
- Dalmasso, Cyril, Broet, Philippe, and Thierry Moreau, 2005, A Simple Procedure for Estimating the False Discovery Rate, *Bioinformatics* 21, 660-68.
- deRoon, Frans A., Nijman, Theo E., and Chris Veld, 2000, Hedging Pressure Effects in Futures Markets, *Journal of Finance* 55(3), 1437-1456.
- Deuskar, Prachi, and Timothy C. Johnson, 2010, Market Liquidity and Order-Driven Risk, working paper, University of Illinois.
- Dewally, Michaël, Louis Ederington and Chitru Fernando, 2010, Determinants of Trading Profits of Individual Futures Traders: Risk Premia, Information, and Luck, Working paper, Department of Finance, University of Oklahoma.
- Easley, D., Kiefer, N., O'Hara, M., and J. Paperman, 1996, Liquidity, Information and Infrequently Traded Stocks, *Journal of Finance* 51, 1405-1436.

- Erb, Claude, and Campbell R. Harvey, 2006, The Tactical and Strategic Value of Commodity Futures, *Financial Analysts Journal*, March/April.
- Evans, M. D. D., and R. K. Lyons, 2002, Order Flow and Exchange Rate Dynamics, *Journal of Political Economy* 110, 170-180.
- \_\_\_\_\_, 2008, How is Macro News Transmitted to Exchange Rates? *Journal of Financial Economics* 88(1), 26-50.
- Fama, Eugene F., and Kenneth R. French, 2009, Luck Versus Skill in the Cross Section of Mutual Fund Returns, Working paper, Booth School of Business, University of Chicago, forthcoming in the *Journal of Finance*.
- Fishe, Raymond P. H., and Michel Robe, 2004, The Impact of Illegal Insider Trading in Dealer and Specialist Markets: Evidence from a Natural Experiment, *Journal of Financial Economics* 71, 461-488.
- Fisher, R. A., 1922, On the Interpretation of  $\chi^2$  from Contingency Tables, and the Calculation of P, *Journal of the Royal Statistical Society* 85(1), 87–94.
- George, Thomas J., Kaul, Gautam, and M. Nimalendran, 1991, Estimation of the Bid-Ask Spread and Its Components: A New Approach, *Review of Financial Studies* 4, 623-656.
- Glosten, Lawrence R, and Lawrence E. Harris, 1988, Estimating the Components of the Bid-Ask Spread, *Journal of Financial Economics* 21, 123-142.
- Green, T. C., 2004, Economic News and the Impact of Trading on Bond Prices, *Journal of Finance* 59, 1201–1234.
- Grinblatt, Mark, Keloharju, Matti. and Juhani Linnainmaa, 2009, IQ and Market Participation, working paper, UCLA Anderson School of Management, forthcoming in the *Journal of Finance*.

- Grossman, Sanford J and Merton H. Miller, 1988, Liquidity and Market Structure, *Journal of Finance* 43(3), 617-37.
- Grundy, B., and Y. Kim, 2002, Stock Market Volatility in a Heterogeneous Agent Economy, Journal of Financial and Quantitative Analysis, 37, 1-27.
- Harris, Jeffrey H., and Bahattin Buyuksahin, The Role of Speculators in the Crude Oil Futures Market, working paper, Office of the Chief Economist, Commodity Futures Trading Commission, 2009.
- Harris, Milton, and Artur Raviv, 1993, Differences of Opinion Make a Horse Race, *Review of Financial Studies* 6(3), 473-506.
- Hartzmark, Michael L., 1987, Returns to Individual Traders of Futures: Aggregate Results, Journal of Political Economy 95 (6), 1292-1306.
- \_\_\_\_\_, 1991, Luck Versus Forecast Ability: Determinants of Trader Performance in Futures Markets, *Journal of Business* 64(1), 49-74.
- Hasbrouck, Joel, 1988, Trades, Quotes, Inventories, and Information, *Journal of Financial Economics* 22, 229-252.
- \_\_\_\_\_, 1991, Measuring the Information Content of Stock Trades, *Journal of Finance* 46(1), 179-207.
- \_\_\_\_\_, 1995, One Security, Many Markets: Measuring the Contributions to Price Discovery, *Journal of Finance* 50, 1175-1199.
- Haung, Roger D., and Hans R. Stoll, 1997, The Components of the Bid-Ask Spread: A General Approach, *Review of Financial Studies* 10, 995-1034.
- Hieronymus, Thomas A., 1977, *Economics of Futures Trading for Commercial and Personal Profit* (2<sup>nd</sup> Ed.), New York: Commodity Research Bureau.

- Houthakker, H. S., 1957, Can Speculators Forecast Prices?, *Review of Economics and Statistics* 39(2), 143-151.
- Ito, Takatoshi, Lyons, Richard K., and Michael Melvin, 1998, Is There Private Information in the FX Market? The Tokyo Experiment, *Journal of Finance* 80, 1111-1130.
- Kandel, Eugene, and N. Pearson, 1995, Differential Interpretation of Public Signals and Trade in Speculative Markets, *Journal of Political Economy*103, 831–872.
- Keim, Donald B., and Ananth Madhavan, 1995, Anatomy of the Trading Process: Empirical Evidence on the Behavior of Institutional Traders, *Journal of Financial Economics* 37, 371-398.
- Kim, Oliver and Robert E. Verrecchia, 1991, Market Reactions to Anticipated Announcements, Journal of Financial Economics 30, 273-309.
- \_\_\_\_\_, 1997, Pre-announcement and Event-period Private Information, *Journal of Accounting and Economics* 24, 395-419.
- Kosowski, R., Timmermann, A., Wermers, R., and H. White, 2006, Can Mutual Fund 'Stars' Really Pick Stocks? New Evidence from a Bootstrap Analysis, *Journal of Finance* 61, 2551-2595.
- Kurov, A., and D. Lasser, 2004, Price Dynamics in the Regular and E-Mini Futures Markets, Journal of Financial and Quantitative Analysis, 39, 365-384.
- Leuthold, Raymond, Garcia, Philip and Richard Lu, 1994, The Returns and Forecasting Ability of Large Traders in the Frozen Pork Bellies Futures Market, *Journal of Business* 67(4), 459-473.
- Li, K-C., 1991, Sliced Inverse Regression for Dimension Reduction, *Journal of the American Statistical Association*, 86, 316–327

- Love, Ryan, and Richard Payne, 2008, Macroeconomic News, Order Flows, and Exchange Rates, *Journal of Financial and Quantitative Analysis* 43(2), 467–488.
- Menkhoff, L. and M. Schmeling, 2010, Whose Trades Convey Information? Evidence From a Cross-Section of Traders, *Journal of Financial Markets*, 13, 101-128.
- Meulbroek, L.K., 1992, An Empirical Analysis of Illegal Insider Trading, *Journal of Finance* 47(5), 1661–1699

Miller, R.G., 1981, Simultaneous Statistical Inference, 2nd ed, Springer Verlag New York.

- Moskowitz, Tobias, Ooi, Yao Hua, and Lasse H. Pedersen, 2010, Trends, working paper, University of Chicago - Booth School of Business.
- Nicolosi, Gina, Peng, Liang and Ning Zhu, 2009, Do Individual Investors Learn from Their Trading Experience?, *Journal of Financial Markets* 12, 317-336.
- Odean, Terrance, 1999, Do Investors Trade Too Much?, American Economic Review 89, 1279-1298.
- Pasquariello, Paolo, and Carlo Vega, 2007, Informed and Strategic Order Flow in the Bond Markets, *Review of Financial Studies* 20(6), 1975-2019.
- Phillips, Gordon M., and Robert J. Weiner, 1994, Information and Normal Backwardation as Determinants of Trading Performance: Evidence from the North Sea Oil Forward Market, *The Economic Journal* 104, 76-95.
- Rockwell, Charles S., 1967, Normal Backwardation, Forecasting, and the Returns to Commodity Futures Traders," reprinted in *Classic Futures: Lessons from the Past for the Electronic Age*, edited by Lester G. Telser, London: Risk Publications, 2000.
- Storey, John D., 2002, A Direct Approach to False Discovery Rates, *Journal of the Royal Statistical Society* 64, 479-498.

- \_\_\_\_\_, 2003, The Positive False Discovery Rate: A Bayesian Interpretation and the q-Value, *Annals of Statistics* 31, 2013-2035.
- \_\_\_\_\_, Taylor, Jonathan E., and David Siegmund, 2004, Strong Control, Conservative Point Estimation and Simultaneous Conservative Consistency of False Discovery Rates: A Unified Approach, *Journal of the Royal Statistical Society* 66, 187-205.
- Tang, Ke and Wei Xiong, 2009, Index Investing and the Financialization of Commodities, working paper, Princeton University.
- Teweles, Richard J., Harlow, Charles V. and Herbert L. Stone, 1977, *The Commodity Futures Game: Who Wins? Who Loses? Why?*, New York: McGraw-Hill.
- Wermers, R., 2000, Mutual Fund Performance: An Empirical Decomposition into Stock-Picking Talent, Style, Transactions Costs, and Expenses, *Journal of Finance* 55, 1655-1695.
- White, Halbert, 1980, A Heteroscedasticity-consistent Covariance Matrix Estimator and a Direct Test for Heteroscedasticity," *Econometrica* 48, 817-838.

#### APPENDIX

#### Implementing the False Discovery Rate (FDR) Method

Implementing the *FDR* procedure requires that we estimate the three terms on the right hand side of equation (5) in the text. Given *c*, the first term,  $\Pr(z_j > c | j \in \{\text{null}\})$ , is the *p*-value for  $z_j$ , which can be retrieved from the standard normal distribution. The second term,  $\Pr(z_j > c)$ , can be estimated from the proportion of null hypotheses that are rejected at critical value *c*.

Formulating an estimator for the third term,  $\pi_0$ , the proportion of null traders in the population requires some assumptions about the distribution of  $z_j$  for the informed and uninformed traders. In their seminal paper, Benjamini and Hochberg (1995) avoid such assumptions by setting  $\pi_0=1$ , which reduces the *FDR* to

$$FDR_{BH}(c) = \frac{\Pr(z_j > c \mid j \in \{\text{null}\})}{\Pr(z_j > c)}$$
(A1)

Because  $\pi_0$  cannot exceed one, equation (A1) provides an upper bound on the true *FDR*. Benjamini and Hochberg prove that this approach generates a conservative test, for which the proportion of falsely rejected hypotheses is less than *FDR*<sub>BH</sub>.

Alternatively, Storey (2002, 2004) provides a method of estimating  $\pi_0$  based on the view that *z*-statistics close to zero are generated by null traders. In terms of *p*-values, he assumes that null traders generate all *z*-statistics for two-sided *p*-values greater than some level,  $\lambda$ . We write Storey's assumption as

$$\Pr\left(j \in \{\operatorname{null}\} \middle| 2\left(1 - \Phi(|z_j|)\right) > \lambda\right) = 1,$$
(A2)

where  $\Phi$  denotes the standard normal CDF and  $2(1-\Phi(|z_j|))$  is the two-sided *p*-value.

Rewriting the condition  $2(1-\Phi(|z_j|)) > \lambda$  as  $|z_j| < \Phi^{-1}(1-0.5\lambda)$  and applying this assumption yields

$$\pi_{0} = \Pr\left(j \in \{\text{null}\}\right)$$

$$= \frac{\Pr\left(j \in \{\text{null}\} \mid z_{j} \mid < \Phi^{-1}(1 - 0.5\lambda)\right) \Pr\left(|z_{j}| < \Phi^{-1}(1 - 0.5\lambda)\right)}{\Pr\left(|z_{j}| < \Phi^{-1}(1 - 0.5\lambda)\right| \ j \in \{\text{null}\}\right)}$$

$$= \frac{\Pr\left(|z_{j}| < \Phi^{-1}(1 - 0.5\lambda)\right)}{1 - \lambda}$$
(A3)

where the denominator in the last line follows from the definition of  $\lambda$ . Thus, for a given  $\lambda$  we can estimate  $\pi_0$  as

$$\hat{\pi}_{0} = \frac{1}{(1-\lambda)n} \sum_{i=1}^{n} \mathbb{1}\left( |z_{i}| < \Phi^{-1} \left( 1 - 0.5\lambda \right) \right)$$
(A4)

where  $1(\bullet)$  denotes the unit indicator function. Combining Storey's approach with estimates for the first two terms gives the follow estimator for equation (5):

$$FDR(c) = \frac{(1 - \Phi(c))\hat{\pi}_0}{n^{-1}\sum_{i=1}^n \mathbf{1}(z_i > c)}$$

$$= \frac{(1 - \Phi(c))\sum_{i=1}^n \mathbf{1}(|z_i| < \Phi^{-1}(1 - 0.5\lambda))}{(1 - \lambda)\sum_{i=1}^n \mathbf{1}(z_i > c)}$$
(A5)

We implement the *FDR* test by choosing the minimum *c* such that  $FDR(c) \le 0.05$ , and judge trader *j* to have predictive ability if  $z_j > c$ .

The important question presented by Storey's approach is how to choose  $\lambda$ . Storey (2002) uses a bootstrapping method, which is also implemented by Barras, Scaillet and Wermers (2010). For our analysis, we develop a chi-square goodness-of-fit test to select an optimal  $\lambda$ . This

method is simple to implement and also provides information on how much the alternative hypothesis vitiates the null population.

To illustrate Storey's and our approach, consider how Dalmasso, Broet and Moreau (2005) construct the general problem of identifying  $\pi_0$  as a mixture of distributions. In any set of multiple tests, the observed *p*-values will be distributed under the null with probability  $\pi_0$  and under the alternative(s) with probability  $(1 - \pi_0)$ . Thus, the expected *p*-value across tests is defined by the following mixture using the mixing parameter,  $\pi_0$ :

$$E[P] = \pi_0 E_0 [P] + (1 - \pi_0) E_1 [P]$$
(A6)

where  $E_0$  and  $E_1$  are expectations taken over the null and alternative distributions, respectively. Under the null, *p*-values are uniformly distributed with the  $n^{\text{th}}$  moment from the origin equal to 1/(n+1). Dividing (A1) by  $E_0[P^n]$  produces an obvious estimator that is expected to bound  $\pi_0$ :

$$\pi_0 \le (n+1) \sum_{i=1}^{M} p_i^n / M$$
(A7)

where *M* is the number of tests and  $p_i$  is the observed *p*-value of each test. Dalmasso, Broet and Moreau (2005) develop a class of estimators based on generalizing (A6) to the expectation over a transformation of *P*. The transformation they develop is denoted by  $\varphi(p)^n = [-\ln(1-p)]^n$ , which has the feature that the bias for  $\pi_0$  is decreasing in *n*. In contrast, the bias of the moments estimator in (A7) is increasing in *n*, so the best choice is n = 1, or two times the sample mean of the *p*-values.

Storey (2002) uses the result that the null distribution is uniform to develop his bootstrapping approach for bounding  $\pi_0$ . This method relies on finding a good cutoff point ( $\lambda$ ) in the

distribution of ordered *p*-values (lowest to highest) beyond which the *p*-values are likely from the null and exhibit uniformity. For a given cutoff, the estimate of  $\pi_0$  is defined as:

$$\pi_0^*(\lambda) = \frac{\#\{p_i > \lambda\}}{M} \frac{1}{(1-\lambda)}$$
(A8)

where the first term represents the fraction of p-values exceeding the cutoff and the second term rescales this fraction to the entire distribution of p-values.

Storey's (2002, 2003) method computes  $FDR(\lambda,\alpha)$  for the sample *p*-values over a discrete range (*R*) of possible  $\lambda$ , given a known rejection rate ( $\alpha$ ). The procedure then bootstraps the sample distribution of *p*-values to compute b = 1,2,3...,B samples of size *M*. In the  $b^{\text{th}}$  sample, let  $FDR_b(\lambda,\alpha)$  define the false discovery rate statistic, then over all *B* samples the mean square error (*MSE*) is computed as:

$$MSE(\lambda) = \frac{1}{B} \sum_{b=1}^{B} \left[ FDR_b(\lambda, \alpha) - \min_{\lambda' \in \mathbb{R}} \{ FDR(\lambda', \alpha) \} \right]^2$$
(A9)

The value of  $\lambda$  that minimizes (A4) is selected as the optimal cutoff and used in (A8) to estimate an upper bound on  $\pi_0$ .

We offer an alternative procedure for estimating the optimal cutoff that avoids the bootstrapping method. To implement this approach, divide the ordered *p*-values into i = 1,2,3...,r discrete groups according to the discrete range of possible  $\lambda$ . Beginning with the last two groups, which contain the highest of the ordered *p*-values, compute the chi-squared goodness-of-fit test statistic for the null that these groups are drawn from a uniform distribution. The last two groups are *r* and *r*-1. Let  $\chi^2(\lambda(r-i), k)$  denote this test statistic, where k (= i) defines the degrees of

freedom and  $\lambda(r-i)$  represents the cutoff value implied by the  $(r-i)^{\text{th}}$  group. Select the cutoff value from the results across all groups:

$$\lambda_{best} = \arg \max_{i=1...(r-1)} \left\{ prob[\chi^2(\lambda(r-i),k)] \right\}$$
(A10)

The solution to (A10) produces in a cutoff value implied by the set of *p*-values that are *least likely* to reject the null hypothesis of a uniform distribution. This solution is substituted into (A8) to estimate the bound on  $\pi_0$ . The intuition for this test is same as Storey (2002) offered for equation (A8), which is that the greatest effect from an alternative hypothesis is likely to be found in the lowest *p*-values. Thus, for some group of high *p*-values, we expect that the uniform distribution is a good fit. The chi-squared statistic offers a functional measure to define that fit, although other measures may also perform well (e.g., the Anderson-Darling statistic). In applying this test to our empirical analyses, we use 10 bins to define the number of groups, although 20 bins gave similar results.

## Table 1Six Trader Types

A trader is characterized by whether her open interest is consistent with before, during or after price changes. Open interest is observed in discrete intervals from yesterday to today's close, and price changes are matched accordingly. The omitted group type is a trader that acts completely at random (i.e., a noise trader) and thus is not statistically significant in any of the six trader types.

| Open Interest (OI)     | OI move is consistent<br>with price changes | OI move is inconsistent<br>with price changes |
|------------------------|---|---|
| OI moves before prices | Informed                                    | Uninformed                                    |
| OI moves with prices   | Large Liquidity<br>Demander                 | Large Liquidity<br>Supplier                   |
| OI moves after prices  | Momentum                                    | Contrarian                                    |

## Table 2Sample Characteristics

This table provides summary information for our sample data, which includes all futures positions held at the end-of-day by traders in who reported to the CFTC. The sample includes all contracts traded between January 2000 and May 2009. Trader-specific characteristics are shown in Panel (A) with an explanation of how each characteristic is measured. Panel (B) shows the distribution of traders across business lines. Panel (C) shows how representative the CFTC sample is compared to all positions in the market and the total number of unique reporting traders by commodity. Excluding the overlap between commodities, the sample include 8,921 unique traders.

| Characteristic  | Crude<br>oil | Copper | Corn         | Cotton        | Gold        | Heating<br>oil | Natural<br>gas | Silver | Soybean<br>oil | Soybeans | Sugar | Wheat |
|---|--------------|--------|--------------|---------------|-------------|----------------|----------------|--------|----------------|----------|-------|-------|
|   |              |        | Panel (A)    | : Trader-Sp   | ecific Cha  | racteristics   |                |        |                |          |       |       |
| <b>Experience:</b> Average number of position days; that is, OI>0   | 359          | 346    | 449          | 156           | 338         | 457            | 437            | 361    | 429            | 343      | 160   | 357   |
| Active: Days with trades divided by total position days   | 0.82         | 0.55   | 0.72         | 0.67          | 0.69        | 0.77           | 0.79           | 0.67   | 0.62           | 0.75     | 0.72  | 0.66  |
| Size: Average number of futures long and short contracts held at the end of daily trading                     | 332          | 175    | 285          | 282           | 357         | 232            | 194            | 237    | 347            | 157      | 938   | 238   |
| Expirations: Average number of contract expirations held  | 3.47         | 1.77   | 2.54         | 1.87          | 1.56        | 3.74           | 6.13           | 1.53   | 2.28           | 2.02     | 2.56  | 1.81  |
| <b>Net Long:</b> Average net long position divided by average of long and short positions; zero if net short  | 0.48         | 0.64   | 0.66         | 0.69          | 0.77        | 0.50           | 0.54           | 0.74   | 0.60           | 0.59     | 0.69  | 0.66  |
| <b>Net Short:</b> Average net short position divided by average of long and short positions; zero if net long | 0.56         | 0.68   | 0.64         | 0.59          | 0.61        | 0.49           | 0.40           | 0.61   | 0.50           | 0.58     | 0.60  | 0.54  |
|   |              | Pa     | nel (B): Dis | stribution of | f Traders b | y Business     | Line           |        |                |          |       |       |
| Commercials (AD,AM,AO,AP)   | 15.4%        | 11.7%  | 28.7%        | 16.2%         | 7.2%        | 29.8%          | 21.2%          | 9.1%   | 19.3%          | 20.0%    | 23.7% | 16.2% |
| Swap/Derivatives Dealer (AS)  | 3.7%         | 2.7%   | 1.1%         | 4.8%          | 2.3%        | 4.7%           | 4.8%           | 2.4%   | 3.4%           | 1.3%     | 6.8%  | 1.9%  |
| Floor Broker/Trader (FBT)   | 9.7%         | 6.7%   | 10.5%        | 12.9%         | 6.9%        | 9.2%           | 8.9%           | 6.7%   | 15.9%          | 11.0%    | 10.8% | 12.1% |
| Managed Money Trader (MMT)  | 24.3%        | 31.1%  | 15.9%        | 35.7%         | 26.5%       | 26.9%          | 28.0%          | 28.1%  | 30.7%          | 16.4%    | 36.1% | 24.3% |
| Other (NRP & specialized)   | 47.0%        | 47.8%  | 43.6%        | 30.4%         | 57.2%       | 29.5%          | 37.1%          | 53.6%  | 30.8%          | 51.5%    | 22.6% | 45.3% |
|   |              | I      | Panel (C): R | Representati  | on of Repa  | orted Positio  | ons            |        |                |          |       |       |
| Percent of All Market OI  | 92.1%        | 89.2%  | 82.1%        | 79.2%         | 91.2%       | 86.9%          | 94.6%          | 90.1%  | 89.0%          | 80.8%    | 75.7% | 85.6% |
| Total Reporting Traders   | 1,830        | 1,065  | 3,345        | 705           | 1,580       | 709            | 1,276          | 940    | 857            | 2,856    | 545   | 1,881 |

# Table 3 Overnight Informed Identified Using the False Discovery Rate (FDR) Method

Overnight informed traders are identified using the lagged trading profits and position profit rules to define forecasting success. This table shows the number of such traders found to be significant by the FDR method when the critical value is set at the 5% level of significance in a one-sided test. The informed are identified using close-to-close prices. The unconditional and HM test counts are reported and show the percent of such informed traders relative to all traders in the sample. Also, the percent of daily informed open interest relative to sample daily open interest is shown for the position profits results. The total counts across commodities are shown at the bottom of the table, which also shows the range of critical values arising from applying the FDR criterion to determine the significance of these tests. The sample includes trading between January 2000 to May 2009 and uses all traders who had 30 or more observations on the respective success variable. A liquidity filter is used to remove days in which there was less than 30 contracts traded in a given expiration.

|                                 |                           | Lagged Tra | ding Profi | ts        | Position Profits |                |         |       |           |         |  |  |
|---------------------------------|---------------------------|------------|------------|-----------|------------------|----------------|---------|-------|-----------|---------|--|--|
|                                 | <b>Unconditional Test</b> |            | HM         | 1 Test    | Uı               | nconditional T | `est    |       | HM Test   |         |  |  |
| Commodity                       | Count                     | % Traders  | Count      | % Traders | Count            | % Traders      | % of OI | Count | % Traders | % of OI |  |  |
| Crude Oil                       | 0                         | 0.0%       | 0          | 0.0%      | 4                | 0.3%           | 0.3%    | 4     | 0.4%      | 0.2%    |  |  |
| Copper                          | 0                         | 0.0%       | 0          | 0.0%      | 72               | 7.3%           | 10.4%   | 51    | 8.1%      | 7.9%    |  |  |
| Corn                            | 0                         | 0.0%       | 0          | 0.0%      | 33               | 1.1%           | 3.2%    | 25    | 1.4%      | 3.9%    |  |  |
| Cotton                          | 0                         | 0.0%       | 0          | 0.0%      | 0                | 0.0%           | 0.0%    | 0     | 0.0%      | 0.0%    |  |  |
| Gold                            | 0                         | 0.0%       | 2          | 0.2%      | 0                | 0.0%           | 0.0%    | 1     | 0.1%      | 0.1%    |  |  |
| Heating oil                     | 1                         | 0.2%       | 1          | 0.2%      | 0                | 0.0%           | 0.0%    | 4     | 0.8%      | 0.3%    |  |  |
| Natural gas                     | 0                         | 0.0%       | 0          | 0.0%      | 4                | 0.4%           | 0.3%    | 4     | 0.5%      | 0.4%    |  |  |
| Silver                          | 0                         | 0.0%       | 0          | 0.0%      | 93               | 12.0%          | 14.3%   | 0     | 0.0%      | 0.0%    |  |  |
| Soybean oil                     | 0                         | 0.0%       | 0          | 0.0%      | 27               | 3.5%           | 7.4%    | 1     | 0.2%      | 0.4%    |  |  |
| Soybeans                        | 0                         | 0.0%       | 1          | 0.1%      | 11               | 0.5%           | 0.9%    | 4     | 0.2%      | 0.5%    |  |  |
| Sugar                           | 1                         | 0.3%       | 1          | 0.4%      | 0                | 0.0%           | 0.0%    | 0     | 0.0%      | 0.0%    |  |  |
| Wheat                           | 0                         | 0.0%       | 0          | 0.0%      | 2                | 0.1%           | 2.3%    | 2     | 0.2%      | 1.2%    |  |  |
| Total                           | 2                         | 0.0%       | 5          | 0.1%      | 246              | 2.8%           | 3.3%    | 96    | 1.1%      | 1.2%    |  |  |
| Total (less duplicates)         | 2                         |            | 5          |           | 230              |                |         | 94    |           |         |  |  |
| Range of FDR<br>Critical Values | 3.58                      | 8 to 4.03  | 3.66       | to 4.07   | 2.28             | 3 to 3.98      |         | 2.59  | to 4.30   |         |  |  |

#### Table 4

#### **Characteristics Inferred from Overnight Informed Traders**

This table uses the overnight informed results in Table 3 to infer how group characteristics affect a trader's success or nonsuccess as judged by the position profits rule. For these regressions, the discrete dependent variable equals one if the trader is identified by the FDR approach as significantly informed; zero otherwise. The independent variables are defined in Table 2. The Experience, Active, Size, and Expirations variables enter as logarithms. For the business line dummy variables. The model is estimated without a constant term, so the coefficient on each business-type dummy variable represents the full effect of the group. The sample percentage representation of each business type is shown for comparison. The price measure used to judge success is shown at the top of the table. Observations are pooled across commodities with fixed effects defined for each commodity. We adjust for clustering between commodities by trader and use White's (1980) estimator to correct for heteroscedasticity. The coefficients reported here are transformed using the projection defined by equation (8) in the text. The significance level of the underlying coefficient is reported next to the transformed coefficient, where an "\*" ("\*\*") indicates significance at the 90% (95%) level of confidence.

|                  | Proportion in | Unc     | onditional ' | Гest     | Proportion in |         | H-M Test |          |
|------------------|---------------|---------|--------------|----------|---------------|---------|----------|----------|
| Variables        | Sample        | (1)     | (2)          | (3)      | Sample        | (4)     | (5)      | (6)      |
|                  |               |         |              |          |               |         |          |          |
| Experience       |               |         | 0.23 **      | 0.25 **  |               |         | 0.10 *   | 0.10 *   |
| Active           |               |         | -0.07 *      | -0.07 *  |               |         | 0.16 **  | 0.12 **  |
| Size             |               |         | 0.08         | 0.11 *   |               |         | 0.17 *   | 0.23 **  |
| Expirations      |               |         | 0.02         | 0.06 **  |               |         | 0.01     | 0.01     |
| Net Long         |               |         | -0.06 **     | -0.02    |               |         | -0.23 ** | -0.19 ** |
| Net Short        |               |         | -0.19 **     | -0.12 ** |               |         | -0.28 ** | -0.22 ** |
| Commercial       | 0.21          | 0.06 ** |              | 0.02     | 0.21          | 0.06 ** |          | 0.07 **  |
| Swaps Dealer     | 0.03          | 0.01 *  |              | -0.01    | 0.03          | 0.01    |          | 0.00     |
| FBT              | 0.10          | 0.24 ** |              | 0.19 **  | 0.14          | 0.40 ** |          | 0.27 **  |
| Managed Money    | 0.25          | 0.36 ** |              | 0.28 **  | 0.27          | 0.24 ** |          | 0.19 **  |
| Other            | 0.41          | 0.33 ** |              | 0.42 **  | 0.36          | 0.29 ** |          | 0.36 **  |
| <b>R-Squared</b> |               | 0.007   | 0.015        | 0.02     |               | 0.006   | 0.012    | 0.018    |
| Sample Size      |               | 14518   | 14518        | 14518    |               | 10037   | 10037    | 10037    |
| FDR counts       |               | 246     | 246          | 246      |               | 96      | 96       | 96       |

# Table 5 Identifying Intraday Liquidity, Informed, Momentum & Contrarian Traders

This table shows the counts of momentum and contrarian traders as identified by the previous-overnight price change rule and informed, liquidity demanders and suppliers as identified by the triple test. To select momentum (contrarian) traders, the test identifies traders whose change in position moves with (against) the close-to-open price change on the previous overnight. To find informed traders, the triple test identifies traders whose change in open interest moves directly with intraday prices, conversely with the prior open-to-close price and whose end-of-day positions are predictive of the subsequent day's price change. The liquidity traders are selected from the groups that fail to predict the subsequent day's price change. The triple test counts include only traders who also pass the intraday price change test.

|                            |             |                  | Tripl                 |         | Previous Overnight Test |         |       |         |            |         |  |
|----------------------------|-------------|------------------|-----------------------|---------|-------------------------|---------|-------|---------|------------|---------|--|
|                            | Liqı<br>Sup | uidity<br>oplier | Liquidity<br>Demander |         | Info                    | ormed   | Mom   | ientum  | Contrarian |         |  |
| Commodity                  | Count       | Percent          | Count                 | Percent | Count                   | Percent | Count | Percent | Count      | Percent |  |
| Crude Oil                  | 9           | 0.8%             | 14                    | 1.3%    | 7                       | 0.7%    | 56    | 4.9%    | 70         | 6.1%    |  |
| Copper                     | 2           | 0.3%             | 0                     | 0.0%    | 9                       | 1.5%    | 93    | 14.1%   | 165        | 25.0%   |  |
| Corn                       | 2           | 0.1%             | 24                    | 1.2%    | 28                      | 1.4%    | 145   | 6.9%    | 322        | 15.4%   |  |
| Cotton                     | 0           | 0.0%             | 16                    | 5.0%    | 0                       | 0.0%    | 14    | 4.1%    | 16         | 4.7%    |  |
| Gold                       | 4           | 0.5%             | 8                     | 1.0%    | 5                       | 0.6%    | 99    | 11.4%   | 108        | 12.4%   |  |
| Heating oil                | 0           | 0.0%             | 6                     | 1.2%    | 3                       | 0.6%    | 28    | 5.4%    | 34         | 6.6%    |  |
| Natural gas                | 6           | 0.7%             | 16                    | 1.9%    | 22                      | 2.7%    | 67    | 7.8%    | 101        | 11.7%   |  |
| Silver                     | 6           | 1.2%             | 11                    | 2.3%    | 7                       | 1.4%    | 56    | 10.7%   | 90         | 17.1%   |  |
| Soybean oil                | 3           | 0.6%             | 13                    | 2.5%    | 21                      | 4.1%    | 55    | 9.9%    | 60         | 10.8%   |  |
| Soybeans                   | 18          | 1.2%             | 17                    | 1.1%    | 36                      | 2.4%    | 124   | 7.5%    | 114        | 6.9%    |  |
| Sugar                      | 0           | 0.0%             | 0                     | 0.0%    | 1                       | 0.3%    | 9     | 3.0%    | 25         | 8.3%    |  |
| Wheat                      | 43          | 4.6%             | 27                    | 2.9%    | 19                      | 2.0%    | 88    | 8.6%    | 65         | 6.4%    |  |
| Total                      | 93          |                  | 152                   |         | 158                     |         | 834   |         | 1170       |         |  |
| Total (less<br>duplicates) | 90          |                  | 115                   |         | 91                      |         | 444   |         | 912        |         |  |

### Table 6 Effects Inferred from Triple and Previous-Overnight Tests

This table reports estimates of the inverse regression coefficients using the counts identified by the pairs and triple tests. Previous-overnight test results are shown for momentum and contrarian traders and triple test results are shown for liquidity demanders, liquidity suppliers and informed traders. For these regressions, the discrete dependent variable equals one if the trader is identified by the FDR approach as significant; zero otherwise. The independent variables are defined in Table 2. The Experience, Active, Size, and Expirations variables enter in logarithms. The constant is omitted in these regressions. Observations are pooled across commodities with fixed effects defined for each commodity. We adjust for clustering between commodities by trader and use White's (1980) estimator to correct for heteroscedasticity. The coefficient reported here are transformed using the projection defined by equation (15) in the text. The significance level of the underlying coefficient is reported next to the transformed coefficient, where an "\*" ("\*\*") indicates significance at the 90% (95%) level of confidence.

|  |                   |                     |                     |                     | Triple               | Test to Idea         | ntify                |                      |                      |                      |                       | Prev                  | vious-Overni          | ght Test to Iden       | tify                   |                        |
|--|-------------------|---------------------|---------------------|---------------------|----------------------|----------------------|----------------------|----------------------|----------------------|----------------------|-----------------------|-----------------------|-----------------------|------------------------|------------------------|------------------------|
|  | Sample<br>Percent | Liq                 | uidity Supp         | lier                | Liqu                 | idity Demar          | ıder                 |                      | Informed             |                      |                       | Momentum              |                       | Contrarian             |                        |                        |
| Variables                              |                   | (1)                 | (2)                 | (3)                 | (4)                  | (5)                  | (6)                  | (7)                  | (8)                  | (9)                  | (10)                  | (11)                  | (12)                  | (13)                   | (14)                   | (15)                   |
| Experience                             |                   |                     | 0.41 **             | 0.38 **             |                      | 0.01                 | 0.04                 |                      | 0.10                 | 0.15 *               |                       | 0.19 **               | 0.22 **               |                        | 0.24 **                | 0.21 **                |
| Active                                 |                   |                     | 0.23 **             | 0.19 **             |                      | 0.26 **              | 0.27 **              |                      | 0.39 **              | 0.39 **              |                       | 0.00                  | 0.04                  |                        | 0.14 **                | 0.12 **                |
| Size                                   |                   |                     | -0.13               | -0.06               |                      | 0.57 **              | 0.44 **              |                      | 0.79 **              | 0.58 **              |                       | 0.10 *                | 0.02                  |                        | 0.01                   | 0.06                   |
| Expirations                            |                   |                     | 0.08                | 0.01                |                      | -0.40 **             | -0.25 **             |                      | -0.52 **             | -0.36 **             |                       | -0.34 **              | -0.20 **              |                        | 0.08 **                | 0.03 *                 |
| Net Long                               |                   |                     | 0.02                | 0.02                |                      | -0.08 **             | -0.06 **             |                      | 0.03                 | 0.02                 |                       | -0.01                 | 0.00                  |                        | 0.03 **                | 0.03 **                |
| Net Short                              |                   |                     | 0.09 **             | 0.08 **             |                      | -0.14 **             | -0.08 **             |                      | -0.08 **             | -0.04 *              |                       | -0.04 **              | -0.02                 |                        | 0.08 **                | 0.05 **                |
| Commercial                             | 0.25              | 0.42 **             |                     | 0.24 **             | 0.02 *               |                      | 0.04 **              | 0.02 *               |                      | 0.02                 | 0.08 **               |                       | 0.08 **               | 0.46 **                |                        | 0.31 **                |
| Swaps Dealer                           | 0.04              | 0.08 **             |                     | 0.05 *              | 0.07 **              |                      | 0.05 **              | 0.17 **              |                      | 0.13 **              | 0.03 **               |                       | 0.03 **               | 0.04 **                |                        | 0.02 **                |
| FBT                                    | 0.13              | 0.24 **             |                     | 0.19 **             | 0.19 **              |                      | 0.13 **              | 0.13 **              |                      | 0.09 **              | 0.05 **               |                       | 0.05 **               | 0.13 **                |                        | 0.11 **                |
| Managed Money                          | 0.26              | 0.07 **             |                     | 0.09 **             | 0.52 **              |                      | 0.37 **              | 0.53 **              |                      | 0.34 **              | 0.56 **               |                       | 0.37 **               | 0.13 **                |                        | 0.13 **                |
| Other                                  | 0.33              | 0.20 **             |                     | 0.30 **             | 0.20 **              |                      | 0.27 **              | 0.15 **              |                      | 0.28 **              | 0.28 **               |                       | 0.33 **               | 0.25 **                |                        | 0.32 **                |
| R-Squared<br>Sample Size<br>FDR counts |                   | 0.004<br>9845<br>93 | 0.011<br>9845<br>93 | 0.012<br>9845<br>93 | 0.009<br>9845<br>152 | 0.017<br>9845<br>152 | 0.021<br>9845<br>152 | 0.016<br>9845<br>158 | 0.032<br>9845<br>158 | 0.043<br>9845<br>158 | 0.046<br>10544<br>834 | 0.060<br>10544<br>834 | 0.076<br>10544<br>834 | 0.032<br>10544<br>1170 | 0.073<br>10544<br>1170 | 0.083<br>10544<br>1170 |

## Table 7 Rank-Sum and t-Tests of Informed and Not-Informed Using Daily Profits per Trader

This table compares the daily profits per trader between those identified as informed and not informed in the FDR tests. The p-value for the Wilcoxon rank-sum is shown along with the mean rank-sums for both informed and not informed participants by commodity. The t-test comparing average daily profits per trader is reported below the Wilcoxon results. The p-value for the t-test is computed assuming unequal variances. Panel (A) shows these results for informed traders identified in Table 3 as found by the HM and unconditional position tests using close-to-close prices. Panel (B) shows these results for the intraday informed traders identified in Table 5 by computing the profits for position changes evaluated using midpoint and closing prices. Excluding the overlap between commodities, the sample include 8,921 unique traders.

|                                  | Crude oil  | Copper     | Corn      | Cotton       | Gold            | Heating oil     | Na    | tural gas | Silver     | So   | ybean oil | So | ybeans | Sugar | Nheat         |
|----------------------------------|------------|------------|-----------|--------------|-----------------|-----------------|-------|-----------|------------|------|-----------|----|--------|-------|---------------|
|                                  |            |            | Pa        | nel (A): Ove | ernight Informe | d Identified i  | n Ta  | ble 3     |            |      |           |    |        |       |               |
| Unconditional Position Test:     |            |            |           |              |                 |                 |       |           |            |      |           |    |        |       |               |
| Informed Mean Rank Sum           | 2,420      | 2,355      | 2,426     | ó n.a.       | n.a.            | n.a.            |       | 2,341     | 2,407      |      | 2,387     |    | 2,379  | n.a.  | 2,368         |
| Not Informed Mean Rank Sum       | 2,189      | 2,244      | 2,193     | 8 n.a.       | n.a.            | n.a.            |       | 2,111     | 2,192      |      | 2,236     |    | 2,238  | n.a.  | 2,251         |
| Wilcoxon p-value                 | 0.000      | 0.002      | 0.000     | ) n.a.       | n.a.            | n.a.            |       | 0.000     | 0.000      |      | 0.000     |    | 0.000  | n.a.  | 0.001         |
| Informed Aver. Daily Profits     | \$ 33,351  | \$ 26,499  | \$ 7,350  | ) n.a.       | n.a.            | n.a.            | \$    | 39,807    | \$ 15,390  | \$   | 16,743    | \$ | 12,432 | n.a.  | \$<br>12,699  |
| Not Informed Aver. Daily Profits | \$ 1,333   | \$ (4,918) | \$ 443    | 8 n.a.       | n.a.            | n.a.            | \$    | (1,596)   | \$ (10,064 | ) \$ | (1,595)   | \$ | (774)  | n.a.  | \$<br>891     |
| t-test p-value                   | 0.008      | 0.012      | 0.000     | ) n.a.       | n.a.            | n.a.            |       | 0.254     | 0.089      |      | 0.100     |    | 0.049  | n.a.  | 0.531         |
| HM Position Test:                |            |            |           |              |                 |                 |       |           |            |      |           |    |        |       |               |
| Informed Mean Rank Sum           | 2,037      | 2,357      | 2,391     | n.a.         | 2,102           | 1,781           |       | 2,299     | n.a.       |      | 1,719     |    | 2,342  | n.a.  | 2,266         |
| Not Informed Mean Rank Sum       | 1,826      | 2,242      | 2,228     | 8 n.a.       | 2,060           | 1,681           |       | 2,151     | n.a.       |      | 1,656     |    | 2,220  | n.a.  | 2,225         |
| Wilcoxon p-value                 | 0.000      | 0.002      | 0.000     | ) n.a.       | 0.130           | 0.003           |       | 0.000     | n.a.       |      | 0.043     |    | 0.001  | n.a.  | 0.141         |
| Informed Aver. Daily Profits     | \$ 45,237  | \$ 40,179  | \$ 7,396  | 5 n.a.       | \$ (13,571)     | \$ 38,651       | \$    | 39,807    | n.a.       | \$   | (548)     | \$ | 4,086  | n.a.  | \$<br>6,555   |
| Not Informed Aver. Daily Profits | \$ (3,401) | \$ (726)   | \$ (156   | 5) n.a.      | \$ (10,366)     | \$ (5,693)      | ) \$  | (8,194)   | n.a.       | \$   | (660)     | \$ | 5,760  | n.a.  | \$<br>(307)   |
| t-test p-value                   | 0.017      | 0.006      | 0.003     | 3 n.a.       | 0.798           | 0.014           |       | 0.185     | n.a.       |      | 0.977     |    | 0.890  | n.a.  | 0.210         |
|                                  |            |            | Pa        | nel (B): Int | traday Informed | l Identified ir | ı Tal | ble 5     |            |      |           |    |        |       |               |
| Informed Mean Rank Sum           | 2.399      | 1.647      | 2,198     | 3 n.a.       | 2.592           | 2.235           |       | 2,743     | 2.504      |      | 2.664     |    | 2.718  | n.a.  | 2.818         |
| Not Informed Mean Rank Sum       | 2.169      | 1.355      | 1.437     | 7 n.a.       | 1,940           | 2.366           |       | 1.858     | 2.032      |      | 1.953     |    | 1.897  | n.a.  | 1.797         |
| Wilcoxon p-value                 | 0.000      | 0.000      | 0.000     | ) n.a.       | 0.000           | 0.000           |       | 0.000     | 0.000      |      | 0.000     |    | 0.000  | n.a.  | 0.000         |
| Informed Aver. Daily Profits     | \$ 24.841  | \$ 4.764   | \$ 10.006 | 5 n.a.       | \$ 26.766       | \$ 7.615        | \$    | 95.885    | \$ 13.689  | \$   | 4.218     | \$ | 13.110 | n.a.  | \$<br>15,488  |
| Not Informed Aver. Daily Profits | \$ 9,849   | \$ 35      | \$ (210   | )) n.a.      | \$ (3,585)      | \$ 9,358        | \$    | (2,693)   | \$ 93      | \$   | (1,542)   | \$ | 2,376  | n.a.  | \$<br>(1,681) |
| t-test p-value                   | 0.000      | 0.000      | 0.000     | ) n.a.       | 0.000           | 0.242           | ·     | 0.000     | 0.000      |      | 0.000     | ·  | 0.149  | n.a.  | 0.000         |

# Table 8Unique Counts of Informed Traders by Commodity

This table shows total unique counts of informed traders by commodity and the fraction of traders found to be significantly informed in the tests shown, which is the overlap of informed traders between tests. The counts are for combining the unconditional and intraday (or triple) tests and the HM and intraday tests. The total counts are shown at the bottom of the table along with the average overlap across commodities.

|                   | Unique Co                 | ounts          | Fraction Overlapping      |                |  |  |  |  |
|-------------------|---------------------------|----------------|---------------------------|----------------|--|--|--|--|
| Commodity         | Unconditional<br>& Triple | HM &<br>Triple | Unconditional<br>& Triple | HM &<br>Triple |  |  |  |  |
| Crude Oil         | 10                        | 10             | 10.0%                     | 10.0%          |  |  |  |  |
| Copper            | 77                        | 58             | 5.2%                      | 3.4%           |  |  |  |  |
| Corn              | 56                        | 50             | 8.9%                      | 6.0%           |  |  |  |  |
| Cotton            | 0                         | 0              | 0.0%                      | 0.0%           |  |  |  |  |
| Gold              | 5                         | 6              | 0.0%                      | 0.0%           |  |  |  |  |
| Heating oil       | 3                         | 7              | 0.0%                      | 0.0%           |  |  |  |  |
| Natural gas       | 26                        | 26             | 0.0%                      | 0.0%           |  |  |  |  |
| Silver            | 97                        | 7              | 3.1%                      | 0.0%           |  |  |  |  |
| Soybean oil       | 46                        | 22             | 4.3%                      | 0.0%           |  |  |  |  |
| Soybeans          | 46                        | 40             | 2.2%                      | 0.0%           |  |  |  |  |
| Sugar             | 1                         | 1              | 0.0%                      | 0.0%           |  |  |  |  |
| Wheat             | 21                        | 21             | 0.0%                      | 0.0%           |  |  |  |  |
| Total and Average | 388                       | 248            | 4.1%                      | 2.4%           |  |  |  |  |

# **AFBORITHMIC FINANCE**

## Discovering the ecosystem of an electronic financial market with a dynamic machine-learning method

Shawn Mankad; George Michailidis; Andrei Kirilenko

Algorithmic Finance (2013), 2:2, 151-165 DOI: 10.3233/AF-13023

Abstract, HTML, and PDF: http://algorithmicfinance.org/2-2/pp151-165

**Aims and Scope** Algorithmic Finance is a high-quality academic research journal that seeks to bridge computer science and finance, including high frequency and algorithmic trading, statistical arbitrage, momentum and other algorithmic portfolio management strategies, machine learning and computational financial intelligence, agent-based finance, complexity and market efficiency, algorithmic analysis on derivatives, behavioral finance and investor heuristics, and news analytics.

#### **Managing Editor**

Philip Maymin, NYU-Polytechnic Institute

#### **Deputy Managing Editor**

Jayaram Muthuswamy, Kent State University

#### **Advisory Board**

Kenneth J. Arrow, Stanford University Herman Chernoff, Harvard University David S. Johnson, AT&T Labs Research Leonid Levin, Boston University Myron Scholes, Stanford University Michael Sipser, Massachusetts Institute of Technology Richard Thaler, University of Chicago Stephen Wolfram, Wolfram Research

#### **Editorial Board**

#### **Associate Editors**

Peter Bossaerts, California Institute of Technology Emanuel Derman, Columbia University Ming-Yang Kao, Northwestern University Pete Kyle, University of Maryland David Leinweber, Lawrence Berkeley National Laboratory Richard J. Lipton, Georgia Tech Avi Silberschatz, Yale University Robert Webb, University of Virginia

#### **Affiliate Editors**

Giovanni Barone-Adesi, University of Lugano Bruce Lehmann, University of California, San Diego



Subscription, submission, and other info: www.iospress.nl/journal/algorithmic-finance

**Unique Features of the Journal** *Open access*: Online articles are freely available to all. *No submission fees*: There is no cost to submit articles for review. There will also be no publication or author fee for at least the first two volumes. *Authors retain copyright*: Authors may repost their versions of the papers on preprint archives, or anywhere else, at any time. *Enhanced content*: Enhanced, interactive, computable content will accompany papers whenever possible. Possibilities include code, datasets, videos, and live calculations. *Comments*: Algorithmic Finance is the first journal in the Financial Economics Network of SSRN to allow comments. *Archives*: The journal is published by <u>IOS Press</u>. In addition, the journal maintains an <u>archive on SSRN.com Legal</u>: While the journal does reserve the right to change these features at any time without notice, the intent will always be to provide the world's most freely and quickly available research on algorithmic finance. *ISSN*: Online ISSN: 2157-6203. Print ISSN: 2158-5571.

## Discovering the ecosystem of an electronic financial market with a dynamic machine-learning method\*

Shawn Mankad<sup>†‡</sup>, George Michailidis<sup>‡</sup>, and Andrei Kirilenko<sup>‡‡</sup> <sup>‡</sup>Department of Statistics, University of Michigan, Michigan, USA <sup>‡‡</sup>MIT Sloan School of Management, Cambridge, MA, USA

Abstract. Not long ago were traded by human traders in face-to-face markets. The ecosystem of an open outcry market was well-known, visible to a human eye, and rigidly prescribed. Now trading is increasingly done in anonymous electronic markets where traders do not have designated functions or mandatory roles. In fact, the traders themselves have been replaced by algorithms (machines) operating with little or no human oversight. While the process of electronic trading is not visible to a human eye, machine-learning methods have been developed to recognize persistent patterns in the data. In this study, we develop a dynamic machine-learning method that designates traders in an anonymous electronic market into five persistent categories: high frequency traders, market makers, opportunistic traders, fundamental traders, and small traders. Our method extends a plaid clustering technique with a smoothing framework that filters out transient patterns. The method is fast, robust, and suitable for a discovering trading ecosystems in a large number of electronic markets.

Keywords: trading strategies, high frequency trading, machine learning, clustering

#### 1. Introduction

The words "stock market", "futures market" or "trading pit" used to elicit a mental picture of a chaotic crowd of agitated people wearing brightlycolored jackets, gesticulating wildly and shouting at each other. Yet, a trained human eye would see a great deal of structure behind this frenzy. Some of the people were market makers who stood at certain posts and "made markets" in securities or derivatives that were designated only to them. Some were floor brokers who formed circles around the market makers to get the best prices for a broad range of their customers - from

\*The views expressed in this paper are our own and do not constitute an official position of any agency, its management or staff. <sup>†</sup>Corresponding author: Electronic address: smankad@umich pension funds investing their assets, to banks hedging exposures on their balance sheets. Others were different types of floor traders from scalpers to spreaders to opportunistic position takers, who wandered around the floor looking for opportunities to exploit. The ecosystem of an open outcry market was well-known, visible to a human eye, and rigidly prescribed: traders had designated functions, used common gestures to trade, wore jackets of certain colors, and could be found in specific locations on a trading floor.

The transition to anonymous electronic trading has obfuscated the prescribed ecosystem of roles, relationships, and designations previously clearly visible on a trading floor. Trading floors have been replaced by server farms, prescribed gestures have been replaced by message protocols, and the traders themselves have been replaced by algorithms often operating with little or no human oversight.

While the process of electronic trading is not visible to a human eye, machine-learning methods have been

<sup>.</sup>edu; Fax Number: 734-763-4676; Phone Number: 734-763-3519 Mailing Address: University of Michigan; 269 West Hall, 1085 South University; Ann Arbor, MI 48109-1107

developed to recognize persistent patterns in the data. Even without a formal, regulatory designation, a trader who follows a particular strategy would leave a distinct footprint in the data.

In this study, we present a novel machine-learning method to parse through the footprints of all traders in a highly liquid, anonymous electronic market and find certain common "paths" that they follow, thus, describing the roles and functions of participants who inhabit the new ecosystem of an electronic financial market.

Our method combines a static plaid clustering technique with a dynamic smoothing framework that filters out transient patterns. The plaid clustering technique - a regression-based method to describe empirical regularities in cross-sectional data - was previously used only for a single, static data matrix. Our method extends the plaid model by making use of a time series of data matrices. Our extension, which we refer to as the smooth plaid model, is able to consistently identify categories of traders and trading outcomes that persist through time.

We utilize synthetic data generated from an agent based model (Paddrik et al., 2011; Hayes et al., 2012) that is calibrated on actual E-mini 500 stock index futures contract (E-mini) data made available to us by the CFTC. We originally employed our method on regulatory, transaction-level data for the E-mini the price discovery vehicle for the broad U.S. stock market. However, due to a Chicago Mercantile Exchange (CME) complaint, there has been a hold by the CFTC on the use and reporting of any regulatory data by the academic community. Using synthetic data allows us to make public our machine-learning methodology and results without concerns associated with the extreme confidentiality of regulatory data. The more methods that are available in the literature for analysis of electronic trading data, the more rigorous the discussion on the role and consequences of electronic markets. We note that results of applying our method on actual regulatory data are broadly similar. Moreover, we were able to designate traders into the same categories that were recovered manually by Kirilenko et al. (2010) in their analysis of the Flash Crash using similar data.

Our results are as follows. Using the smooth plaid model, we assign 6387 traders in the simulated data into five distinct categories: high frequency traders (7 traders), market makers (73), opportunistic traders (2405), fundamental buyers and sellers (1281), and small traders (2849). These traders occupy quite

distinct, albeit sometimes overlapping, positions in the ecosystem of the market. High frequency traders, whose data footprint sometimes resembles scalpers on steroids, occupy a very special position in the market. They trade through an enormous number of contracts each day, but carry very little inventory at any point in time. Market makers are in the market all the time; they quickly buy and sell on demand and manage their inventory very tightly. Fundamental traders accumulate directional inventory over long periods of time, often days, presumably to take a longer-term investment view or to hedge their other exposures. Opportunistic traders take on and manage directional bets for minutes or hours at a time, in search of opportunities to profit from the perceived imbalances. Small traders do not exhibit any persistent pattern; they enter the market very infrequently at seemingly random times and trade in trivial quantities.

We believe that the smooth plaid model in particular, and machine-learning methods, more generally, can be effectively used for the analysis of traders and their strategies in electronic financial markets. In an environment where traders do not have formal designations, the smooth plaid model forms a useful first step to separate tens of thousands of trading accounts into manageable trader categories for subsequent analysis - be it a market event like a Flash Crash<sup>1</sup>, co-movement of asset prices<sup>2</sup>, or the impact of trading strategies on market quality<sup>3</sup>.

The paper proceeds as follows. In the next section, we briefly summarize the plaid model (2.1), and discuss our modifications to create the smooth plaid algorithm (2.2). In Section 3, we illustrate the proposed model using a simple set of simulated data. In Section 4, we apply our method to simulated data generated by an agent-based simulation model of an electronic market calibrated to the E-mini. We conclude with a discussion and review of this study (Section 5).

<sup>&</sup>lt;sup>1</sup>See, Kirilenko et al. (2010). The authors separate their traders into categories manually. They arrive at similar categories as the ones presented in this study.

<sup>&</sup>lt;sup>2</sup>See, Huang (2011) for an investigation of co-movement of exchange rates. The authors develop a variant of a machine-learning technique with a parametric way to deal with the time-series dimension.

<sup>&</sup>lt;sup>3</sup>See, Chaboud et al. (2011) and Hendershott et al. (2010). The authors rely on designations given to them by a trading venue. They do not use a machine-learning method to cluster traders into categories.

#### 2. Methods

#### 2.1. Biclustering and the plaid model

Suppose we observe a data matrix  $X \in \mathbb{R}^{n \times p}$ , where  $X_{ij}$  represents the *i*th sample (i = 1, ..., n)and *j*th variable (j = 1, ..., p). The plaid model, first introduced by Lazzeroni and Owen (2000), aims to decompose the data to reveal the underlying structure. In our setting, the model is trained to discover groups of traders that have similar trading behaviors.

The term biclustering was first used by Cheng and Church (2000) to refer to grouping procedures appropriate when both the samples and variables are of scientific interest. In contrast, clustering methods belong to a closely related topic in machine learning and are concerned with discovering the structure of samples only. Hence, biclustering methods extract groups of samples (rows) and variables (columns) to find homogeneous submatrices in a static data matrix. These methods typically allow samples to be in more than one cluster, or in none at all. This flexibility is also given to variable groups, that is, variables can be defined with respect to only a subset of samples, not necessarily with respect to all of them. Moreover, these flexible models allow for overlapping biclusters.

In our application setting, samples are individual traders and the variables are measures of trading activity for each trader: trading volume, net position, change in inventory, trades per second, and median intertrade duration. A bicluster is then a group of traders and measurements of their trading activity that are similar. With respect to the biclustered variables, a trader is more similar to other traders in the same bicluster than traders outside of the bicluster.

Next, we introduce an important concept to the plaid model, namely that of an additive "layer". A layer is a canonical matrix matching the dimensions of the given data matrix, with zeros everywhere except the biclustered elements. In the plaid model, the data is decomposed into a series of additive layers that capture the underlying structure of the data. As a consequence, layers combine to provide a reconstruction that highlights the main features of the given data matrix.

The plaid model first includes a background layer that consists of *all* traders and variables to account for global effects in the data. In our application setting, the background layer accounts for market trends that affect trading behavior of all traders, such as for example, a major liquidity event. There are in principle many ways to construct the background layer. The simplest approach is to set each element of the background layer to be equal to the global average of the given data matrix. One could also estimate a parametric model that incorporates a priori information about the traders and variables. In our analysis, we set each column of the background layer to the corresponding variable's mean. This is equivalent to standardizing the data, which is necessary since a variable like volume is strictly non-negative, while others like net position can be negative. Subsequent layers represent additional effects corresponding to specific traders and variables that exhibit a strong pattern not explained by the background layer.

Formally the data matrix  $X \in \mathbb{R}^{n \times p}$  can be represented as

$$X_{ij} = \mu_0 + \sum_{k=1}^{K} \theta_{ijk} r_{ik} c_{jk},$$
 (1)

where i = 1, ..., n indexes samples and j = 1, ..., pindexes variables,  $\mu_0$  captures the background layer and  $\theta_{ijk}$  describes the bicluster effect; k is a layer index running to the number of biclusters K. The parameters  $r_{ik}$  and  $c_{jk}$  are indicator variables that combine to identify the bicluster, that is, they denote bicluster membership for, respectively, the traders and variables.

There are several modeling choices for the form of  $\theta_{ijk}$ , the most common being

$$\theta_{ijk} = \mu_k + \alpha_{ik} + \beta_{jk},\tag{2}$$

where i = 1, ..., n, j = 1, ..., p and k = 1, ..., K. Each bicluster has a mean, trader, and variable effect. Hence, each bicluster is expressed as a two-way analysis of variance (ANOVA) model. In other words, each trader in a bicluster can be interpreted as following similar strategies ( $\mu_k$ ). Yet, traders in the bicluster may differ slightly due to differences in preference, amount of available capital, and so on. This traderspecific effect is captured by the effect { $\alpha_{ik}$ }. Similarly, biclustered measures of trading activity can differ by trading strategies { $\beta_{jk}$ }.

The biclusters are discovered in a sequential fashion. Suppose K-1 layers have been estimated in addition to the background layer. The residual data matrix is given by

$$\hat{Z}_{ij} = X_{ij} - \hat{\mu}_0 - \sum_{k=1}^{K-1} \hat{\theta}_{ijk} \hat{r}_{ik} \hat{c}_{jk}.$$
(3)

The *K*th bicluster is found by minimizing the usual residual sum of squares over all parameters of interest

$$\min_{\{\theta_{ijK}, r_{iK}, c_{jK}\}} \sum_{i=1}^{n} \sum_{j=1}^{p} (\hat{Z}_{ij} - \theta_{ijK} r_{iK} c_{jK})^2.$$
(4)

Estimates of the bicluster memberships  $(\hat{r}_{iK}, \hat{c}_{jK})$  are obtained with a numerical search. A simple search procedure based on k-means clustering (see Hastie et al., 2001) is presented throughout this paper. More complex strategies and comprehensive discussion can be found in Lazzeroni and Owen (2000) and Turner et al. (2005). When given bicluster memberships, estimates of the bicluster-specific effects ( $\hat{\theta}_{ijK}$ ) are easy to compute, as one can use the usual two-way ANOVA estimators (Turner et al., 2005).

The plaid model estimates the behavior of each trader and then seeks groups of traders that have similar behavior over the biclustered variables. The estimation procedure is an iterative one based on minimizing sum of squares of the data minus estimated layers (pseudocode is given in Algorithm 1). First the background layer is estimated, then biclusterspecific layers are added one at a time. The statistical significance of layers are determined by a permutation test. The algorithm terminates when a significant layer cannot be found.

Next, we provide a brief review of the permutation test discussed starting on page 8 of Lazzeroni and Owen (2000). A comprehensive review can also be found in Turner et al. (2005). The permutation test is intuitively similar to bootstrapping, and relies on resampling of the data to approximate significance of the bicluster. The basic idea is that the data values are independent of biclusters after permuting the rows and columns. Thus, comparing the candidate bicluster against (noise) biclusters obtained after randomizing the data matrix allows one to accept a bicluster only if it is significantly larger than what one would find in noise.

The importance of each bicluster is measured with

$$\sigma_k^2 = \sum_{i=1}^n \sum_{j=1}^p \hat{r}_{ik} \hat{c}_{jk} \hat{\theta}_{ijk}^2,$$
(5)

where k = 1, ..., K. Let  $\pi_r$  be the permutation of the index set  $\{1, ..., n\}$ , and  $\pi_c$  be the permutation of the index set  $\{1, ..., p\}$ . Then  $\tilde{Z}_l = \hat{Z}(\pi_r, \pi_c)$  is the matrix after permuting every row of the residual data matrix  $\hat{Z}$  and then permuting every column of the result. The importance of a bicluster obtained from  $\tilde{Z}_l$ is measured with

$$\hat{\sigma}_{n_l}^2 = \sum_{i=1}^n \sum_{j=1}^p \tilde{r}_{il} \tilde{c}_{jl} \tilde{\theta}_{ijl}^2,$$
(6)

where  $\tilde{r}_{il}, \tilde{c}_{jl}$  are bicluster memberships estimated from  $\tilde{Z}_l$ . The candidate bicluster is rejected if any of the noise biclusters are more important. The selection of the total number of noise biclusters *L* is discussed further in Section 2.3. The permutation test is given in Algorithm 3.

Next, we will discuss an extension of plaid models to detect persistent patterns when given a sequence of data matrices.

| Algorithm 1 The plaid model estimation procedure for static data.   |
|---|
| <b>Input:</b> Matrix $X \in \mathbb{R}^{n \times p}$  |
| <b>Output:</b> Sequentially discovered biclusters $\{\hat{r}_{ik}, \hat{c}_{jk}, \hat{\theta}_{ijk}\}_{k=1}^{K}$          |
| 1: $\hat{\mu}_0 = \frac{1}{n} 1_{n \times n} X$   |
| 2: $\hat{Z} = X - \hat{\mu}_0$  |
| 3: $K = 1$ (bicluster counter)  |
| 4: repeat   |
| 5: $\{\hat{r}_{iK}, \hat{c}_{jK}, \hat{\theta}_{ijK}\} = estimateBicluster(\hat{Z}) \text{ (see Algorithm 2)}$            |
| 6: $b = permuteTest(\{\hat{r}_{iK}, \hat{c}_{jK}, \hat{\theta}_{ijK}\})$ (see Algorithm 3)                                |
| 7: <b>if</b> $b = 0$ <b>then</b>  |
| 8: $\hat{Z}_{ij} = X_{ij} - \hat{\mu}_0 - \sum_{k=1}^{K} \hat{\theta}_{ijk} \hat{r}_{ik} \hat{c}_{jk}$                    |
| 9: $K = K + 1$  |
| 10: <b>end if</b>   |
| 11: <b>until</b> $b = 1$  |
| 12: <b>return</b> $\{\hat{r}_{ik}, \hat{c}_{jk}, \hat{\theta}_{ijk}\}, i = 1, \dots, n, j = 1, \dots, p, k = 1, \dots, K$ |
|   |

#### Algorithm 2 estimateBicluster

**Input:** Matrix  $\hat{Z} \in \mathbb{R}^{n \times p}$ 

**Output:** Bicluster  $\{\hat{r}_{iK}, \hat{c}_{jK}, \hat{\theta}_{ijK}\}$ 1: Apply k-means (k=2) to rows of  $\hat{Z}$ . Set  $\hat{r}_{iK}$  to the smaller cluster. 2: Apply k-means (k=2) to columns of  $\hat{Z}$ . Set  $\hat{c}_{iK}$  to the smaller cluster. repeat 3:  $\hat{\theta}_{ijK} = \operatorname{argmin}_{\{\theta_{ijK}\}} \sum_{i=1}^{n} \sum_{j=1}^{p} (\hat{Z}_{ij} - \theta_{ijK} \hat{r}_{iK} \hat{c}_{jK})^2$ 4: for i=1,...,n do 5: if  $\sum_{j=1}^{p} (\hat{Z}_{ij} - \hat{\theta}_{ijK} \hat{c}_{jK})^2 < \sum_{j=1}^{p} \hat{Z}_{ij}^2$  then 6:  $\hat{r}_{iK} = 1$ 7: 8: else 9:  $\hat{r}_{iK} = 0$ end if 10: end for 11: for j=1,...,p do 12: if  $\sum_{i=1}^{n} (\hat{Z}_{ij} - \hat{\theta}_{ijK} \hat{r}_{iK})^2 < \sum_{i=1}^{n} \hat{Z}_{ij}^2$  then  $\hat{c}_{jK} = 1$ 13:  $14 \cdot$ else 15:  $\hat{c}_{jK} = 0$ 16: end if 17: end for 18: **until**  $\hat{r}_{iK}, \hat{c}_{jK}$  converge or maximum iteration number attained 19: **return**  $\{\hat{r}_{iK}, \hat{c}_{jK}, \hat{\theta}_{ijK}\}, i = 1, ..., n, j = 1, ..., p$ 20:

#### Algorithm 3 permuteTest

**Input:** Bicluster  $\{\hat{r}_{iK}, \hat{c}_{iK}, \hat{\theta}_{ijK}\}$ **Output:** {0, 1} 1:  $\sigma_K^2 = \sum_{i=1}^n \sum_{j=1}^p \hat{r}_{iK} \hat{c}_{jK} \hat{\theta}_{ijK}^2$ 2:  $\hat{Z}_{ij} = X_{ij} - \hat{\mu}_0 - \sum_{k=1}^K \hat{\theta}_{ijk} \hat{r}_{ik} \hat{c}_{jk}$ 3: for l=1,...,L do  $\pi_r = permutation(\{1, \ldots, n\})$ 4:  $\pi_{c} = permutation(\{1, \dots, p\})$   $\tilde{Z}_{l} = \hat{Z}(\pi_{r}, \pi_{c})$   $\tilde{r}_{il}, \tilde{c}_{jl}, \tilde{\theta}_{ijl} = estimateBicluster(\tilde{Z}_{l})$ 5: 6: 7: 8:  $\hat{\sigma}_{n_l}^2 = \sum_{i=1}^n \sum_{j=1}^p \tilde{r}_{il} \tilde{c}_{jl} \tilde{\theta}_{ijl}^2$ 9: end for 10: if  $\hat{\sigma}_{K}^{2} > \max\{\hat{\sigma}_{n_{1}}^{2}, ..., \hat{\sigma}_{n_{L}}^{2}\}$  then return 0 11: 12: else return 1 13: 14: end if

## 2.2. Smooth plaid models for multidimensional time-series

Suppose we observe a time series of matrices with the rows consisting of individual traders and the columns representing various measures of their trading activity at different points in time. Formally, we have  $\{X_{ij}^{(t)}\}_{t=1}^{T}$ , where t is a time index and  $i = 1, \ldots, n, j = 1, \ldots, p$ . Since each row corresponds to a trader, we have a total of n traders that transact at least once in the data. For each trader, at each point in time, we observe the same p variables that measure different aspects of trader behavior. Then we can represent the data matrix  $X^{(t)}$  at time t as

$$X_{ij}^{(t)} = \mu_0^{(t)} + \sum_{k=1}^{K} \theta_{ijk}^{(t)} r_{ik}^{(t)} c_{jk}^{(t)},$$
(7)

where t = 1, ..., T. The expression and indicator parameters  $(\{\hat{r}_{ik}^{(t)}, \hat{c}_{jk}^{(t)}, \hat{\theta}_{ijk}^{(t)}\})$  reflect whether a bicluster is active in a given time period t. The total number of biclusters K is fixed for all time periods, since additionally allowing K to vary with time creates identifiability and implementation challenges.

Visually, the data can be organized as a three dimensional array, shown in Figure 1, with traders arranged in the rows, trading features in the columns, and time as the 3rd 'depth' dimension.

Note that a direct analysis may proceed by collapsing the temporal dimension and working with a



Fig. 1. For our study, we construct a series of matrices, one for each period of time, consisting of approximately six thousand rows (one for each trader) and several columns of trading measures. Once the data is organized, we are working with a three dimensional array with traders arranged in the rows, trading measures in the columns, and time as the 3rd dimension.

single two dimensional matrix with traders arranged in the rows and variables that have been combined over time in the columns. However, collapsing the time dimension would aggregate away a significant amount of valuable time-series information present in the data. To illustrate, Figure 2 shows the general structure that remains after integrating over time. As shown in the stylized plot<sup>4</sup> of net position vs. volume/number of trades, a distinct group of high frequency traders emerges holding a very small open position at the end of the trading day, together with fundamental traders holding large positive or negative positions. However, there is a very large number of significant traders that are not allocated to an interpretable group.

A potential remedy for this is to analyze each data set at each point in time separately. However, this would ignore the potentially important time component of trading strategies, while becoming dominated by transient patterns. Moreover, if we repeatedly and directly apply a method like the plaid model, the estimated groupings can change for each data set when a temporally stable structure is more appropriate.

In this study, we design and employ a dynamic method that is in between these two direct approaches. A penalized optimization framework accounts for auto-correlation by effectively averaging the groups over a rolling window of time. Such an approach helps mitigate the effects of transient patterns, while enhancing structural regularities in the data.



Fig. 2. The general structure that remains after collapsing the temporal dimension. High frequency traders, buyers, and sellers are prominent and easily detectable. However, there are a large number of residual traders that cannot be easily interpreted.

Formally, we consider the search for the Kth layer over the interval  $t = T - W, \ldots, T$ . The search is initialized with starting values for  $\{\hat{r}_{iK}^{(t)}\}_{t=T-W}^{T}$  and  $\{\hat{c}_{jK}^{(t)}\}_{t=T-W}^{T}$ , which denote whether the candidate bicluster was detected in the previous W time periods. The objective function is given next.

$$\min_{\{\theta^{(T-W)},\dots,\theta^{(T)}\}} \sum_{t=T-W}^{T} \sum_{i=1}^{n} \sum_{j=1}^{p} (\hat{Z}_{ij}^{(t)} - \theta_{ijK}^{(t)} \hat{r}_{iK}^{(t)} \hat{c}_{jK}^{(t)})^{2} \\
+ \lambda \sum_{t=T-W+1}^{T} \sum_{i=1}^{n} \sum_{j=1}^{p} (\theta_{ijK}^{(t)} \hat{r}_{iK}^{(t)} \hat{c}_{jK}^{(t)} \\
- \theta_{ijK}^{(t-1)} \hat{r}_{iK}^{(t-1)} \hat{c}_{jK}^{(t-1)})^{2},$$
(8)

<sup>&</sup>lt;sup>4</sup>Plots of aggregate data are used throughout this work to protect confidentiality.

where  $\lambda$  is a tuning parameter and W is a parameter determining the number of previous time periods to consider. Following the notation in Equation 3,  $\hat{Z}^{(t)}$  is the residual data matrix at time t

$$\hat{Z}_{ij}^{(t)} = X_{ij}^{(t)} - \hat{\mu}_0^{(t)} - \sum_{k=1}^{K-1} \hat{\theta}_{ijk}^{(t)} \hat{r}_{ik}^{(t)} \hat{c}_{jk}^{(t)}, \qquad (9)$$

where i = 1, ..., n, j = 1, ..., p and t = 1, ..., T.

For given  $\lambda$  and W, we solve the optimization problem through coordinate descent, which has been developed and implemented for such objective functions by Friedman et al. (2010). This optimization approach achieves large improvements in computational efficiency over other minimization approaches, and allows our framework to be feasible for large-size data problems.

The basic steps of the algorithm are given below in Algorithm 4 and illustrated on a toy example in Figure 3. The main idea behind the algorithm is to:

- 1. Use results from previous time steps to form candidate biclusters for the current time period, then apply the penalization framework and permutation test to discover significant and stable biclusters (see Algorithm 5).
- 2. After candidate biclusters from previous times have been exhausted, a final search is performed with the penalization framework and permutation test for new biclusters that were not captured in the previous results (see Algorithm 6).
- 3. Return all significant biclusters discovered in previous steps 1 and 2.

We note that when t = T,  $\{\theta_{ijk}^{(t)}\}_{t=T-W}^{T}$  are simultaneously estimated. After that, t = T + 1, and  $\{\theta_{ijk}^{(t)}\}_{t=T-W+1}^{T+1}$  are estimated independently. As easily seen, these two sets are highly overlapping. We use the most recent estimate as the final estimator due mainly to its simplicity. This strategy is similar to using a rolling window smoother. In principle, other methods that combine the overlapping estimates could be employed. Though in practice, more complex strategies can sometimes complicate implementation without fundamentally changing the final estimator.

#### 2.3. Implementation issues

Our implementation is performed in R (version 2.15), with all auxiliary functions supported in the basic distribution (R Core Team, 2012). R code is available at www.stat.lsa.umich.edu/~smankad/. Numerical results presented in the following sections are obtained using the code specification above on a Linux platform. Next, we discuss the permutation test used in the stopping criterion for the smooth plaid algorithm, and selection of the parameters  $\lambda$  and W.

**Permutation Test.** The permutation test utilized in Algorithm 3 is modified to accommodate the additional structure between data matrices. Specifically, matrix observations at different times should be permuted separately, so that global time effects are maintained. Also, the importance of bicluster k is measured over the time interval, instead of at a single time:  $\sigma_k^2 = \sum_{t=T-W}^T \sum_{i=1}^n \sum_{j=1}^p \hat{r}_{ik}^{(t)} \hat{c}_{jk}^{(t)} \hat{\theta}_{ijk}^{(t)2}$ . Algorithm 8 shows the smooth plaid permutation test.

#### Algorithm 4 Smooth plaid models estimation procedure.

 $\begin{aligned} & \text{Input: Matrices } \{X^{(t)} \in \mathbb{R}^{n \times p}\}_{t=T-W}^{T}, M \text{ biclusters from previous times } \{\hat{r}_{im}^{(t)}, \hat{c}_{jm}^{(t)}\}_{m=1}^{M} \\ & \text{Output: Biclusters } \{\hat{r}_{ik}^{(t)}, \hat{c}_{jk}^{(t)}, \hat{\theta}_{ijk}^{(t)}\}_{k=1, t=T-W}^{K,T} \\ & \text{ i: } \hat{\mu}_{0}^{(t)} = \frac{1}{n} \mathbf{1}_{n \times n} X^{(t)}, t = T - W, \dots, T \\ & \text{ : } \hat{Z}^{(t)} = X^{(t)} - \hat{\mu}_{0}^{(t)}, t = T - W, \dots, T \\ & \text{ : } \hat{Z}^{(t)} = X^{(t)} - \hat{\mu}_{0}^{(t)}, t = T - W, \dots, T \\ & \text{ : } \{\hat{r}_{ik}^{(t)}, \hat{c}_{jk}^{(t)}, \hat{\theta}_{ijk}^{(t)}\}_{k=1, t=T-W}^{K',T} = searchPrevBCResults(\{\hat{Z}^{(t)}\}_{t=T-W}^{T}, \{\hat{r}_{im}^{(t)}, \hat{c}_{jm}^{(t)}\}_{m=1}^{M}) \\ & \text{ (see Algorithm 5)} \\ & \text{ : } \hat{Z}_{ij}^{(t)} = \hat{Z}_{ij}^{(t)} - \sum_{k=1}^{K'} \hat{r}_{ik}^{(t)} \hat{c}_{jk}^{(t)} \hat{\theta}_{ijk}^{(t)} \text{ for } t = T - W, \dots, T, i = 1, \dots, n, j = 1, \dots, p. \\ & \text{ : } \{\hat{r}_{ik}^{(t)}, \hat{c}_{jk}^{(t)}, \hat{\theta}_{ijk}^{(t)}\}_{k=K'+1, t=T-W}^{K,T} = searchNewBC(\{\hat{Z}^{(t)}\}_{t=T-W}^{T}) \text{ (see Algorithm 6)} \\ & \text{ : return } \{\hat{r}_{ik}^{(t)}, \hat{c}_{jk}^{(t)}, \hat{\theta}_{ijk}^{(t)}\}_{k=1, t=T-W}^{K,T} \\ \end{aligned}$ 



Fig. 3. The smooth plaid algorithm on a toy example. The first row contains raw data and the second row shows the fitted values. We are interested in estimating t = 3. We combine the imperfectly detected bicluster from time steps t = 1 and t = 2 to form the initial condition or candidate bicluster for t = 3. The double arrows denote the penalty.

#### Algorithm 5 searchPrevBCResults

It is argued in Lazzeroni and Owen (2000) that, since after permuting rows and columns the data values are independent of row and column labels, the approximate probability of accepting k or more false biclusters is  $(L + 1)^{-k}$ , where L is the total number of noise biclusters. The authors suggest four or fewer noise biclusters for each permutation test. Though, this is highly dependent on the size of the data and available computing power (costs are proportional to the number of noise biclusters). With the large sized
Algorithm 6 searchNewBCInput: Matrices  $\{Z^{(t)} \in \mathbb{R}^{n \times p}\}_{t=T-W}^{T}$ Output: Biclusters  $\{\hat{r}_{ik}^{(t)}, \hat{c}_{jk}^{(t)}, \hat{\theta}_{ijk}^{(t)}\}_{k=1,t=T-W}^{K,T}$ 1: K = 1 (bicluster counter)2: repeat3:  $\{\hat{r}_{iK}^{(t)}, \hat{c}_{jK}^{(t)}, \hat{\theta}_{ijK}^{(t)}\}_{t=T-W}^{T} = estimateSmoothBicluster(\{\hat{Z}^{(t)}\}_{t=T-W}^{T})$  (see Algorithm 7)4:  $b = permuteTest2(\{\hat{r}_{iK}^{(t)}, \hat{c}_{jK}^{(t)}, \hat{\theta}_{ijK}^{(t)}\}_{t=T-W}^{T}, \{\hat{Z}^{(t)}\}_{t=T-W}^{T})$  (see Algorithm 8)5: if b = 0 then6:  $\hat{Z}_{ij}^{(t)} = X_{ij}^{(t)} - \hat{\theta}_{ijk}^{(t)} \hat{r}_{ik}^{(t)} \hat{c}_{jk}^{(t)}, t = T - W, \dots, T$ 7: Set K = K + 18: end if9: until b = 110: return  $\{\hat{r}_{ik}^{(t)}, \hat{c}_{jk}^{(t)}, \hat{\theta}_{ijk}^{(t)}\}_{k=1,t=T-W}^{K,T}$ 

#### Algorithm 7 estimateSmoothBicluster

**Input:** Matrices  $\{\hat{Z}^{(t)} \in \mathbb{R}^{n \times p}\}_{t=T-W}^{T}$ **Output:** Biclusters  $\{\hat{r}_{iK}^{(t)}, \hat{c}_{jK}^{(t)}, \hat{\theta}_{ijK}^{(t)}\}_{t=T-W}^{T}$ 1: Apply k-means (k=2) to rows of  $\hat{Z}^{(T)}$ . Set  $\{\hat{r}_{iK}^{(t)}\}_{t=T-W}^{T}$  to the smaller cluster. 2: Apply k-means (k=2) to columns of  $\hat{Z}^{(T)}$ . Set  $\{\hat{c}_{jK}^{(t)}\}_{t=T-W}^{T}$  to the smaller cluster. 3: repeat  $\{\theta_{ijK}^{(t)}\}_{t=T-W}^{T} = \operatorname{argmin} \text{ of Equation 8}$  for t=T-W,...,T do 4: 5: for i=1,...,n do 6: if  $\sum_{t=T-W}^{T} \sum_{j=1}^{p} (\hat{Z}_{ij}^{(t)} - \hat{\theta}_{ijK}^{(t)} \hat{c}_{jK}^{(t)})^2 < \sum_{t=T-W}^{T} \sum_{j=1}^{p} \hat{Z}_{ij}^{(t)2}$  then 7: 8. else 9:  $\hat{r}_{iK}^{(t)} = 0$ 10: end if 11: end for 12: **for** j=1,...,p **do** 13: if  $\sum_{t=T-W}^{T} \sum_{i=1}^{n} (\hat{Z}_{ij}^{(t)} - \hat{\theta}_{ijk}^{(t)} \hat{r}_{ik}^{(t)})^2 < \sum_{t=T-W}^{T} \sum_{i=1}^{n} \hat{Z}_{ij}^{(t)2}$  then 14:  $\hat{c}_{jK}^{(t)} = 1$ else 15: 16:  $\hat{c}_{iK}^{(t)} = 0$ 17: end if 18: end for 19: end for 20: 21: **until**  $\{\hat{r}_{iK}^{(t)}, \hat{c}_{jK}^{(t)}\}$  converge or maximum iteration number attained 22: return  $\{\hat{r}_{iK}^{(t)}, \hat{c}_{jK}^{(t)}, \hat{\theta}_{ijK}^{(t)}\}_{t=T-W}^{T}$ 

data encountered in our application, we set L=3. This parameter can be adjusted by the user to balance accuracy and computational ease.

**Choosing**  $\lambda$ . The effect of  $\lambda$  is to create smoother paths over time for each bicluster. Specifically, larger

penalization levels force biclusters to have similar estimated values as in neighboring time steps.

A systematic way to choose  $\lambda$  is through crossvalidation. The idea behind cross validation is to use a random subset of the data to fit the model, and the rest

#### Algorithm 8 permuteTest2

**Input:** Biclusters  $\{\hat{r}_{iK}^{(t)}, \hat{c}_{jK}^{(t)}, \hat{\theta}_{ijK}^{(t)}\}_{t=T-W}^T$ , Matrices  $\{Z^{(t)} \in \mathbb{R}^{n \times p}\}_{t=T-W}^T$ **Output:** {0, 1} 1:  $\sigma_K^2 = \sum_{t=T-W}^T \sum_{i=1}^n \sum_{j=1}^p \hat{r}_{iK}^{(t)} \hat{c}_{jK}^{(t)} \hat{\theta}_{ijK}^{(t)2}$ 2: **for** l=1,...,L (number of noise layers) **do** for t=1,...,T do  $\pi_r^{(t)} = permutation(\{1, \dots, n\})$ 4:  $\begin{aligned} \pi_{c}^{(t)} &= permutation(\{1, \dots, p\}) \\ \tilde{Z}_{l}^{(t)} &= \hat{Z}^{(t)}(\pi_{r}^{(t)}, \pi_{c}^{(t)}) \end{aligned}$ 5: 
$$\begin{split} \boldsymbol{\omega}_{l} &= \boldsymbol{\omega}^{(v)}(\boldsymbol{\pi}_{r}^{(v)}, \boldsymbol{\pi}_{c}^{(v)}) \\ \text{end for} \\ \tilde{r}_{il}^{(t)}, \tilde{c}_{jl}^{(t)}, \tilde{\theta}_{ijl}^{(t)} &= estimateSmoothBicluster(\{\tilde{Z}_{l}^{(t)}\}) \text{ (see Algorithm 7)} \\ \hat{\sigma}_{n_{l}} &= \sum_{t=T-W}^{T} \sum_{i=1}^{n} \sum_{j=1}^{p} \tilde{r}_{il}^{(t)} \tilde{c}_{jl}^{(t)} \tilde{\theta}_{ijl}^{(t)2}. \\ \text{end for} \\ \boldsymbol{v} \in \boldsymbol{\mathcal{O}} \end{split}$$
6: 7: 8: 9: 10: 11: if  $\hat{\sigma}_{K}^{2} > \max\{\hat{\sigma}_{n_{1}}^{2}, ..., \hat{\sigma}_{n_{L}}^{2}\}$  then return 0 12: 13: else return 1 14: 15: end if

of the data to assess model accuracy. Different values of  $\lambda$  are cycled over and the one that corresponds to the lowest test error is chosen.

In particular, suppose one is given a sequence of potential  $\lambda$ s. Cross-validation divides the samples into G groups. Then for each potential  $\lambda$ , Equation 8 is minimized G times, once with each of the groups omitted. The coefficients from each estimation are used to predict the omitted group. The error is accumulated, and average error and standard deviation over the G groups is computed. Finally, the  $\lambda$  corresponding to the lowest mean squared error is chosen. A full algorithmic description, including selecting the initial  $\lambda$  sequence, with code can be found in Friedman et al. (2010).

Under this framework, the 'optimal' value of  $\lambda$ (denoted by  $\lambda^*$ ) can change each time we minimize the penalized objective function. Thus, for the same bicluster,  $\lambda^*$  over time periods  $T - W, \ldots, T$  may be different than the  $\lambda^*$  chosen for  $T - W + 1, \ldots, T + 1$ . Further,  $\lambda^*$  can change for different biclusters over the same time window. This flexibility is ideal given that different trader groups may follow different dynamics, and those dynamics may be time-varying.

**Choosing** W. The parameter, W, controls the window width for smoothing, e.g., the number of previous time steps to include in the smoothing. Larger values of W mean that the model has more memory so it incorporates more observations for estimation. This risks missing sharper changes in the data and

only detecting the most persistent patterns. On the other hand, small values of W make the fitting more sensitive to sharp changes, but increase variance due to smaller number of observations. We find setting W = 1 (penalizing over adjacent matrices) is sufficient for filtering out most noisy expressions. Other values could be used if external information is known, or if additional smoothing is needed.

#### 3. An illustrative example

Before applying our model to the E-mini S&P 500 Futures Contract, we illustrate and validate our methodology with simulated data.

We generate a time-series of data matrices, where each matrix has 100 rows and columns, with embedded biclusters that evolve through time. In particular, we have  $X^{(t)} \in \mathbb{R}^{100 \times 100}$ , where  $X_{ij}^{(t)} \sim N(\mu_k^{(t)}, 1)$ . The background layer has mean  $\mu_0 = 0$ , and there are two biclusters with means shown in Figure 4. One bicluster has constant mean, while the other oscillates with time. The size of each bicluster is 16% of the size of X, and is constant throughout time.

We use only a mean effect for the bicluster effect for simplicity, that is,  $\theta_{ijk} = \mu_k$ . Figure 5 shows the estimated values for different levels of penalization. We see that if  $\lambda$  is too large, only the most persistent pattern is detected. On the other hand, the  $\lambda$  from cross-validation yields a nearly complete



Fig. 4. True and estimated bicluster means in simulated data. One bicluster has constant mean, while the other oscillates. The estimates are obtained with  $\lambda$  chosen through cross validation.



Fig. 5. The left panel shows an example of the raw data, unfolded (repeatedly concatenated) to be a 2-D array  $[X^{(1)}|X^{(2)}|...|X^{(T)}]$ . The central panel shows the estimated values with  $\lambda = \infty$ . The right panel shows the estimated values when choosing  $\lambda$  via cross validation.

perfect detection of the biclusters, while smoothing the expression patterns over time.

We compare the proposed smooth plaid model with the more direct approach of applying the static plaid model to each time step separately. The percentage of biclusters detected and false positive results are presented in Table 1. We see that the smooth plaid procedures perform favorably, since it does a significantly better job at detecting the dynamic bicluster, while maintaining a negligible number of false positives and the detection of the stable bicluster.

Panel B of Table 1 shows that the recovered bicluster contributions to explained variance are relatively small. This highlights that the proposed approach can be advantageous when finding 'needles in a haystack', and is closely related to anomaly detection.

 $\begin{array}{l} \mbox{Table 1} \\ \mbox{Simulation Results from the Illustrative Example. \% Variance} \\ \mbox{Explains is defined as } 1 - \sum_t ||\hat{X}^{(t)} - X^{(t)}||_F^2 / \sum_t ||X^{(t)}||_F^2. \end{array}$ 

| Panel A: Detection Accuracy                          |                                   |                                    |                     |  |  |
|--|-----------------------------------|------------------------------------|---------------------|--|--|
| Algorithm  | % Stable<br>Bicluster<br>Detected | % Dynamic<br>Bicluster<br>Detected | % False<br>Positive |  |  |
| Smooth Plaid   | 93.7                              | 87.9                               | 1.6                 |  |  |
| Static Plaid   | 89.1                              | 40.6                               | 1.7                 |  |  |
| Panel B: Estimated Smooth Plaid Bicluster Statistics |                                   |                                    |                     |  |  |
| Bicluster  | Number                            | Number                             | % Variance          |  |  |
| Bielustei  | Rows                              | Columns                            | Explained           |  |  |
| Bicluster 1  | 16                                | 16                                 | 16.6                |  |  |
| Bicluster 2  | 16                                | 16                                 | 2.9                 |  |  |
| Overall  |                                   |                                    | 19.5                |  |  |

In summary, the smooth plaid procedures perform favorably in this synthetic setting by discovering the true, underlying biclustering structure and evolution.

# 4. Applying smooth plaid models to transaction-level data

We now apply the smooth plaid model to transaction-level data generated by an agent-based simulation model calibrated to the E-mini S&P 500 futures contract. The E-mini trades electronically on the CME Globex trading platform, a fully electronic limit order market. Limit orders are submitted by traders wishing to buy and sell a certain number of contracts up to a certain price (or at the market price for market orders). Submitted orders are matched by a matching algorithm. The number of outstanding E-mini contracts is created directly by buying and selling interests. There is no limit on how many contracts can be outstanding at any given time. The CME Globex matching algorithm for the E-mini offers strict price and time priority. Specifically, orders to buy at higher prices or sell at lower prices are placed in queues ahead orders to buy at lower prices or sell at higher prices. Orders that offer to buy or sell at the same price arranged in the order that they have arrived into the Globex matching engine.

We use simulated, transaction-level data generated by an agent-based model (Paddrik et al., 2011; Hayes et al., 2012) that retains key attributes of actual E-mini 500 stock index futures data observed by the exchanges and regulators. The model of Paddrik et al. (2011) and Hayes et al. (2012) simulates traders who submit orders to a limit order book according to different combinations of trading styles. Orders are matched according to strict price and time priority like with the CME Globex matching algorithm. The six different trading styles are based on the findings of Kirilenko et al. (2010) and consistent of fundamental buyers and sellers, market makers, opportunistic and high frequency traders. The different trader types are calibrated to match key aspects of behavior found in real E-mini 500 stock index futures market data, such as trading speed, market volume share, and position limits. For each simulated transaction, we know the buyer and the seller, the price and quantity at which they traded, and the time of execution.

We use the following variables to cluster traders into groups: trades per second, trading volume (total number of contracts traded), cumulative inventory/net position (reset to zero for each trading account at the end of each trading day), change in inventory, and median duration for each trader. Each of the variables is calculated for a preset time period. We define the time period as 600 transactions (trades). Our results are robust with respect to different sampling schemes. Though, if too small of a time period is used, then most traders will not have participated in any transactions and the data matrices will contain many zeros. Such sparsity can mask the slower groups of traders. A brief description of each variable follows.

Trades per second are computed by dividing the total number of transactions that a trader makes in a given time period by the total number of seconds in that time period. Trades per second is indicative of the decision horizons and execution strategies for different traders.

Trading volume is computed for each trader by summing up the total number of contracts transacted in each time period. Trading volume is indicative of the overall trading activity of a particular trader.

Change in inventory is computed by subtracting the number of contracts sold during a particular period from the number of contracts bought during the period. Change in inventory is indicative of a risk exposure of a particular trader accumulated during a period of time.

Cumulative net inventory is calculated by accumulating a trader's inventory from the beginning of the day to the end of the current time period. Cumulative net inventory indicates the direction (long or short) and size of the risk exposure of a trader accumulated from the beginning of the day.

Lastly, intertrade duration is defined as the time (in seconds) until the next trade. Specifically, for each transaction involving a given trader, we compute the time, in seconds, until the next transaction between any two traders. We then compute the median intertrade duration for a trader during a sample period.

Each trading variable measures different aspects of how much, in which direction, and how quickly each trader transacts. Once organized, the data contains 6387 rows (traders), 5 columns, and 792 time periods. Each day contains between approximately 30 to 50 time periods, depending on the number of transactions per day.

After we apply our algorithm to the data, we use additional filtering on the fitted values to separate traders into five broad groups. The additional grouping consists of variable thresholds that separate traders into groups and is based on characteristics that measure a trader's strategic profile, such as, among others, the rate of a trader's mean reversion of fitted accumulated inventory. Thus, we employ the smooth plaid method as a temporal filter to facilitate trader classifications. Our method improves on the direct approaches by cleaning the temporal noise, allowing the additional classification of thousands of traders into interpretable categories.

The main benefit of employing the smooth plaid method is in separating market makers, opportunistic, and small traders. The more direct approaches struggle with these traders since they have strategies that can appear very similar statistically when the time dimension is aggregated.

Altogether, we find 7 high frequency traders (HFT's), 73 market makers, 2405 opportunistic traders, 1281 fundamental position traders, and 2849 small/residual traders.

HFTs occupy a distinctive niche in the ecosystem of modern electronic markets. They trade through an enormous number of contracts each day, but carry very little inventory at any point in time. Market makers have a footprint qualitatively similar to HFTs, but significantly smaller volume-wise. Fundamental traders are primarily characterized by large positive or negative cumulative net positions at the end of a trading day. The use of temporal information is quite important in identifying this group, since some of its members accumulate directional positions by executing many small-size orders, while others execute a few larger-size orders, thus trying to disguise their behavior so as not to be taken advantage by the market. This feature is to a large extent lost to an analysis ignoring the temporal dimension. The same holds true for the group of small traders that trade infrequently at random points in time, hence lacking any persistent pattern. The final group consists of opportunistic traders that have a persistent presence in the market, but their trading behavior bifurcates between fundamental positioning and market making.

Figures 6, 7, and 8 illustrate that the five groups exhibit different trading signatures. Fundamental traders accumulate either a large positive or negative net imbalance. On the other hand, all other groups have on average zero net position. Opportunistic traders net positions vary more than market makers, which vary more than high frequency traders. These trader groups are conceptually similar to the bicluster in the illustrative example with oscillating mean structure, and as we saw, the smooth plaid model has superior performance for such dynamic behavior by conditioning on previous time points.



Fig. 6. Stylized representation of the net position (x-axis) versus volume/number of trades. After applying the smooth plaid procedure, we can additionally classify market makers, opportunists and small traders.

#### 5. Conclusion

In this study, we present a dynamic machinelearning method that designates traders in a liquid financial market into five persistent categories based on their footprint in the data. Our method is based on a plaid clustering technique enhanced by a smoothing framework that filters out transient patterns. The method performs extremely well on regulatory, transaction-level data for the E-mini S&P 500 stock index futures contract, the price discovery vehicle for the broad U.S. stock market. However, in order to preserve confidentiality of the regulatory data, the results we present employ simulated data generated by an agent-based simulation model of an electronic market calibrated to the E-mini.

For comparison, Table 2 shows that our classification of traders is consistent with the study by Kirilenko et al. (2010), which classified trader behavior using similar E-mini futures data three days before and during the Flash Crash of May 6, 2010. While investigating the triggering event of the Flash Crash, Kirilenko et al. (2010) manually designated trading accounts that traded in the E-mini on May 6, 2010 into the same six distinct categories. The categorization was essentially based on the dynamics of two characteristics: end of day holdings and intraday trading volume for each trading account.

The similarity in our groupings validates and demonstrates the usefulness of our method, since these previous reports manually classified each trader through an exhaustive and labor intensive procedure. Our biclustering algorithm was able to detect similar groups and



Fig. 7. Volume-based market share for each group of trader over each day of data. The high frequency traders persistently trade a significantly large number of contracts.



Fig. 8. This heatmap portrays average net position for each group of trader over each day of the month. The different trader types have different signatures in the data.

Table 2 Grouping traders in the E-mini S&P 500 futures contract. Our analysis is performed on simulated E-mini data, while Kirilenko et al. (2010) analyze regulatory data for the May 2010 E-mini S&P 500 futures contract.

| Trader Type                | Smooth Plaid | Kirilenko et al. (2010) |
|----------------------------|--------------|-------------------------|
| HFT                        | 7            | 16                      |
| Market maker               | 73           | 179                     |
| Opportunistic              | 2405         | 5808                    |
| Fundamental Buyers/Sellers | 1281         | 2539                    |
| Small                      | 2849         | 6880                    |

the relevant variables that consistently separate them over time using a novel machine-learning methodology.

We argue that the smooth plaid model can be effectively used for the analysis of traders and their strategies in electronic financial markets. In an environment where traders do not have formal designations, the smooth plaid model forms a useful first step to separate tens of thousands of trading accounts into manageable trader categories for subsequent academic, policy and regulatory analysis.

We also expect our method to be useful in other applications where one is given a time-series of matrices, such as examining traders across different markets or analyzing macroeconomic variables for different entities over time.

#### 6. Acknowledgments

Shawn Mankad and George Michailidis were supported by grant DMS 1228164.

The authors also thank an anonymous referee and the Editor Philip Maymin for their valuable comments and suggestions to improve the quality of the paper. The authors are also extremely grateful to Mark Paddrik for sharing the transaction-level data generated by the agent-based simulation model that he and his co-authors have developed.

#### References

- Chaboud, A., Hjalmarsson, E., Vega, C., Chiquoine, B., 2011. Rise of the machines: Algorithmic trading in the foreign exchange market. SSRN eLibrary http://ssrn.com/abstract=1501135 [Accessed July 11, 2013].
- Cheng, Y., Church, G.M., 2000. Biclustering of expression data. In: Proceedings of the Eighth International Conference on Intelligent Systems for

Molecular Biology, AAAI Press, Menlo Park, CA, pp. 93–103. ISBN 1-57735-115-0.

- Friedman, J., Hastie, T., Tibshirani, R., 2010. Regularization paths for generalized linear models via coordinate descent. J. Stat. Softw. 33 (1), 1–22. http://www.jstatsoft.org/v33/i01/ [Accessed July 11, 2013].
- Hastie, T., Tibshirani, R., Friedman, J.H., 2001. The elements of statistical learning: data mining, inference, and prediction: with 200 full-color illustrations. Springer-Verlag, New York.
- Hayes, R.L., Paddrik, M.E., Todd, A., Yang, S.Y., Beling, P., Scherer, W., 2012. Agent based model of the E-Mini S&P 500 future: Application for policy making. Proceedings of the 2012 Winter Simulation Conference. http://ssrn.com/abstract=2112139 [Accessed July 11, 2013].
- Hendershott, T., Jones, C.M., Menkveld, A.J., 2010.
  Does algorithmic trading improve liquidity?
  J. Finance. 66, 1–33. http://ssrn.com/abstract= 1100635 [Accessed July 11, 2013].
- Huang, Q.-H., 2011. Discovery of time-inconsecutive co-movement patterns of foreign currencies using an evolutionary biclustering method. App. Math. Comput. 218 (8), 4353–4364.
  ISSN 0096-3003. doi: 10.1016/j.amc.2011.10.011. http://www.sciencedirect.com/science/article/pii/ S0096300311012410.
- Kirilenko, A.A., Kyle, A.S., Samadi, M., Tuzun, T., 2010. The Flash Crash: The impact of high frequency trading on an electronic market. http:// ssrn.com/abstract=1686004 [Accessed July 11, 2013].
- Lazzeroni, L., Owen, A., 2000. Plaid models for gene expression data. *Stat. Sin.* 12, 61–86.
- Paddrik, M.E., Hayes, R.L., Todd, A., Yang, S.Y., Scherer, W., Beling, P., 2011. An agent based model of the E-Mini S&P 500 and the Flash Crash. SSRN eLibrary. http://ssrn.com/abstract=1932152 [Accessed July 11, 2013].
- R Core Team. 2012. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. http://www.Rproject.org/. ISBN 3-900051-07-0 [Accessed July 11, 2013].
- Turner, H., Bailey, T., Krzanowski, W., 2005. Improved biclustering of microarray data demonstrated through systematic performance tests. *Comput. Stat. Data Anal.* 48 (2), 235–254. ISSN 0167-9473. doi: 10.1016/j.csda.2004.02.003. http://www.sciencedirect.com/science/article/pii/ S0167947304000295.

# Gaussian Process-Based Algorithmic Trading Strategy Identification

Steve Y. Yang<sup>†</sup>, Qifeng Qiao<sup>†</sup>, Peter A. Beling<sup>†</sup>,

William T. Scherer<sup>†</sup>, and Andrei A. Kirilenko<sup>‡</sup> <sup>‡</sup>

<sup>†</sup>The University of Virginia, 151 Engineer's Way, P.O. Box 400747, Charlottesville, VA 22904

‡ MIT Sloan School of Management, 50 Memorial Drive,

Cambridge, Massachusetts 02142

Corresponding author. Email: yy6a@viginia.edu; and ak67@mit.edu

The research presented in this paper was co-authored by Andrei Kirilenko, a former full-time CFTC employee, Peter Beling, a former CFTC contractor who performed work under CFTC OCE contract (CFCE-11-CO-0233), William Scherer, a former CFTC contractor who performed work under CFTC OCE contract (CFCE-11-CO-0234), Steve Yang, a former CFTC contractor who performed work under CFTC DOE contract (CFEN-12-CO-0092), and Qifeng Qiao, a Ph.D. student at the University of Virginia who had no official affiliation with the CFTC. The Office of the Chief Economist and CFTC economists produce original research on a broad range of topics relevant to the CFTCs 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 Commission staff, or the Commission itself.

#### Abstract

Many market participants now employ algorithmic trading, commonly defined as the use of computer algorithms to automatically make certain trading decisions, submit orders, and manage those orders after submission. Identifying and understanding the impact of algorithmic trading on financial markets has become a critical issue for market operators and regulators. Advanced data feed and audit trail information from market operators now make the full observation of market participants' actions possible. A key question is the extent to which it is possible to understand and characterize the behavior of individual participants from observations of trading actions.

In this paper, we consider the basic problems of categorizing and recognizing traders (or, equivalently, trading algorithms) on the basis of observed limit orders. These problems are of interest to regulators engaged in strategy identification for the purposes of fraud detection and policy development. Methods have been suggested in the literature for describing trader behavior using classification rules defined over a feature space consisting of the summary trading statistics of volume and inventory and derived variables that reflect consistency of buying or selling behavior. Our principal contribution is to suggest an entirely different feature space that is constructed by inferring key parameters of a sequential optimization model that we take as a surrogate for the decision making process of the traders. In particular, we model trader behavior in terms of a Markov decision process (MDP). We infer the reward (or objective) function for this process from observation of trading actions using a process from machine learning known as inverse reinforcement learning (IRL). The reward functions learned through IRL then constitute a feature space that can be the basis for supervised learning (for classification or recognition of traders) or unsupervised learning (for categorization of traders). Making use of a real-world data set from the E-Mini futures contract, we compare two principal IRL variants, linear IRL (LIRL) and Gaussian Process IRL (GPIRL), against a method based on summary trading statistics. Results suggest that IRL-based feature spaces support accurate classification and meaningful clustering. Further, we argue that, because they attempt to learn traders' underlying value propositions under different market conditions, the IRL methods are more informative and robust than the summary statistic-based approach and are well suited for discovering new behavior patterns of market participants.

#### **Index Terms**

Inverse Reinforcement Learning; Gaussian Process; High Frequency Trading; Algorithmic Trading; Behavioral Finance; Markov Decision Process; Support Vector Machine

#### I. INTRODUCTION

Financial markets have changed dramatically over the past 10 years or so. These changes reflect the culmination of a decade-long trend from a market structure with primarily manual floor

trading to a market structure dominated by automated computer trading. This rapid transformation has been driven by the evolution of technologies for generating, routing, and executing orders, which have dramatically improved the speed, capacity, and sophistication of the trading functions that are available to market participants.

High-quality trading markets promote capital formation and allocation by establishing prices for securities and by enabling investors to enter and exit their positions in securities wherever and whenever they wish to do so. The one important feature of all types of algorithmic trading strategies is to discover the underlying persistent tradable phenomena and generate trading opportunities. These trading opportunities include microsecond price movements that allow a trader to benefit from market-making trades, several minute-long strategies that trade on momentum forecasted by market microstructure theories, and several hour-long market movements that surround recurring events and deviations from statistical relationship (Aldridge (2010)). Algorithmic traders then design their trading algorithms and systems with the aim of generating signals that result in consistent positive outcomes under different market conditions. Different strategies may target different frequencies, and the profitability of a trading strategy is often measured by a certain return metric. The most commonly used measure is the Sharpe ratio, a risk-adjusted return metric first proposed by Sharpe (Edwards (1966)).

In particular, there is a subgroup within the algorithmic trading strategies called High Frequency Trading (HFT) strategies that have attracted a lot of attention from investors, regulators, policy makers, and academics broadly. According to the U.S. Securities and Exchange Commission, high-frequency traders are "professional traders acting in a proprietary capacity that engage in strategies that generate a large number trades on daily basis." (The SEC Concept Release on Equity Market Structure, 75 Fed. Reg. 3603, January 21, 2010). The SEC characterized HFT as (1) the use of extraordinary high-speed and sophisticated computer programs for generating, routing, and executing orders; (2) use of co-location services and individual data feeds offered by exchanges and others to minimize network and other types of latencies; (3) very short timeframes for establishing and liquidating positions; (4) the submission of numerous orders that are canceled shortly after submission; and (5) ending the trading day in as close to a flat position as possible (that is, not carrying significant, unhedged positions over night). Although many HFT strategies exist today and they are largely unknown to public, researchers have shed lights on their general characteristics recently. Several illustrative HFT strategies include: (1) acting as an informal or formal market-maker, (2) high-frequency relative-value trading, and (3) directional

trading on news releases, order flow, or other high-frequency signals (Jones (2012)).

In the past few years, there have been a number of studies of HFT and algorithmic trading more generally. Their primary objective is to understand the economic impact of these algorithmic trading practices to the market quality including liquidity, price discovery process, trading costs, etc. On the empirical side, some researchers have been able to identify a specific HFT in data, and others are able to identify whether a trade is from algorithmic traders. Given the amount of information provided by exchanges and data vendors, it is possible to describe patterns in algorithmic order submission, order cancellation, and trading behavior. It is also possible to see whether algorithmic or HFT activities are correlated with bid-ask spreads, temporary and/or permanent volatility, trading volume, and other market activity and quality measures. Hendershott et al. (2011) study the implementation of an automated quote at the New York Exchange. They find that the implementation of auto-quote is associated with an increase in electronic message traffic and an improvement in market quality including narrowed effective spreads, reduced adverse selection, and increase price discovery. These effects are concentrated in large-cap firms, and there is little effect in small-cap stocks. Menkveld (2012) studies the July 2007 entry of a high-frequency market-maker into the trading of Dutch stocks. He argues that competition between trading venues facilitated the arrival of this high-frequency market-maker and HFT more generally, and he shows that high-frequency market-maker entry is associated with 23% less adverse selection. The volatility measured using 20 minutes realized volatility is unaffected by the entry of the high-frequency market-maker. Riordan et al. (2012) examine the effect of a technological upgrade on the market quality of 98 actively traded German stocks. They conclude that the ability to update quotes faster helps liquidity providers minimize their losses to liquidity demanders, and there are more price discovery take place. Boehmer et al. (2012) examine international evidence on electronic message traffic and market quality across 39 stock exchanges over the 2001-2009 period. They find that co-location increases algorithmic trading and HFT, and introduction of co-location improves liquidity and the information efficiency of prices. However, they claim volatility does not decline as much as it would be based on the observed narrower bid-ask spreads. Gai et al. (2012) study the effect of two recent 2010 Nasdaq technology upgrades that reduce the minimum time between messages from 950 nanoseconds to 200 nanoseconds. These technological changes lead to substantial increase in the number of canceled orders without much change in overall trading volume. There is so little change in bid-ask spreads and depths. Overall, these studies have focused on empirical evidence that an increase in algorithmic trading has positive influence on market quality in general.

On the theoretical side, there are a number of models developed to understand the economic impact of these algorithmic trading practices. Biais et al. (2012) conclude HFT can trade on new information more quickly, generating adverse selection costs, and they also find multiple equilibrium in their model, and some which exhibit socially inefficient over investment in HFT. The model from Jovanovic et al. (2010) shows that HFT can avoid some adverse selection, and can provide some that benefit to uninformed investors who need to trade. Martinez et al. (2012) conclude from their model that HFT obtains and trades on information an instant before it is available to others, and it imposes adverse selection on market-makers. Therefore liquidity is worse, and prices are no longer efficient. Martinez et al. (2012) focus on HFTs that demand liquidity, and suggest that HFT makes market prices extremely efficient by incorporating information as soon as it becomes available. Markets are not destabilized, as long as there is a population of market-makers standing ready to provide liquidity at competitive prices. Other related theoretical models include Pagnotta et al. (2012), who focus on the investment in speed made by exchanges in order to attract trading volume from speed sensitive investors. Moallemi et al. (2012) argue that a reduction in latency allows limit order submitters to update their orders more quickly, thereby reducing the value of the trading option that a limit order grants to a liquidity demander. The common theme in these models is that HFT may increase adverse selection, and it is harmful for liquidity. However, the ability to intermediate traders who arrive at different times is generally good for liquidity.

Moreover, there have been a number of studies focused on algorithmic traders' behaviors. These studies examine the trading activities of different types of traders and try to distinguish their behavioral differences. Hendershott *et al.* (2012) use exchange classifications distinguish algorithmic traders from orders managed by humans. They find that algorithmic traders concentrates in smaller trade sizes, while large block trades of 5,000 shares or more are predominantly originated by human traders. Algorithmic traders consume liquidity when bid-ask spreads are relatively narrow, they supply liquidity when bid-ask spreads are relatively wide. This suggests that algorithmic traders provide a more consistent level of liquidity through time. Brogaard (2012) and Hendershott *et al.* (2011) work with Nasdaq data that flag whether trades involves HFT.

Hendershott et al. (2011) find that HFT accounts for about 42% of (double-counted) Nasdaq volume in large-cap stocks but only about 17% of volume in small-cap stocks. They estimate a state-space model that decomposes price changes into permanent and temporary components, and measures the contribution of HFT and non-HFT liquidity supply and liquidity demand to each of these price change components. They find that when HFTs initiate trades, they trade in the opposite direction to the transitory component of prices. Thus, HFTs contribute to price discovery and contribute to efficient stock prices. Brogaard (2012) similarly finds that 68% of trades have an HFT on at least one side of the transaction, and he also finds that HFT participation rates are higher for stocks with high share prices, large market caps, narrow bid-ask spreads, or low stock-specific volatility. He estimates a vector autoregressive permanent price impact model and finds that HFT liquidity suppliers face less adverse selection than non-HFT liquidity suppliers, suggesting that they are somewhat judicious in supplying liquidity. Kirilenko et al. (2011) use account-level tick-by-tick data on the E-Mini S&P 500 futures contract, and they classify traders into various categories, including HFTs, opportunistic traders, fundamental traders and noise traders. Benos et al. (2012) conduct a similar analysis using UK equity data. These different datasets provide considerable insight into overall HFT trading behavior.

One of the important goals of learning traders trading strategies is to be able to categorize and identify the market participants, and be able to further understand their influences related to such important economic issues as multiple characterizations of price formation processes, market liquidity, and order flow, etc. (Hasbrouchk *et al.* (2001), Gabaix *et al.* (2003), Gatheral (2010), Hasbrouck (1991), and Jones *et al.* (1994)). We assert that enhanced understanding of the economic implication of these different algorithmic trading strategies will yield quantitative evidence of value to market policy makers and regulators seeking to maintain transparency, fairness and overall health in the financial markets.

In particular, traders deploy different trading strategies where each strategy has a unique value proposition under a particular market condition. In other words, we can cast this problem as a sequential decision problem under different conditions. Traders aim to optimize their decisions overtime and consequently maximize their reward under different market conditions. We can theoretically use reward functions to represent the value system that are encapsulated in the various different trading strategies. It is possible to derive new policies based on the reward functions learned and apply them in a new environment to govern a new autonomous process.

This process is defined as reward learning under the framework of inverse reinforcement learning (Ng et al. (2000), Abbeel et al. (2004) and Ramachandran et al. (2007). For example, a simple keep-or-cancel strategy for buying one unit, the trader has to decide when to place the order and when to cancel the order based on the market condition which may likely be characterized as a stochastic process. However the value proposition for the trader is to buy one unit of the security at a lowest price possible. This could be realized in a number of ways. It could be described as a reward function meaning when the system is in a particular state, the trader is always looking for a fixed reward. This notion of value proposition drives the trader to take corresponding actions according to the market conditions. This ultimately constitutes trader's policies or strategies. Therefore a strategy under certain value proposition can be consistently programmed in algorithms to achieve its goal of buy-one-unit in an optimal way. Consequently, strategies developed under certain value frameworks can be observed, learned and even reproduced in a different environment (such as a simulated financial market where impact of these strategies can be readily assessed). As documented by Yang et al. (2012), Hayes et al. (2012) and Paddrik et al. (2012), the manipulative or disruptive algorithmic strategies can be studied and monitored by market operators and regulators to prevent unfair trading practices. Furthermore, new emerging algorithmic trading practices can be assessed and new regulations and policies can be evaluated to maintain the overall health of the financial markets.

In this study, we model the trading behavior of different market participants by the solution to the inverse Markov decision process (MDP). We try to describe how traders are able to take actions in a highly uncertain environment to reach return goals on different horizons. This task can be solved using dynamic programming (DP) and reinforcement learning (RL) based on MDP. The model accounts for traders' preferences and expectations of uncertain state variables. In a general MDP modeling setting, we describe these variables in two spaces: the state space and the action space. From the trading decision perspective, we can parameterize learning agents using reward functions that depend on state and action. We consider the market dynamics in view of the learning agents' subjective beliefs. The agents perform DP/RL through a sense, trial and learn cycle. First, the agents gain state information from sensory input. Based on the current state, knowledge and goals, the agents find and choose the best action. Upon receiving new feedback, the agents learn to update their knowledge with a goal of maximizing their cumulative expected reward. In the discrete-valued state and action problem space, DP and RL methods use

similar techniques involving policy iteration and value iteration algorithms (Bertsekas (1996), and Sutton *et al.* (1998)) to solve MDP problems. Formalisms for solving forward problems of RL are often divided into model-based and model-free approaches (Daw *et al.* (2005), and Sutton *et al.* (1998)).

As framed by Abbeel *et al.* (2004) under the Inverse Reinforcement Learning (IRL) framework, the entire field of reinforcement learning is founded on the presupposition that the reward function, rather than policy, is the most succinct, robust, and transferable definition of the task. However, the reward function is often difficult to know in advance for some real-world tasks, so the following difficulties may arise: 1) We have no experience to tackle the problem; 2) We have experience but can not interpret the reward function explicitly; 3) The problem we solve may be interacting with the adversarial decision makers who make all their effort to keep the reward function secret. Rather than accessing the true reward function, it is easier to observe the behavior of some other agents (teacher/expert) to determine how to solve the problem. Hence, we have motivation to learn from observations. Technical approaches to learning from observations generally fall into two broad categories Ratliff *et al.* (2009). The first category, called imitation learning, attempts to use supervised learning to predict actions directly from observations of features of the environments, which is unstable and vulnerable to highly uncertain environment. The second category is concerned with how to learn the reward function that characterizes the agent's objectives and preferences in MDP (Ng *et al.* (2000)).

IRL was first introduced in machine learning literature by Ng *et al.* (2000) in formulating it as an optimization problem to maximize the sum of differences between the quality of the optimal action and the quality of the next-best action. Other algorithms have been developed or integrated into apprenticeship learning based on this linear approximation of the reward function. The principal idea of apprenticeship learning using IRL is to search mixed solutions in a space of learned policies with the goal that the cumulative feature expectation is near that of the expert (Abbeel *et al.* (2004) and Syed *et al.* (2008)).

Other algorithms have also been developed under the IRL framework. A game-theoretic approach to apprenticeship learning using IRL was developed in the context of a two-player zero-sum game in which the apprentice chooses a policy and the environment chooses a reward function (Syed *et al.* (2007)). Another algorithm for IRL is policy matching, in which the loss function penalizing deviations from the expert's policy that is minimized by tuning the

parameters of the reward functions (Neu *et al.* (2007)). The maximum entropy IRL is proposed in the context of modeling real-world navigation and driving behaviors (Ziebart *et al.* (2008)). The algorithms for apprenticeship learning using IRL do not actually aim to recover the reward function but instead are only concerned with the optimal policy. Ramachandran and Amir consider IRL from a Bayesian perspective without assuming the linear approximation of the reward function (Ramachandran *et al.* (2007)). Their model interprets the observations from the expert as the evidence that is used to obtain a posterior distribution over reward using Markov Chain Monte Carlo simulation. Recent theoretical works on IRL such as the framework of the linear-solvable MDP (Dvijotham *et al.* (2010)), have also improved the learning performance. IRL has also been successfully applied to many real-world problems, such as the automatic control of helicopter flight (Abbeel *et al.* (2010)).

We apply a Gaussian process-based IRL (GPIRL) model proposed by Qiao and Beling (2011) to learn the trading behaviors under different market conditions. In this GPIRL, a Gaussian prior is assigned on the reward function and the reward function is treated as a Gaussian process. This approach is similar to that of Ramachandran et al. (2007), who view the state-action samples from agents as the evidence that will be used to update a prior value in the reward function, under a Bayesian framework. The solution (Ramachandran et al. (2007)) depends on non-convex optimization using Markov Chain Monte Carlo simulation. Moreover, the ill-posed nature of the inverse learning problem also presents difficulties. Multiple reward functions may yield the same optimal policy, and there may be multiple observations at a single state given the true reward function. The GPIRL model aims to address the ill-posed nature of this problem by applying Bayesian inference and preference graphs. Here we are faced with the challenge of modeling traders' action as non-deterministic policies. In general, agent's policies range from deterministic Markovian to randomized history dependent, depending on how traders incorporate past information and how traders select actions. Due to the uncertainty of the environment and the random error of the measurement in the observations, a deterministic policy could very likely be perceived as a non-deterministic one. Modeling traders' reward function using a Gaussian process is well suited to address these issues. One of the main novel features of this approach is that it not only represents a probabilistic view but is also computationally tractable.

The dynamic nature of financial markets makes it possible to postulate a priori a relationship between the market variables we observe and those we wish to predict. The main contributions of this study can be summarized as follows:

- We propose an entirely different feature space that is constructed by inferring key parameters of a sequential optimization model that we take as a surrogate for the decision making process of the traders. We infer the reward (or objective) function for this process from observation of trading actions using a process from machine learning known as inverse reinforcement learning (IRL).
- 2) We model traders' reward functions using a Gaussian process. We also apply preference graphs to address the non-deterministic nature of the observed trading behaviors, reducing the uncertainty and computational burden caused by the ill-posed nature of the inverse learning problem.
- 3) We suggest a quantitative behavioral approach to categorizing algorithmic trading strategies using weighted scores over time in the reward space, and we conclude that it performs consistently better than the existing summary statistic-based trader classification approach (Kirilenko *et al.* (2011)).

The remainder of this paper is organized as follows: First we discuss summary statistics approach to trader classification in section II We then discuss the framework of which we use to model market dynamics and the traders' decisions in sectionII We extend the MDP and introduce IRL formulation in section IV We review the original linear IRL formulation and provide a Bayesian probabilistic model to infer the reward function using Gaussian processes. We apply the GPIRL algorithm to the E-Mini S&P 500 Futures market as experiments in section V. We show that the GPIRL algorithm can accurately capture algorithmic trading behavior based on observations of the high frequency data. We also compare our behavior-based classification results with the results from Kirilenko *et al.* (2011), and show that our behavioral approach represents a consistent improvement. Finally we provide concluding remarks about the GPIRL and its applications in section VI

#### **II. SUMMARY STATISTICS APPROACH TO TRADER CLASSIFICATION**

Kirilenko *et al.* (2011) suggest an approach to classifying individual trading accounts based on the summary statistics of trading volume and inventory and consistency of buying or selling behavior. Six categories are used to describe individual trading accounts:

- 1) High Frequency Traders high volume and low inventory;
- 2) Intermediaries low inventory;
- 3) Fundamental Buyers consistent intraday net buyers;
- 4) Fundamental Sellers consistent intraday net sellers;
- 5) Opportunistic Traders all other traders not classified;
- 6) Small Traders low volume.

In this section, we develop the details for classification rules corresponding to the six categories from Kirilenko *et al.* (2011), and then apply these rules to a real-world futures contract data set.

#### A. E-Mini Market Data Description

The E-Mini S&P 500 is a stock market index of futures contracts traded on the Chicago Mercantile Exchange's (CME) Globex electronic trading platform. The notional value of one contract is \$50 times the value of the S&P 500 stock index. The tick size for the E-Mini S&P 500 is 0.25 index points or \$12.50. For example, the S&P 500 Index futures contract is trading at \$1,400.00, then the value of one contract is \$70,000. The advantages to trading E-mini S&P 500 contracts include liquidity, greater affordability for individual investors and around-the-clock trading.

Trading takes place 24 hours a day with the exception of a short technical maintenance shutdown period from 4:30 p.m. to 5:00 p.m. The E-Mini S&P 500 expiration months are March, June, September, and December. On any given day, the contract with the nearest expiration date is called the front-month contract. The E-Mini S&P 500 is cash-settled against the value of the underlying index and the last trading day is the third Friday of the contract expiration month. The initial margin for speculators and hedgers are \$5,625 and \$4,500, respectively. Maintenance margins for both speculators and hedgers are \$4,500. There is no limit on how many contracts can be outstanding at any given time.

The CME Globex matching algorithm for the E-Mini S&P 500 offers strict price and time priority. Specifically, limit orders that offer more favorable terms of trade (sell at lower prices and buy at higher prices) are executed prior to pre-existing orders. Orders that arrived earlier are matched against the orders from the other side of the book before other orders at the same price. This market operates under complete price transparency. This straight forward matching

algorithm allows us to reconstruct the order book using audit trail messages archived by the exchanges and allows us to replay the market dynamics at any given moment.

In this paper, empirical work is based on a month of E-Mini order book audit trail data. The audit trail data includes all the order book events timestamped at a millisecond time resolution, and contains the following data fields: date, time (the time when the client submits the order to the exchange), conf\_time (the time when the order is confirmed by the matching engine), customer account, tag 50 (trader identification number), buy or sell flag, price, quantity, order ID, order type (market or limit), and func\_code (message type, e.g. order, modification, cancellation, trade, etc.).

# B. Summary Statistic-based Classification of E-Mini Market Data

We apply the set of the statistics based trader classification rules documented by Kirilenko *et al.* (2011) on our E-Mini dataset. For Fundamental traders, we calculate their end of the day net position. If the end-of-the-day net position is more than 15% of their total trading volume on that day, we categorize them either as Fundamental Buyers or Fundamental Sellers depending on their trading directions. We also identify Small Traders as those accounts with a trading volume of 9 contracts or less. We apply the criteria (Kirilenko *et al.* (2011)) for Intermediaries, Opportunistic Traders and High Frequency Traders, and obtain consistent results based on the one-month data. There are two steps involved in this process. First, we ensure that the account's net holdings fluctuate within 1.5% of its end of day level, and second, we ensure the account's end of the day net position is no more than 5% of its daily trading volume. Then if we define HFTs as a subset of Intermediaries (top 7% in daily trading volume), we find that there is a significant amount of overlap between HFTs and Opportunistic Traders. The problem is that the first criterion is not well defined, as the fluctuation of net holdings is vaguely defined. Net holdings could be measured in different ways.

In consultation with the authors of Kirilenko *et al.* (2011), we choose the standard deviation of an account's net position measured on the event clock as a measure of an account's holding fluctuation. With this definition, we find that a 1.5% fluctuation is too stringent for HFTs, because many high trading volume accounts are classified as Opportunistic Traders, while in reality their end of day positions are still very low compared with other Opportunistic Traders. Therefore, it is adequate to relax the first criterion requiring that the standard deviation of

the account's net holdings throughout the day is less than its end of day holding level. We find that the newly adjusted criteria classify most high volume trading accounts as HFTs, and this classification rule is validated from the registration information we can acquire. Without this adjustment, almost all the top trading accounts are incorrectly classified as Opportunistic Traders. Table I summarizes the results after applying the new classification rule demonstrating that the modified classification criteria identified more HFTs. On average, there are 38 HTF accounts, 118 Intermediary accounts, 2,658 Opportunistic accounts, 906 Fundamental Buyer accounts, 775 Fundamental Seller accounts, and 5,127 Small Trader accounts. Over the 4-week period, only 36% of the 120 accounts that consistently classified as the same type of traders. If we rank these accounts by their daily trading volume, we find that only 40% of the top 10 accounts are consistently classified as the same trader types. The variation occurs among the HFTs, Intermediaries, and Opportunistic Traders.

### **III. MARKOV DECISION PROCESS MODEL OF MARKET DYNAMICS**

In this section, we develop a Markov decision process (MDP) model of trader behavior. This model will then serve as the basis for the inverse reinforcement learning process described in section IV.

# A. MDP Background and Notation

The primary aim of our trading behavior-based learning approach is to uncover decision makers' policies and reward functions through the observations of an expert whose decision process is modeled as an MDP. In this paper, we restrict our attention to a finite countable MDP for easy exposition, but our approach can be extended to continuous problems if desired. A discounted finite MDP is defined as a tuple  $\mathbb{M} = (S, \mathcal{A}, \mathcal{P}, \gamma, r)$ , where

- $\mathcal{S} = \{s_n\}_{n=1}^N$  is a set of N states. Let  $\mathcal{N} = \{1, 2, \cdots, N\}.$
- $\mathcal{A} = \{a_m\}_{m=1}^M$  is a set of M actions. Let  $\mathcal{M} = \{1, 2, \cdots, M\}$ .
- P = {P<sub>am</sub>}<sup>M</sup><sub>m=1</sub> is a set of state transition probabilities (here P<sub>am</sub> is a N×N matrix where each row, denoted as P<sub>am</sub>(s<sub>n</sub>,:), contains the transition probabilities upon taking action a<sub>m</sub> in state s<sub>n</sub>. The entry P<sub>am</sub>(s<sub>n</sub>, s<sub>n'</sub>) is the probability of moving to state s<sub>n'</sub>, n' ∈ N in the next stage.).
- $\gamma \in [0,1]$  is a discount factor.

|            |      |        | listic          | lten   | lten   | nber.                 | in a       |
|------------|------|--------|-----------------|--|--|-----------------------|------------|
|            |      | GC .   | X <sup>10</sup> | and the second s | and the second s | A.M.                  | A CONTRACT |
|            |      | A A A  | 000             | 100  | 100  | (Pa)                  | પૈયા       |
|            |      | Ź      | 07              | E.   | 14<br>14   | <i>R</i> <sup>2</sup> | R          |
| Date       | HFTs | Makers | Traders         | Buyers   | Sellers  | of Accounts           | Volume     |
| 10/04/2012 | 39   | 193    | 2,833           | 940  | 818  | 10,425                | 3,261,852  |
| 10/05/2012 | 38   | 162    | 2,598           | 1191   | 1055   | 11,495                | 3,875,232  |
| 10/06/2012 | 38   | 167    | 2,401           | 895  | 712  | 9,065                 | 2,852,244  |
| 10/07/2012 | 39   | 196    | 2,726           | 919  | 747  | 9,841                 | 3,424,768  |
| 10/08/2012 | 32   | 162    | 2,511           | 847  | 812  | 9,210                 | 3,096,800  |
| 10/11/2012 | 21   | 118    | 1,428           | 636  | 573  | 6,230                 | 1,765,254  |
| 10/12/2012 | 38   | 186    | 2,687           | 896  | 745  | 9,771                 | 3,236,904  |
| 10/13/2012 | 38   | 187    | 2,582           | 1020   | 840  | 10,297                | 3,699,108  |
| 10/14/2012 | 30   | 198    | 3,001           | 1070   | 795  | 10,591                | 4,057,824  |
| 10/15/2012 | 46   | 210    | 3,109           | 890  | 773  | 9,918                 | 4,437,826  |
| 10/18/2012 | 37   | 173    | 2,126           | 869  | 724  | 8,735                 | 2,458,510  |
| 10/19/2012 | 52   | 216    | 3,651           | 1030   | 974  | 11,600                | 5,272,672  |
| 10/20/2012 | 39   | 176    | 2,949           | 951  | 877  | 10,745                | 3,956,790  |
| 10/21/2012 | 43   | 240    | 3,370           | 952  | 771  | 10,980                | 4,230,194  |
| 10/22/2012 | 32   | 143    | 1,837           | 676  | 629  | 7,370                 | 2,026,234  |
| 10/25/2012 | 38   | 181    | 2,533           | 888  | 684  | 9,228                 | 3,074,558  |
| 10/26/2012 | 37   | 175    | 2,726           | 816  | 709  | 9,568                 | 3,000,628  |
| 10/27/2012 | 45   | 186    | 2,973           | 919  | 820  | 10,472                | 3,850,556  |
| 10/28/2012 | 39   | 185    | 2,873           | 914  | 705  | 9,777                 | 3,485,910  |
| 10/29/2012 | 37   | 160    | 2,247           | 794  | 744  | 8,369                 | 3,012,860  |

 TABLE I

 The E-Mini S&P 500 Futures Market Data Summary

• r denotes the reward function, mapping from  $\mathcal{S} \times \mathcal{A}$  to  $\Re$  with the property that

$$r(s_n, a_m) \triangleq \sum_{n' \in \mathcal{N}} \mathbf{P}_{a_m}(s_n, s_{n'}) r(s_n, a_m, s_{n'})$$

where  $r(s_n, a_m, s_{n'})$  denotes the function giving the reward of moving to the next state  $s_{n'}$  after taking action  $a_m$  in current state  $s_n$ . The reward function  $r(s_n, a_m)$  may be further reduced to  $r(s_n)$ , if we neglect the influence of the action. We use **r** for reward vector through out this paper. If the reward only depends on state, we have  $\mathbf{r} = (\mathbf{r}(s_1), \ldots, \mathbf{r}(s_N))$ . If we let **r** be the vector of the reward depending on both state and action. We have

$$\mathbf{r} = (\mathbf{r}_1(s_1), \dots, \mathbf{r}_1(s_N), \dots, \mathbf{r}_M(s_1), \dots, \mathbf{r}_M(s_N))$$
$$= (\mathbf{r}_1, \cdots, \mathbf{r}_M).$$

In an MDP, the agent selects an action at each sequential stage, and we define a *policy* (*behavior*) as the way that the actions are selected by a decision maker/agent. Hence this process can be described as a mapping between state and action, i.e., a random state-action sequence  $(s^0, a^0, s^1, a^1, \dots s^t, a^t, \dots)$ , <sup>1</sup> where  $s^{t+1}$  is connected to  $(s^t, a^t)$  by  $\mathbf{P}_{a^t}(s^t, s^{t+1})$ .

We also define rational agents as those that behave according to the optimal decision rule where each action selected at any stage maximizes the value function. The value function for a policy  $\pi$  evaluated at any state  $s^0$  is given as  $V^{\pi}(s^0) = E[\sum_{t=0}^{\infty} \gamma^t r(s^t, a^t) | \pi]$ . This expectation is over the distribution of the state sequence  $\{s^0, s^1, ...\}$  given the policy  $\pi = \{\mu^0, \mu^1, \cdots\}$ , where  $a^t = \mu^t(s^t), \ \mu^t(s^t) \in U(s^t)$  and  $U(s^t) \subset \mathcal{A}$ . The objective at state s is to choose a policy that maximizes the value of  $V^{\pi}(s)$ . The optimal policy is then  $V^*(s^0) = \sup_{\pi} E[\sum_{t=0}^{\infty} \gamma^t r(s^t, a^t) | \pi]$ . Similarly, there is another function called the *Q*-function (or *Q*-factor) that judges how well an action is performed in a given state. The notation  $Q^{\pi}(s, a)$  represents the expected return from state s when action a is taken and thereafter policy  $\pi$  is followed.

In the infinite-horizon case, the stationary Markovian structure of the problem implies that the only variable that affects the agent's decision rule and the corresponding value function should be time invariant. We then have the essential theory of MDPs (Bellman R. (1957)) as follows:

Theorem 1 (Bellman Equations): Given a stationary policy  $\pi$ ,  $\forall n \in \mathcal{N}, m \in \mathcal{M}, V^{\pi}(s_n)$  and

<sup>&</sup>lt;sup>1</sup>Superscripts represent time indices. For example  $s^t$  and  $a^t$ , with the upper-index  $t \in \{1, 2, \dots\}$ , denote state and action at the t-th horizon stage, while  $s_n$  (or  $a_m$ ) represents the n-th state (or m-th action) in S (or A).

 $Q^{\pi}(s_n, a_m)$  satisfy

$$V^{\pi}(s_n) = r(s_n, \pi(s_n)) + \gamma \sum_{n' \in \mathcal{N}} \mathbf{P}_{\pi(s_n)}(s_n, s_{n'}) V^{\pi}(s_{n'}),$$
  
$$Q^{\pi}(s_n, a_m) = r(s_n, a_m) + \gamma \sum_{n' \in \mathcal{N}} \mathbf{P}_{a_m}(s_n, s_{n'}) V^{\pi}(s_{n'}).$$

Theorem 2 (Bellman Optimality):  $\pi$  is optimal if and only if,  $\forall n \in \mathcal{N}, \pi(s_n) \in \arg \max_{a \in \mathcal{A}} Q^{\pi}(s_n, a)$ .

Based on the above theorem of MDPs, we have the following equations to represent the Q-function as a the reward function.

$$Q^{\pi}(s_n, a_m) = \mathbf{r}_m(s_n) + \gamma \mathbf{P}_{a_m}(s_n, :)(\mathbf{I} - \gamma \mathbf{P}_{\pi})^{-1} \mathbf{r}_m,$$

where  $\mathbf{P}_{\pi}$  represents the state transition probability matrix for following policy  $\pi$  at every state, and  $\mathbf{r}_{\mathbf{m}}$  represents the reward vector under action  $a_m$ .

# B. Constructing an MDP Model from Order Book Data

Figure 1 shows the entire life-cycle of an order initiated by a client of an exchange. The order book audit trail data contains these messages, and the entire order history (i.e. order creation, order modifications, fills, cancellation, etc.) can be retrieved and analyzed. To construct an MDP model of trader behavior, we first reconstruct the limit order book using the audit trail messages. The order book then contains bid/ask prices, market depth, liquidity, etc. During this process on the E-Mini data described in section 2.1, we processed billions of messages for each trading date, and built price queues using the price and time priority rule.

Once we have the order book at any given event tick, we take the market depth at five different levels as our base variables and then discretize these variables to generate an MDP model state space. This study extends the MDP model documented by Yang *et al.* (2012) to obtain five variables, i.e., order volume imbalance between the best bid and the best ask prices, order volume imbalance between the 2nd best bid and the 2nd best ask prices, order volume imbalance between the 3rd best bid and the 3rd best ask prices, the order book imbalance at the 5th best bid and the 5th ask prices, and the inventory level/holding position (see Figure 2 (b)). Then we discretize the values of the five variables into three levels defined as high (above  $\mu + 1.96\sigma$ ), neutral ( $\mu \pm 1.96$ ), and low (below  $\mu - 1.96\sigma$ ). Based on our observation that the first 3 best bid and ask prices change the most, we select the first 3 level order book imbalance



Fig. 1. **CME Globex Order Lifecycle.** T1: Trader submits a new order; T2: The state of an order is changed, if a stop is activated; T3: A trader may choose to cancel an order, and the state of an order can be modified multiple times; T4: When an order is partially filled, the quantity remaining decreases; T5: Order elimination is similar to order cancellation except it is initiated by the trading engine; T7: An order may be filled completely; T6: Trades can be eliminated after the fact by the exchanges.

variables in modeling the limit order book dynamics. As argued by Yang *et al.* (2012), these volume-related variables reflect the market dynamics on which the traders/algorithms depend to place their orders at different prices.

As the volume imbalance at the best bid/ask prices is the most sensitive indicator of the trading behavior of HFTs, Intermediaries and some of the Opportunistic traders, we also hypothesize that the volume imbalance at other prices close to the book prices will be useful in inferring trader behavior. As demonstrated in previous work (Yang *et al.* (2012)), the private variable of a trader's inventory level provides critical information about trader's behavior. Traders in high frequency environments strive to control their inventory levels as a critical measure of controlling the risk of their position (Kirilenko *et al.* (2011), Easley *et al.* (2010) and Brogaard *et al.* (2010)). HFTs and Market Makers tend to turn over their inventory level five or more times a day and to hold very small or even zero inventory positions at the end of the trading session. These observations provide strong support for the introduction of a position variable



Limit order book MDP model with 5 different levels of volume imbalances, and 10 buckets of price placement.

Fig. 2. Order Book MDP Model: This graph shows the state variables used in the MDP model.

to characterize trader behavior in our model. Therefore, together with the volume imbalance variables, we propose a computational model with  $3^5 = 243$  states.

Next, we need to define the action space. In general, there are three types of actions: placing a new order, canceling an existing order, or placing a market order. We divide the limit order book into 10 buckets at any given point of time by the following price markers: the best bid price, the 2nd best bid price, the 3rd best bid price, between the 4th and 5th bid prices, below the 5th best bid price, the best ask price, the 2nd best ask price, the 3rd best ask price, the 3rd best ask price, the 4th and 5th ask prices, and above the 5th best ask price. Then, at any given point of time, a trader can take 22 actions. The price markers used to define the price ranges are illustrated in Figure (2). We use unit size for all the actions. Orders other than unit size are treated as repeated actions without state transition.

#### **IV. INVERSE REINFORCEMENT LEARNING**

Given an MDP  $\mathbb{M} = (S, \mathcal{A}, \mathcal{P}, \gamma, r)$ , let us define the inverse Markov decision process (IMDP)  $\mathbb{M}_{\mathbb{I}} = (S, \mathcal{A}, \mathcal{P}, \gamma, \mathcal{O})$ . The process  $\mathbb{M}_{\mathbb{I}}$  includes the states, actions, and dynamics of  $\mathbb{M}$ , but lacks a specification of the reward function, r. By way of compensation,  $\mathbb{M}_{\mathbb{I}}$  includes a set of observations  $\mathcal{O}$  that consists of state-action pairs generated through the observation of a decision maker. We can define the *inverse reinforcement learning* (IRL) problem associated with  $\mathbb{M}_{\mathbb{I}} = (S, \mathcal{A}, \mathcal{P}, \gamma, \mathcal{O})$  to be that of finding a reward function such that the observations  $\mathcal{O}$  could have come from an optimal policy for  $\mathbb{M} = (S, \mathcal{A}, \mathcal{P}, \gamma, r)$ . The IRL problem is, in general, highly under-specified, which has led researchers to consider various models for restricting the set of reward functions under consideration. Ng *et al.* (2000), in a seminal consideration of IMDPs and associated IRL problems, observed that, by the optimality equations, the only reward vectors consistent with an optimal policy  $\pi$  are those that satisfy the set of inequalities

$$(\mathbf{P}_{\pi} - \mathbf{P}_{a})(\mathbf{I} - \gamma \mathbf{P}_{\pi})^{-1}\mathbf{r} \ge \mathbf{0}, \forall a \in \mathcal{A},$$
(1)

where  $\mathbf{P}_{\pi}$  is the transition probability matrix relating to observed policy  $\pi$  and  $\mathbf{P}_a$  denotes the transition probability matrix for other actions. Note that the trivial solution  $\mathbf{r} = \mathbf{0}$  satisfies the constraints (1), which highlights the under-specified nature of the problem and the need for reward selection mechanisms.

In the machine learning and artificial intelligence literature, a principal motivation for considering IRL problems is the idea of apprenticeship learning, in which observations of state-action pairs are used to learn the policies followed by experts for the purpose of mimicking or cloning behavior. By its nature, apprenticeship learning problems arise in situations where it is not possible or desirable to observe all state-action pairs for the decision maker's policy. The basic idea of apprenticeship learning through IRL is to first use IRL techniques to learn the reward function (vector) and then use that function to define an MDP problem, which can then be solved for an optimal policy. Our process is quite different. We learn the reward function with IRL and then directly use the rewards as features for classifying and clustering traders or trading algorithms.

#### A. Linear IRL

Ng *et al.* (2000) advance the idea choosing the reward function to maximize the difference between the optimal and suboptimal policies, which can be done using a linear programming formulation.

Most of the existing IRL algorithms make some assumption about the form of the reward function. Prominent examples include the model in Ng *et al.* (2000), which we term linear IRL (LIRL) because of its linear nature. In LIRL, the reward function is written linearly in terms of basis functions, and effort is made to maximize the quantity

$$\sum_{s \in \mathcal{S}} \left[ Q^{\pi}(s, a') - \max_{a \in \mathcal{A} \setminus a'} Q^{\pi}(s, a) \right], \forall a \in \mathcal{A}.$$
 (2)

The optimization problem in Ng *et al.* (2000) is equivalent to the following optimization program:

$$\begin{split} \max_{\mathbf{r}} \sum_{s \in \mathcal{S}} \beta(s) &- \lambda \sum_{s \in \mathcal{S}} |\mathbf{r}(s)| \\ \text{s.t.} \\ (\mathbf{P}_{\pi} - \mathbf{P}_{a}) (I - \gamma \mathbf{P}_{\pi})^{-1} \mathbf{r} & \geq \beta(s), \quad \forall a \in \mathcal{A}, \forall s \in \mathcal{S} \\ \beta(s) & \geq \mathbf{0}, \quad \forall s \in \mathcal{S}, \end{split}$$

where  $\lambda$  is a regularization parameter included to encourage sparse solution vectors. Yang *et al.* (2012) use this approach to find a feature space that can be used to classify and cluster simulated trading agents.

#### B. Bayesian IRL Framework

Ramachandran *et al.* (2007) originally proposed a Bayesian Framework for IRL. The posterior over reward is written as

$$p(r|\mathcal{O}) = p(\mathcal{O}|r)p(r) \propto \prod_{(s,a)\in\mathcal{O}} p(a|s,r).$$

Then, the IRL problem is written as  $\max_r \log p(\mathcal{O}|r) + \log p(r)$ . For many problems, however, the computation of  $p(r|\mathcal{O})$  may be complicated and some algorithms use Markov chain Monte Carlo (MCMC) to sample the posterior probability. Below we adopt a different approach that uses the idea of selecting reward on the basis of *maximum a posteriori* (MAP) estimate computed using convex optimization.

#### C. Gaussian Process IRL

We now turn to an IRL problem that addresses observations from a decision making process in which the reward function has been contaminated by Gaussian noise. In particular, we assume



Fig. 3. Action Preference Graph Examples: (a). This graph shows an example action preference graph at state 158. (b). This graph shows an example action preference graph at state 14.

that the reward vector can be modeled as  $r + \mathcal{N}(0, \sigma^2)$ , where  $\mathcal{N}(0, \sigma^2)$  is Gaussian noise. In the financial trading problem setting, we may observe certain trading behavior over a period of time, but we may not observe the complete polices behind a particular trading strategy. As discussed earlier, different trading strategies tend to look at different time horizons. Therefore, the observation period becomes critical to the learning process. Furthermore, two types of errors may be introduced into our observations: The first type of error may be introduced during our modeling process. Resolution of these discrete models will introduce errors into our observations. The second potential source of error is the strategy execution process. Execution errors will occur due to the uncertainty of market movements and will eventually appear in our observations, confounding our efforts to determine the true policy. Overall, there are two types of challenges in this learning problem: the uncertainty about reward functions given the observation of decision behavior and the ambiguity involved in observing multiple actions at a single state.

Qiao and Beling (2011) argue for two different modeling techniques in learning reward functions. To lessen the ambiguity of observing multiple actions at a state, they argue that Bayesian inference should be the basis for understanding the agent's preferences over the action space. This argument is reasonable because the goal of IRL is to learn the reward subjectively perceived by the decision maker from whom we have collected the observation data. The intuition is that decision makers will select some actions at a given state because they prefer these actions to others. These preferences are among the countable actions that can be used to represent multiple observations at one state.

Here we use two examples to demonstrate how the action preference relationships have been constructed based on the MDP model and observed actions. Table II shows two example states

| State | Action | Frequency Observed | State | Action | Frequency Observed |
|-------|--------|--------------------|-------|--------|--------------------|
| 14    | 1      | 0.23               | 158   | 1      | 0.30               |
| 14    | 2      | 0.14               | 158   | 3      | 0.07               |
| 14    | 7      | 0.06               | 158   | 7      | 0.11               |
| 14    | 11     | 0.26               | 158   | 11     | 0.30               |
| 14    | 12     | 0.09               | 158   | 17     | 0.07               |
| 14    | 16     | 0.17               | 158   | 18     | 0.07               |
| 14    | 26     | 0.06               | 158   | 20     | 0.07               |

TABLE II Action Preference Graph Examples



Fig. 4. **Examples Preference Graphs:** (a) An example of observing two actions at a state. (b) An example of a unique observation at a state.

with multiple observed actions. We then sort the frequency in descending order and construct a two-layer graph: the top layer has the most frequently observed actions and the bottom layer holds all the other actions. Based on this preference observation, we can construct two preference graphs as shown in Figure (3). The state transition matrix can be constructed for the entire market for the observation period. In our MDP model, we have a 243x243 matrix for every single action.

In the following, we first introduce the preference theory for the IMDP model, and then we formalize the idea of modeling the reward function as a Gaussian process under the Bayesian inference framework.

1) Action Preference Learning: In this section, we first define the action preference relationship and the action preference graph. At state  $s_n$ ,  $\forall \hat{a}, \check{a} \in \mathcal{A}$ , we define the action preference relation as:

- 1) Action  $\hat{a}$  is weakly preferred to  $\check{a}$ , denoted as  $\hat{a} \succeq_{s_n} \check{a}$ , if  $Q(s_n, \hat{a}) \ge Q(s_n, \check{a})$ ;
- 2) Action  $\hat{a}$  is strictly preferred to  $\check{a}$ , denoted as  $\hat{a} \succ_{s_n} \check{a}$ , if  $Q(s_n, \hat{a}) > Q(s_n, \check{a})$ ;
- 3) Action  $\hat{a}$  is equivalent to  $\check{a}$ , denoted as  $\hat{a} \sim_{s_n} \check{a}$ , if and only if  $\hat{a} \succeq_{s_n} \check{a}$  and  $\check{a} \succeq_{s_n} \hat{a}$ .

An action preference graph is a simple directed graph showing preference relations among the countable actions at a given state. At state  $s_n$ , the action preference graph  $G_n = (\mathcal{V}_n, \mathcal{E}_n)$ comprises a set  $\mathcal{V}_n$  of nodes together with a set  $\mathcal{E}_n$  of edges. For the nodes and edges in graph  $G_n$ , let us define

- 1) Each node represents an action in  $\mathcal{A}$ . Define a one-to-one mapping  $\varphi: \mathcal{V}_n \to \mathcal{A}$ .
- 2) Each edge indicates a preference relation.

Furthermore, we make the following assumption as a rule to build the preference graph, and then we show how to draw a preference graph at state  $s_n$ :

At state  $s_n$ , if action  $\hat{a}$  is observed, we have the following preference relations:  $\hat{a} \succeq_{s_n} \check{a}, \forall \check{a} \in \mathcal{A} \setminus \{\hat{a}\}.$ 

It is therefore straightforward to show the following according to Bellman optimality. The variable  $\hat{a}$  is observed if and only if  $\hat{a} \in \arg \max_{a \in \mathcal{A}} Q(s_n, a)$ . Therefore, we have

$$Q(s_n, \hat{a}) > Q(s_n, \check{a}), \ \forall \check{a} \in \mathcal{A} \setminus \{\hat{a}\}$$

According to the definition on preference relations, it follows that if  $Q(s_n, \hat{a}) > Q(s_n, \check{a})$ , we have  $\hat{a} \succ_{s_n} \check{a}$ . Hence, we can show that the preference relationship has the following properties:

- 1) If  $\hat{a}, \check{a} \in \mathcal{A}$ , then at state  $s_n$  either  $\hat{a} \succeq_{s_n} \check{a}$  or  $\check{a} \succeq_{s_n} \hat{a}$ .
- 2) If  $\hat{a} \succeq_{s_n} \check{a}$  or  $\check{a} \succeq_{s_n} \tilde{a}$ , then  $\hat{a} \succeq_{s_n} \tilde{a}$ .

At this point, we have a simple representation of the action preference graph that is constructed by a two-layer directed graph. We may have either multiple actions at  $s_n$  as in Figure (4) (a) or a unique action at  $s_n$  as in Figure (4) (b). In this two-layer directed graph, the top layer  $\mathcal{V}_n^+$ is a set of nodes representing the observed actions and the bottom layer  $\mathcal{V}_n^-$  contains the nodes denoting the other actions. The edge in the edge set  $\mathcal{E}_n$  can be represented by a formulation of its beginning node u and ending node v. We write the k-th edge as  $(u \to v)_k$  if  $u \in \mathcal{V}_n^+, v \in \mathcal{V}_n^-$ , or the l-th edge  $(u \leftrightarrow v)_l$  if  $u \in \mathcal{V}_n^-, v \in \mathcal{V}_n^-$ . Recall the mapping between  $\mathcal{V}_n$  and  $\mathcal{A}$ , the representation  $u \to v$  indicates that action  $\varphi(u)$  is preferred over  $\varphi(v)$ . Similarly,  $u \leftrightarrow v$  means that action  $\varphi(u)$  is equivalent to  $\varphi(v)$ .

In the context of financial trading decision process, we may observe multiple actions from one particular trader under certain market conditions. That is to say that the observation data  $\mathcal{O}$ may represent multiple decision trajectories generated by non-deterministic policies. To address IRL problems in those cases, Qiao and Beling (2011) propose to process  $\mathcal{O}$  into the form of pairs of state and preference graphs similar to the representation shown in Figure (5), and then we apply Bayesian inference using the new formulation.

According to Qiao and Beling (2011), we can represent  $\mathcal{O}$  as shown in Figure (5). At state  $s_n$ , its action preference graph is constructed by a two-layer directed graph: a set of nodes  $\mathcal{V}_n^+$  in the top layer and a set of nodes  $\mathcal{V}_n^-$  in the bottom layer. Under the non-deterministic policy assumption, we adopt a reward structure depending on both state and action.



Fig. 5. Observation structure for MDP.

2) Gaussian Reward Process: Recall that the reward depends on both state and action, and consider  $r_m$ , the reward related to action  $a_m$ , as a Gaussian process. We denote by  $k_m(s_i, s_j)$ the function generating the value of entry (i, j) for covariance matrix  $\mathbf{K}_m$ , which leads to  $\mathbf{r}_m \sim N(0, \mathbf{K}_m)$ . Then the joint prior probability of the reward is a product of multivariate Gaussian, namely  $p(\mathbf{r}|S) = \prod_{m=1}^{M} p(\mathbf{r}_m|S)$  and  $\mathbf{r} \sim N(0, \mathbf{K})$ . Note that  $\mathbf{r}$  is completely specified by the positive definite covariance matrix  $\mathbf{K}$ , which is block diagonal in the covariance matrices  $\{\mathbf{K}_1, \mathbf{K}_2, \dots, \mathbf{K}_M\}$  based on the assumption that the reward latent processes are uncorrelated. In practice, we use a squared exponential kernel function, written as:

$$k_m(s_i, s_j) = e^{\frac{1}{2}(s_i - s_j)\mathbf{T}_m(s_i - s_j)} + \sigma_m^2 \delta(s_i, s_j),$$

where  $\mathbf{T}_m = \kappa_m \mathbf{I}$  and  $\mathbf{I}$  is an identity matrix. The function  $\delta(s_i, s_j) = 1$ , when  $s_i = s_j$ ; otherwise  $\delta(s_i, s_j) = 0$ . Under this definition the covariance is almost unity between variables whose inputs are very close in the Euclidean space, and decreases as their distance increases.

Then, the GPIRL algorithm estimates the reward function by iteratively conducting the following two main steps:

1) Get estimation of  $\mathbf{r}_{MAP}$  by maximizing the posterior  $p(\mathbf{r}|\mathcal{O})$ , which is equal to minimize  $-\log p(\mathcal{O}|\mathbf{r}) - \log p(\mathbf{r}|\boldsymbol{\theta})$ , where  $\boldsymbol{\theta}$  denotes the vector of hyper-parameters including  $\kappa_m$  and  $\sigma_m$  that control the Gaussian process.

2) Optimize the hyper-parameters by using gradient decent method to maximize  $\log p(\mathcal{O}|\boldsymbol{\theta}, \mathbf{r}_{MAP})$ , which is the Laplace approximation of  $p(\boldsymbol{\theta}|\mathcal{O})$ .

3) Likelihood Function and MAP Optimization: GPIRL adopts the following likelihood functions to capture the strict preference and equivalent preference respectively.

$$p((\hat{a} \succ_{s_n} \check{a})_k | \mathbf{r}) = \Phi(\frac{Q(s_n, \hat{a}) - Q(s_n, \check{a})}{\sqrt{2}\sigma})$$
(3)

$$p((\hat{a} \sim_{s_n} \hat{a}')_l | \mathbf{r}) \propto e^{-\frac{1}{2}(Q(s_n, \hat{a}) - Q(s_n, \hat{a}'))^2}$$
(4)

In Eq. 3, the function  $\Phi(x) = \int_{-\infty}^{x} N(v|0, 1) dv$ , where N(v|0, 1) denotes a standard Gaussian variable.

As we stated earlier, if we model the reward functions as being contaminated with Gaussian noise that has a mean of zero and an unknown variance  $\sigma^2$ , we can then define the likelihood function for both the k-th strict preference relation and the l-th equivalent preference relation. Finally, we can formulate the following proposition:

*Proposition 3:* The likelihood function, given the evidence of the observed data ( $\mathcal{O}$  in the form of pairs of state and action preference graph ( $\mathcal{G}$ ), is calculated by

$$p(\mathcal{O}|\mathbf{r}) \propto p(\mathcal{G}|\mathcal{S}, \mathbf{r}) = \prod_{n=1}^{N} p(G_n|s_n, \mathbf{r}) = \prod_{n=1}^{N} \prod_{k=1}^{n_n} p((\hat{a} \succ_{s_n} \check{a})_k |\mathbf{r}) \prod_{l=1}^{m_n} p((\hat{a} \sim_{s_n} \hat{a}')_l |\mathbf{r}), \quad (5)$$

where  $n_n$  denotes the number of edges for strict preference and  $m_n$  means the number of edges for equivalent preference at state  $s_n$ .

In conclusion, the probabilistic IRL model is controlled by the kernel parameters  $\kappa_m$  and  $\sigma_m$  which compute the covariance matrix of reward realizations, and by  $\sigma$  which tunes the noise level in the likelihood function. We put these parameters into the hyper-parameter vector  $\boldsymbol{\theta} = (\kappa_m, \sigma_m, \sigma)$ . More often than not, we do not have prior knowledge about the hyper-parameters. And then we can apply maximum a posterior estimate to evaluate the hyper-parameters.

Essentially, we now have a hierarchical model. At the lowest level, we have reward function values encoded as a parameter vector **r**. At the top level, we have hyper-parameters in  $\theta$  controlling the distribution of the parameters. Inference takes place one level at a time. At

the bottom level, the posterior over function values is given by Bayes' rule:

$$p(\mathbf{r}|\mathcal{S}, \mathcal{G}, \boldsymbol{\theta}) = \frac{p(\mathcal{G}|\mathcal{S}, \boldsymbol{\theta}, \mathbf{r})p(\mathbf{r}|\mathcal{S}, \boldsymbol{\theta})}{p(\mathcal{G}|\mathcal{S}, \boldsymbol{\theta})}.$$
(6)

The posterior combines the prior information with the data, reflecting the updated belief about  $\mathbf{r}$  after observing the decision behavior. We can calculate the denominator in Eq.6 by integrating  $p(\mathcal{G}|\mathcal{S}, \boldsymbol{\theta}, \mathbf{r})$  over the function space with respect to  $\mathbf{r}$ , which requires a high computational capacity. Fortunately, we are able to maximize the non-normalized posterior density of  $\mathbf{r}$  without calculating the normalizing denominator, as the denominator  $p(\mathcal{G}|\mathcal{S}, \boldsymbol{\theta})$  is independent of the values of  $\mathbf{r}$ . In practice, we obtain the maximum posterior by minimizing the negative log posterior, which is written as

$$U(\mathbf{r}) \triangleq \frac{1}{2} \sum_{m=1}^{M} \mathbf{r}_{m}^{T} \mathbf{K}_{m}^{-1} \mathbf{r}_{m} - \sum_{n=1}^{N} \sum_{k=1}^{n_{n}} \ln \Phi(\frac{Q(s_{n}, \hat{a}) - Q(s_{n}, \check{a})}{\sqrt{2}\sigma}) + \sum_{n=1}^{N} \sum_{l=1}^{m_{n}} \frac{1}{2} (Q(s_{n}, \hat{a}) - Q(s_{n}, \hat{a}'))^{2}$$
(7)

Qiao and Beling (2011) present a proof that Proposition (7) is a convex optimization problem. At the minimum of  $U(\mathbf{r})$  we have

$$\frac{\partial U(\mathbf{r})}{\partial \mathbf{r}_m} = 0 \Rightarrow \hat{\mathbf{r}}_m = K_m(\nabla \log P(\mathcal{G}|\mathcal{S}, \boldsymbol{\theta}, \hat{\mathbf{r}}))$$
(8)

where  $\hat{\mathbf{r}} = (\hat{\mathbf{r}}_1, \cdots, \hat{\mathbf{r}}_{a_m}, \cdots, \hat{\mathbf{r}}_m)$ . In Eq.8, we can use Newton's method to find the maximum of U with the iteration,

$$\hat{\mathbf{r}}_m^{\mathrm{new}} = \hat{\mathbf{r}}_m - (\frac{\partial^2 U(\mathbf{r})}{\partial \mathbf{r}_m \partial \mathbf{r}_m})^{-1} \frac{\partial U(\mathbf{r})}{\partial \mathbf{r}_m}$$

# V. EXPERIMENT WITH THE E-MINI S&P 500 EQUITY INDEX FUTURES MARKET

In this section, we conduct two experiments using the MDP model defined earlier to identify algorithmic trading strategies. We consider the six trader classes defined by Kirilenko *et al.* (2011), namely High Frequency Traders, Market Makers, Opportunistic Traders, Fundamental Buyers, Fundamental Sellers and Small Traders. As we argue earlier, the focus of our study will be more on HFTs and Market Makers due to the large daily volume and their potential

impact to the financial markets. In Kirilenko *et al.* (2011)'s paper, there are only about from 16 to 20 HFTs on the S&P500 Emini market. Although this is a small population, their impact to the market has drawn increased attention from policy makers, regulators and academia. That is why we focus our attention on this small population. Among the roughly 10,000 trading accounts for the S&P500 Emini market, we narrow down to about 120 accounts based their high daily trading volume. In the first experiment, we select the top 10 trading accounts by their volume and end-of-the-day positions. In this we guarantee our subjects are HTFs. In the second experiment, we randomly select 10 out of the 120 accounts. This selection criterion ensures that our subjects are of either HTF or market making strategies. With these two experimentation, we show the performance of our IRL based approach to identify the high impact population of the Algorithmic trading strategies.

### A. Trader Behavior Identification

Yang *et al.* (2012) examine different trading behaviors using a linear IRL (LNIRL) algorithm with the simulated E-Mini S&P 500 market data. That MDP model contains three variables: the volume imbalance at the bid/ask prices, the volume imbalance at the 3rd best bid/ask prices, and the position level. Although this MDP model is relatively simple, it is evident from the experimental results that the IRL reward space is effective in identifying trading strategies with a relatively high accuracy rate.

This paper tries to address two important issues during the modeling process to solve a realistic market strategy learning problem using real market data. The first issue is that in reality, we often do not have a complete set of observations of a trader's policies. As the market presents itself as a random process in terms of both prices and volume, it is unlikely that we will be able to capture all possible states during our observation window. In contrast, the study performed by Yang *et al.* (2012) assumes complete observation of a trader's decision policies for the simulated trading strategies. In other words, the policies simulated by a distribution can be completely captured when the simulation is run long enough. The convergence of these simulated policies and the testing results are consistent with their assumptions. However, when we use real market data to learn about trading strategies, it is necessary to address the incomplete observation problem. The second issue is to the assumption of deterministic policy vs. non-deterministic policy. Yang *et al.* (2012) make a deterministic policy assumption. Under the linear feature optimization framework,

non-deterministic policies can be represented by a single maximum deterministic policy. In this study, we relax the deterministic policy assumption and allow non-deterministic policies under a Gaussian process framework. As we argue earlier, Gaussian process learning allows us to infer policies even when we have a very limited set of observations. At the same time, we incorporate Gaussian preference learning into our inference process. This approach helps us to incorporate less frequently observed policies into our reward learning process. Together, the proposed GPIRL approach results in a model that relies less on observations and makes fewer assumptions on the polices we are to learn.

# B. Multi-class SVM Trader Classifier using GPIRL vs. LNIRL

This section uses the support vector machine (SVM) classification method to identify traders based on reward functions that we recover from the observations of the trader's behaviors. We select a group of traders whose behaviors are consistently observed during the period we study. The primary reason for choosing the SVM classification method is its flexibility that allows us to explore feature separation in different high dimensional spaces using kernel functions. We aim to compare the performance of the two behavior learning algorithms LNIRL and GPIRL, and to show that GPIRL perform better in real world trading strategy identification.

We constructed 80 sample trajectories for each of the top 10 trading accounts. While there are 121 trading accounts consistently traded over the 4-week period, this study focuses on the top 10 trading accounts. We apply both the LNIRL (Ng *et al.* (2000) and Yang *et al.* (2012)), and GPIRL (Qiao and Beling (2011)) to these 800 samples. And then we apply the SVM algorithm to the 10 traders using pair-wise classification. For each pair, we first train a SVM classifier (with Gaussian kernel) with 60 randomly selected samples, and test the classification on the remaining 20 samples. We repeat the sampling 100 times and then take the average classification accuracy. We list both LNIRL classification results in Table III, and GPIRL results in Table IV. On average, LNIRL gives a classification accuracy of 0.6039, while GPIRL achieves a classification accuracy of 0.9650. This result confirms our earlier assumption that GPIRL performs better when we have incomplete observations, and incorporate non-deterministic policies through Gaussian preference learning.
|  | [,1]   | [,2]   | [,3]   | [,4]   | [,5]   | [,6]   | [,7]   | [,8]   | [,9]   | [,10]  |
|--|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|
| [1,]   | 0.0000 | 0.5437 | 0.5187 | 0.4812 | 0.6375 | 0.4812 | 0.5312 | 0.5750 | 0.7750 | 0.5937 |
| [2,]   | 0.5437 | 0.0000 | 0.5250 | 0.5125 | 0.7437 | 0.5562 | 0.4937 | 0.4250 | 0.7625 | 0.6812 |
| [3,]   | 0.5187 | 0.5250 | 0.0000 | 0.4687 | 0.6875 | 0.5250 | 0.5187 | 0.5250 | 0.7312 | 0.6250 |
| [4,]   | 0.4812 | 0.5125 | 0.4687 | 0.0000 | 0.6937 | 0.5000 | 0.4937 | 0.5062 | 0.6562 | 0.6625 |
| [5,]   | 0.6375 | 0.7437 | 0.6875 | 0.6937 | 0.0000 | 0.6625 | 0.7375 | 0.6875 | 0.7750 | 0.5437 |
| [6,]   | 0.4812 | 0.5562 | 0.5250 | 0.5000 | 0.6625 | 0.0000 | 0.5500 | 0.5500 | 0.6500 | 0.6375 |
| [7,]   | 0.5312 | 0.4937 | 0.5187 | 0.4937 | 0.7375 | 0.5500 | 0.0000 | 0.4937 | 0.8000 | 0.6125 |
| [8,]   | 0.5750 | 0.4250 | 0.5250 | 0.5062 | 0.6875 | 0.5500 | 0.4937 | 0.0000 | 0.6437 | 0.6562 |
| [9,]   | 0.7750 | 0.7625 | 0.7312 | 0.6562 | 0.7750 | 0.6500 | 0.8000 | 0.6437 | 0.0000 | 0.7437 |
| [10,]  | 0.5937 | 0.6812 | 0.6250 | 0.6625 | 0.5437 | 0.6375 | 0.6125 | 0.6562 | 0.7437 | 0.0000 |
| Notes: The columns and rows of this table represent anonymous traders. |        |        |        |        |        |        |        |        |        |        |

TABLE III

#### PAIR-WISE TRADER CLASSIFICATION USING SVM BINARY CLASSIFICATION USING LNIRL

|  | [,1]   | [,2]   | [,3]   | [,4]   | [,5]   | [,6]   | [,7]   | [,8]   | [,9]   | [,10]  |
|--|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|
| [1,]   | 0.0000 | 1.0000 | 0.9875 | 0.9750 | 0.9500 | 0.9750 | 0.9625 | 1.0000 | 0.9750 | 1.0000 |
| [2,]   | 1.0000 | 0.0000 | 0.9750 | 0.9375 | 0.9875 | 0.9750 | 0.9625 | 0.9625 | 0.9875 | 1.0000 |
| [3,]   | 0.9875 | 0.9750 | 0.0000 | 0.9750 | 0.9625 | 0.9875 | 1.0000 | 0.9750 | 0.9750 | 0.9875 |
| [4,]   | 0.9750 | 0.9375 | 0.9750 | 0.0000 | 0.9750 | 0.9500 | 0.9375 | 0.9875 | 0.9875 | 0.9750 |
| [5,]   | 0.9500 | 0.9875 | 0.9625 | 0.9750 | 0.0000 | 1.0000 | 1.0000 | 0.9625 | 0.9875 | 1.0000 |
| [6,]   | 0.9750 | 0.9750 | 0.9875 | 0.9500 | 1.0000 | 0.0000 | 0.9625 | 0.8750 | 0.9125 | 0.9750 |
| [7,]   | 0.9625 | 0.9625 | 1.0000 | 0.9375 | 1.0000 | 0.9625 | 0.0000 | 0.8625 | 0.9625 | 0.9875 |
| [8,]   | 1.0000 | 0.9625 | 0.9750 | 0.9875 | 0.9625 | 0.8750 | 0.8625 | 0.0000 | 0.8000 | 1.0000 |
| [9,]   | 0.9750 | 0.9875 | 0.9750 | 0.9875 | 0.9875 | 0.9125 | 0.9625 | 0.8000 | 0.0000 | 0.9625 |
| [10,]  | 1.0000 | 1.0000 | 0.9875 | 0.9750 | 1.0000 | 0.9750 | 0.9875 | 1.0000 | 0.9625 | 0.0000 |
| Notes: The columns and rows of this table represent anonymous traders. |        |        |        |        |        |        |        |        |        |        |

TABLE IV

#### PAIR-WISE TRADER CLASSIFICATION USING SVM BINARY CLASSIFICATION USING GPIRL

#### C. Trading Strategy Clustering and Comparison with the Summary Statistic-Based Approach

Next, we will show that our IRL based behavior identification approach is far superior to the statistic-based approach. We will use the top 10 trading accounts as examples to demonstrate improvement of behavior-based trading strategy identification achieved using the Gaussian Preference IRL model.

In the previous section, we discovered that using reward functions we can reliably identify a particular trading strategy over a period of time with a relatively high accuracy. In this section, we want to study the similarity of reward characterization among the different trading strategies. This problem can be characterized as an unstructured learning problem - clustering. We have the characterization of rewards over the state space and action space, and we aim to group trading strategies based on their similarity over the Cartesian product of the state and action spaces. We also attempt to establish connections between these trading strategy classification definitions established by Kirilenko et al. (Kirilenko *et al.* (2011)) and our behavior-based trading strategy clustering.

The first problem we have to address is the dimensionality of the feature space. We essentially have a reward structure over a large set of feature sets. This feature set is a product of the state space and the action space in our computational model. Fortunately, under the LNIRL algorithm, we reduce the feature space to only the state space because in this linear feature expectation optimization problem we only consider reward at a particular state. Under the deterministic policy assumption, we assume that the value function converges at a particular state. In other words, the reward function is not a function of actions. In this case, we have 243 features that must be considered during the clustering. However, under the GPIRL framework we do not assume deterministic policy, and we treat reward as a function of both states and actions. Therefore we have 243x30 features for the latter approach. We also observe that the reward matrix is relatively sparse where there are zero values at many states. To consider computational tractability and efficiency, we first examine the data structure through Principal Component Analysis.

In the LNIRL case, the first two Principal Components (PCs) explain 79.78% of the data variation, and from the upper left plot in Figure (6) (a) we see that the first 200 PCs provide nearly 100% explanatory power. In the GPIRL case, the first two PCs only explain 38.98% of the data variation. Looking at the upper left plot in Figure (6) (b), we see that more PCs are

needed to have better represent the data. To balance the accuracy and computational efficiency, we choose the first 200 PCs for the LNIRL and the first 400 PCs for the GPIRL case. This reduction leads to significant gain in computational efficiency and sacrifices less than 2% data variation (lower left figure in both Figure (6) (a) and (b)). From the upper right plots in both the LNIRL and the GPIRL spaces, we see that the first two PCs give a good representation along the first PC and that in the LNIRL case, the feature vector representation is evenly distributed between the first two PCs. The LNIRL space includes distinctly separated observations. On the other hand, the GPIRL space contains concentrations of observations, but unclear boundaries. In both cases, we would expect the PC dimension reduction technique to achieve relatively good representation of the data variation.

Now we apply unsupervised learning method to group the trading behavior observed on a selected group of trading accounts over the observation period. We select 10 trading accounts with the highest average daily trading volume over a period of 4 weeks (20 days) in our first experiment. We define an observation instance as a continuous period covering two hours over which we take all the activities of a particular trader including placing new orders, modifying and canceling existing orders, and placing market orders. For each trader, we collect four observation instances on each trading date: two observation instances during the morning trading and two observation instances during the afternoon trading. The two observation periods in the morning and in the afternoon have an hour overlap time, but the observations in the morning and the afternoon do not overlap. This observation distribution is selected based on the general theory of intraday U-shaped patterns in volume - namely, that trading is heavy at the beginning and the end of the trading day and relatively light in the middle of the day (Ekman et al. (2006), Admati et al. (1988), Lee et al. (2001), and Chordia et al. (2001)). We also examined traders' actions throughout the entire trading day. We found that the two-hour observation time is a good cut-off, and with the overlapping instances in both the morning and the afternoon we expect to capture the U-shaped pattern of the market.

We then perform hierarchical clustering and generate a heat map and dendrogram of the observations in both the LNIRL reward space and the GPIRL reward space. The simplest form of matrix clustering clusters the rows and columns of a dataset using Euclidean distance metric and average linkage. For both Figure (7) (a) and (b), the left dendrogram shows the clustering of the observations (rows), and the top dendrogram shows the clustering of the PCs (columns).





Fig. 6. **Principal Component Representation of the Reward Data:** (a). Data representation under the first two PCs in the LNIRL reward space; The upper left figure shows cumulative percentage of the data variance explained by the PCs. The lower left figure plots the loadings of all the observations onto the first two PCs; The upper right figure shows the projection of the observation and feature vector onto the first two PCs. The lower right shows the projection of the observation onto the first two PCs with boundary point markers. (b). Data representation under the first two Principal Components in the GPIRL reward space. The upper left figure shows the cumulative percentage of the data variance explained by the PCs. The lower left figure plots the loadings of all the observations on to the first two PCs. The upper right figure shows the projection of the observation and feature vectors onto the first two PCs. The upper right figure shows the projection of the observation and feature vectors onto the first two PCs. The upper right figure shows the projection of the observation and feature vectors onto the first two PCs. The lower right figure shows the projection of the observation and feature vectors onto the first two PCs. The lower right figure shows the projection of the observation onto the first two PCs with boundary point markers.

It is evident that there is a clear division of the observations (rows) in both cases. Upon closer examination, the left dendrogram contains two clusters: the top cluster and the bottom cluster with a black dividing strip in the middle of the second small cluster. We then zoomed into the small cluster and look for the sources of these observations<sup>2</sup>. In the LNIRL reward space, we find that the small cluster consists mostly of observations from trader 1 (observations numbered from 1 to 80) and trader 2 (observations numbered from 81 to 160). Observations from trader 9 (observations numbered from 641 to 720) form the black division between these two groups. In the GPIRL reward space, we find that the small cluster consists of three traders: trader 1 (observations numbered from 1 to 80), trader 2 (observations numbered from 81 to 160), and trader 5 (observations numbered from 321 to 400) with the observations from the rest of the traders lying on the other side of the divide. Moreover, we find that the observations from trader 9 (observations numbered from 641 to 720) form the black division between these two groups. These observations show that the majority of the top 10 traders form one group with 2 or 3 traders behaving a little differently. Furthermore, we observe that the clustering has less than perfect purity. In other words, individual observations from the top cluster occasionally lie in the small cluster at the bottom indicating that behavior changes over time. The interpretation of this observation is that the HFTs may behave like Opportunistic Traders for a short period of time. We also occasionally observe Opportunistic Traders behaving like HFTs. In this case, observations cross the divide into the top cluster.

Next we propose a continuous measure of clustering using the hierarchical clustering method. We use the summary statistic-based trader classification method proposed by Kirilenko et al. (Kirilenko *et al.* (2011)) to create reference labels. For this market data, we do not have true labels on those trading accounts. We aim to improve the labeling methods documented by Kirilenko et al. The motivation for creating a continuous measure of clustering is to address the potential changes in trading behavior over time. As we mentioned earlier, we applied the summary statistic-based classification rule on the 200 observations over the 4-week period and found we can only consistently label the traders as a single type 40% of the time. We now define a weighted scoring system to evaluate both the rule-based classifier and the behavior-based classifier. Among

<sup>&</sup>lt;sup>2</sup>Note: In both Figure (6) (a) and (b), we group observations from the same trader together in our data matrix. We have 10 traders and each has 80 observations. From the lower left graph in both (a) and (b), observations are ordered sequentially by trader IDs. For example, observations 1 through 80 come from trader 1, and observations 81 through 160 come from trader 2. This continues along the X-axis up to observations 721 through 800 from trader 10.

the 6 types of traders defined in the data section, we only concerned with labeling HFTs, Intermediaries, and Opportunistic Traders. The other three types of traders, e.g., Fundamental Buyers, Fundamental Sellers, and Small Traders, can be reliably identified by their daily volume and their end of day positions. Here, we assign score 2 if a trader is classified as a HFT; we assign score 1 if a trader is classified as an Opportunistic Trader; and we assign 0 if a trader is classified as an Intermediary. Labels for clustering are assigned using the majority voting rule based on the summary statistic classification rule. We then combine the scores using a weight defined as the frequency with which a particular score is assigned to a particular trader. Here, we want to compare the summary statistic-based trader type classification with the behavior-based trader type classification in an effort to find connections between these two methods.

The visual representations in Figure (8) (a) show that trader 1 and trader 5 have a wide range of end of day positions, but their daily trading volumes remains at relatively the same levels. These traders will likely be classified sometimes as HFTs and sometimes as Opportunistic Traders. While trader 2 exhibits a smaller range of end of day positions than trader 1 and trader 5, the general pattern is very similar to that of traders 1 and 5, and we should classify trader 2 as an Opportunistic Trader. Based on this manual examination, traders 1, 2 and 5 should be classified as Opportunistic Traders and the rest should be classified as HFTs. Now, we compare the results of the summary statistic-based classification rule with those of the behavior-based classification. Figure (9) (a) shows that two groups of traders exist in the LNIRL reward space. Eight out of ten are identified as HFTs and only trader 1 and trader 2 are classified as Opportunistic Traders. This result is consistent with our observation from the dendrogram in Figure (7) (a). When we compare this result with the GPIRL reward space, we can locate all three traders (1, 2, and 5) that we identified through the manual process. This result is also consistent with our observation from the heat map in Figure (7) (b). The statistic-based classification method misclassified trader 2 because the cut-off in the statistic-based approach is based on a simple ratio between the trading volume and the end position. We can see that trader 2 has a relatively small spread of end position. However, the behavior-based approach can identify this pattern and is able to cluster this trader with other traders with similar patterns.

We run another experiment using 10 randomly selected traders out of the traders with the top 30 trading volumes. We know this selection will only result in three types of traders, i.e., HFTs, Intermediaries and Opportunistic Traders. We feed these 800 observations to both LNIRL and

GPIRL algorithms to obtain reward representations of their trading behaviors. Based on visual examination (see Figure (8) (b)), we see that trader 1, trader 2 and trader 7 are Opportunistic Traders and the rest are HFTs. We apply the same techniques as before and we use the same cutoff scores (1.85 in the LNIRL reward space, and 1.75 in the GPIRL reward space). As a result, we can accurately identify the two classes of traders using the same cut-off score we used for the top 10 case (see Figure (10)). The classification in the LNIRL reward space gives the same result as that in the GPIRL reward space, while the statistic-based classification method misclassified trader 3 as an Opportunistic Trader. Based on the daily end position, daily total trading volume and inventory variance, trader 3 should be classified as a HFT. Again, this misclassification is due to the aggregate cut-off ratio. However, the behavior-based approach can identify this pattern and is able to cluster trader 3 with other traders with similar behavioral patterns. Overall, we argue that the GPIRL reward space score-based classification rule provides an advantage over the summary statistic-based approach in that it is based on similarity in behavior and it can be clearly interpreted. Because the GPIRL rule a better reflects traders' choices of actions under different market conditions than the summary statistics, it is well suited for the discovery of new behavioral patterns of the market participants. We also conclude that the GPIRL reward space is more informative and is a superior measure of trading behavior in terms of the LNIRL reward space.

#### VI. CONCLUSION

We assume incomplete observation of algorithmic trading strategies and model traders' reward functions as a Gaussian process. We also incorporate traders' action preferences under different market conditions through preference graph learning. The aim of this study is to quantify trader behavior-based on IRL reward learning under a Bayesian inference framework. We apply both a linear approach (a linear combination of known features) (Abbeel *et al.* (2004)) and GPIRL (Qiao and Beling (2011)) to a real market dataset (The E-Mini S&P 500 Futures), and we conclude that GPIRL is superior to the LNIRL methods, with a 6% greater rate of identification accuracy. Furthermore, we establish a connection between the summary statistic-based classification (Kirilenko *et al.* (2011)) and our behavior-based classification. We propose a score-based method to classify trader types, and because of the transferable property of the reward structure the cut-off score for classifying a group of traders can be applied to different

The implication of this research is that reward/utility-based trading behavior identification can be applied to real market data to accurately identify specific trading strategies. As documented by Abbeel et al. (Abbeel *et al.* (2004)) and confirmed by many other researchers, the reward function is the most succinct, robust, and transferable definition of a control task. Therefore, the behavior learned under the reward space has much broader applicability than observed policies. Furthermore, these learned reward functions will allow us to replicate a particular trading behavior in a different environment to understand their impact on the market price movement and market quality in general.

We also want to note some future research on improving the identification accuracy and discuss applications of this behavioral characterization:

- During our preference learning inference phase, we only considered a simple two layer preference graph. However, traders' preferences can be further distinguished with multi-layer graphs or other preference learning techniques.
- Our study focused on the sampled algorithmic traders on a market. Future studies can extend these results to a large scale experiment to include market participants (specifically Opportunistic traders), and study their behavioral similarity through clustering. We can then associate the group behavior with market quality measures.
- Under the GPIRL framework, we are able to recover a detailed reward structure. These reward functions can be used to generate new policies under a simulated market condition to understand the full behavior of certain trading strategies. This framework provides a particularly interesting way for market regulators to see how the various trading strategies will interact during stressed market conditions.

#### REFERENCES

- Abbeel, P., Coates, A., and Ng, A. Y., Autonomous Helicopter Aerobatics through Apprenticeship Learning. *The International Journal of Robotics Research*, 2010, 1-31.
- Abbeel, P., and Ng, A.Y., Apprenticeship learning via inverse reinforcement learning. *Proceedings* of the twenty-first international conference on Machine learning, 2004.
- Admati, A.R., and Pfleiderer, P., A Theory of Intraday Patterns: Volume and Price Variability. *The Review of Financial Studies*, 1988, **1**, 3–40.

- Aiolli, F., and Sperduti, A., Learning Preferences for Multiclass Problems. Advances in Neural Information Processing Systems 17, 2004, 17–24.
- Aldridge, I., A Practical Guide to Algorithmic Strategies and Trading Systems High Frequency Trading. *John Wiley & Sons, Inc,* 2010, 339.

Bellman, R., Dynamic programming. Princeton University Press, 1957.

- Benos, E., and Sagade, S. High-frequency trading behavior and its impact on market quality: evidence from the UK equity market *Bank of England Working Paper No.* 469 2012.
- Bertsekas, D.P., Neuro-Dynamic Programming. Athena Scientific, 2007.
- Biais, B., Foucault, T., and Moinas, S. Equilibrium high-frequency trading. *SSRN Working Paper* 2012.
- Boehmer, E., Fong, K., and Wu, J. International evidence on algorithmic trading. *EDHEC Working Paper* 2012.
- Boularias, A., and Chaib, B., Bootstrapping Apprenticeship Learning. Advances in Neural Information Processing Systems 24, 2010.
- Brogaard, J.A. High Frequency Trading and its impact on market quality. *Ph.D. Dissertation, Kellogg School of Management, Northwestern University,* 2010.
- Brogaard, J. High frequency trading and market quality SSRN Working Paper 2012.
- Chordia, T., Roll, R., and Subrahmanyam, A., Market Liquidity and Trading Activity, *Journal* of Finance, 2001, 56, 501–530.
- Daw, N.D., Niv, Y., and Dayan, P., Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, 2005, **8**, 1704–1711.
- Dekel, O., Manning, C.D., and Singer, Y., Log-Linear Models for Label Ranking. 21st International Conference on Machine Learning, 2004.
- Dvijotham, K., and Todorov, E., Inverse Optimal Control with Linearly-solvable MDPs. *Proc.* 27th International Conf. on Machine learning, 2010 (ACM).
- Easley, D., Lopez, M.M., and O'Hara, M., The microstructure of the "Flash Crash". *Cornell* University Working Paper, 2010, 6, 7–14.
- Sharpe, W.F., Mutual Fund Performance. Journal of Business, 1966, 39, 119-138.
- Ekman, P., Intraday patterns in the S&P 500 index futures market. *The Journal of Futures Markets*, 2004, **12**, 365–381.

Furnkranz, J., and Hullermeier, E., Preference Learning. Kunstliche Intelligenz, 2005.

- Gabaix, X., Gopikrishnan, P., Plerou, V. and Stanley, H.E., A theory of power-law distributions in financial market fluctuations. *Nature*, 2003, **423**, 267–270.
- Gai, J., Yao, C., and Ye, M., The externalities of high-frequency trading, *University of Illinois Working Paper* 2012.
- Gatheral, J., No-dynamic-arbitrage and market impact. *Quantitative Finance*, 2010, 10, 749–759.
- Hasbrouck, J., Measuring the Information Content of Stock Trades. *The Journal of Finance*, 1991, **46**, 179–207.
- Hasbrouchk, J., and Seppi, D.J., Common factors in prices, order flows and liquidity. *Journal of Financial Economics*, 2001, **59**, 383–411.
- Hayes, R., Paddrik, M. E., Todd, A., Yang, S. Y., Scherer, W., and Beling, P. Agent Based Model of the E-MINI Future Market: Applied to Policy Decisions. *Proceedings of 2012 Winter Simulation Conference, Berlin, Germany* 2012.
- Hendershott, T., Jones C.M., and Menkveld, A.J., Does Algorithmic Trading Improve Liquidity?. *Journal of Finance*, 2004, **66**, 1–33.
- Hendershott, T., and Riordan, R. Algorithmic trading and the market for liquidity. *Journal of Financial and Quantitative Analysis,* forthcoming 2013.
- Hida, T., and Hitsuda, M., Gaussian Processes. 1993.
- Jovanovic, B., and Menkveld, A.J., Middlemen in Limit-Order Markets. NYU Working Paper, New York, 2010.
- Jones, C.M., Kaul, G. and Lipson, M.L., Transactions, volume, and volatility. *Review of Financial Studies*, 1994, 7, 631–651.
- Jones, C.M., What do we know about high-frequency trading? SSRN Working Paper 2012.
- Kirilenko, A., Kyle, A.S., Samadi, M., and Tuzun, T., The Flash Crash: The Impact of High Frequency Trading on an Electronic Market, *SSR Working Paper Series*, 2011.
- Kolmogorov, A.N., Foundations of the Theory of Probability. 1956 (AMS Chelsea).
- Lee, S.J., and Zoran, P., Learning behavior styles with inverse reinforcement learning. ACM SIGGRAPH 2010 papers, 2010, 1–7.
- Lee, Y. T., Fox, C. R., and Liu, Y., Explaining Intraday Pattern of Trading Volume from the Order Flow Data. *Journal of Business Finance and Accounting*, 2001, **28**, 306–686.
- Levine, S., Popovic, Z., and Koltun, V., Feature Construction for Inverse Reinforcement Learning.

Advances in Neural Information Processing Systems 24, 2010, 6, 7–14.

- Martinez, V.H., and Rosu, I., High frequency traders, news and volatility. SSRN Working Paper, 2012.
- Menkveld, A., High frequency trading and the new market makers. SSRN Working Paper 2012.
- Moallemi, C.C., and Saglam, M., The cost of latency. SSRN Working Paper, 2012.
- Neu, G., and Szepesvari, C., Apprenticeship learning using inverse reinforcement learning and gradient methods. *Proc. Uncertainty in Artificial Intelligence*, 2007.
- Ng, A.Y., and Russel, S., Algorithms for inverse reinforcement learning. *In Proc. ICML*, 2000, 663–670.
- Paddrik, M. E., Hayes, R., Todd, A., Yang, S. Y., Scherer, W., and Beling, P, An Agent Based Model of the E-Mini S&P 500 and the Flash Crash. *Proceedings of IEEE Computational Intelligence in Financial Engineering and Economics*, 2012, 2012.
- Pagnotta, E., and Philippon, T., Competing on speed. SSRN Working Paper, 2012.
- Qiao, Q., and Beling, P., Inverse Reinforcement Learning with Gaussian Process. *Proceedings* of 2011 American Control Conference, 2011.
- Ramachandran, D., and Amir, E., Bayesian Inverse Reinforcement Learning. *In Proc. IJCAI*, 2007, 2586-2591.
- Rasmussen, C.E., and Williams, C.K., Gaussian Processes for Machine Learning. *MIT Press*, 2006.
- Ratliff, N., Ziebart, B., Peterson, K., Bagnell, J.A., Hebert, M., Dey, A.K., and Srinivasa, S., Inverse Optimal Heuristic Control for Imitation Learning. *Proc. AISTATS*, 2009, 424–431.
- Riordan, R., and Storkenmaier, A., Latency, liquidity and price discovery. *Journal of Financial Market*, forthcoming 2013.
- Seeger, M., Gaussian Processes for Machine Learning. *International Journal of Neural Systems*, 2004, **14**, 2004–2016.
- Sutton, R.S., and Barto, A.G., Reinforcement Learning: An Introduction. The MIT Press, 1998.
- Syed, U., and Schapire, R.E., A game-theoretic approach to apprenticeship learning. *In Advances in Neural Information Processing Systems*, 2008, 1449–1456.
- Syed, U., Bowling, M., and Schapire, R.E., Apprenticeship learning using linear programming. *Proc. 25th international Conf. on Machine learning*, 2008, 1032–1039.
- Wei, W., and Zoubin, G., Preference learning with Gaussian processes. Proc. 22th International

Conf. on Machine learning, 2005, 137-144.

- Yang, S.Y., Paddrik, M.E., Hayes, R.J., Todd, A., Kirilenko, A.A., Beling, P., and Scherer, W., Behavior Based Learning in Identifying High Frequency Trading Strategies. *Proceedings of IEEE Computational Intelligence in Financial Engineering and Economics*, 2012, 2012.
- Ziebart, B.D., Mass, A., Bagnell, J.A., and Dey, A.K., Maximum Entropy Inverse Reinforcement Learning. *In Proceedings of the Twenty-Third AAAI on Artifical Intelligence*, 2008, 1433-1438.



(a)

(b)

Fig. 7. **Hierarchical Clustering of Data Matrix:** (a). Heat map of 800 observations of the LNIRL Rewards in the first 200 PCs. (b). Heat map of 800 observations in GPIRL Rewards in the first 400 PCs.



Fig. 8. Top 10 Traders: Trader's Daily Trading Volume vs. Daily End Position during a 20 Day Period. (a) Randomly Selected 10 Traders: Traders 1, 2 and 5 have varying end positions. (b) Traders 1, 2, and 7 have varying end positions.



Fig. 9. Trader Type Classification Compared with the Summary Statistic Based Rule Classification for the Top 10 Traders. (a) Hierarchical clustering in the LNIRL reward space. (b) Hierarchical clustering in the GPIRL reward space.



Fig. 10. Trader Type Classification Compared with the Summary Statistic-Based Rule Classification for the 10 Randomly Selected Traders: (a) Hierarchical clustering in LNIRL reward space. (b) Hierarchical clustering in GPIRL reward space.

# The Flash Crash: High-Frequency Trading in an Electronic Market

#### ANDREI KIRILENKO, ALBERT S. KYLE, MEHRDAD SAMADI, and TUGKAN TUZUN\*

#### ABSTRACT

We study intraday market intermediation in an electronic market before and during a period of large and temporary selling pressure. On May 6, 2010, U.S. financial markets experienced a systemic intraday event – the Flash Crash – where a large automated selling program was rapidly executed in the E-mini S&P 500 stock index futures market. Using audit trail transaction-level data for the E-mini on May 6 and the previous three days, we find that the trading pattern of the most active nondesignated intraday intermediaries (classified as High Frequency Traders) did not change when prices fell during the Flash Crash.

<sup>\*</sup>Kirilenko is with Imperial College London, Kyle is with the University of Maryland, Samadi is with Southern Methodist University, and Tuzun is with the Federal Reserve Board of Governors. We thank Robert Engle, Chester Spatt, Larry Harris, Cam Harvey, Bruno Biais, Simon Gervais, participants at the Western Finance Association Meeting, NBER Market Microstructure Meeting, Centre for Economic Policy Research Meeting, Q-Group Seminar, Wharton Conference in Honor of Marshall Blume, Princeton University Quant Trading Conference, University of Chicago Conference on Market Microstructure and High-Frequency Data, NYU-Courant Mathematical Finance Seminar, Columbia Conference on Quantitative Trading and Asset Management, and seminar participants at Columbia University, MIT, Boston University, Brandeis University, Boston College, UMass-Amherst, Oxford University, Cambridge University, the University of Maryland, Bank for International Settlements, Commodity Futures Trading Commission, Federal Reserve Board, and the International Monetary Fund, among others. The research presented in this paper was coauthored by Andrei Kirilenko, a former full-time CFTC employee, Albert Kyle, a former CFTC contractor who performed work under CFTC OCE contract (CFCE-09-CO-0147), Mehrdad Samadi, a former full-time CFTC employee and former CFTC contractor who performed work under CFTC OCE contracts (CFCE-11-CO-0122 and CFCE-13-CO-0061), and Tugkan Tuzun, a former CFTC contractor who performed work under CFTC OCE contract (CFCE-10-CO-0175). The Office of the Chief Economist and CFTC economists produce original research on a broad range of topics relevant to the CFTC's mandate to regulate commodity futures markets and commodity options markets, and its expanded mandate to regulate the swaps markets pursuant to the Dodd-Frank Wall Street Reform and Consumer Protection Act. The analyses and conclusions expressed in this paper are those of the authors and do not reflect the views of the Federal Reserve System, the members of the Office of the Chief Economist, other CFTC staff, or the CFTC itself. The Appendix can be found in the online version of the article on the Journal of Finance website.

On May 6, 2010, U.S. financial markets experienced a systemic intraday event known as the "Flash Crash." The CFTC-SEC (2010b) joint report describes the Flash Crash as follows:

"At 2:32 [CT] p.m., against [a] backdrop of unusually high volatility and thinning liquidity, a large fundamental trader (a mutual fund complex) initiated a sell program to sell a total of 75,000 E-mini [S&P 500 futures] contracts (valued at approximately \$4.1 billion) as a hedge to an existing equity position.  $[\ldots]$  This large fundamental trader chose to execute this sell program via an automated execution algorithm ("Sell Algorithm") that was programmed to feed orders into the June 2010 E-mini market to target an execution rate set to 9% of the trading volume calculated over the previous minute, but without regard to price or time. The execution of this sell program resulted in the largest net change in daily position of any trader in the E-mini since the beginning of the year (from January 1, 2010 through May 6, 2010). [...] This sell pressure was initially absorbed by: high frequency traders ("HFTs") and other intermediaries in the futures market; fundamental buyers in the futures market; and cross-market arbitrageurs who transferred this sell pressure to the equities markets by opportunistically buying E-mini contracts and simultaneously selling products like [the] SPY [(S&P 500 exchange-traded fund ("ETF"))], or selling individual equities in the S&P 500 Index. [...] Between 2:32 p.m. and 2:45 p.m., as prices of the E-mini rapidly declined, the Sell Algorithm sold about 35,000 Emini contracts (valued at approximately \$1.9 billion) of the 75,000 intended. [...] By 2:45:28 there were less than 1,050 contracts of buy-side resting orders in the E-mini, representing less than 1% of buy-side market depth observed at the beginning of the day. [...] At 2:45:28 p.m., trading on the E-mini was paused for five seconds when the Chicago Mercantile Exchange ("CME") Stop Logic Functionality was triggered in order to prevent a cascade of further price declines.<sup>1</sup> [...] When trading resumed at 2:45:33 p.m., prices stabilized and shortly thereafter, the E-mini began to recover, followed by the SPY. [...] Even though after 2:45 p.m. prices in the E-mini and SPY were recovering from their severe declines, sell orders placed for some individual securities and ETFs (including many retail stop-loss orders, triggered by declines in prices of those securities) found reduced

<sup>&</sup>lt;sup>1</sup>The CME's Globex Stop Logic Functionality is an automated pre-trade safeguard procedure designed to prevent the execution of cascading stop orders that would cause "excessive" declines or increases in prices due to lack of sufficient depth in the central limit order book. In the context of this functionality, "excessive" is defined as being outside of a predetermined "no bust" range. The no bust range varies from contract to contract; for the E-mini, it was set at 6 index points (24 ticks) in either direction. After Stop Logic Functionality is triggered, trading is paused for a certain period of time as the matching engine goes into what is called a "reserve state." The length of the trading pause varies between 5 and 20 seconds from contract to contract; it was set at 5 seconds for the E-mini. During the reserve state, orders can be submitted, modified, or cancelled, but no executions can take place. The matching engine exits the reserve state by initiating the same auction opening procedure as it does at the beginning of each trading day. After the starting price is determined by the re-opening auction, the matching engine returns to the standard continuous matching protocol.

buying interest, which led to further price declines in those securities. [...] [B]etween 2:40 p.m. and 3:00 p.m., over 20,000 trades (many based on retail-customer orders) across more than 300 separate securities, including many ETFs, were executed at prices 60% or more away from their 2:40 p.m. prices. [...] By 3:08 p.m., [...] the E-mini prices [were] back to nearly their pre-drop level [... and] most securities had reverted back to trading at prices reflecting true consensus values."

To illustrate the large and temporary decline in prices and the corresponding increase in trading volume on May 6, Figure 1 presents end-of-minute transaction prices (solid line) and minute-by-minute trading volume (dashed line) in the E-mini on May 6.

#### <Insert Figure 1>

The accumulation of the largest daily net short position of the year by a single trader over a matter of minutes can be thought of as a period of large and temporary selling pressure. Theory suggests that a period of large and temporary selling pressure can trigger a market crash even in the absence of a fundamental shock. Building on the Grossman and Miller (1988) framework, Huang and Wang (2008) develop an equilibrium model that links the cost of maintaining continuous market presence with market crashes even in the absence of fundamental shocks and with perfectly offsetting idiosyncratic shocks. In their model, market crashes emerge endogenously when a sudden excess of sell orders overwhelms the insufficient risk-bearing capacity of market makers. Because the provision of continuous market presence is costly, market makers choose to maintain equilibrium risk exposures that are too low to offset large but temporary liquidity imbalances. In the event of a large enough sell order, the liquidity on the buy side can only be obtained after a price drop that is large enough to compensate increasingly reluctant market makers for taking on additional risky inventory.

Weill (2007) presents an equilibrium model of optimal dynamic inventory adjustment of competitive capital-constrained intermediaries faced with large and temporary selling pressure. This framework begins with an exogenous negative aggregate shock to outside investors' marginal utility of holding the asset, which leads to a sharp price drop. During and immediately following the price drop, there is no change in intermediaries' inventories. As intermediaries anticipate that the marginal utilities of some outside investors' will begin to increase and the selling pressure will subside, they find it optimal to dynamically accumulate a long position, during which time market prices rise. They then unwind their inventory just as market prices reach their initial level. As shown in Figure 1 of Weill (2007), the co-movement between intermediary inventories and prices varies over time, suggesting that this relationship is dynamic. More generally, Nagel (2012) shows that return reversals are related to the risk-bearing capacity of intermediaries.

Intermediation is an essential function in markets in which buyers and sellers do not arrive simultaneously. As technology has transformed the way financial assets are traded, intermediation has been increasingly provided by market participants without formal obligations. An important question is how nondesignated intraday intermediaries behave during periods of large and temporary buying or selling pressure in automated financial markets.

In this paper, we empirically examine intraday market intermediation in an electronic market before and during a period of large and temporary selling pressure.<sup>2</sup> We use audit trail account-level transaction data in the E-mini S&P 500 stock index futures

<sup>&</sup>lt;sup>2</sup>We use the term intraday intermediation instead of market making or liquidity provision because the two latter terms have become associated with affirmative obligations to provide two-sided quotes, serve a customer base, and maintain "fair and orderly markets." Market making has also been formally recognized in a plethora of government regulations, regulations by self-regulatory organizations, and court decisions. Intraday intermediation, in contrast, does not necessarily entail designated market making or mandatory liquidity provision. Intraday intermediation can be provided by not only designated market makers, but also by proprietary traders trading exclusively for their own trading accounts without acting in any agency capacity such as, for example, routing customer order flow or providing customer advice (see Committee on the Global Financial System (2014)). The term intraday intermediation is also distinct from the notion of financial intermediation, which refers to the process of asset transformation "by purchasing assets and selling liabilities" (see Madhavan (2000)).

market over the period May 3 through 6, 2010.<sup>3</sup> Guided by the literature on inventory management by intermediaries (see O'Hara (1995) and Hasbrouck (2006), among others), we classify trading accounts that do not accumulate large directional positions and whose inventories display mean-reversion during May 3 through 5 as intraday intermediaries. If an account is classified as an intermediary on any of these three days, we keep it in the same category on May 6, 2010. Importantly, this approach does not require that an intermediary maintain low inventory on the day of the Flash Crash. We further separate intraday intermediaries into High Frequency Traders and Market Makers.<sup>4</sup> As their category name suggests, High Frequency Traders participate in a markedly larger proportion of trading than Market Makers.<sup>5</sup>

Theory suggests that intermediaries optimally adjust inventory in relation to falling prices. If the intermediaries' risk-bearing capacity is overwhelmed, they become unwilling to accumulate more inventory without large price concessions. Consistent with the theory of limited risk-bearing capacity of intermediaries, the combined net inventories of the accounts classified as intraday intermediaries over the four days of our sample, including May 6, did not exceed 6,000 E-mini contracts – a sum that is an order of magnitude smaller than the large sell program of 75,000 contracts documented in CFTC-SEC (2010b). In contrast to Weill (2007), during the period of large and temporary selling pressure on May 6, we find that both categories of intraday intermediaries also accumulate net long inventory positions as prices decline.

To examine the dynamic risk-bearing capacity of intermediaries before and during

<sup>&</sup>lt;sup>3</sup>The CFTC-SEC report's narrative of the Flash Crash in the E-mini was based in part on the preliminary analysis contained in the original version of this paper (see footnote 22 of CFTC-SEC (2010b)).

<sup>&</sup>lt;sup>4</sup>Throughout the paper we employ the following convention: we use upper case letters whenever we refer to the categories that we define, e.g., Market Makers and High Frequency Traders and lower case letters whenever we refer to general type of activity, e.g., market making and high frequency trading.

<sup>&</sup>lt;sup>5</sup>Accounts classified as High Frequency Traders based on inventory and volume patterns might be representative of a subset of all high frequency trading strategies.

the Flash Crash, we empirically study the second-by-second co-movement of their inventory changes and price changes over May 3 through 6. We find that inventory changes of High Frequency Traders exhibit a statistically significant relationship with both contemporaneous and lagged price changes and that this relationship did not change when prices fell during the Flash Crash. However, the statistical relationship between Market Maker inventory changes and price changes did change during the Flash Crash compared with the previous three days.

Moreover, we find that inventory changes of Market Makers are negatively related to contemporaneous price changes, consistent with theories of traditional market making (see Hendershott and Seasholes (2007), among others). In contrast, inventory changes of High Frequency Traders are positively related to contemporaneous price changes. Foucault, Roell, and Sandas (2003), Menkveld and Zoican (2016), and Budish, Cramton, and Shim (2015) provide theoretical mechanisms through which the inventories of intermediaries may positively co-move with price changes at high frequencies. These studies suggest that if certain traders can react marginally faster to a signal, they can adversely select stale quotes of marginally slower market makers, engaging in "stale quote sniping" or "latency arbitrage." Consequently, faster traders are able to trade ahead of price changes at short time horizons.

Consistent with the theory of "stale quote sniping," we find that over May 3 through 5, when High Frequency Traders are net buyers in a given second, prices increase in the following second and remain higher over the subsequent 20 seconds. We examine the extent to which High Frequency Traders' trading activity precedes price changes and find that High Frequency Traders lift a disproportionate amount of the final best ask depth before an increase in the best ask level and provide a disproportionate proportion of depth first transacted against at the new price level.

Our main contribution is empirically studying theories of intermediation during a pe-

riod of large and temporary selling pressure. The closest studies to ours are Brogaard, Hendershott, and Riordan (2016), who study high frequency traders as classified by NAS-DAQ during the 2008 short-sale ban and Brogaard et al. (2016), who study the activity of high frequency traders as classified by NASDAQ around extreme price movements.<sup>6</sup> In contrast, we focus on trading during the Flash Crash in the inclusive, centralized Emini market with individual account IDs and use the entire universe of trading accounts. Our analysis makes use of a detailed and comprehensive set of transaction-level data in the E-mini three days before and on the day of the Flash Crash. Focusing on trading in the E-mini during the Flash Crash provides two additional advantages. Unlike the U.S. equity markets, there are no market maker obligations in the fully electronic E-mini. Thus, a focus on trading in the E-mini during the Flash Crash may help us understand the potential implications of not imposing market making obligations as markets become more automated, especially during periods of market stress. Furthermore, all of the trading in the E-mini takes place in one venue. Consequently, our results are not affected by the fragmentation of trading, and we are able to study the entire universe of

<sup>&</sup>lt;sup>6</sup>Since the release of CFTC-SEC (2010b), a number of studies have examined the Flash Crash. For example, Madhavan (2012) studies the propagation of the Flash Crash to ETFs where trades were disproportionately broken and finds that ETFs that traded at stub quote price levels were characterized by a relatively high degree of trading fragmentation. Menkveld and Yueshen (2016) study the trading of the large sell program during the Flash Crash and argue that the arbitrage relationship between the E-mini and the S&P 500 ETF (SPY) may have broken down during the Flash Crash and subsequent recovery. Easley, Lopez, and O'Hara (2011) apply the Volume Synchronized Probability of Informed Trading (VPIN) measure to the day of the Flash Crash and find abnormal levels of "order-flow toxicity" in the hours leading up to the crash. Market data vendor and commentator Nanex also analyzes trading during the Flash Crash and argues that the large fundamental seller never submitted marketable orders. In contrast, Menkveld and Yueshen (2016) document that "half of the sell orders were limit orders, the other half market orders." While these studies contribute to our overall understanding of how the Flash Crash became a systemic financial marketwide event, we focus on the trading of intraday intermediaries in the stock index futures market, where, according to the CFTC–SEC (2010b) report, the triggering event occurred.

trading of a given account in the E-mini June 2010 contract.<sup>7</sup>

The rest of the paper proceeds as follows. In Section I, we discuss the market structure of the E-mini and the data employed in this paper. In Section II, we present our empirical methodology and results. In Section III, we conclude.

# I. Institutional Background and Data

# A. The E-mini S&P 500 Futures Market

The CME introduced the E-mini contract in 1997. The E-mini owes its name to the fact that it is traded electronically and in denominations five times smaller than the original S&P 500 futures contract. Since its introduction, the E-mini has become a popular instrument to hedge exposures to baskets of U.S. stocks or to speculate on the direction of the entire stock market. The E-mini contract attracts the highest dollar volume among U.S. equity index products (futures, options, or exchange-traded funds). Hasbrouck (2003) shows that of all U.S. equity index products, the E-mini contributes the most to the price discovery of the U.S. stock market. The contracts are cashsettled against the value of the underlying S&P 500 equity index at expiration dates in March, June, September, and December of each year. The contract with the nearest expiration date, which attracts the majority of trading activity, is called the "frontmonth" contract. In May 2010, the front-month contract was the contract expiring in

<sup>&</sup>lt;sup>7</sup>A number of studies have analyzed the behavior of high frequency traders as classified by NASDAQ using data from NASDAQ exchanges only (see Brogaard, Hendershott, and Riordan (2014, 2016), Carrion (2013), Hirschey (2016) and Brogaard et al. (2016), inter alia). However, as of the end of Q3, 2010, trading on NASDAQ exchanges represented approximately a third of Tape C (the tape for NASDAQ stocks) trading volume. Our approach also differs from studies that attempt to infer the behavior of high frequency traders from aggregate market data (see Hendershott, Jones, and Menkveld (2011), Hasbrouck and Saar (2013), and Conrad, Wahal, and Xiang (2015), inter alia). We are also able to study the trading of all accounts active in the E-mini rather than the trading of one high frequency trader or institutional investor (see Menkveld (2013) and Menkveld, and Yueshen (2016), respectively).

June 2010. The notional value of one E-mini contract is \$50 multiplied by the S&P 500 stock index. During May 3 - 6, 2010, the S&P 500 index fluctuated slightly above 1,000 points, making each E-mini contract worth about \$50,000. The minimum price increment, or "tick" size, of the E-mini is 0.25 index points, or \$12.50; a price move of one tick represents a fluctuation of about 2.5 basis points. The E-mini trades exclusively on the CME Globex trading platform, a fully electronic limit order market. Trading takes place 24 hours a day with the exception of one 15-minute technical maintenance break each day. The CME Globex matching algorithm for the E-mini follows a "price-time priority" rule in that orders offering more favorable prices are executed ahead of orders with less favorable prices, and orders with the same prices are executed in the order they were received by Globex. The market for the E-mini features both pre- and post-trade transparency. Pre-trade transparency is provided by transmitting to the public in real time the quantities and prices for buy and sell orders resting in the central limit order book up or down 10 tick levels from the last transaction price. Post-trade transparency is provided by transmitting to the public prices and quantities of executed transactions. The identities of individual traders submitting, canceling, or modifying bids and offers, as well as those whose orders have been executed, are not made available to the public.

## B. Data

Our sample consists of intraday audit trail transaction-level data for the E-mini S&P 500 June 2010 futures contract for the sample period spanning May 3 - 6, 2010. These data come from the Trade Capture Report (TCR), which the CME provides to the CFTC.<sup>8</sup> For each of the four days, we examine all regular transactions occurring during

<sup>&</sup>lt;sup>8</sup>Due to the highly confidential nature of these data and commonality across certain trading accounts, we aggregate trading accounts into trader categories. Prior to the release of this paper, all matters related to the aggregation of data, presentation of results, and sharing of the results with the public were reviewed by the CFTC.

the 405-minute period starting at the opening of the market for the underlying stocks at 8:30 a.m. CT (CME Globex is in the Central Time Zone) or 9:30 a.m. ET and ending at the time of the technical maintenance break at 3:15 p.m. CT, 15 minutes after the close of trading in the underlying stocks. For each transaction, we use fields with the account identifiers for the buyer and the seller, the price and quantity transacted, the date and time (to the nearest second), a sequence ID number that sorts trades into chronological order within one second, a field indicating whether the trade resulted from a limit (both marketable and nonmarketable) or market order, an order ID that assigns multiple trade executions to the original order, and an "aggressiveness" indicator stamped by the CME Globex matching engine as "N" for a resting order and "Y" for an order that executed against a resting order. We do not study message-level data and, thus, do not observe activity for orders that did not execute.

# C. Descriptive Statistics

Market-level descriptive statistics are presented in Table I. We report statistics separately for May 3 to 5 and May 6. Statistics in the May 3 to 5 column represent three-day averages.

# <Insert Table I>

Trading volume and the number of trades on May 6 were more than double the average daily trading volume over the previous three days. Volatility measured as the log of the intraday price range was also significantly larger on May  $6.^9$  The average trade size on both May 3 - 5 and May 6 was approximately five contracts. Over 90%

 $<sup>^{9}</sup>$ In the Internet Appendix, we present the daily five-minute realized variance of the SPY for 2004 to 2013 and find that the daily realized variances observed on May 3 - 5 were not abnormal.

of trading and trading volume were executed via limit orders (both marketable and non-marketable).

# II. Methodology and Results

We classify over 15,000 unique accounts trading in the E-mini into intraday intermediaries and other categories of traders to provide an empirical analysis of intraday intermediation before and during the Flash Crash. We then study the behavior of the most active intermediaries defined as High Frequency Traders in more detail.

# A. Trader Categories

Over 15,000 unique accounts traded in the E-mini during our sample period. Traders in the E-mini, including those that buy and sell throughout a trading day, do not have formal designations such as market makers, dealers, or specialists. To classify accounts as intraday intermediaries, we adopt a data-driven approach based on trading activity and inventory patterns. Our definition of intraday intermediaries is designed to capture traders who follow a strategy of consistently buying and selling throughout a trading day while maintaining low levels of inventory.<sup>10</sup>

Market intermediaries can be broadly defined as "traders who can fill gaps arising from imperfect synchronization between the arrivals of buyers and sellers" (see Grossman and Miller (1988)). This definition implies that intermediaries often participate in a significant proportion of transactions (see Glosten and Milgrom (1985) and Kyle (1985)) and that intermediaries' inventories are mean-reverting at a relatively high frequency (see Garman (1976), Amihud and Mendelson (1980), and Ho and Stoll (1983), among

<sup>&</sup>lt;sup>10</sup>We use a broad definition of intermediation to classify accounts as intraday intermediaries that does not use the relationship between intermediary trading and prices or price fluctuations.

others). Empirically, intraday mean-reversion in inventories and relatively high trading volume are salient characteristics of intermediation (see Hasbrouck and Sofianos (1993), and Madhavan and Smidt (1993)). A growing literature on the most active intermediaries variously defines them as fast traders, high frequency traders, or high frequency market makers (see Ait-Sahalia and Saglam (2016), Jovanovic and Menkveld (2016), Biais, Foucault, and Moinas (2015), as well as empirical studies by Menkveld (2013), Brogaard, Hendershott, and Riordan (2014), and Carrion (2013), and a survey by Jones (2013)).

A trader is classified as an intraday intermediary if it holds small intraday and endof-day net positions relative to its daily trading volume over May 3 - 5, 2010, irrespective of its trading behavior on May 6. To be classified as an intraday intermediary, a trader denoted by j must meet criteria (i) with respect to its daily trading volume  $(Vol_{j,d})$ , where d denotes a trading day, (ii) with respect to its end-of-day position  $(NP_{j,d,t=405})$ relative to its daily trading volume, where t denotes each minute during a trading day, and (iii) with respect to its intraday minute-by-minute inventory  $(NP_{j,d,t})$  pattern.

We set the following specific levels for each criterion (to simplify notation, we suppress the subscript j and set beginning-of-day inventories for all trading accounts to zero  $(NP_{j,d,t=0} = 0)$ ):

(i) An account must trade 10 or more contracts on at least one of the three days prior to the Flash Crash (May 3, 4, and 5, 2010).

$$Vol_d \ge 10$$
,

According to the data, this volume cutoff is a conservative way to first remove accounts that do not trade an economically significant amount before categorizing intraday intermediaries.<sup>11</sup>

(ii) The three-day average of the absolute value of the ratio of the account's end-ofday net position to its daily trading volume must not exceed 5%.

$$\frac{\sum_{d=1}^{3} \frac{|NP_{d,t=405}|}{Vol_d}}{3} \le 5\%$$

Specifically, we compute the daily ratio of a trader's end-of-day position to its daily trading volume on May 3, 4, and 5, compute the absolute value of the ratios for each day, and calculate the three-day average of the absolute values of the ratio.

(iii) The three-day average of the square root of the account's daily mean of squared end-of-minute net position deviations from its end-of-day net position over its daily trading volume must not exceed 0.5%.

$$\frac{\sum_{d=1}^{3} \sqrt{\frac{1}{405} \sum_{t=1}^{405} \left(\frac{NP_{d,t} - NP_{d,t=405}}{Vol_d}\right)^2}}{3} \le 0.5\%.$$

These cutoff levels are specific to our sample and may need to be adjusted if applied in other markets.<sup>12</sup>

Of the accounts that are classified as intraday intermediaries, we further classify the 16 most active accounts, that is, those with the highest number of trades over May 3 -

<sup>&</sup>lt;sup>11</sup>In setting the volume cutoff, there is a tradeoff. On the one hand, the number of contracts traded needs to be large enough to ensure that economically small traders are not mistakenly categorized as intraday intermediaries, but not so high that accounts characterized by consistent buying and selling are mistakenly categorized as Small Traders. Using a back-of-the-envelope approximation from Table II, the average number of contracts traded per day by an average Small Trader is 1.98 ((2,397,639 × 0.005)/6,065 ≈ 1.98). The corresponding approximation for intraday intermediaries is 5,255 contracts ((2,397,639 × 0.4471)/204 ≈ 5,255). There is a significant difference between these different types of categories in the data. However, rather than making the volume cutoff larger, we apply two additional criteria that also link to the theory and empirical evidence of intermediation.

<sup>&</sup>lt;sup>12</sup>Kirilenko, Mankad, and Michailidis (2013) confirm the qualitative intuition of our classification using a dynamic unsupervised machine learning method that does not rely on user-specified cutoffs.

5, as High Frequency Traders.<sup>13</sup> The other intraday intermediary accounts are classified as Market Makers. A High Frequency Trader is thus similar to a Market Maker in all respects, except that High Frequency Traders participate in a significantly greater number of trades.<sup>14</sup> If an account is classified as a High Frequency Trader or a Market Maker over May 3 - 5, 2010, it remains in the same category for May 6, 2010, as well. As previously mentioned, this restriction does not require that a High Frequency Trader or a Marker Maker maintain low inventory relative to volume on the day of the Flash Crash.<sup>15</sup>

We classify all other traders as Small Traders, Fundamental Buyers, and Fundamental Sellers. We call the remaining accounts Opportunistic Traders. Unlike High Frequency Traders and Market Makers, these trader categories are classified separately for each of the four trading days, including May 6, 2010.<sup>16</sup>

On each day, an account is classified as a Small Trader if it trades fewer than 10 contracts. Over 6,000 out of the 15,000 accounts are classified as Small Traders. The

<sup>&</sup>lt;sup>13</sup>Results are qualitatively similar when we classify the most active accounts based on trading volume. According to Figure 2 below, there is also a large difference in the trading volume between the 16th and 17th ranked intraday intermediaries in terms of daily trading volume.

<sup>&</sup>lt;sup>14</sup>High Frequency Traders trade significantly more frequently than any other trader category, including Market Makers. Over May 3 - 5, 15 High Frequency Traders were active on average. The three-day average of the High Frequency Traders' daily number of trades per second is 5.98. In contrast, over May 3 - 5, 189 Market Makers were active on average and the three-day average of the Market Makers' daily number of trades per second is 2.14. These estimates suggest that on average a High Frequency Trader trades about 30 times more often than a Market Maker. While we do not observe the messages or latency of traders with our data, Clark-Joseph (2014) applies our classification methodology to message-level data and confirms that High Frequency Traders submit messages in the millisecond environment. Hayes et al. (2012) confirm our classification with simulated data calibrated on the E-mini.

<sup>&</sup>lt;sup>15</sup>Sixteen unique accounts were classified as High Frequency Traders over May 3 - 6, of which, 14 of the 16 accounts traded on May 3, all 16 accounts traded on May 4, and 15 of the 16 accounts traded on May 5. No new accounts that satisfy the criteria of High Frequency Traders enter the E-mini on May 6. The accounts classified as High Frequency Traders based on inventory and volume patterns may be representative of a subset of all high frequency trading strategies as defined by the SEC (2014) concept release on market structure.

<sup>&</sup>lt;sup>16</sup>The rationale for classifying Small, Fundamental, or Opportunistic traders separately each day is that they may trade only on one day. It is also possible that the same account can be classified differently on different days. For example, an account can be a Fundamental Buyer on one day, a Small Trader on another day, and a Fundamental Seller or Opportunistic Trader on yet another day.

Small Traders category likely captures retail traders (see Kaniel, Saar, and Titman (2008), and Seasholes and Zhu (2010) among others). Small Traders account for less than 1% of the total trading volume in our sample.

On each day, an account is classified as a Fundamental Buyer if it trades 10 contracts or more and accumulates a net long end-of-day position equal to at least 15% of its total trading volume for the day. Similarly, an account is classified as a Fundamental Seller if it trades 10 contracts or more and the absolute value of its net short position at the end of the day is at least 15% of its total trading volume for the day. This category is meant to capture accounts that accumulate significant directional positions on a given day and most likely reflects trading patterns of institutional investors with longer holding horizons (see Anand et al. (2013), and Puckett and Yan (2011), among others).

The remaining accounts are categorized as Opportunistic Traders. Opportunistic Traders move in and out of positions throughout the day but adjust their net holdings with significantly larger fluctuations and lower frequency than intraday intermediaries. Opportunistic Traders may follow a variety of arbitrage trading strategies, including cross-market arbitrage (for example, long futures/short securities), statistical arbitrage, and news arbitrage (buy if the news indicators are positive/sell if the news indicators are negative). Opportunistic Traders may also engage in providing intermediation across days or weeks rather than intraday.

Our classification methodology is based entirely on directly observed individual inventory and trading volume patterns of traders. Unlike many other markets, traders in our data set do not have designations due to regulatory, reporting, or other mandatory or voluntary disclosure requirements. In that regard, our classification differs from papers that use NASDAQ data, which classify high frequency traders using a variety of qualitative and quantitative criteria, or the approach of Biais, Declerck, and Moinas (2016) which uses a combination of a proprietary/agency flag along with quantitative criteria. Our approach also differs from those that use only qualitative criteria to identify traders such as Kurov and Lasser (2004), who use a proprietary/agency code, Joint Staff Report (2015) on the October 15 "Flash Rally" in U.S. Treasuries, which classifies accounts based on their organizational structure or Chaboud et al. (2014), who use a flag provided by a trading platform.

Figure 2 provides a visual representation of two of our classification dimensions: trading activity and end-of-day positions for all but the Small Traders, whose activity is negligible. The four panels correspond to each of the four trading days. The shaded areas are stylistically drawn to cover the areas populated by the individual trading accounts that fall into each of the categories based on their trading volume (vertical axis) and end-of-day position scaled by market trading volume (horizontal axis).<sup>17</sup>

#### <Insert Figure 2>

According to Figure 2, the ecosystem of the E-mini market consists of five fairly distinct clusters of traders: Fundamental Buyers, Fundamental Sellers, High Frequency Traders, Opportunistic Traders, and Market Makers. In terms of their trading activity, High Frequency Traders stand out from all the other trading categories and are clearly distinct from Market Makers. By accumulating a significant negative inventory, the cloud of Fundamental Sellers spreads out to the left of the origin, while the cloud of Fundamental Buyers spreads out to the right. Opportunistic Traders overlap to some extent with all of the other categories of traders.

Average indicators of trading activity for all categories of traders are presented in Table II. Panel A presents averages for the three days prior to the Flash Crash (May 3

<sup>&</sup>lt;sup>17</sup>For confidentiality reasons, we do not present trading volume or net position of individual accounts.

to 5, 2010), while Panel B presents indicators for the day of the Flash Crash (May 6, 2010).

#### <Insert Table II>

According to Table II, during the three days prior to the Flash Crash, 15 High Frequency Traders on average accounted for an average of 34.22% of the total trading volume and 189 Market Makers, on average accounted for an additional 10.49% of total trading volume. On the day of the Flash Crash, their respective shares of total trading volume dropped to 28.57% and 9.00%, respectively.<sup>18</sup>

Table II also presents average trade-weighted and volume-weighted "Aggressiveness Ratios," defined as the percentage of trades or contracts in which a side of the trade was the marketable side as opposed to a nonexecutable (that is, passive or resting). Over May 3 to 5, 2010, the three-day average of the volume-weighted proportions of aggressive trade executions by High Frequency Traders and Market Makers are 49.86% and 34.99%, respectively. On May 6, 2010, the proportions are only slightly different at 46.59% and 32.49%, respectively.<sup>19</sup> On May 6, trades of Fundamental Sellers resulted from markedly larger portions of orders that were executed than the other trader categories. Over 99% of High Frequency Traders' and Market Makers' trades result from limit orders, while only 65% of Small Traders' trades result from limit orders.

# B. Intermediation and the Flash Crash

Theory links liquidity crashes to the risk-bearing capacity of intermediaries. Huang and Wang (2008, 2010) develop an equilibrium framework in which market crashes

<sup>&</sup>lt;sup>18</sup>Some accounts classified as Market Makers for May 3 to 5 did not trade on May 6.

<sup>&</sup>lt;sup>19</sup>During the re-opening auction after the triggering of the Stop Logic Functionality on May 6, 2010, both sides of transactions were marked as passive.

emerge endogenously when a sudden excess of sell orders overwhelms the insufficient risk-bearing capacity of market makers. Further, Ait-Sahalia and Saglam (2016) link elevated price volatility with tighter inventory bounds for "high frequency" intermediaries, reflecting their capacity to bear risk associated with increased volatility.

The risk-bearing capacity of intermediaries can be identified by the observed bounds of their net positions.<sup>20</sup> Figure 3 presents the end-of-minute net inventories of Market Makers and High Frequency Traders alongside the price level of the E-mini. The dashed lines plot Market Makers' and High Frequency Traders' net positions, while the solid lines plot the price level of the E-mini. The top four panels present the net position of Market Makers over May 3 to 6, while the bottom four panels present the net positions of High Frequency Traders.

# <Insert Figure 3>

On each of the four days in our sample, High Frequency Traders never accumulated inventories greater than approximately 4,000 contracts, which is much smaller than the size of the 75,000-contract order of the large sell program documented in CFTC-SEC (2010b).<sup>21</sup> Similarly, Market Makers do not take on net inventories that exceed 1,500 contracts in either direction. These findings are consistent with the theory of the limited

<sup>&</sup>lt;sup>20</sup>See, for example, the inventory control models such as those in Amihud and Mendelson (1980) and Ho and Stoll (1983), among others. For empirical analysis, see Madhavan and Smidt (1993) and Hasbrouck and Sofianos (1993), among others.

<sup>&</sup>lt;sup>21</sup>In the Internet Appendix, we also document an approximately 30,000-contract trade imbalance between Fundamental Sellers and Fundamental Buyers in the minutes leading up to the Flash Crash. This imbalance is nearly an order of magnitude larger than the documented inventory capacity of High Frequency Traders. In addition, we show that the majority of the Fundamental Trader trade imbalance was picked up by Opportunistic Traders, who may be able to take on larger inventories in the E-mini because they are simultaneously selling shares in equity markets in order to conduct crossmarket arbitrage. The most active Opportunistic Traders in our sample also took on significant long inventories during the Flash Crash, likely while engaging in cross-market arbitrage. We present their net inventories under the title "High Frequency Arbitragers" in the Internet Appendix. Our results are consistent with the notion that the imbalance between Fundamental Sellers and Buyers was larger than the risk-bearing capacity of both High Frequency Traders and Market Makers.

risk-bearing capacity of intermediaries during a liquidity crash, as intraday intermediaries did not take on larger inventories compared with their pre-May 6 inventories. In contrast to Weill (2007), during the period of large and temporary selling pressure on May 6, we find that both categories of intraday intermediaries also accumulate net long inventory positions as prices decline.<sup>22</sup>

On May 6, as discussed in CFTC-SEC (2010b), shortly before the Stop Logic Functionality was triggered during the Flash Crash, High Frequency Traders aggressively liquidated approximately 2,000 contracts accumulated earlier, which coincided with significant additional price declines. In contrast, Market Makers did not liquidate the inventories that they had accumulated in the early minutes of the Flash Crash until after the Stop Logic Functionality was activated.<sup>23</sup>

To empirically examine the risk-bearing capacity of intraday intermediaries before and during the Flash Crash, we examine the second-by-second co-movement between the inventory changes of High Frequency Traders and Market Makers and market prices. Hasbrouck and Sofianos (1993) estimate vector autoregressions that include price changes, signed orders, and NYSE specialist inventory positions. More recently, Hendershott and Menkveld (2014) examine dynamics between the NYSE specialist inventories and prices, and Brogaard, Hendershott, and Riordan (2014) examine co-movements between high frequency traders as defined by NASDAQ and price changes, further decomposing price changes into permanent and temporary price changes.

We employ an empirically similar approach to establish a baseline statistical relationship between changes in inventories and changes in prices over May 3 to 5, 2010. With this baseline analysis, we simply examine the co-movement of intraday intermediary in-

 $<sup>^{22}</sup>$ The partial consistency of our empirical results with Weill (2007) could be due to the fact the Flash Crash takes place in an automated central limit order market, while Weill (2007) studies a market in which outside investors must be connected to each other by intermediaries.

<sup>&</sup>lt;sup>23</sup>For additional description of the trading activity during the seconds prior to the activation of the Stop Logic Functionality, see the Internet Appendix.

ventories and price changes without making causal inferences, as prices and inventories are jointly determined. We employ this baseline analysis separately for High Frequency Traders and Market Makers to account for possible differences in statistical relationships. Our baseline inventory and price regression is given as<sup>24</sup>

$$\Delta y_t = \alpha + \phi \cdot \Delta y_{t-1} + \delta \cdot y_{t-1} + \sum_{i=0}^{20} [\beta_i \cdot \Delta p_{t-i}/0.25] + \epsilon_t, \tag{1}$$

where  $y_t$  and  $\Delta y_t$  denote the inventories and changes in inventories of High Frequency Traders or Market Makers for each second of a trading day, t = 0 corresponds to the opening of stock trading on the NYSE at 8:30:00 a.m. CT (9:30:00 ET) and t = 24,300denotes the close of Globex at 15:15:00 CT (4:15:00 p.m. ET), and  $\Delta p_t$  denotes the price change in index point units between the high-low midpoint of second t - 1 and the high-low midpoint of second t to account for bid-ask bounce. To convert price changes into the number of ticks, we divide  $\Delta p$  by 0.25.<sup>25</sup> We present *t*-statistics obtained from White (1980) standard errors.<sup>26</sup>

# <Insert Table III>

In all baseline specifications, the regression coefficient on the lagged intermediary inventory level is negative, reflecting the mean-reversion of High Frequency Trader and Market Maker inventories. High Frequency Trader inventory changes are positively related to contemporaneous and lagged price changes in both specifications up to four lags. By the 10th lagged price change, High Frequency Traders inventory changes become negatively related to price changes. In contrast, Market Maker inventory changes

 $<sup>^{24}\</sup>mathrm{To}$  allay concerns of nonstationarity, we first-difference intraday intermediary inventories and market prices.

 $<sup>^{25}</sup>$ For reference, we also estimate the same regressions without the contemporaneous price change. See the Internet Appendix.

<sup>&</sup>lt;sup>26</sup>In Augmented Dickey Fuller tests, we reject the null of a unit root for all variables.
are negatively related to contemporaneous price changes but are generally positively related to lagged price changes.<sup>27</sup> Hendershott and Seasholes (2007) argue that market makers are willing to accommodate trades to less patient investors only if they are able to buy (sell) at a discount (premium) relative to future prices. Thus, the inventories of intermediaries should coincide with buying and selling pressure, which causes price movements that subsequently reverse themselves, implying a negative contemporaneous relationship between market maker inventories and prices. Although the co-movement between Market Maker inventory changes and price changes fits this paradigm, its High Frequency Trader counterpart does not. The fact that the regression coefficients of High Frequency Traders lagged inventory levels are larger than their Market Maker counterparts may speak to the difference in holding horizon and inventory mean-reversion of these two categories.<sup>28</sup>

To test whether the statistical relationship between intermediary inventory changes and price changes significantly changed during the Flash Crash, we estimate the following regressions:

$$\begin{split} \Delta y_t &= \alpha + \phi \Delta y_{t-1} + \delta y_{t-1} + \Sigma_{i=0}^{20} [\beta_i \times p_{t-i}/0.25] \\ &+ D_t^D \{ \alpha^D + \phi^D \Delta y_{t-1} + \delta^D y_{t-1} + \Sigma_{i=0}^{20} [\beta_i^D \times p_{t-i}/0.25] \} \\ &+ D_t^U \{ \alpha^U + \phi^U \Delta y_{t-1} + \delta^U y_{t-1} + \Sigma_{i=0}^{20} [\beta_i^U \times p_{t-i}/0.25] \} + \epsilon_t . \end{split}$$

In these regressions, we stack observations from May 3, May 4, May 5, and May 6 and include two sets of interaction terms,  $D_t^D$  and  $D_t^U$ . where  $D_t^D$  corresponds to the

 $<sup>^{27}{\</sup>rm The}$  contemporaneous price change coefficient for High Frequency Traders is statistically distinguishable from its Market Maker counterpart at the 1% level.

<sup>&</sup>lt;sup>28</sup>Results are qualitatively similar when we when we incorporate lead price changes in these regressions and when we include more price change and inventory lags. See the Internet Appendix.

"down" period of the Flash Crash and  $D_t^U$  corresponds to the "up" period (between 13:32:00 and 13:45:28 CT and between 13:45:33 and 14:08:00 CT, respectively).<sup>29</sup> The interaction coefficients measure differences between the coefficient estimates for the respective periods of the Flash Crash and for the non-Flash Crash periods. Results are presented in Table IV.

#### <Insert Table IV>

For High Frequency Traders, during the "down" phase of the Flash Crash, all interaction coefficients except for the fourth lagged price change are statistically insignificant - that is, the statistical relationship between High Frequency Traders' inventory changes and price changes did not significantly change in the seconds during which the price of the E-mini fell. During the "up" phase, which commenced after a five-second pause in trading, seven coefficients changed - notably, the coefficients on the interaction terms of the contemporaneous price change and the first two lagged price change interaction coefficients are negative and significant. We construct an F-test from the  $R^2$  estimated from the baseline regression presented in Table III and fail to reject the null that the interaction coefficients are jointly distinguishable from zero, lending little evidence to the view that High Frequency Traders' trading pattern changed.<sup>30</sup>

In contrast to High Frequency Traders, the contemporaneous and lagged price change interaction coefficients are statistically significant for Market Makers during both the "down" and "up" phases of the Flash Crash. During the "down" and "up" phases, the correlation between the Market Maker inventory changes and contemporaneous prices

 $<sup>^{29}</sup>$ Since we study intraday intermediation before and during the Flash Crash, we exclude the observations after 14:08:00 (CT) on May 6.

<sup>&</sup>lt;sup>30</sup>In the Internet Appendix, we document that it is the "down" phase of the Flash Crash that best corresponds to the period of large and temporary selling pressure, as the net selling of Fundamental Sellers exceeds the net buying of Fundamental Buyers by 33,944 contracts. Only one interaction coefficient is statistically significant for High Frequency Traders during this period

increased, while the correlation between the lagged prices and inventories decreased. The net effect of these positive and negative changes (sum of the significant coefficients) is close to zero, suggesting that the co-movement between Market Maker inventories and prices appears to have shifted down the lag structure. We construct an F-test from the  $R^2$  estimated from the baseline regression presented in Table III and reject the null that the interaction coefficients are jointly distinguishable from zero at the 1% level.

The change in co-movement between Market Makers' inventories and prices during the Flash Crash is consistent with the theory of liquidity crashes when intermediation is endogenous. CFTC-SEC (2010b) also indicates a reduced number of Market Makers during periods of the Flash Crash. In contrast, the mean-reversion of High Frequency Traders' inventory, as well as the co-movement between the inventories of High Frequency Traders and market prices did not significantly change during the Flash Crash.

### C. High Frequency Traders

To better understand High Frequency Traders' responses to the Flash Crash, we conduct additional empirical analyses of their intraday trading behavior. A developing theoretical literature models the behavior of High Frequency Traders differently than that of a traditional market marker. Broadly speaking, in these models faster intraday traders are able to "snipe" stale orders of slower market participants (see, for example, Foucault, Roell, and Sandas (2003), Cvitanic and Kirilenko (2010), Budish, Cramton, and Shim (2015), and Menkveld and Zoican (2016)). Quote sniping provides an economic rationale through which the inventories of faster intraday traders may positively co-move with price changes at high frequencies. Empirically, Harris and Schultz (1998) study the trading of the so-called SOES Bandits who picked off stale dealer quotes in NASDAQ stocks. A testable empirical pattern consistent with these predictions entails certain

traders regularly trading ahead of price changes at short time horizons.

We conduct two sets of tests that further analyze the statistical relationship between changes in the net positions of High Frequency Traders and market prices at very short time horizons. In the first set of tests, we analyze how prices change up to 20 seconds after High Frequency Traders trade. Figure 4 illustrates the results. The upper-left panel presents results for High Frequency Traders buy events over May 3 to 5, the upper-right panel presents results for High Frequency Traders buy events on May 6, and the lowerleft and lower-right panels present corresponding results for High Frequency Traders sell events.<sup>31</sup>

### <Insert Figure 4>

When High Frequency Traders are net buyers over May 3 to 5, prices rise by 17% of a tick in the next second then begin to gradually fall; 20 seconds after a net buy by High Frequency Traders, prices remain 15% of a tick higher. The total effect of net buy High Frequency Traders' trades can be separated into net aggressive and net passive buy trades. When High Frequency Traders buy aggressively, prices rise by 20% of a tick in the next second, continue rising into the next second, stabilize at about 23% of a tick during seconds 2 to 11, and then begin to gradually fall; 20 seconds after a net aggressive buy by High Frequency Traders, prices remain 15% of a tick higher. When High Frequency Traders buy passively, prices remain 15% of a tick higher. When High Frequency Traders buy passively, prices rise by 2% of a tick in the next second, prices then slowly trend downward to about negative 3% of a tick at the 20th second. The results are nearly the same for High Frequency Traders' sell trades, with the notable

<sup>&</sup>lt;sup>31</sup>For an "event-second" in which High Frequency Traders are net buyers, net aggressive buyers, and net passive buyers, value-weighted average prices paid by the High Frequency Traders in that second are subtracted from the value-weighted average prices for all trades in the same second and each of the following 20 seconds. The results are averaged across event-seconds, weighted by the magnitude of the High Frequency Traders' net position change in the event-second. Price differences on the vertical axis are given in the number of ticks (\$12.50 per one E-mini contract).

exception that while prices do decrease in the first second for passive High Frequency Traders' sell trades, they never cross into negative territory and, in fact, drift upward to about 12% of a tick 19 seconds later. The results are qualitatively similar on May 6, though prices appear to have a larger and more persistent response after sales by High Frequency Traders. It is important to note that prices increase in the second after High Frequency Traders complete their position change and not during the second that High Frequency Traders change their position, consistent with timing ability and not just the mechanical result of the price impact of marketable orders. This finding also cannot be explained by persistence in High Frequency Traders' inventory changes as one-second High Frequency Traders inventory changes are not autocorrelated, as can be seen in Table III.

In the second set of tests, we directly study the theory of quote sniping by analyzing how High Frequency Traders trade before and after decreases in the best bid or increases in the best offer. In a centralized limit order book market, a pattern consistent with stale quote sniping involves traders lifting posted depth just prior to a price change and then offsetting their position immediately at the new price level. Despite not directly examining limit order book data, the exact sequence of transactions and an aggressive/passive flag allows us to infer trader activity around price change events in the centralized E-mini order book in trade and volume time. We define a price increase (decrease) event as the best ask (bid) price increasing (decreasing). This definition ensures that we do not consider bid-ask bounces as price change events. When the best ask (bid) price increases (decreases), we count backwards the number of contracts traded at the "old" best ask price preceding the price change event. When we get to 100 contracts, we stop and attribute each side of the 100 contracts traded to one of the six categories of traders, add up the contract sides for each category, and calculate for each category the percentage share of trading volume for the last 100 contracts traded at the "old" best ask (bid) price before the price increase (decrease). We then compute averages for each category's overall price increase and decrease events over May 3 to 5 and May 6. Similarly, we calculate the percentage shares of aggressive and passive volume for the 100 contracts at the "new" best price after the price change. Our choice of examining the last and first 100 contracts traded around a price change event is motivated by the fact that the average posted depth at the best bid and offer during our sample was roughly 500 contracts. Furthermore, over May 3 to 5, 98.67 contracts were traded per second on average.<sup>32</sup> Results are presented in Table V.

#### <Insert Table V>

Table V has four panels. Panel A presents price increase events over May 3 to 5. Panel B presents price decrease events over May 3 to 5. Panels C and D present price increase and decrease events on May 6, respectively. In each panel, the last column presents the volume participation of different categories of traders without conditioning on price changes.

As shown in the last column of Panel A, when not conditioning on price changes, High Frequency Traders account for 34.04% of aggressive trading at the best ask price level. The share of High Frequency Traders' aggressive volume rises to 57.70% at the best ask price level before price increase events and falls to 14.84% at the new best ask price level after price increase events.<sup>33</sup> On the passive side, High Frequency Traders account for 34.33% of total passive volume at the best ask price level. However, the share of High Frequency Traders' passive volume at the best ask falls to 28.72% before

 $<sup>^{32}</sup>$ Results are qualitatively similar for 20 and 50 contracts, as well as for 10, 25, and 50 transactions. See the Internet Appendix. *p*-values associated with the averages and differences between High Frequency Traders and Market Makers are always less than 1%.

 $<sup>^{33}\</sup>mathrm{These}$  differences are statistically significant at the 1% level.

a price increase event and rises to 37.93% at the new best ask price level after a price increase event.<sup>34</sup>

For price decrease events, as shown in Panel B, the results are essentially symmetric. High Frequency Traders account for 55.20% of aggressive sell volume for the last 100 contracts traded before a price decrease event, compared with 34.17% when not conditioning on price changes. The share of High Frequency Traders' aggressive volume decreases to 15.04% at the new bid price after a price decrease event. On the passive side, High Frequency Traders account for 34.45% of the total passive volume at the best bid price level. The share of High Frequency Traders' passive volume falls to 27.41% before a price increase event and rises to 38.31% traded at the new best bid price level after a price decrease event. Panels C and D show that the behavior of High Frequency Traders is qualitatively similar on May  $6.^{35}$ 

In contrast, Market Makers follow a noticeably more passive trading strategy than High Frequency Traders. According to Panel A, Market Makers account for 13.48% of passive volume at the ask and only 7.27% of the aggressive volume at the ask. For the last 100 contracts at the old ask, Market Makers' share of volume increases relatively modestly, from 7.27% to 8.78% of aggressive volume at the old best ask price level. However, Market Makers' share of passive volume at the old best ask price also increases, from 13.48% to 15.80%.

These results suggest that High Frequency Traders behave differently than traditional market makers. The behavior of High Frequency Traders is empirically more consistent

<sup>&</sup>lt;sup>34</sup>The respective pre-price change event and post-price change event for the High Frequency Traders' aggressive and passive participation rates over May 3 to 5 are statistically distinguishable at the 1% level.

 $<sup>^{35}</sup>$ There were increases in the participation rate by Opportunistic Traders on the aggressive side of trades on May 6. For example, Opportunistic Traders' share of the aggressive volume at the ask price before a price increase increases from 19.21% over May 3 to 5 to 34.26% on May 6. Similarly, their share of aggressive volume at the bid price before a price decrease increases from 20.99% over May 3 to 5 to 33.86% on May 6.

with the theories of quote sniping or latency arbitrage than theories of traditional market making (see Glosten and Milgrom (1985)).<sup>36</sup> Our results, which are based on all trading in the E-mini, strengthen partial-sample results based on equity trading on NASDAQ (see, for example, Brogaard, Hendershott, and Riordan (2014)). We also directly link our empirical design and results to the theory of quote sniping.

## **III.** Conclusion

In this paper, we study intraday intermediation in the fully automated E-mini S&P 500 futures market before and during the Flash Crash, which was a period of large and temporary selling pressure. Our results suggest that the behavior of nondesignated intraday intermediaries is consistent with the theory of limited risk-bearing capacity: they did not take on large risky inventories relative to the large and temporary selling pressure on May 6. However, unlike textbook market makers, the most active intraday intermediaries (classified as High Frequency Traders) did not significantly alter their inventory dynamics when faced with large liquidity imbalances.

For a period of time, the Flash Crash seemed like an isolated event. However, flash events in the U.S. Treasury markets on October 15, 2014 reignited discussion about the vulnerability of liquid automated markets to severe dislocations and disruptive trading. Our empirical approach provides a framework to study intraday market dynamics before and during such systemic events, which may be a feature of the "new normal."

<sup>&</sup>lt;sup>36</sup>Without a sample of message-level data, we cannot determine whether High Frequency Traders become more aggressive in response to other traders changing their orders or private information, though the resulting trade patterns of either are consistent with quote sniping. Empirical patterns consistent with quote sniping that we document at a higher frequency do not preclude negative correlations of inventories and price changes at lower frequencies, as High Frequency Traders may employ heterogeneous strategies.

Initial submission: December 30, 2010; Accepted: June 8, 2016

Editors: Bruno Biais, Michael R. Roberts, and Kenneth J. Singleton

## REFERENCES

- Ait-Sahalia, Yacine, and Mehmet Saglam, 2016, High frequency market making, Working paper, Princeton University.
- Amihud, Yakov, and Haim Mendelson, 1980, Dealership market: Market making with inventory, Journal of Financial Economics 8, 31-53.
- Anand, Amber, Paul Irvine, Andy Puckett, and Kumar Venkataraman, 2013, Institutional trading and stock resiliency: Evidence from the 2007-2009 financial crisis, *Journal of Financial Economics* 108, 773-797.
- Biais, Bruno, Fany Declerck, and Sophie Moinas, 2016, Who supplies liquidity, how and when?, Working paper, University of Toulouse.
- Biais, Bruno, Thierry Foucault, and Sophie Moinas, 2015, Equilibrium fast trading, Journal of Financial Economics 116, 292-313.
- Brogaard, Jonathan, Allen Carrion, Thibaut Moyaert, Ryan Riordan, Andriy Shkilko, and Konstantin Sokolov, 2016, High-frequency trading and extreme price movements, Working paper, University of Washington.
- Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, 2014, High-frequency trading and price discovery, *Review of Financial Studies* 27, 2267-2306.
- Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan, 2016, High-frequency trading and the 2008 short sale ban, Working paper, University of Washington.
- Budish, Eric, Peter Cramton, and John Shim, 2015, The high-frequency trading arms race: Frequent batch auctions as a market design response, *Quarterly Journal of Economics* 130, 1547-1621.
- Carrion, Allen, 2013, Very fast money: High-frequency trading on the NASDAQ, Journal of Financial Markets 16, 680-711.
- CFTC-SEC Staff Report, 2010a, Preliminary findings regarding the market events of May 6.
- CFTC-SEC Staff Report, 2010b, Findings regarding the market events of May 6.
- Chaboud, Alain, Benjamin Chiquoine, Erik Hjalmarsson, and Clara Vega, 2014, Rise of the machines: Algorithmic trading in the foreign exchange market, *Journal of Finance* 69, 2045-2084.
- Clark-Joseph, Adam, 2014, Exploratory trading, Working paper, University of Illinois.

- Committee on the Global Financial System, 2014, Market-making and proprietary trading: Industry trends, drivers and policy implications.
- Conrad, Jennifer, Sunil Wahal, and Jin Xiang, 2015, High-frequency quoting, trading, and the efficiency of prices, *Journal of Financial Economics* 116, 271-291.
- Cvitanic, Jaksa, and Andrei Kirilenko, 2010, High frequency traders and asset prices, Working paper, California Institute of Technology.
- Easley, David, Marcos Lopez, and Maureen O'Hara, 2011, The microstructure of the "Flash Crash": Flow toxicity, liquidity crashes, and the probability of informed trading, *Journal of Portfolio Management* 37, 118-128.
- Foucault, Thierry, Alisa Roell, and Patrik Sandas, 2003, Market making with costly monitoring: An analysis of the SOES controversy, *Review of Financial Studies* 16, 345-384.
- Garman, Mark, 1976, Market microstructure, *Journal of Financial Economics* 3, 257-275.
- Glosten, Lawrence, and Paul Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71-100.
- Grossman, Sanford, and Merton Miller, 1988, Liquidity and market structure, *Journal* of Finance 43, 617-633.
- Harris, Jeffrey, and Paul Schultz, 1998, The trading profits of SOES bandits, *Journal* of Financial Economics 50, 39-62.
- Hasbrouck, Joel, 2003, Intraday price formation in U.S. equity index markets, *Journal* of Finance 58, 2375-2400.
- Hasbrouck, Joel, 2006, Empirical Market Microstructure: The Institutions, Economics, and Econometrics of Securities Trading (Oxford University Press, Oxford).
- Hasbrouck, Joel, and Gideon Saar, 2013, Low-latency trading, Journal of Financial Markets 16, 646-679.
- Hasbrouck, Joel, and George Sofianos, 1993, The trades of market makers: An empirical analysis of NYSE specialists, *Journal of Finance* 48, 1565-1593.
- Hayes, Roy, Mark Paddrik, Andrew Todd, Steve Yang, Peter Beling, and William Scherer, 2012, An agent based model of the E-mini S&P 500 applied to Flash Crash analysis, *IEEE Conference on Computational Intelligence for Financial En*gineering & Economics, 1-8.

- Hendershott, Terrence, Charles Jones, and Albert Menkveld, 2011, Does algorithmic trading improve liquidity? *Journal of Finance* 66, 1-33.
- Hendershott, Terrence, and Albert Menkveld, 2014, Price pressures, Journal of Financial Economics 114, 405-423.
- Hendershott, Terrence, and Mark Seasholes, 2007, Market maker inventories and stock prices, American Economic Review 97, 210-214.
- Hirschey, Nicholas, 2016, Do high-frequency traders anticipate buying and selling pressure? Working paper, London Business School.
- Ho, Thomas, and Hans Stoll, 1983, The dynamics of dealer markets under competition, Journal of Finance 38, 1053-1074.
- Huang, Jennifer, and Jiang Wang, 2008, Liquidity and market crashes, Review of Financial Studies 22, 2607-2643.
- Huang, Jennifer, and Jiang Wang, 2010, Market liquidity, asset prices, and welfare, Journal of Financial Economics 95, 107-127.
- Joint Staff Report, 2015, The U.S. Treasury Market on October 15, 2014.
- Jones, Charles, 2013, What do we know about high-frequency trading? Working paper, Columbia University.
- Jovanovic, Boyan, and Albert Menkveld, 2016, Middlemen in limit order markets, Working paper, New York University.
- Kaniel, Ron, Gideon Saar, and Sheridan Titman, 2008, Individual investor trading and stock returns, *Journal of Finance* 63, 273-310.
- Kirilenko, Andrei, Shawn Mankad, and George Michailidis, 2013, Discovering the ecosystem of an electronic financial market with a dynamic machine-learning method, *Algorithmic Finance* 2, 151-165.
- Kurov, Alexander, and Dennis Lasser, 2004, Price dynamics in the regular and E-mini futures market, *Journal of Financial and Quantitative Analysis* 39, 365-384.
- Kyle, Albert, 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315-1336.
- Madhavan, Ananth, 2000, Market microstructure: A survey, *Journal of Financial Markets* 3, 205-258.
- Madhavan, Ananth, 2012, Exchange-traded funds, market structure, and the Flash Crash, *Financial Analysts Journal* 68, 20-35.

- Madhavan, Ananth, and Seymour Smidt, 1993, An analysis of changes in specialist inventories and quotations, *Journal of Finance* 48, 1595-1628.
- Menkveld, Albert, 2013, High frequency trading and the new market makers, *Journal* of Financial Markets 16, 712-740.
- Menkveld, Albert, and Bart Yueshen, 2016, The Flash Crash: A cautionary tale about highly fragmented markets, Working paper, VU University Amsterdam.
- Menkveld, Albert, and Marius Zoican, 2016, Need for speed? Exchange latency and liquidity, Working paper, VU University Amsterdam.
- Nagel, Stefan, 2012, Evaporating liquidity, Review of Financial Studies 25, 2005-2039.
- O'Hara, Maureen, 1995, Market Microstructure Theory (Blackwell, Cambridge, MA).
- Puckett, Andy, and Xuemin Yan, 2011, The interim trading skills of institutional investors, *Journal of Finance* 66, 601-633.
- Seasholes, Mark, and Ning Zhu, 2010, Individual investors and local bias, Journal of Finance 65, 1987-2010.
- SEC, 2014, Equity Market Structure Literature Review Part II: High Frequency Trading.
- Weill, Pierre-Olivier, 2007, Leaning against the wind, *Review of Economic Studies* 74, 1329-1354.
- White, Halbert, 1980, A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* 48, 817-838.

# Table IMarket Descriptive Statistics

This table presents summary statistics for the June 2010 E-mini S&P 500 futures contract. The first column presents averages calculated for May 3 through 5, 2010, between 8:30 and 15:15 CT. The second column presents statistics for May 6, 2010 between 8:30 and 15:15 CT. Volume is the number of contracts traded. The number of traders is the number of trading accounts that traded at least once during a trading day. Order size and trade size are measured in number of contracts. The use of limit orders is presented in both percent of the number of transactions and trading volume. Volatility is calculated as the natural logarithm of maximum price over minimum price within a trading day.

|                                       | May 3–5         | May 6      |
|---------------------------------------|-----------------|------------|
| Daily Trading Volume                  | $2,\!397,\!639$ | 5,094,703  |
| # of Trades                           | 446,340         | 1,030,204  |
| # of Traders                          | $11,\!875$      | $15,\!422$ |
| Trade Size                            | 5.41            | 4.99       |
| Limit Orders % Volume                 | 95.45%          | 92.44%     |
| Limit Orders % Trades                 | 94.36%          | 91.75%     |
| Volatility (Log High-Low Price Range) | 1.54%           | 9.82%      |
| Return                                | -0.02%          | -3.05%     |

# Table II Summary Statistics of Trader Categories

This table presents summary statistics for trader categories. Panel A presents three-day average statistics for May 3 through 5, 2010 from 8:30 to 15:15 CT. Percentage of trading volume is a three-day average of the daily percentage of total trading volume for each trader category. Percentage of trades is a three-day average of the daily percentage of total trades for each trader category. Trade Size (Avg.) is a three-day average of the daily account-level average trade size within each trader category. Order Size (Avg.) is a three-day average of the daily account-level average size of the executed portion of an order within each trader category. Limit Orders % of volume is a three-day average of the percentage of trader category trading volume that resulted from marketable and nonmarketable limit orders. Agg Ratio Vol-Weighted is a three-day average of the percentage of trader category trading volume that resulted from 8:30 to 15:15 CT.

| Tallel A. May 5-5 Three-Day Average |  |   |   |  |  |   |
|-------------------------------------|--|---|---|--|--|---|
| % Volume                            | % of Trades  | # Traders   | Trade Size  | Order Size   | Limit Orders   | Agg Ratio   |
|                                     |  |   | (Avg.)  | (Avg.)   | %Volume  | Vol-Weighted  |
| 34.22%                              | 32.56%   | 15  | 5.69  | 14.75  | 100.00%  | 49.86%  |
| 10.49%                              | 11.63%   | 189   | 4.88  | 7.92   | 99.61%   | 34.99%  |
| 11.89%                              | 10.15%   | 1,013   | 6.34  | 14.09  | 91.26%   | 58.40%  |
| 12.11%                              | 10.10%   | 1,088   | 6.50  | 14.20  | 92.18%   | 54.98%  |
| 30.79%                              | 33.34%   | 3,504   | 4.98  | 8.80   | 92.14%   | 50.49%  |
| 0.50%                               | 2.22%  | 6,065   | 1.22  | 1.25   | 70.09%   | 58.54%  |
| Volume                              | # of Trades  | # Traders   | Trade Size  | Order Size   | Limit Orders   |   |
|                                     |  |   | (Avg.)  | (Avg.)   | %Volume  |   |
| $2,\!397,\!639$                     | 446,340  | 11,875  | 5.41  | 10.83  | 95.45%   |   |
|                                     | % Volume<br>34.22%<br>10.49%<br>11.89%<br>12.11%<br>30.79%<br>0.50%<br>Volume<br>2,397,639 | % Volume         % of Trades           34.22%         32.56%           10.49%         11.63%           11.89%         10.15%           12.11%         10.10%           30.79%         33.34%           0.50%         2.22%           Volume         # of Trades           2,397,639         446,340 | % Volume         % of Trades         # Traders           34.22%         32.56%         15           10.49%         11.63%         189           11.89%         10.15%         1,013           12.11%         10.10%         1,088           30.79%         33.34%         3,504           0.50%         2.22%         6,065           Volume         # of Trades         # Traders           2,397,639         446,340         11,875 | % Volume         % of Trades         # Traders         Trade Size (Avg.)           34.22%         32.56%         15         5.69           10.49%         11.63%         189         4.88           11.89%         10.15%         1,013         6.34           12.11%         10.10%         1,088         6.50           30.79%         33.34%         3,504         4.98           0.50%         2.22%         6,065         1.22           Volume         # of Trades         # Traders         Trade Size (Avg.)           2,397,639         446,340         11,875         5.41 | % Volume         % of Trades         # Traders         Trade Size<br>(Avg.)         Order Size<br>(Avg.)           34.22%         32.56%         15         5.69         14.75           10.49%         11.63%         189         4.88         7.92           11.89%         10.15%         1,013         6.34         14.09           12.11%         10.10%         1,088         6.50         14.20           30.79%         33.34%         3,504         4.98         8.80           0.50%         2.22%         6,065         1.22         1.25           Volume         # of Trades         # Traders         Trade Size         Order Size           (Avg.)         2,397,639         446,340         11,875         5.41         10.83 | % Volume         % of Trades         # Traders         Trade Size         Order Size         Limit Orders           34.22%         32.56%         15         5.69         14.75         100.00%           10.49%         11.63%         189         4.88         7.92         99.61%           11.89%         10.15%         1,013         6.34         14.09         91.26%           12.11%         10.10%         1,088         6.50         14.20         92.18%           30.79%         33.34%         3,504         4.98         8.80         92.14%           0.50%         2.22%         6,065         1.22         1.25         70.09%           Volume         # of Trades         # Traders         Trade Size         Order Size         Limit Orders           (Avg.)         (Avg.)         (Avg.)         % Volume         54.46,340         11.875         5.41         10.83         95.45% |

Panel A: May 3–5 Three-Day Average

Panel B: May 6

|                        |           |             |           | -          |            |              |              |  |
|------------------------|-----------|-------------|-----------|------------|------------|--------------|--------------|--|
| Trader Type            | % Volume  | % of Trades | # Traders | Trade Size | Order Size | Limit Orders | Agg Ratio    |  |
|                        |           |             |           | (Avg.)     | (Avg.)     | % Volume     | Vol-Weighted |  |
| High Frequency Traders | 28.57%    | 29.35%      | 16        | 4.85       | 9.86       | 100.00%      | 46.59%       |  |
| Market Makers          | 9.00%     | 11.48%      | 179       | 3.89       | 5.88       | 99.64%       | 32.49%       |  |
| Fundamental Buyers     | 12.01%    | 11.54%      | 1,263     | 5.15       | 10.43      | 88.84%       | 56.43%       |  |
| Fundamental Sellers    | 10.04%    | 6.95%       | 1,276     | 7.19       | 21.29      | 89.99%       | 55.30%       |  |
| Opportunistic Traders  | 40.13%    | 39.64%      | 5,808     | 5.05       | 10.06      | 87.39%       | 52.98%       |  |
| Small Traders          | 0.25%     | 1.04%       | $6,\!880$ | 1.20       | 1.24       | 63.61%       | 63.63%       |  |
|                        | Volume    | # of Trades | # Traders | Trade Size | Order Size | Limit Orders |              |  |
|                        |           |             |           | (Avg.)     | (Avg.)     | % Volume     |              |  |
| All                    | 5,094,703 | 1,030,204   | 15,422    | 4.99       | 9.76       | 92.443%      |              |  |
|                        |           |             |           |            |            |              |              |  |

# Table IIIBaseline Regression: Net Holdings and Prices

This table presents estimated coefficients for the regression:  $\Delta y_t = \alpha + \phi \Delta y_{t-1} + \sum_{i=0}^{20} [\beta_i \times \Delta p_{t-i}/0.25] + \epsilon_t$ . The dependent variable is the change in holdings of High Frequency Traders or Market Makers, as indicated. Both changes in holdings,  $\Delta y_t$ , and lagged holdings,  $y_{t-1}$ , are in contracts. Price changes,  $\Delta p_{t-i}$ , are in ticks. The sampling frequency is one second. *t*-statistics, calculated using the White (1980) estimator are reported in parentheses. Observations are stacked for May 3 through 5.

|                       | $\Delta$ NP HFT | $\Delta$ NP MM |
|-----------------------|-----------------|----------------|
| Intercept             | -1.64           | -0.53          |
|                       | (-3.54)         | (-3.33)        |
| $\Delta NPHFT_{t-1}$  | -0.01           |                |
| NPHET.                | (-0.69)         |                |
| 1 t 1 1 1 t t - 1     | (-11.77)        |                |
| $\Delta NPMM_{t-1}$   |                 |                |
| NDMM                  |                 | (-0.79)        |
| M P M M t - 1         |                 | (-8.93)        |
| $\Delta P_t$          | 32.09           | -13.54         |
| 0                     | (18.44)         | (-23.83)       |
| $\Delta P_{t-1}$      | 17.18           | -1.22          |
|                       | (12.58)         | (-2.71)        |
| $\Delta P_{t-2}$      | 8.36            | 2.16           |
|                       | (7.15)          | (4.99)         |
| $\Delta P_{t-3}$      | 5.09            | 2.53           |
| A D                   | (4.93)          | (5.97)         |
| $\Delta r_{t-4}$      | 3.91            | 2.00           |
| Δ <i>P.</i> -         | (3.02)          | 2.50           |
| $\Delta t t = 5$      | (1.56)          | (5.91)         |
| $\Delta P_{\star}$ c  | -0.08           | 2.16           |
| 1-0                   | (-0.07)         | (5.42)         |
| $\Delta P_{t-7}$      | -1.00           | 1.84           |
|                       | (-0.97)         | (4.96)         |
| $\Delta P_{t-8}$      | -1.76           | 1.47           |
|                       | (-1.56)         | (3.83)         |
| $\Delta P_{t-9}$      | -1.81           | 0.45           |
|                       | (-1.70)         | (1.19)         |
| $\Delta P_{t-10}$     | -3.90           | 0.52           |
| 4.17                  | (-3.78)         | (1.37)         |
| $\Delta P_{t-11}$     | -4.73           | -0.03          |
| Δ P                   | (-4.70)         | (-0.07)        |
| $\Delta t t - 12$     | (-3.33)         | (0.41)         |
| $\Delta P_{\star}$ 12 | -3.80           | 0.27           |
| 1-13                  | (-3.74)         | (0.72)         |
| $\Delta P_{t-14}$     | -4.77           | 0.32           |
|                       | (-4.70)         | (0.86)         |
| $\Delta P_{t-15}$     | -2.74           | -0.19          |
|                       | (-2.63)         | (-0.53)        |
| $\Delta P_{t-16}$     | -2.21           | -0.64          |
|                       | (-2.09)         | (-1.72)        |
| $\Delta P_{t-17}$     | -2.52           | -0.10          |
| A D                   | (-2.45)         | (-0.26)        |
| $\Delta r_{t-18}$     | -4.30           | 0.04           |
| $\Delta P_{t-10}$     | -4.21           | 0.57           |
| $\rightarrow t - 19$  | (-4.16)         | (1.51)         |
| $\Delta P_{t=20}$     | -5.86           | -0.12          |
| $\iota = 20$          | (-5.86)         | (-0.33)        |
| #obs                  | 72837           | 72837          |
| $Adj - R^2$           | 0.019           | 0.026          |

### Table IV

#### High Frequency Traders and Market Makers: The Flash Crash

This table presents estimated coefficients for the regression:  $\Delta y_t = \alpha + \phi \Delta y_{t-1} + \Delta y_{t-1} + \sum_{i=0}^{20} [\beta_i \times p_{t-i}/0.25] + D_t^D \{\alpha^D + \phi^D \Delta y_{t-1} + \delta^D y_{t-1} + \sum_{i=0}^{20} [\beta_i^D \times p_{t-i}/0.25]\} + D_t^U \{\alpha^U + \phi^U \Delta y_{t-1} + \delta^U y_{t-1} + \sum_{i=0}^{20} [\beta_i^U \times p_{t-i}/0.25]\} + \epsilon_t \text{ over May 3 through 6, 2010 with dummy variables } D_t^D$  and  $D_t^U$  included to interact with observations during the "Down" (from 13:32:00 CT to 13:45:28 CT) and "Up" (from 13:45:33 CT to 14:08:00 CT) phases of the Flash Crash. The period between 13:45:28 CT and 13:45:33 CT corresponds to the five-second pause in trading; there are no changes in prices or inventory during the five-second pause. The cutoff for observations on May 6, 2010 is 14:08:00 CT. The dependent variable the change in holdings of High Frequency Traders or Market Makers, as indicated. Both changes in holdings,  $\Delta y_t$ , and lagged holdings,  $y_{t-1}$ , are in contracts. Price changes,  $\Delta p_{t-i}$ , are in ticks. Estimates are computed for second-by-second observations. t-statistics, calculated using the White (1980) estimator are reported in parentheses.

| Variable           | $\Delta$ NP HFT | $\Delta$ NP MM | Variable (cont)         | $\Delta$ NP HFT | $\Delta$ NP MM | Variable (cont)           | $\Delta$ NP HFT | $\Delta$ NP MM |
|--------------------|-----------------|----------------|-------------------------|-----------------|----------------|---------------------------|-----------------|----------------|
| Intercept          | -2.04           | -0.48          | $Intercept^D$           | 9.22            | 9.15           | $Intercept^U$             | 2.27            | 0.49           |
|                    | (-4.78)         | (-3.34)        | D                       | (1.19)          | (2.41)         |                           | (0.55)          | (0.33)         |
| $\Delta N P_{t-1}$ | -0.005          | -0.024         | $\Delta NP_{t-1}^D$     | -0.031          | -0.024         | $\Delta NP_{t-1}^{U}$     | 0.004           | 0.085          |
|                    | (-0.69)         | (-3.31)        | D                       | (-0.80)         | (-0.63)        |                           | (0.10)          | (2.74)         |
| $NP_{t-1}$         | -0.005          | -0.005         | $NP_{t-1}^D$            | -0.002          | -0.007         | $NP_{t-1}^U$              | -0.001          | 0.000          |
|                    | (-12.95)        | (-10.78)       | P                       | (-0.38)         | (-1.62)        |                           | (-0.21)         | (-0.17)        |
| $\Delta P_t$       | 31.47           | -15.48         | $\Delta P_t^D$          | 1.29            | 14.13          | $\Delta P_t^U$            | -40.83          | 14.29          |
|                    | (16.89)         | (-21.96)       |                         | (0.18)          | (6.73)         |                           | (-12.18)        | (13.68)        |
| $\Delta P_{t-1}$   | 14.96           | -0.54          | $\Delta P_{t-1}^D$      | -3.02           | 11.44          | $\Delta P_{t-1}^0$        | -9.60           | 5.63           |
|                    | (12.17)         | (-1.23)        | · - D                   | (-0.57)         | (5.11)         | 11                        | (-3.44)         | (7.12)         |
| $\Delta P_{t-2}$   | 6.24            | 2.69           | $\Delta P_{t-2}^{D}$    | -6.84           | 1.87           | $\Delta P_{t-2}^{\circ}$  | -9.72           | -1.83          |
|                    | (5.36)          | (5.99)         | 1 D D                   | (-1.26)         | (0.81)         | 1 DU                      | (-3.57)         | (-2.20)        |
| $\Delta P_{t-3}$   | 3.02            | 2.65           | $\Delta P_{t-3}^D$      | -4.16           | -2.03          | $\Delta P_{t-3}^0$        | -3.97           | -2.47          |
|                    | (3.31)          | (7.14)         | 1 D D                   | (-0.69)         | (-1.22)        | 1 DU                      | (-1.61)         | (-3.75)        |
| $\Delta P_{t-4}$   | 1.92            | 2.74           | $\Delta P_{t-4}^D$      | -9.74           | -4.91          | $\Delta P_{t-4}^0$        | -1.12           | -2.51          |
|                    | (2.04)          | (7.78)         | · - D                   | (-1.98)         | (-3.11)        | 11                        | (-0.49)         | (-3.70)        |
| $\Delta P_{t-5}$   | 0.63            | 2.21           | $\Delta P_{t-5}^D$      | -10.94          | -3.45          | $\Delta P_{t-5}^0$        | 1.86            | -2.86          |
|                    | (0.64)          | (5.99)         | . – D                   | (-1.57)         | (-2.25)        | 11                        | (0.75)          | (-4.36)        |
| $\Delta P_{t-6}$   | -1.89           | 1.99           | $\Delta P_{t-6}^D$      | 0.59            | -2.91          | $\Delta P_{t-6}^{\circ}$  | 4.27            | -2.45          |
|                    | (-2.03)         | (5.72)         | . – D                   | (0.11)          | (-1.86)        | 11                        | (1.78)          | (-3.71)        |
| $\Delta P_{t-7}$   | -2.85           | 1.92           | $\Delta P_{t-7}^D$      | -1.66           | -2.71          | $\Delta P_{t-7}^{\circ}$  | -4.54           | -3.38          |
|                    | (-2.89)         | (5.18)         | . – D                   | (-0.31)         | (-1.59)        | 11                        | (-1.73)         | (-5.05)        |
| $\Delta P_{t-8}$   | -2.52           | 1.43           | $\Delta P_{t-8}^D$      | 2.45            | -2.97          | $\Delta P_{t-8}^{\circ}$  | 1.79            | -1.65          |
|                    | (-2.68)         | (4.33)         | • <b>D</b>              | (0.44)          | (-1.92)        | 1 DU                      | (0.76)          | (-2.76)        |
| $\Delta P_{t-9}$   | -2.59           | 0.48           | $\Delta P_{t-9}^D$      | -4.32           | -2.98          | $\Delta P_{t-9}^{\circ}$  | 2.69            | -1.64          |
|                    | (-2.76)         | (1.44)         | . – D                   | (-0.61)         | (-1.70)        | 11                        | (1.12)          | (-2.54)        |
| $\Delta P_{t-10}$  | -5.18           | 0.91           | $\Delta P_{t-10}^{D}$   | 3.93            | -3.40          | $\Delta P_{t-10}^{\circ}$ | 4.41            | -1.52          |
|                    | (-4.66)         | (2.12)         | . – D                   | (0.50)          | (-1.78)        | 11                        | (1.76)          | (-2.22)        |
| $\Delta P_{t-11}$  | -5.07           | -0.05          | $\Delta P_{t-11}^{D}$   | 9.84            | -6.35          | $\Delta P_{t-11}^{o}$     | 6.01            | -0.36          |
|                    | (-5.76)         | (-0.16)        | . – D                   | (1.30)          | (-2.96)        | 11                        | (2.27)          | (-0.51)        |
| $\Delta P_{t-12}$  | -4.05           | -0.10          | $\Delta P_{t-12}^{D}$   | 8.38            | -0.73          | $\Delta P_{t-12}^{\circ}$ | 4.37            | -0.79          |
|                    | (-4.46)         | (-0.31)        | 1 D D                   | (1.07)          | (-0.37)        | 1 DU                      | (1.34)          | (-1.26)        |
| $\Delta P_{t-13}$  | -3.86           | -0.07          | $\Delta P_{t-13}^{D}$   | 11.92           | -4.69          | $\Delta P_{t-13}^{\circ}$ | 10.02           | 0.28           |
|                    | (-4.27)         | (-0.20)        | 1 D D                   | (1.64)          | (-2.10)        | 1 DU                      | (3.34)          | (0.43)         |
| $\Delta P_{t-14}$  | -4.36           | 0.28           | $\Delta P_{t-14}^{D}$   | -8.56           | 0.79           | $\Delta P_{t-14}^{o}$     | 1.64            | -0.59          |
| A D                | (-5.01)         | (0.84)         | 1 DD                    | (-1.29)         | (0.41)         | A DU                      | (0.62)          | (-0.98)        |
| $\Delta P_{t-15}$  | -2.05           | -0.17          | $\Delta P_{t-15}^{D}$   | 8.46            | -5.41          | $\Delta P_{t-15}^{\circ}$ | 1.47            | -0.09          |
| A D                | (-2.27)         | (-0.50)        | 1 DD                    | (1.17)          | (-2.55)        | A DU                      | (0.64)          | (-0.15)        |
| $\Delta P_{t-16}$  | -2.01           | -0.39          | $\Delta P_{t-16}^{D}$   | -3.25           | 3.92           | $\Delta P_{t-16}^{\circ}$ | 1.07            | 0.99           |
| A D                | (-2.10)         | (-1.11)        | A DD                    | (-0.41)         | (1.80)         | 1 DU                      | (0.37)          | (1.56)         |
| $\Delta P_{t-17}$  | -2.67           | 0.01           | $\Delta P_{t-17}^{-}$   | 6.24            | -1.57          | $\Delta P_{t-17}$         | 5.19            | 0.48           |
| A D                | (-3.05)         | (0.02)         | A DD                    | (0.81)          | (-0.69)        | $\Lambda DU$              | (2.13)          | (0.75)         |
| $\Delta P_{t-18}$  | -3.89           | 0.19           | $\Delta P_{t-18}$       | -8.62           | 0.86           | $\Delta P_{t-18}$         | (0.59)          | -0.69          |
| A D                | (-4.10)         | (0.58)         | A DD                    | (-1.05)         | (0.42)         | $\Lambda DU$              | (2.58)          | (-1.12)        |
| $\Delta P_{t-19}$  | -3.50           | 0.70           | $\Delta P_{t-19}^{-}$   | -1.05           | -3.07          | $\Delta P_{t-19}$         | -0.75           | -0.88          |
| A D                | (-3.88)         | (2.08)         | A DD                    | (-0.12)         | (-1.39)        | $\Lambda DU$              | (-0.30)         | (-1.44)        |
| $\Delta P_{t-20}$  | -5.30           | -0.33          | $\Delta P_{t-20}^{-}$   | -2.32           | 3.13           | $\Delta P_{t-20}$         | 4.88            | -0.06          |
|                    | (-3.82)         | (-1.00)        | # of Obs                | (-0.30)         | (1.30)         |                           | (2.14)          | (-0.09)        |
|                    |                 |                | # 01 005<br>Adjusted R? | 0.021           | 0.036          |                           |                 |                |
|                    |                 |                | inajustea 112           | 0.021           | 0.000          |                           |                 |                |

### Table V

### Shares of Passive and Aggressive Trading Volume Around Price Increase and Price Decrease Events

This table presents each trader category's share of aggressive and passive trading volume for the last 100 contracts traded before a price increase event or price decrease event and the first 100 contracts traded at the new higher (lower) price after a price increase (decrease) event. For comparison purposes, this table also presents the unconditional share of aggressive and passive trading volume of each trader category. Trading categories are High Frequency Traders (Hft), Market Makers (Mm), Fundamental Buyers (Buyer), Fundamental Sellers (Seller), Opportunistic Traders (Opp), and Small Traders (Small). To emphasize the symmetry between buying and selling, the rows for Buyer and Seller in Panels B and D have been reversed relative to Panels A and C.

| Panel A: Trading at the Best Ask Around Price Increase Events, May 3–5, 2010 |                    |                |            |                     |                  |                        |  |  |
|--|--------------------|----------------|------------|---------------------|------------------|------------------------|--|--|
|  |                    |                |            |                     |                  |                        |  |  |
|  | Last 100           | 0 Contracts    | First 10   | First 100 Contracts |                  | at Best Ask            |  |  |
|  | Passive            | Aggressive     | Passive    | Aggressive          | Passive          | Aggressive             |  |  |
| Hft  | 28.72%             | 57.70%         | 37.93%     | 14.84%              | 34.33%           | 34.04%                 |  |  |
| Mm   | 15.80%             | 8.78%          | 19.58%     | 7.04%               | 13.48%           | 7.27%                  |  |  |
| Buyer  | 6.70%              | 11.61%         | 4.38%      | 26.17%              | 4.57%            | 21.53%                 |  |  |
| Seller   | 16.00%             | 2.65%          | 11.82%     | 7.09%               | 16.29%           | 5.50%                  |  |  |
| Opp  | 32.27%             | 19.21%         | 25.95%     | 43.39%              | 30.90%           | 31.08%                 |  |  |
| Small  | 0.51%              | 0.04%          | 0.34%      | 1.46%               | 0.44%            | 0.58%                  |  |  |
| Panel  | B: Tradir          | ng at the Best | Bid Arou   | nd Price Dec        | crease Events, l | May 3–5, 2010          |  |  |
|  | T ( 10)            |                | D: ( 10    |                     |                  |                        |  |  |
|  | Last 100           | U Contracts    | First 10   | 0 Contracts         | All Volum        | e at Best Bid          |  |  |
| TTC  | Passive            | Aggressive     | Passive    | Aggressive          | Passive          | Aggressive             |  |  |
| Hft  | 27.41%             | 55.20%         | 38.31%     | 15.04%              | 34.45%           | 34.17%                 |  |  |
| Mm   | 15.49%             | 8.57%          | 20.64%     | 6.58%               | 13.79%           | 7.45%                  |  |  |
| Seller   | 5.88%              | 11.96%         | 3.83%      | 24.87%              | 5.67%            | 20.91%                 |  |  |
| Buyer  | 17.98%             | 3.22%          | 12.71%     | 8.78%               | 15.40%           | 6.00%                  |  |  |
| Opp  | 32.77%             | 20.99%         | 24.18%     | 43.41%              | 30.30%           | 30.89%                 |  |  |
| Small  | 0.47%              | 0.06%          | 0.34%      | 1.32%               | 0.39%            | 0.58%                  |  |  |
| Pane   | l C: Trad          | ing at the Be  | st Ask Aro | und Price In        | crease Events,   | May 6, 2010            |  |  |
|  | Last 100 Contracts |                | First 10   | 0 Contracts         | All Volum        | All Volume at Best Ask |  |  |
|  | Passive            | Aggressive     | Passive    | Aggressive          | Passive          | Aggressive             |  |  |
| Hft  | 28.46%             | 38.86%         | 30.55%     | 14.84%              | 30.94%           | 26.98%                 |  |  |
| Mm   | 12.95%             | 5.50%          | 13.88%     | 5.45%               | 12.26%           | 5.82%                  |  |  |
| Buver  | 6.31%              | 17.49%         | 5.19%      | 21.76%              | 5.45%            | 20.12%                 |  |  |
| Seller   | 13.84%             | 3.84%          | 14.30%     | 5.71%               | 14.34%           | 4.40%                  |  |  |
| OPP  | 38.26%             | 34.26%         | 35.94%     | 51.87%              | 36.86%           | 42.37%                 |  |  |
| Small  | 0.19%              | 0.06%          | 0.16%      | 0.37%               | 0.16%            | 0.31%                  |  |  |
| Panel  | D: Tradi           | ing at the Bes | st Bid Aro | und Price De        | ecrease Events,  | May 6, 2010            |  |  |
|  |                    |                |            |                     |                  |                        |  |  |
|  | Last 100 Contracts |                |            | 0 Contracts         | All Volum        | All Volume at Best Bid |  |  |
|  | Passive            | Aggressive     | Passive    | Aggressive          | Passive          | Aggressive             |  |  |
| Hft  | 28.38%             | 38.67%         | 30.13%     | 14.59%              | 30.09%           | 26.29%                 |  |  |
| Mm   | 12.27%             | 5.04%          | 14.85%     | 5.64%               | 12.05%           | 5.88%                  |  |  |
| Seller   | 4.19%              | 16.46%         | 3.77%      | 21.21%              | 3.82%            | 17.55%                 |  |  |
| Buyer  | 15.83%             | 5.90%          | 13.89%     | 6.97%               | 15.27%           | 7.26%                  |  |  |
| Opp  | 39.12%             | 33.86%         | 37.15%     | 51.10%              | 38.56%           | 42.68%                 |  |  |
| Small  | 0.21%              | 0.08%          | 0.21%      | 0.48%               | 0.21%            | 0.34%                  |  |  |



Figure 1: Prices and trading volume of the E-mini S&P 500 stock index futures contract. Source: "Preliminary Findings Regarding the Market Events of May 6, 2010." This figure presents minute-by-minute transaction prices and trading volume of the June 2010 E-mini S&P futures contract on May 6, 2010 between 8:30 and 15:15 CT. Trading volume is calculated as the number of contracts traded during each minute. Transaction price is the last transaction price of each minute.



Figure 2: Trading accounts, trading volume, and net position scaled by market trading volume. This figure presents trader categories superimposed (as shaded areas) over all individual trading accounts ranked by their trading volume and net position scaled by market trading volume. The panels reflect trading activity in the June 2010 E-mini S&P 500 futures contract for May 3 through 6, 2010.





Figure 3: Net position of Market Makers and High Frequency Traders. This figure presents the net position (left vertical axis) of Market Makers and High Frequency Traders and market transaction prices (right vertical axis) in the June 2010 E-mini S&P 500 futures contract over one-minute intervals during May 3, 4, 5, and 6 between 8:30 and 15:15 CT. Net position is calculated as the difference between the total open long and total open short positions of Market Makers and High Frequency Traders at the end of each minute. Transaction price is the last market transaction price each minute. The top panel presents the net positions of Market Makers and the bottom panel presents the net positions of High Frequency Traders.



Figure 4: High Frequency Traders' trading and prices. This figure illustrates how prices change after HFT trading activity in a given second. The upper-left panel presents results for buy trades for May 3 through 5, 2010, the upper-right panel presents results for buy trades on May 6, 2010 and the lower-left and lower-right panels present corresponding results for sell trades. For an "event-second" in which High Frequency Traders are net buyers, net Aggressive Buyers, and net Passive Buyers, value-weighted average prices paid by the High Frequency Traders in that second are subtracted from the value-weighted average prices for all trades in the same second and each of the following 20 seconds. The results are averaged across event-seconds, weighted by the magnitude of High Frequency Traders' net position change in the event-second. The upper-left panel presents results for May 6, and the lower two panels present results for sell trades calculated analogously. Price differences on the vertical axis are scaled so that one unit equals one tick (\$12.50 per contract).