

Singapore Management University

## Institutional Knowledge at Singapore Management University

---

Research Collection Lee Kong Chian School Of  
Business

Lee Kong Chian School of Business

---

6-2018

### The competitive landscape of high-frequency trading firms

Ekkehart BOEHMER

*Singapore Management University, EBOEHMER@smu.edu.sg*

Dan LI

*University of Hong Kong*

Gideon SAAR

*Cornell University*

Follow this and additional works at: [https://ink.library.smu.edu.sg/lkcsb\\_research](https://ink.library.smu.edu.sg/lkcsb_research)



Part of the [Corporate Finance Commons](#), [Finance and Financial Management Commons](#), and the [Portfolio and Security Analysis Commons](#)

---

#### Citation

BOEHMER, Ekkehart; LI, Dan; and SAAR, Gideon. The competitive landscape of high-frequency trading firms. (2018). *Review of Financial Studies*. 31, (6), 2227-2276.

Available at: [https://ink.library.smu.edu.sg/lkcsb\\_research/5328](https://ink.library.smu.edu.sg/lkcsb_research/5328)

This Journal Article is brought to you for free and open access by the Lee Kong Chian School of Business at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection Lee Kong Chian School Of Business by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email [cherylds@smu.edu.sg](mailto:cherylds@smu.edu.sg).

## The Competitive Landscape of High-Frequency Trading Firms

Ekkehart Boehmer, Dan Li, and Gideon Saar\*

This version: June 2017

### Abstract

We examine product differentiation in the high-frequency trading (HFT) industry, where the “products” are secretive proprietary trading strategies. We demonstrate how principal component analysis can be used to detect underlying strategies that are common to multiple HFT firms, and show that there are three product categories with distinct attributes. We study how HFT competition in each product category impacts the market environment, presenting evidence that indicates how it influences the short-horizon volatility of stocks as well as the viability of trading venues.

\*Ekkehart Boehmer ([eboehmer@smu.edu.sg](mailto:eboehmer@smu.edu.sg)) is from the Lee Kong Chian School of Business, Singapore Management University; Dan Li ([lidan@hku.hk](mailto:lidan@hku.hk)) is from the Faculty of Business and Economics, University of Hong Kong; and Gideon Saar ([gs25@cornell.edu](mailto:gs25@cornell.edu)) is from the Johnson Graduate School of Management, Cornell University. We thank Matthew Baron, Evangelos Benos, Gregg Berman, James Brugler, Bjorn Hagstromer, Johan Hombert, Adam Joseph-Clark, Elvis Jarnećić, Boyan Jovanovic, Andrew Karolyi, Albert Menkveld, Sophie Moinas, Dermot Murphy, Maureen O’Hara, Mike Prior, Ioanid Rosu, Mehmet Saglam, Mehrdad Samadi, and Giorgio Valente for helpful comments, as well as conference and seminar participants at the 2<sup>nd</sup> Dauphine Microstructure Workshop, the 3<sup>rd</sup> Annual SEC/CFP Conference on Financial Market Regulation, the 13<sup>th</sup> Annual IDC Conference, the FINRA/Columbia Market Structure Conference, Cornell University, Indiana University, Monash University, the PBC School of Finance at Tsinghua University, Tel Aviv University, the University of Victoria Wellington, and VU Amsterdam. Dan Li thanks the Investment Industry Regulatory Organization of Canada (IIROC) for providing her with the data used in this study during her work at IIROC as part of a research agreement, and Maureen Jensen for her support and many helpful discussions. All conclusions are those of the authors and do not represent the views of IIROC. Part of this research was conducted while Gideon Saar was visiting the School of Management at Tel Aviv University. An earlier version of this work was circulated under the title “Correlated High-Frequency Trading.”

Industry competitiveness is an important theme in the economics of industrial organization. It is of particular interest when the industry in question plays a central role in the performance of our securities markets, which are essential to information production and allocation of capital. In this paper, we closely examine product differentiation and intra-industry competition in the high-frequency trading (HFT) industry using unique regulatory data from Canada. HFT firms use computerized algorithms for proprietary trading, and they engage in electronic market making, cross-trading venue price arbitrage, short-term statistical arbitrage, and various other opportunistic strategies.<sup>1</sup> This subset of financial service firms is responsible for most of the order flow and about half of the trading in equity markets in both the United States and Canada. The dominant role played by HFT firms has generated concerns over both the fairness and stability of equity markets, and these concerns are fueled in part by the secrecy that surrounds the operations of these firms.

As in Hoberg and Philips (2016), our study is based on the idea that product similarity is germane to industry boundaries. Given that the “product” of an HFT firm is a proprietary (and secretive) trading strategy, we first grapple with how to define product categories and assess the number of firms competing in each one. Besanko et al (2010) suggest identifying close substitutes by asking what a product does and when, where, and how it is used. We therefore identify HFT firms that are close substitutes by looking for strategies that do the same things at the same time in the same stocks. The greater the similarity in their activity, the more likely it is that they respond to the same stimuli and pursue the same profit opportunities. The industrial organization literature has long established that the more similar the products, the more competitive the environment. We invoke this “principle of differentiation” (Tirole (1988)) to motivate measures of similarity in HFT strategies as proxies for the intensity of competition between HFT firms. We use these measures to study whether competition between HFT firms affects the short-horizon volatility of individual stocks, and to examine the relationship between HFT competition and market concentration.

Our investigation of these issues is facilitated by data from the Investment Industry Regulatory Organization of Canada (IIROC), which is the national self-regulatory organization that oversees all equity trading venues in Canada. These data enable us to observe HFT activity coming from dealers’

---

<sup>1</sup> High-frequency trading therefore does not include agency algorithms that execute orders on behalf of investors; the term “HFT” is used only to denote proprietary trading operations conducted by standalone firms or the trading desks of larger financial firms.

proprietary trading desks as well as from HFT firms that use direct market access (DMA) arrangements with dealers. Moreover, we can track the activity of each HFT firm across all trading venues in Canada. Our sample consists of S&P/TSX 60 Index stocks, and we analyze activity from 30 days that represent bullish, bearish, and neutral trading environments during a period that runs from June 2010 through March 2011.

We use unmasked user and firm identities to characterize 31 trading firms as “high-frequency traders.” The strategies of these HFT firms are secretive insofar as we do not know their objectives, but we can measure their output: orders, cancellations, and trades. Each strategy creates a unique profile of orders, cancellations, and trades, but it is impossible to understand the nature behind any one strategy simply by observing millions of data points. We use principal component analysis (PCA) to classify these profiles into product categories by decomposing the correlation matrix of the occurrences of orders, cancellations, and trades, and construct the principal components that capture the meta-profiles that are common to multiple HFT firms. Our analysis suggests that there are at least three distinct underlying common strategies (or product categories), each followed by several HFT firms. While there are 17 HFT firms that do not appear to pursue one of these common strategies, the 14 firms that follow the common strategies represent most of the HFT activity, accounting for 96.21% of the messages that HFT firms send to the market and 78.97% of the volume they trade. Therefore, we find that competing HFT firms in three differentiated product categories generate most of the HFT activity in the Canadian market.

The PCA does not attach any economic interpretation to the underlying common strategies it identifies. We therefore analyze the principal component scores by regressing them on various attributes of the HFT firms and the market environment. Our results could be interpreted to suggest that one of the principal components represents a cross-venue arbitrage strategy, another one involves market making, and a third is related to short-horizon directional speculation. While we stress that these labels are merely suggested by patterns of regression coefficients (they do not constitute definitive characterizations of these strategies), the important takeaway from our analysis is that there is clear product heterogeneity in this industry. Therefore, we believe it may be more constructive to discuss the impact of HFT on the market in the context of specific strategies rather than in the aggregate.

Given the aforementioned “principle of differentiation,” we construct a measure of similarity in trading strategies within each product category as a proxy for the competitiveness of the product market.

While competition has desirable properties, if it means that HFT firms engage in strategies that are highly correlated across stocks, the central role these firms play in our markets could exacerbate return movements and increase short-horizon stock volatility. We study whether this is the case, and our investigation uncovers the opposite result: the short-horizon volatility of most stocks loads negatively on the extent of similarity in strategies of HFT firms that follow the underlying common strategy that we associate with market making. We investigate potential drivers of this relationship, and find that portions of short-horizon volatility related to both the permanent and temporary trade price impacts drive the negative relationship with HFT competition. In other words, our results suggest that competition between HFT intermediaries in the spirit of Jovanovic and Menkveld (2016) could lead to faster revelation of hard-information signals, lowering adverse selection costs and hence reducing volatility induced by the permanent price impact of trades. Competition in market making may also lower the compensation HFT market makers earn, which in turn explains the reduction in volatility that stems from the temporary price impact. These results indicate that HFT competition may bring dual benefits to the market in the form of potential lower rents for the services HFT firms provide as well as lower volatility. We stress, however, that our analysis reflects normal market conditions and this conclusion may not hold during an intense market breakdown, such as the Flash Crash in May of 2010 in the United States.

We also investigate how competition between HFT firms relates to competition between trading venues in the market. Brokers are required to provide their customers with best execution and to search all venues on which a security is traded to find the most beneficial terms of trading. Therefore, the trading venues displaying the best prices or smallest spreads attract more order flow. HFT firms play a central role in setting prices and spreads, and hence could affect competition between trading venues. We show that concentration of trading in the market is negatively related to similarity in strategies of HFT firms that follow two of the three underlying common strategies: cross-venue arbitrage and market making. We study one channel through which HFT competition affects market concentration by considering whether greater HFT competition on each trading venue increases the percentage of time that the venue features the best prices or the narrowest spreads, which we view as measures of the viability or competitiveness of the venue. We find that smaller trading venues are more viable while the dominant trading venue is less viable when HFT competition is more intense, which is likely one of the drivers behind the negative relationship we document between HFT competition and market concentration.

We contribute to the literature by presenting new insights into the nature of competition in the HFT industry and its effects on the market. Our most important contribution is the ability to define product categories in the HFT industry despite the complete lack of information about the nature of each firm's strategy, and identify firms that compete directly with one another in each category. Our use of PCA to investigate product differentiation is novel, and could potentially be utilized by others in similar contexts. We also provide new insights into the impact of HFT competition within each product category on the market environment, documenting evidence that it lowers the short-horizon volatility of stocks and influences market concentration by enhancing the viability of smaller trading venues.

## **1. Literature Review**

Our paper joins a rapidly expanding body of literature on HFT in financial markets. For recent surveys on the topic of HFT, see Jones (2013) and Goldstein, Kumar, and Graves (2014). Among the theoretical contributions are those of Han, Khapko, and Kyle (2014), Hoffmann (2014), Biais, Foucault, and Moinas (2015), Ait-Sahalia and Saglam (2016), Foucault, Hombert, and Rosu (2016), Jovanovic and Menkveld (2016), and Rosu (2016). In particular, Jarrow and Protter (2012) show that, when HFT firms respond to common signals, their correlated activity affects market prices, thereby increasing market volatility and generating abnormal profits for these firms. While we indeed show that multiple HFT firms pursue correlated strategies, we find that the short-horizon volatility of most stocks loads negatively, not positively, on a measure of similarity between HFT strategies. Budish, Cramton, and Shim (2015) and Menkveld and Zoican (2016) model HFT firms that pursue heterogeneous strategies in the market, which we also document empirically.

Many empirical contributions focus on intraday analysis of aggregate HFT behavior (e.g., Carrion 2013; Hasbrouck and Saar 2013; Brogaard, Hendershott, and Riordan 2014; Jarnecic and Snape 2014; Hirschey 2016; Kirilenko et al. 2016). Several papers use data on trading by individual HFT firms, rather than aggregate behavior, to investigate HFT strategies (Baron, Brogaard, and Kirilenko 2012; Hagstromer and Norden 2013; Hagstromer, Norden, and Zhang 2014; Benos and Sagade 2016), although they focus predominantly on whether HFT firms demand or supply liquidity. Clark-Joseph (2014) uses index futures data from the CMA to examine exploratory trading by HFT firms.

In studies that are more closely related to our analysis, Breckenfelder (2013) and Brogaard and

Garriott (2016) examine one aspect of competition: the entry and exit of HFT firms. Specifically, Brogaard and Garriott (2016) analyze data from one alternative trading system in Canada and show that new entrants take volume away from incumbents even as they increase the overall market share of HFT firms, and these effects decline with each successive entry. They also find that market liquidity improves after the entry of an HFT firm and deteriorates after an exit (especially when there are only one or two HFT firms trading in a given stock). Breckenfelder (2013) uses data from the Stockholm Stock Exchange and finds the opposite result: deterioration of liquidity for entries of HFT firms and improvement for exits. Menkveld (2013) examines the trading of one HFT firm and makes the case that this particular firm enhances the viability of a new trading venue, which is related to our result that competition between HFT firms is positively related to the viability of smaller trading venues.<sup>2</sup> Hagstromer and Norden (2013) find that HFT firms that predominantly supply liquidity appear to mitigate intraday volatility, which complements our finding pertaining to the relationship between HFT competition and volatility, although they do not address the competition dimension.<sup>3</sup>

Lastly, our paper joins several other papers that use Canadian order-level data to investigate HFT-related issues. Malinova, Park, and Riordan (2014) study the cost of trading around changes in market structure that affect mainly HFT firms and other algorithmic traders. Comerton-Forde, Malinova, and Park (2016) study the nature of liquidity provision around changes in dark trading regulation, while Korajczyk and Murphy (2016) examine HFT liquidity provision to large institutional trades. Brogaard, Hendershott, and Riordan (2016) study how limit orders and trades of high-frequency traders contribute to price discovery.

---

<sup>2</sup> In constructing measures of similarity in strategies, we use correlations between the activities of HFT firms. Several papers have examined correlations in the context of algorithmic trading or HFT. Chaboud et al. (2014) investigate algorithmic trading on the foreign-exchange EBS platform. Benos et al. (2015) use transactions data from the London Stock Exchange to study ten HFT firms. They use vector autoregressions to show that there is a positive dynamic relationship between HFT firms, and construct a measure of concurrent directional trading that is meant to capture correlated behavior. Anand and Venkataraman (2016) investigate the correlated activity of market makers on the Toronto Stock Exchange, some of which are HFT firms that operate as Electronic Liquidity Providers (ELPs).

<sup>3</sup> Brogaard, Carrion, Moyaert, Riordan, Shkilko, and Sokolov (2017) identify short intervals with large price movements and show that NASDAQ HFT firms in the aggregate supply rather than demand liquidity during these intervals, hence possibly dampening volatility. Hasbrouck and Saar (2013) also find that aggregate HFT activity appears to lower the intraday volatility of NASDAQ stocks.

## 2. The Canadian Market: Sample, Data, and HFT Firms

Our data come from the Investment Industry Regulatory Organization of Canada (IIROC). All trading venues in Canada are required to provide data feeds to IIROC, which performs both real-time and post-trade market surveillance of trading activities. Directly connecting to trading venues in Canada requires an IIROC membership, and IIROC admits only security dealers as members. Other financial firms, such as asset managers, banks, insurance companies, and proprietary trading firms, can trade through dealers' brokerage arms or via direct market access (DMA) arrangements provided by dealers. DMA arrangements allow non-dealer trading firms to access trading venues without having to hand over their orders to brokers for execution.

During our sample period (June 2010 through March 2011), Canada has five trading venues organized as electronic limit-order books that trade stocks listed on the Toronto Stock Exchange: Alpha ATS Limited Partnership (ALF), Chi-X Canada ATS (CHX), Omega ATS (OMG), Pure Trading (PTX), and the Toronto Stock Exchange (TSX).<sup>4</sup> Trading on crossing networks ("dark trading") in Canada during our sample period is limited essentially to one dark pool (MATCH Now) with no more than a 3% market share.<sup>5</sup>

### 2.1 Sample

We carry out the empirical work on 30 trading days that are selected to capture variation across market conditions. We rank the daily returns of the S&P/TSX Composite Index from June 2010 through March 2011, and select the 10 worst days (down days), the 10 best days (up days), and the 10 days closest to (and centered on) zero return (flat days). This design allows us to examine whether the similarity in the strategies of HFT firms depends on market conditions (as summarized by the daily return on a broad index). Our sample stocks consist of 52 constituents of the S&P/TSX 60 Index, which accounts for approximately 73% of Canada's total equity market capitalization. Eight stocks are excluded from the Index because they converted from income trusts to corporate structures (five), listed for less than one year (one), delisted (one), or changed symbols (one). Panel A of Table 1 presents summary statistics for

---

<sup>4</sup> Alpha became a stock exchange on April 1, 2012. In July 2012, Alpha was acquired by the TMX Group, which also owns TSX. During our sample period, however, Alpha and TSX were independent trading venues.

<sup>5</sup> MATCH Now also provides a real-time data feed to IIROC and is included in our data. Liquidnet Canada, another dark pool, executed only a few trades each day during our sample period. As a result, IIROC did not require it to participate in the real-time centralized data feed, instead requiring it to submit trade information manually at the end of the trading day.



the sample.<sup>6</sup> The mean market capitalization is 19.4 billion Canadian dollars (CAD), with an average stock price of 39.1 CAD, and average daily volume of 78.2 million CAD. Panel A clearly shows that our sample period encompasses three distinct market conditions: the average daily returns of stocks on down, flat, and up days are  $-1.72\%$ ,  $-0.08\%$ , and  $1.66\%$ , respectively.

## 2.2 Data

The order-level data we obtain from IIROC cover all Canadian trading venues, and contain information about order submissions, cancellations, modifications, and executions with 10-millisecond time stamps. The time stamps from all trading venues are synchronized with the regulator's time stamps and reflect the local time at which a message (a general term used for submissions, cancellations, and executions of orders) is processed. The record of each message contains the following information: ticker symbol, order side (buy or sell), trading venue, price, total quantity, non-displayed quantity, broker ID, user ID, order type (e.g., client orders, inventory trading), various flags (e.g., short sell, market on close, opening trade), and order/trade ID. Trade messages are identified as buyer-initiated or seller-initiated. Events in an order's life, including modification, partial fill, full fill, and cancellation have the same order ID. We provide a more complete description of the data in the Online Appendix.

One advantage of our data is that the same user IDs are used on all trading venues in Canada. Most HFT firms in Canada do not operate as licensed dealers but rather gain access through DMA arrangements with one or multiple dealers. We obtain tables that identify the user IDs of all trading firms that use DMA arrangements. Hence, we can accurately detect the activity of each HFT firm on all trading venues irrespective of whether it trades via multiple DMA arrangements with dealers or uses multiple user IDs.

## 2.3 Identifying HFT Firms

Researchers have used a variety of methods to identify HFT activity, dictated mostly by the quality and attributes of available data. Some papers use publicly available data to identify aggregate HFT activity (e.g., Hasbrouck and Saar (2013)), while others have used data with masked user IDs in conjunction with quantitative criteria (e.g., end-of-day inventory that is less than 20% of daily volume, see Brogaard, Hendershott, and Riordan (2016)). The method considered most accurate for identifying HFT firms

---

<sup>6</sup> Data on stock characteristics are obtained from the Summary Information Database of the Canadian Financial Markets Research Centre (CFMRC).

requires knowledge of the trading entities. The advantage of our data is that they include unmasked user and firm identities, hence providing us with an unparalleled ability to observe the HFT firms.<sup>7</sup>

Our data contain hundreds of trading firms, which makes it difficult to investigate every one of them at length. We therefore rank trading firms along several dimensions simply as a way to flag candidates for more thorough analysis. The dimensions we use are number of trades, number of messages, orders-to-trades ratio, the number of times a firm's intraday inventory positions cross the end-of-day position (or zero), and cross-trading-venue activity in the same or neighboring time-stamps. For example, if a firm has almost no proprietary trading and executes only a handful of trades every day, we believe that it is not engaged in what is generally perceived as HFT and therefore does not merit additional attention. We do not set a formal cutoff for any of these dimensions because our classification of trading firms is not quantitative in nature but rather manual; the quantitative analysis simply eliminates firms with little relevant activity.<sup>8</sup>

It is important to stress that we manually examine many firms—those that were eventually identified as engaging in HFT as well as other firms—to ensure that we capture all relevant entities. For each of those firms, we use its name to search for information about the firm's business in media news reports and on the web. For example, some firms are mentioned in newspaper articles as engaging in HFT. Other firms have websites on which they explicitly state that they pursue high-frequency strategies. Still others are listed as participants in TSX's Electronic Liquidity Providers (ELP) program, which offers fee incentives to HFT firms that provide liquidity.<sup>9</sup> We use this information to ensure that we identify HFT firms even if they do not rank highly along some of the quantitative dimensions. Our procedure enables us to identify 31 firms as HFT firms. Some of these firms have DMA arrangements with dealers while others are dealers engaged in proprietary trading. Our goal in this identification procedure is to

---

<sup>7</sup> The NASDAQ HFT dataset was constructed by NASDAQ with knowledge of the trading entities, although NASDAQ released only a general HFT flag rather than a separate flag for each HFT firm.

<sup>8</sup> Furthermore, the quantitative variables are computed using data from all trading days in September 2010, while we carry out the empirical work on the 30 down, flat, and up days. This out-of-sample property is meant to ensure that the identification is exogenous to our empirical analysis, even though the quantitative variables are used simply for rough initial screening and our classification procedure utilizes the identities of the entities to manually select the HFT firms.

<sup>9</sup> This is a program that offers fee incentives to firms that use proprietary capital and high-frequency electronic trading algorithms to provide liquidity on the exchange. ELPs are either independent proprietary trading firms or proprietary trading desks within large banks or financial firms, and the program requires that they passively trade at least 65% of their volume.

ensure that we study as complete a set of HFT firms as possible.<sup>10</sup>

Since we have both user IDs and a mapping of the user IDs to the firms, the first choice we need to make is whether to operate at the user-ID level or aggregate all user IDs that belong to the same firm. There are no rules (to the best of our knowledge) that guide how firms utilize user IDs.<sup>11</sup> One possible concern is that firms assign multiple user IDs to the order flow of an algorithm and send orders via DMA arrangements with multiple dealers to make it more difficult for outsiders to detect their activity. Routing via multiple dealers could also be driven by the desire to limit dependence on one dealer. We choose to work at the HFT-firm level (i.e., to aggregate all user IDs belonging to the same firm) to make our analysis robust to whatever gaming could be going on in terms of the firm's discretionary assignment of user IDs to its orders. It is important to note that, for dealers who may have brokerage arms in addition to proprietary trading operations, we exclude any order or trade made in the capacity of an agent, and include in our measurement of their HFT activity only those orders and trades identified in the order type field as proprietary activity.

The 31 HFT firms we identify send orders via 572 user IDs during our sample period.<sup>12</sup> It is important to stress that, for broker-dealers, these IDs are used solely for proprietary trading. Figure 1 plots the number of user IDs for each of the 31 HFT firms (sorted from the largest to the smallest). The mean number of user IDs per HFT firm is 18.45 and the median is 4. The large dispersion in the number of user IDs per firm is clearly seen in the figure: one firm sends orders through 137 IDs, while there are 6 HFT firms that operate through a single ID. The correlation between the HFT firms' market share in terms of

---

<sup>10</sup> While speed is an essential characteristic of HFT firms, we do not use speed per se to identify such firms. An IIROC (2014) report that investigated the use of machine-learning methods to identify HFT firms used speed measures such as order amendment speed or the percent of "simultaneous" orders. The submission of simultaneous orders to multiple venues is indeed one of the dimensions we examine, but we are unclear whether order amendment speed would identify the majority of HFT firms correctly. Some of the HFT firms may not use AMEND orders but prefer operating with other order types (ENTER and DELET) and/or utilizing IOC and FoK orders. Ultimately, we believe that a manual classification using unmasked user and firm identities is the most accurate way to identify HFT firms and hence is the method we pursue in the paper.

<sup>11</sup> According to IIROC (2012), "one user ID may identify a single approved trader (for example a registered market maker or inventory trader), a business stream (for example, orders originated by a Participant's online discount brokerage system), or a client who accesses the markets directly through a DMA relationship with a Participant." User IDs are therefore utilized at the complete discretion of trading firms.

<sup>12</sup> The aforementioned IIROC (2014) report that examined the feasibility of classifying user IDs using machine learning methods manually labeled 49 user IDs based on knowledge of the trading entities, which is considered the most accurate classification method, to train their algorithm. The trained algorithm was then used on all user IDs, and ended up classifying 98 user IDs as engaging in HFT. While it is possible that these 98 IDs are the most active ones among the 572 user IDs we identify, these results highlight the difficulty involved in classifying HFT firms using quantitative methods.

volume traded and the number of IDs they use is 0.362 (i.e., firms with a larger market share operate through more IDs). However, the number of IDs is essentially uncorrelated (0.015) with market share in terms of the number of messages a firm sends to the market. If an HFT firm operates multiple algorithms and actually designates a separate user ID to each algorithm, our procedure will lump them together, although the principal component analysis we implement in Section 3 could potentially differentiate these separate strategies.<sup>13</sup>

Panel B of Table 1 provides summary statistics for the 31 HFT firms as well as for categories formed on market share of volume. Specifically, MS1 consists of four firms with market shares greater than 4%, MS2 consists of six firms with market shares between 1% and 4%, and MS3 consists of the rest of the firms. Overall, these 31 HFT firms have 46.4% of the market share in terms of volume. The average number of daily trades of a firm in our sample is 19,445, but it ranges from 102,035 for firms in MS1 to 2,489 for firms in MS3. Similarly, the number of daily messages a firm sends to the market (submissions and cancellations of non-marketable limit orders as well as the number of marketable orders) is 1,063,974, ranging from 5,496,423 for firms in MS1 to 239,157 on average for firms in MS3. The mean (median) number of times the intraday inventory position of a firm in our sample crosses its end-of-day inventory increases with market share: 4.3 (1.8) for MS3, 13.4 (4.3) for MS2, and 73.0 (70.1) for MS1. Similarly, the number of 10-millisecond time stamps in which an HFT firm is active both on the primary market (TSX) and on another trading venue increases with market share. However, it is interesting to note that the Orders/Trades Ratio does not share the same pattern, with the highest mean and median (116.5 and 71.4, respectively) characterizing the smaller HFT firms.

### **3. HFT Competition and Product Differentiation**

#### *3.1 Differentiation, Competition, and Correlated Activity*

The HFT industry is shrouded in secrecy. Most HFT firms are private and hence reveal no financial or operating information; detailed information about proprietary trading desks of larger, publicly listed firms are seldom disclosed to the public. In general, HFT firms do not reveal information about the objectives of their algorithms beyond speaking generally about concepts such as “liquidity provision” and “arbitrage.” How can one investigate the extent of competition in such an industry? Examining the

---

<sup>13</sup> To ensure the robustness of our results, we also carry out a principal component analysis in terms of user IDs, and discuss the results in the Online Appendix.

economic profits of firms would be ideal, but the costs associated with hiring and retaining the individuals who develop the algorithms as well as other operating costs are simply unavailable. Researchers can only ascertain firms' net trading revenues (or what is left when shares are bought and sold) from datasets like ours, but trading revenues without costs cannot be used to make a correct inference about economic profits or the competitiveness of an industry.<sup>14</sup>

Our first goal is to study whether there are product categories in the HFT industry such that each product category is characterized by multiple HFT firms that pursue similar strategies. Hoberg and Phillips (2016), for example, design a new classification scheme to define industries using text-based analysis of product descriptions. Besanko et al. (2010) define products that are close substitutes as having the same or similar product performance characteristics (i.e., what they do) and having the same or similar occasions for use (i.e., when, where, and how they are used). These criteria motivate the manner in which we identify HFT strategies that are close substitutes: these strategies submit, cancel, or execute orders at the same time in the same stocks. The activity of each HFT firm consists of the messages (e.g., submission and cancellation of limit orders) it sends to the market to implement its strategy. The higher the correlation between the activities of two HFT firms, the greater is the similarity (or the lower is the differentiation) in their proprietary trading strategies. In other words, the higher the correlation, the more likely it is that the firms pursue the same profit opportunities and respond to the same trading signals. Hence, we use the correlation between the manifestations of strategies (e.g., the messages HFT firms send to the market) as a measure of product similarity. In particular, we decompose the correlation matrix in the activities of the HFT firms using principal component analysis to establish the main product categories and determine how many HFT firms operate in each one.

According to Aldridge (2013), the secrecy in the HFT industry stems from purely competitive business considerations. She notes that every strategy has a finite capacity, and a competitor is bound to significantly diminish the profitability of an HFT strategy. As a result, when multiple firms carry out the same strategy, each firm's profits go down. The industrial organization literature has long recognized the "principle of differentiation," according to which firms want differentiated products to soften competition

---

<sup>14</sup> Even estimating net trading revenues requires assumptions about valuing inventory that can dramatically influence the results (see Carion (2013)). For completeness, the Online Appendix contains an analysis of net trading revenues (including fees and liquidity rebates) in the context of the underlying common strategies of HFT firms that we identify in this paper.

and increase their profits (Tirole (1988)). Hotelling (1929) showed in the context of two firms that the more similar the products, the closer one gets to Bertrand competition, and this result has been generalized in many other papers (see, for example, Salop (1979) for the case of  $n$  firms). In general, Shy (1995) states that in both Bertrand and Cournot games with differentiated products, the firms' profits increase when the products become more differentiated (Propositions 7.1 and 7.2). In both cases, product differentiation softens competition; the more similar the products, the more competitive the environment.

Strategy differentiation for HFT firms helps alleviate the capacity problem and increases profitability in the same way that product differentiation in the aforementioned industrial organization models helps alleviate price competition and increases profits. The more similar the strategies, the more intensely HFT firms compete. In Section 4 we use measures of strategy similarity as proxies for competition to investigate how HFT competition affects the market environment.

One conceptual issue we want to clarify is that we use the terms “strategy” and “activity” somewhat interchangeably. While a strategy is typically a plan of action, we do not know the specifics of the algorithms employed by each HFT firm. We can measure only the outcomes of a strategy, although it is natural to recognize that actions we observe are manifestations of a strategy, and hence to treat measures of these actions as representing the strategy. Are strategy and activity necessarily the same? One can presumably construct examples in which they diverge, such as when two distinct strategies produce the same action in one particular state of the world. For example, market making and aggressive news trading could both be associated with a high level of activity when there is an outburst of news (Foucault, Hombert, and Rosu (2016)). However, these two strategies would have completely different patterns of activity in other states of the world (e.g., when there is no news). The strength of our methodology is that we draw on over 36 million observations to estimate the similarities in strategies: the PCA uses every 1-second interval during the market's opening hours, for 30 days (which represent three distinct market conditions), in 52 stocks. Only if the strategies of multiple HFT firms produce similar actions across the many intervals and stocks would our methodology identify an underlying common strategy and point to the firms that follow it. Therefore, while our empirical methodologies exploit correlated activity, our inference is about similarity in the strategies of HFT firms.

### *3.2 Describing HFT Activity*

The main variable we use to describe HFT firm activity in an interval, MSG, emphasizes actions initiated

by an HFT firm, and is defined as the sum of three components: the number of submissions of non-marketable limit orders, the number of cancellations of non-marketable limit orders, and the number of marketable order executions. Hence, MSG for an interval (say one second) describes the total number of messages that an HFT firm sends to the market to initiate a change in its position (either in terms of presence in the limit order book or to transact immediately).<sup>15</sup> We investigate two additional variables that summarize HFT activity to ensure the robustness of our conclusions. The first, TRD, is the number of trades made by an HFT firm in an interval. These trades could occur as the result of marketable orders' submissions or the execution of previously submitted limit orders that rested in the book. The second, LMT, comprises all submissions and cancellations of non-marketable limit orders, and hence describes the actions a firm takes in the interval solely to change its presence in the limit order book.<sup>16</sup>

To economize on the presentation of the results of the PCA, we report only the results for MSG. In Section 4, where the regression and partial correlations analyses enable an efficient presentation of the results for the three measures, we report the results side by side. Our preference for MSG as the main measure is based in part on the fact that most HFT activity comprises submissions and cancellations of orders, not actual trading, and hence a measure that includes these submissions and cancellations would offer a more complete portrayal of their activity. This idea is also highlighted in Brogaard, Hendershott, and Riordan (2016) and Subrahmanyam and Zheng (2016), who stress the importance of HFT limit order behavior.

Before proceeding to the main analysis, however, we establish two stylized facts about similarity in the activity of HFT firms that are directly relevant to the economics of HFT strategies and hence are

---

<sup>15</sup> An AMEND order type is considered two messages, a cancellation and a resubmission, for the purpose of our measures. Our measures include both non-displayed and displayed orders. A refresh of an iceberg order (when the displayed part is executed and shares from the non-displayed part become displayed) does not lead to a change in price or quantity; hence only the initial order is counted. This standardizes the treatment for the various order types according to their economic functions, and enables us to summarize all activity in terms of submissions and cancellations of non-marketable orders and executions of marketable orders.

<sup>16</sup> We also construct a variable to represent HFT activity that consists only of aggressive orders: marketable orders, Immediate-or-Cancel (IOC) orders, and Fill-or-Kill (FoK) orders. It is important to note that these types of orders make up only a small fraction of the total messages of HFT firms. However, there are eight HFT firms for which the ratio of aggressive orders to all orders is higher than 10%, with one firm submitting 54.93% aggressive orders, but these are not the large HFT firms. The PCA results indicate that these eight firms carry out unique strategies that are not related to one another in terms of the patterns of when or in which stocks they submit the orders (they do not follow the underlying common strategies we identify). Our analysis suggests that trying to differentiate strategies based on aggressive orders is less useful than including all messages HFT firms send to the market, perhaps because they reveal only a partial picture of the HFT algorithm. We provide the analysis using aggressive orders in the Online Appendix.

relevant to the design of our tests in the rest of the paper. First, we consider whether the predominant similarity in HFT strategies involves directional activity or total activity. Second, we examine whether similarity in HFT strategies changes on days in which the market experiences large positive or negative returns. We construct two similarity measures: market-wide and stock-specific. Our market-wide similarity measure provides information pertaining to whether the strategies of HFT firms are correlated across stocks at a given time. For each one-second time interval, we compute the correlation coefficient between the activities of each pair of HFT firms across the stocks in the sample, and average the correlations for all pairs of firms in a certain group (e.g., our market share subgroups MS1, MS2, and MS3). The stock-specific similarity measure provides information pertaining to whether the strategies of HFT firms are correlated over time for a given stock. For each stock, we compute the correlation coefficient between the activities of each pair of HFT firms over all time intervals, and average the correlations of all pairs of firms in a certain group.<sup>17</sup>

Our analysis of directional activity is motivated by empirical studies on herding in financial markets (e.g., Wermers 1999; Khandani and Lo 2007, 2011; Choi and Sias 2009; Pedersen 2009; Brown, Wei, and Wermers 2014) that recognize the destabilizing influence that simultaneous actions in one direction (buying or selling) by institutional investors can have on asset prices. In Figure 2, we compare the similarity in total HFT activity (buy plus sell orders) and directional HFT activity (buy minus sell orders).<sup>18</sup> We observe a striking result: the similarity measures involving total activity are four to twelve times the magnitude of those involving directional activity. For example, Panel A of Figure 2 shows that when we use all messages a firm sends to the market to represent HFT activity, the market-wide similarity

---

<sup>17</sup> Specifically, the market-wide similarity measure is computed as follows. Let  $Y_{ik}$  be the activity of HFT firm  $k$  in stock  $i$  and interval  $t$ . If  $k$  and  $l$  are two HFT firms, then  $Cor(Y_k, Y_l)_t$  is the correlation for a specific interval  $t$  over the stocks in the sample. There are  $0.5 * N * (N - 1)$  such correlations for a group of  $N$  HFT firms in each time interval  $t$ , and our market-wide similarity measure for each interval  $t$ ,  $MktSim(Y)_t$ , is computed as the average of these correlations for all pairs of HFT firms. The stock-specific similarity measure is constructed in an analogous fashion. If  $k$  and  $l$  are two HFT firms, then  $Cor(Y_k, Y_l)_i$  is the correlation between their activities in stock  $i$  over all intervals in the sample period. There are  $0.5 * N * (N - 1)$  such correlations for a group of  $N$  HFT firms, and our stock-specific similarity measure for stock  $i$ ,  $StockSim(Y)_i$ , is computed as the average of these correlations for all pairs of HFT firms.

<sup>18</sup> Since limit-order submissions and cancellations have opposing economic implications, we construct the directional activity measure such that we count the cancellation of a limit buy (sell) order by adding to the number of limit sell (buy) orders. For robustness, we also compute a directional activity measure such that submissions and cancellations of buy (sell) orders are both counted as buy (sell) orders, and the resulting correlations are very close. In particular, the result we present in Figure 2, which indicates that the similarity measures for total activity are much larger in magnitude than the similarity measures for directional activity, holds for both measures of directional activity.



measure of total activity for the largest HFT firms (MS1) is 0.359 versus 0.073 for the directional message activity. A similar pattern is observed for the stock-specific similarity measure reported in Panel B. The picture that emerges is consistent with HFT firms operating simultaneously on both sides of the market or buying and selling very rapidly within the same one-second interval (rather than pursuing very strong directional trading for long periods of time). Many HFT firms design strategies to interact with uninformed order flow. Market microstructure theory often specifies that uninformed order flow is non-directional (in contrast to informed order flow), which could potentially explain our finding of low directional similarity of HFT strategies. Given these findings, we focus in the rest of the paper on analyzing similarity in total activity.

Our analysis of similarity in the activity of HFT firms in three distinct market conditions is motivated by the concern that HFT firms react to adverse market conditions (in terms of declining prices) by changing their strategies and hence cause greater fragility. Our study is designed to address this issue because we include in our sample period the 10 worst days in terms of index returns from June 2010 through March 2011. Figure 3 presents the similarity measures computed separately for the down, flat, and up days, alongside the similarity measures for the entire 30-day period. We see, for example, that the market-wide similarity measure for MSG of the largest HFT firms (MS1) are almost identical on down, flat, and up days (0.354, 0.361, and 0.362, respectively). We do not observe that these three distinct market environments result in correspondingly distinct similarity measures. This stylized fact, like the previous one concerning lack of similarity in directional trading, is reassuring in terms of market fragility, although it does not preclude the possibility that similarity in HFT strategies would increase during times of extreme stress, such as the Flash Crash in the U.S. The absence of such extreme episodes in Canadian markets, however, prevents us from examining this possibility empirically. Given these findings, results in the rest of the paper are presented for the entire sample period rather than by market condition.<sup>19</sup>

### *3.3 Principal Components Analysis*

Our first goal is to ascertain whether multiple HFT firms pursue similar strategies that define a particular product category, and to determine how many such product categories can be identified in the HFT

---

<sup>19</sup> We verified that the results of all our tests do not differ materially on down, flat, or up days. We also classified all 10-minute intervals into three categories based on volatility and three categories based on volume, and examined whether similarity in HFT strategies differs when we focus on intraday periods distinguished by higher volume or volatility. We could not discern any clear patterns across the categories.

industry in Canada. We use a data-driven methodology, Principal Component Analysis (henceforth, PCA), to help us answer these questions. We use such a data-driven methodology because of the difficulty of explicitly characterizing the strategy of an HFT firm just by looking at millions of messages that each HFT firm is sending to the market every day. The space of possible signals to which an HFT firm could respond is enormous, and interaction between the strategies of multiple HFT firms means that a complete characterization may require examining the activities of all HFT firms simultaneously. Unfortunately, this is simply not something that a person can do by glancing at the data. Furthermore, for a researcher to study a strategy by hypothesizing a particular economic goal and devising an algorithm to carry out a successful strategy would likely require the same effort an HFT firm invests in developing an algorithm. Even if one were to develop such an algorithm, the activity implied by the algorithm may not correspond exactly to any of the strategies currently employed by HFT firms that are active in the market. Hence, a more promising approach is to use a data-driven methodology to characterize the similarity between the strategies of the HFT firms and categorize them into subgroups.

In general, PCA may be used to discover and summarize the pattern of intercorrelations between variables. It helps in grouping variables that are correlated with one another, presumably because they are driven by the same underlying phenomenon. In our application, we think about the HFT firms' activities as the "variables" that we seek to summarize, and the underlying phenomenon as the strategies that have the same objective and hence belong to the same product category. PCA can also be used to create synthetic variables that represent the underlying common strategies in each of the product categories. This analysis helps us understand HFT strategies in several ways. First, it tells us whether there are underlying common strategies that HFT firms follow. Second, it identifies the firms in each product category and indicates how close each firm's strategy is to the underlying common strategies. Third, it enables analysis of the empirical representations of these underlying common strategies (i.e., the principal component scores) and exploration of how they relate to the market environment as a means of learning about the economic nature of the strategies.

The input for the PCA is a matrix of HFT firms' activities with 31 columns (one for each HFT firm) and 36,504,000 rows (the number of stocks times the number of 1-second intervals in the sample period). In other words, the stocks are stacked one on top of another and we analyze the time-series and cross-sectional sources of variation in HFT activity together. To describe the activity of HFT firms in this

section we concentrate on MSG, which comprises all messages an HFT firm sends to the market in an interval. While the choice of activity representation to use for the PCA could matter, we find that the HFT firms we identify as following underlying common strategies using one representation are also identified as such using others. For example, the identification of firms that follow the three underlying common strategies is identical whether we use MSG or LMT, and using TRD we identify only one fewer firm (depending on the cutoff for the loadings) than we do with MSG.<sup>20</sup> All these measures that one can use as an input for the PCA are simply different representations of the activity generated by the HFT firms' strategies, and we find that they lead to a similar inference regarding which HFT firms follow common strategies. Therefore, to economize on the exposition in terms of tables and discussion, we present in this section the analysis using MSG.

As with any data-driven methodology, it is important to note the economic meaning of various implementation choices. For example, the PCA decomposes the correlation matrix (rather than the covariance matrix) of the strategies. In other words, the variables that describe the activity (or strategy) of each HFT firm are standardized to have zero mean and unit variance. Standardizing the variables eliminates the possibility that one of them would dominate the procedure because it has a much higher scale or range. From an economic perspective, this means that our procedure gives the very active HFT firms the same weight as any other HFT firm (each contributing one unit of variance to the total variance). This helps us focus on similarity in strategies even if some firms are larger than others. The manner in which we pool all stocks for the estimation of the PCA means that it can take into account various strategies that have a "portfolio" flavor to them. For example, a strategy that involves market-wide signals (e.g., from an ETF) that generates activity in multiple stocks in the same intervals could easily be picked up by the procedure.

When conducting the PCA, we must choose how many principal components to retain for subsequent analysis. The corresponding economic question is how many separate underlying strategies

---

<sup>20</sup> We also examine several additional measures for robustness. For example, when we implement the PCA for the buys component of MSG only (or the sells component), we obtain as an output precisely the same HFT firms following the three underlying common strategies. Similarly, the same HFT firms that we find using the MSG measure also follow the principal components when we analyze just the removal of HFT firms' liquidity from the book, which may be of interest to market regulators concerned about market stability. We also implement the PCA with depth in the book (up to 10 price levels from the best prices), and find a similar inference except that most of the firms with large loadings on the first principal component when using MSG load on the second principal component in the depth analysis, and vice versa.

are common to a significant number of the HFT firms in our sample. We conduct a Scree test by plotting the eigenvalues of the principal components and looking for a natural break.<sup>21</sup> Each of the 31 firms contributes  $100/31=3.2\%$  of the variance. However, meaningful principal components would naturally explain more of the variance. We find that the first principal component explains 11.66% of the variance, the second 4.5%, and the third 4.02%, together accounting for over 20% of total variation. Additional principal components account for less than 4% of the variance each, and there seems to be a natural break after three components. Hence, we extract three principal components for subsequent analysis.

The finding that about 20% of the variance is explained by the first three principal components reflects the economic nature of this industry. There are applications of PCA in which researchers drop some of the variables in order for the procedure to explain more of the variance in the remaining set. In our implementation, this is counterproductive. It is important to stress that we do not use PCA to identify HFT firms, but rather we use our knowledge of the trading entities to determine the sample of HFT firms. We are therefore not interested in maximizing the variance explained; we want the methodology to reflect the actual intercorrelations in the set of all HFT firms that operate in Canada. We use the PCA to ascertain the number of underlying common strategies and identify the firms that follow each common strategy. Part of our contribution is in demonstrating the extent of heterogeneity in HFT strategies and the fact that there are many firms with more unique strategies.<sup>22</sup>

Another choice we need to make when implementing the methodology involves determining the rotation of the principal components. The rotation is meant to help us interpret the loadings, which are the coefficients of each HFT firm on each of the principal components. We utilize the commonly-used varimax orthogonal rotation. The loading of an HFT firm on a particular principal component using this rotation has a simple interpretation: it is equivalent to the bivariate correlation between the HFT firm's activity and the underlying common strategy represented by the principal component. Our conclusions are

---

<sup>21</sup> The question whether the activity of the HFT firms lends itself to study using a PCA can be assessed with Bartlett's test of sphericity, which is highly significant in our sample ( $<0.0001$ ). This indicates that a meaningful intercorrelated structure exists. In addition, the Kaiser's measure of sampling adequacy is greater than 0.8 (MSA=0.836), which suggests that the PCA can produce reliable results. However, these tests do not provide information on the number of principal components to retain.

<sup>22</sup> One way to judge how well the principal components represent the activity of the HFT firms is to look at the "commonality" of each HFT firm, which is the percent of the variance of the HFT firm's activity that is accounted for by the principal components that are retained. We note that for approximately one-third of the HFT firms, the retained principal components explain more than 30% of the variation in the activity of each firm. A more detailed exposition of the commonalities is provided in the Online Appendix.

robust to using other rotations.<sup>23</sup>

Table 2 presents the loadings from the PCA using 1-second intervals. Each principal component can be viewed as representing an underlying strategy that is common to multiple HFT firms and which results in the correlated behavior of the firms, and therefore we use the terms “underlying common strategy” and “principal component” interchangeably. The larger (i.e., closer to 1) the loading of a particular HFT firm on a principal component, the greater is the similarity of the firm’s strategy to that underlying common strategy. For each principal component, the loadings are sorted from the most positive to the most negative, and the 31 firms are represented by F01 through F31. The largest loading on the first principal component, for example, belongs to firm F14 and signifies 0.76 correlation between the strategy of this firm and the underlying common strategy represented by the principal component.

While the choice of a cutoff for the magnitude above which a loading is considered economically significant is somewhat arbitrary, it helps to have some cutoff in mind when examining the results. We consider loadings to be significant from an economic standpoint if they are greater than or equal to 0.35, and mark them with an asterisk in the table. Eight HFT firms have significant loadings on the first principal component, suggesting that it represents a strategy that is common to these eight firms. Similarly, six firms exhibit significant loadings on the second principal component, ranging from 0.71 to 0.35, while four HFT firms exhibit significant loadings on the third principal component. Four firms appear to follow more than one strategy: F31 and F28 (significant loadings on the first and second principal components) and F17 and F20 (significant loadings on the first and third principal components).<sup>24</sup>

To summarize the discussion so far, the first important takeaway from the PCA is that there are

---

<sup>23</sup> We also conducted the analysis using an oblique rotation (promax) that allows the principal components to be correlated. The results are basically identical to those with the orthogonal rotation: the loadings are similar in magnitude and the same HFT firms load on the same principal components. Furthermore, the regressions on component scores that we present in Section 3.3 yield the same results irrespective of whether an orthogonal or an oblique rotation is used.

<sup>24</sup> Our conclusions are not sensitive to the exact cutoff we use to denote an economically significant loading. For example, using a cutoff of 0.3 results in 4 additional firms identified as following the underlying common strategies. The market share of this set of firms is slightly higher than that obtained with the 0.35 cutoff, but our conclusions in the rest of the paper are not affected. In particular, the regressions in the next section that help us decipher the economic meaning of the principal components use all HFT firms and hence our interpretation of the principal components is unaffected by the choice of a cutoff. In the Online Appendix, we discuss additional robustness analysis using Sparse Principal Component Analysis (SPCA), which alleviates the need to specify an explicit cutoff. We find that the set of HFT firms that follow the underlying common strategies identified by the SPCA is identical to the one identified by the standard PCA methodology. Given the wide familiarity with the standard procedure and the nearly identical results, we present in the tables the results of the standard PCA methodology.

14 HFT firms competing in three underlying common strategies, with each underlying common strategy followed by multiple firms.<sup>25</sup> On the other hand, 17 HFT firms do not appear to pursue any of these common strategies but rather pursue more unique strategies. How important are the common strategies relative to the unique strategies? The 14 firms that compete in one of the three common strategies represent most of the HFT activity: 96.21% of the messages that HFT firms send to the market and 78.97% of the volume they trade. Therefore, we find that the HFT firms that concentrate their activity in three distinct product categories generate most of the HFT activity. To obtain a deeper understanding of these product categories, we turn to characterizing how the underlying common strategies represented by the principal components are related to the market environment.

### *3.4 Regressions on Principal Component Scores*

The challenge in interpreting the results of a PCA lies in understanding the economic nature of the principal components. To gain additional insights, we use another output of the methodology: the component scores. Each principal component is essentially a linear combination of the observed variables (the 31 HFT firms' activity variables), and the component scores are computed from these variables using the estimated loadings in Table 2 as weights. Since there is a separate score for each principal component in each stock and time interval, we can use regressions to examine the relationships between these component scores (as dependent variables) and the market environment. One advantage of using the principal component scores for this purpose is that they are constructed from the estimated loadings and activity variables of all HFT firms. Hence, our analysis of the economic meaning of the underlying common strategies is robust to whatever choice we make on the cutoff that defines which loadings are considered economically significant.

The first three explanatory variables in our regression specification provide information about the HFT firms. We are interested in characterizing the extent to which HFT firms operate simultaneously across trading venues. The first variable, *HFTcrossmsg*, counts messages sent only by HFT firms that submit messages to multiple trading venues in the same interval. Similarly, we are interested in the degree to which HFT firms trade passively by supplying liquidity. The second variable, *HFTliqsupply*, counts the

---

<sup>25</sup> These findings are robust to estimating the PCA separately on the 10-day up, down, and flat days. For example, we obtain 14 firms following the three principal components when we estimate the PCA on up days, and we get the same 14 firms plus one additional firm when we estimate the PCA on down days. The same set of HFT firms is also identified if we eliminate trading towards the end of the day, or the 10-minute intervals at the open and before the close. These results are available in the Online Appendix.

executed limit orders of all HFT firms in each interval. The third variable,  $|HFTinventory|$ , is the absolute value of the aggregate inventory position of all HFT firms in Canadian dollars (cumulative from the beginning of the day and assuming that all of them start the day with zero inventory).

The next two explanatory variables represent the degree of integration of the Canadian market: *PriceAlign* is the percentage of time that the three trading venues with the highest market share post the same bid or ask price, while *SpreadAlign* is the percentage of time that these three trading venues have the same bid-ask spread. The next three explanatory variables represent the state of liquidity in the market (aggregated across all trading venues): total depth up to 10 cents from the market-wide best bid or offer (MWBBO), the magnitude of depth imbalance at the MWBBO (defined as the absolute value of the difference between the number of shares at the bid and at the ask), and percentage MWBBO spread.<sup>26</sup> The last five explanatory variables represent market conditions in the interval: return (computed from the last transaction price in each interval), volatility (computed as the absolute value of return), the average time between trades in the interval, the return on an ETF that tracks the S&P/TSX Capped Composite Index (XIC), and the ETF return volatility.<sup>27</sup> The ETF return and volatility variables can also be viewed as potential signals for short-term speculation by HFT firms. All explanatory variables are standardized to have zero mean and unit standard deviation to make the coefficients comparable across variables.

Table 3 shows the regression coefficients together with *t*-statistics computed from double-clustered standard errors (along both the stocks and intervals dimensions) to focus our attention on the most significant relationships. Looking at the regression on the first principal component scores, we observe that the coefficient on HFT cross-venue activity (*HFTcrossmsg*) is positive and highly statistically significant. HFT firms that load on this principal component are also more active when prices are more volatile (a positive and significant coefficient on  $|Return|$ ), and there is less depth in the book (*Depth10*) and greater imbalance at the top of the book ( $|TopDepthImb|$ ). These are times at which some

---

<sup>26</sup> All measures are time-weighted. Since messages are time-stamped to 10 milliseconds, there are orders that are submitted and cancelled (or executed) within the same time stamp. When we consider orders that stay in the book to provide liquidity (e.g., for our measures of depth), we assume that a submitted order stays in the book for 10 milliseconds. The exceptions are the following special order types: immediate-or-cancel orders (IOC), fill-or-kill orders (FOK), all-or-nothing orders (AON), dealer's AG orders (that are generated to fulfill their market-making obligations and execute against an incoming order), odd lot orders (OL), and marketable orders, which are assumed to be executed or canceled upon arrival to the market and hence are not added to the limit-order book.

<sup>27</sup> XIC tracks approximately 200 large-cap stocks traded on the TSX. The correlation in daily returns between the S&P/TSX60 and XIC is 0.99 for our 30-day sample period. We compute 1-second interval return and volatility (the absolute value of return) for XIC in exactly the same manner as for the sample stocks.

trading venues post better spreads than others (a negative and significant coefficient on *SpreadAlign*) and there appears to be a need to move liquidity across trading venues, which the HFT firms indeed do with their more intense cross-venue message activity. This suggests that the first principal component could represent a cross-venue arbitrage strategy.

We want to be clear that, as is always the case with PCA, the principal components do not possess an inherent economic interpretation. In every application of this methodology, however, researchers are tasked with giving them one based on the output. We evaluate the regression coefficients of the principal component scores on attributes of HFT activity and market conditions, and our interpretation simply reflects our reading of these coefficients. Whether these labels (e.g., cross-venue arbitrage) are accurate or not does not detract from our use of principal components as data-driven representations of underlying common strategies in a way that can help us understand the competitive landscape of HFT firms. Our use of the principal components in the rest of the paper does not depend on these labels.

Turning to the second principal component, we find that the coefficient on trading passively by providing liquidity (*HFTliqsupply*) is large and highly statistically significant. This underlying common strategy is more active when prices and spreads are more tightly aligned across the trading venues (positive coefficients on *PriceAlign* and *SpreadAlign*), and may represent times when market-making activity is most profitable (i.e., when market-making firms earn the spread rather than lose to changing prices). This strategy is more active when there is greater depth in the book and a smaller imbalance in the top prices, creating ample opportunities for market making without excessive risk.<sup>28</sup> This combination of coefficients, and especially the large coefficient on passive trading that is traditionally associated with market makers, could indicate that this underlying common strategy is related to market making.

The underlying common strategy represented by the third principal component has a combination of statistically significant coefficients that suggests a distinct strategy, although it shares some attributes

---

<sup>28</sup> Note that the coefficients on spread and volatility are not statistically different from zero. When spreads are wider, market makers' profit potential may be higher, and as a result they increase their activity. By supplying more liquidity, however, they work to narrow the spread. The contemporaneous relationship we estimate could incorporate both of these contrasting effects, resulting in a coefficient that is indistinguishable from zero. As for volatility, it is well known that market makers profit from transitory volatility (e.g., they buy when prices are pushed down too much and sell when they reverse). At the same time, market makers lose to informed traders (or simply faster traders) when the source of volatility is informational. Our volatility measure combines both transitory and permanent volatility, which could explain why we do not see a strong effect as these two would have opposite implications for the intensity of market makers' activity.



with the first principal component. HFT firms that load on the third principal component are also more active when prices are more volatile, when depth in the book is lower, and there is greater imbalance at the top prices. Two of the economic relationships that make the third principal component different are a significant negative coefficient on *PriceAlign*, which means that the best prices are not the same across the trading venues, and a negative and significant coefficient on *HFTliqsupply*, which indicates more active (rather than passive) HFT trading. This pattern of relationships with the market environment could represent a short-horizon directional strategy (e.g., momentum) that requires quick trading using marketable orders and is more profitable when there is less depth throughout the book and prices are fragmented across the trading venues. Furthermore, this is the only principal component with a significant coefficient on an ETF variable. The positive relationship with ETF volatility indicates that the short-horizon directional speculation could involve signals on market movements.

There could certainly be other labels for the underlying common strategies that would fit such patterns in regression coefficients. The important takeaway from the regressions, however, is not the labels. Rather, it is the recognition that there is heterogeneity in the underlying common strategies, as evidenced by the distinct relationships we document between each principal component and the variables that represent various aspects of the market environment. Furthermore, there appears to be a significant number of HFT firms that pursue unique strategies, although they represent a small portion of HFT activity. The heterogeneity in HFT strategies is important insofar as the insights generated by thinking about HFT as a single “entity” could be rather limited because aggregating HFT activity hides diversity that is important for understanding how HFT firms interact with markets and ultimately affect them.

### *3.5 Robustness*

In this section we provide additional discussion concerning the nature of the variables that represent HFT activity as well as the robustness of the PCA results.

#### *3.5.1 HFT Activity Variables<sup>29</sup>*

The data we obtain are unique in terms of our ability to identify HFT firms, and we use these data to create three variables that describe their activity (MSG, TRD, and LMT). Since most researchers have access only to publicly available data, an interesting question is whether our analysis can provide guidance on which publicly available measures are best for capturing general HFT activity as well as the

---

<sup>29</sup> We thank a referee for suggesting this analysis.

specific underlying common strategies we identify. Specifically, we look at four measures that can be constructed from publicly available order-level data:

- (i) Ord2Vol: Number of Limit-Order Submissions / (1+Value of Trades (in CAD\$ 100))
- (ii) Canc2Trd: Number of Limit-Order Cancellations / (1+Number of Trades)
- (iii) CancRatio: Number of Limit-Order Cancellations / (1+Number of Limit-Order Submissions and Cancellations)
- (iv) RunsInProgress: Time-Weighted Number of Strategic Runs in an Interval

The first two measures reflect order (or cancellation) intensity normalized by trading. Related measures have been used widely in the literature to describe algorithmic trading and HFT (e.g., Hendershott, Jones, and Menkveld (2011), Hagstromer and Norden (2013), and Boehmer, Fong, and Wu (2015)). Unlike the first two measures, CancRatio does not utilize trade executions but rather focuses on the intensity of cancellations normalized by limit-order submission. This measure is meant to capture the characterization of HFT firms in the SEC (2014) concept release as engaging in “the submission of numerous orders that are cancelled shortly after submission.” The last measure was proposed by Hasbrouck and Saar (2013) and is based on strategic runs, which are linked submissions, cancellations, and executions that are likely to be parts of a dynamic algorithmic strategy and that can be identified with publicly available order-level data. RunsInProgress is constructed for each interval by time-weighting the number of strategic runs that are longer than 10 messages.

We use our order-level data, without taking advantage of the special user IDs or various other flags, to compute these four publicly available measures for the 52 sample stocks and all 1-second intervals in the sample period. Panel A of Table 4 presents Pearson and Spearman correlations showing that all four measures are positively correlated, with the highest correlations observed between the two measures that have executed trades in the denominator (Ord2Vol and Canc2Trd). Panel B of Table 4 presents the correlations of the publicly available measures with the three representations of HFT activity we use in the paper (MSG, LMT, and TRD) aggregated across the 31 HFT firms and the component scores representing the three underlying common strategies. All four publicly available measures are positively correlated with MSG, our main measure representing the activities of HFT firms, with Pearson correlations ranging from 0.178 to 0.675 and Spearman correlations from 0.424 to 0.906. Canc2Trd (the ratio of limit order cancellation to the number of trades) has the highest correlations with MSG. Ord2Trd

is highly correlated with MSG, but is negatively correlated with TRD, likely because it has volume in the denominator. In terms of underlying common strategies, the three ratios (Ord2Vol, Canc2Trd, and CancelRatio) are correlated the most highly with the first principal component that we associate with cross-venue arbitrage. RunsInProgress has the highest correlation with the second principal component that we associate with market making.

Two points are important to note with respect to these correlations. First, the magnitude of the Pearson correlations is sensitive to the specific functional form used. For example, taking a log transformation of the variables increases the correlation between MSG and RunsInProgress from 0.178 to 0.401, bringing it closer to the Spearman correlation (0.424). Second, which publicly available measure is most highly correlated with our HFT variables seems to depend on the length of the interval. In Panel C of Table 4 we present the Spearman correlations for the 1-second, 10-second, and 60-second intervals (I1, I2, and I3, respectively). As we lengthen the interval, RunsInProgress becomes the most highly correlated measure with the three HFT measures. It is also the only one that remains positively correlated with HFT trading activity and with all three underlying common strategies. Canc2Trd remains highly correlated with MSG, but becomes negatively correlated with HFT trading and with the underlying common strategy we associate with the third principal component.

We conclude that RunsInProgress appears to be a reasonable proxy for HFT activity, and becomes an even better proxy as one lengthens the time intervals. The two ratios involving cancellations (normalized by the number of trades or the number of limit orders submitted) also perform rather well, although they are best when used over very short intervals.

### *3.5.2 Event-Time Analysis*

HFT strategies respond to events in the market: submissions, cancellations, and executions of orders (see, for example, Hasbrouck and Saar (2013)). A natural concern is whether having many intervals without any events affects our PCA inference about the HFT strategies. We use several modifications of our methodology to assess this concern. The first, and most straightforward, modification is to eliminate all intervals without events. If the only intervals used in the estimation are those with events, the analysis de-facto captures only occurrences that are concurrent with events. The second methodology is to aggregate events into buckets, in the spirit of Easley, Lopez de Prado, and O'Hara (2012). One advantage of this methodology is that the bucket size can be made stock-specific and hence reflect the unique pace of

activity in each stock. We use two definitions of events to examine the robustness of our results, one using all messages and the other focusing only on trades. Detailed descriptions of the methodologies as well as tables with regression results are provided in the Online Appendix.

We find that our results are very robust to working in event time. In terms of identifying the HFT firms that follow the underlying common strategies, we find that exactly the same 14 HFT firms identified in Table 2 have significant loadings when we eliminate the intervals with no activity. Sixteen HFT firms have significant loadings when we aggregate messages into buckets: the same 14 firms from our regular analysis and two firms that in our regular analysis are just below the cutoff (with 0.33 and 0.32 loadings). Thirteen firms have significant loadings when we aggregate trades into buckets: twelve of the same 14 firms that we identified before and one additional firm that had a 0.33 loading in the regular analysis. The bottom line is that the event-time PCA identifies almost an identical set of firms.

Furthermore, we use the principal component scores regressions to investigate whether the principal components that come out of the event-time analysis represent the same underlying common strategies that we identify in the regular PCA. We find that taking out intervals without activity has no impact on the direction and significance of the coefficients in the regressions. Moreover, statistically significant variables that help us interpret the underlying common strategies in the original analysis are also significant and point to the same interpretation in the event-time regressions. This is a very reassuring result: the inference about the underlying common strategies is the same irrespective of whether we use regular time intervals or event time.

### *3.5.3 Additional Regression Specifications*

Of the 52 stocks in our sample, 37 are cross-listed in the United States. A natural question is whether HFT strategies are different in some way in these stocks, and in particular whether they respond to price discrepancies between the Canadian and U.S. markets. We use the TAQ database to compute 1-second returns on these stocks in the U.S. market, and create two additional variables: the return differential between the Canadian and U.S. markets and the absolute value of the return differential. We then add these two new variables to the principal component scores regressions for the 37 stocks. We find that the absolute value of the return differential is statistically significant in the regression on the first principal component, which we associated with cross-venue arbitrage. This is an intuitive result: cross-listing in the U.S. provides yet another trading venue, and the same HFT firms that perform cross-venue arbitrage in

Canada appear to serve the same function with respect to the U.S. market. The other coefficients show no major differences in signs or significance between the full sample results and the regressions on the cross-listed stocks, and hence we provide these robustness results in the Online Appendix.

We also run additional specifications in an effort to sharpen the distinction between the first two principal components by considering an interaction term between *HFTcrossmsg* and *HFTliqsupply*. Unfortunately, the interaction is highly correlated with *HFTliqsupply*, and therefore we first regress the interaction on both variables and then incorporate the residuals into the full regression (see Lance (1988)). We observe a positive and significant coefficient on this orthogonalized interaction in the regression on the first principal component, suggesting that when cross-venue activity is particularly high, this underlying common strategy tends to experience an increase in passive trading. This is consistent with the negative coefficient on *SpreadAlign* in the sense that some cross-venue messages could be used to move liquidity across trading venues. The complete results are presented in the Online Appendix.

#### 3.5.4 Net Trading Revenues

Given that our PCA methodology identifies three underlying common strategies, we examine whether the firms that follow these strategies have higher or lower net trading revenues than firms that deploy more idiosyncratic algorithms.<sup>30</sup> We find that firms that follow the first and third principal components have net trading revenues that are significantly higher than the net trading revenues of the firms that do not follow any of the underlying common strategies, and also higher than the firms that follow the second principal component. The higher net trading revenues of firms that follow at least two of the underlying common strategies could explain why we observe multiple firms in each product category: these are strategies with a higher revenue potential, and therefore they elicit more intense competition. A detailed exposition of this analysis and results can be found in the Online Appendix.

### 4. HFT Competition and the Market Environment

Having established in Section 3 that there are three distinct product categories, we proceed to investigate how competition between HFT firms in each product category affects the market environment. As

---

<sup>30</sup> To compute net trading revenues (or what is left when shares are bought and sold by a firm), we incorporate trading venue fees and liquidity rebates and assume that shares left (short positions) at the end of the day are “liquidated” (closed) using either the end-of-day midquote or the closing price (Brogaard, Hendershott, and Riordan 2014). It is important to note that while we can calculate an estimate of net trading revenues, a direct analysis of the economic profits of HFT firms or trading desks cannot be conducted as their operating and other costs are unavailable.

discussed in Section 3.1, we rely on results from the industrial organization literature to motivate using the extent of similarity in strategies as a measure of the intensity of competition. The greater the similarity in the strategies of HFT firms, the more intensely these firms compete among themselves. We use these similarity measures to examine whether HFT competition could exacerbate volatility and study how HFT competition interacts with competition between trading venues.

#### *4.1 HFT Competition and Volatility*

Greater similarity in strategies brings about greater competition but also means that multiple HFT firms engage in activities that are highly correlated across stocks. The dominance of HFT firms in the trading environment coupled with this highly correlated activity could potentially amplify idiosyncratic occurrences at the microstructure level and spread them across the market. While the question of whether the magnitude of HFT activity affects volatility has been studied elsewhere (see, for example, Hagstromer and Norden (2013), Hasbrouck and Saar (2013) and Brogaard et al. (2017)), our focus is on the question of how competition between HFT firms relates to volatility. Specifically, we study how the short-horizon volatility of stocks loads on a time-varying market-wide measure that captures the similarity in HFT strategies across stocks in each of the product categories.

For each 1-second time interval, we construct the similarity measure by computing the correlation coefficient between the activities of pairs of HFT firms across the stocks in the sample and averaging the correlations for all pairs of HFT firms in each product category.<sup>31</sup> Specifically, let  $MSG_{ik}$  be the message activity of HFT firm  $k$  in stock  $i$  and interval  $t$ . If  $k$  and  $l$  are two HFT firms that follow the underlying common strategy we associate with the first principal component, then  $Cor(MSG_k, MSG_l)_t$  is the cross-stock correlation between the activities of these two firms for interval  $t$ . There are  $0.5 \cdot 8 \cdot (8-1) = 28$  such correlations for the group of eight HFT firms that follow this underlying common strategy, and the market-wide similarity measure for interval  $t$ ,  $MktSimPC1_t$ , is computed as the average of these correlations for the 28 pairs. We compute  $MktSimPC2_t$  and  $MktSimPC3_t$  in like fashion as the similarity measures for the groups of firms that follow the underlying common strategies we associate with the

---

<sup>31</sup> We also carried out the analysis considering competition between all HFT firms in the market (not just in each product category). The insights we obtain are similar. In the Online Appendix, we provide tables with summary statistics for the similarity measures computed from MSG, TRD, and LMT, and the analysis using all HFT firms and all interval lengths (I1, I2, and I3).

second and third principal components, respectively.<sup>32</sup>

While MSG is a comprehensive representation of HFT activity in that it incorporates every action a firm initiates over an interval, TRD and LMT represent other aspects of an HFT strategy.<sup>33</sup> A priori, it is unclear which one is more important in explaining volatility. On the one hand, TRD focuses on trade executions, and as such it could be more tightly related to price changes. On the other hand, most of the activity of HFT firms involves order flow that does not culminate in trades (limit order submission and cancellation). Including such orders in the representation of the strategy may result in a better description of HFT activity and therefore could be more appropriate when analyzing how HFT activity affects volatility. We use similarity measures computed from MSG, TRD, and LMT to examine whether the results are sensitive to the specific representation of the strategies.<sup>34</sup>

We investigate whether HFT competition affects the short-horizon volatility of individual stocks by running the following regression:

$$|r_{it}| = a_{0i} + a_{1i}MktSimPC1_t + a_{2i}MktSimPC2_t + a_{3i}MktSimPC3_t + a_{4i}MktPC1_t + a_{5i}MktPC2_t + a_{6i}MktPC3_t + a_{7i}MktRet_t + a_{8i}|MktRet_t| + a_{9i}MktVol_t + error_{it} \quad (1)$$

where  $|r_{it}|$  is the absolute value of the interval return for stock  $i$  in interval  $t$  that we use as our measure of short-horizon (or interval) volatility.<sup>35</sup> The right-hand side of equation (1) contains only market-wide variables. Since we are interested in studying the effect of similarity in HFT strategies, not the magnitude of HFT activity, we add as control variables the aggregated activity of the HFT firms that follow the three principal components ( $MktPC1_t$ ,  $MktPC2_t$ , and  $MktPC3_t$ ). As in a typical market model, we include the value-weighted return of all stocks in our sample ( $MktRet_t$ ), and we also control for market volatility (the absolute value of the market return) and market volume (aggregated for all stocks).<sup>36</sup>

We do not believe that the model in equation (1) suffers from endogeneity in the form of either reverse causality or omitted variables. With respect to the former, it is generally accepted that there is no

---

<sup>32</sup> If an HFT firm follows two underlying common strategies, we include its activity in the similarity measures of both product categories. This applies to HFT firms 28 and 31 (PC1 and PC2) as well as 17 and 20 (PC1 and PC3).

<sup>33</sup> The similarity measures computed from MSG, TRD, or LMT have comparable magnitudes. For example, the time-series means of  $MktSimPC1$ ,  $MktSimPC2$ , and  $MktSimPC3$  are 0.320, 0.325, and 0.196, respectively, when we use MSG to represent HFT activity. These summary statistics for TRD (LMT) are 0.386, 0.342, and 0.151 (0.319, 0.324, and 0.196).

<sup>34</sup> Identifying the HFT firms that follow the three underlying common strategies, though, is done based on the principal component analysis of MSG from Section 3.

<sup>35</sup> The returns are computed from trade prices (using the trade closest to the end of an interval).

<sup>36</sup> The results are similar when we use the equal-weighted return rather than the value-weighted return as a proxy for the market portfolio.

reverse causality when one regresses a stock attribute on market attributes. When we run a market model, for example, we normally assume that the stock return does not “cause” the market return. Similarly, the volatility of a single stock in equation (1) does not cause market volatility, nor does it cause the market-wide similarity in HFT strategies that is computed using HFT activity across all stocks in the sample. As for omitted variables, conventional wisdom suggests that the activity of HFT firms is driven by volume and volatility.<sup>37</sup> It is certainly possible that a market volatility shock, for example, could affect both stock volatility (if stocks respond to market conditions) and the extent of similarity in HFT strategies. Once market volatility is controlled for in the regression, however, the shock will not be captured by the similarity measures.

Panel A of Table 5 presents the correlations between the similarity measures and the rest of the market-wide variables used in the regressions. The highest correlation between two similarity measures is 0.308, and the correlations between the similarity measures and the control variables range from -0.215 to 0.130, which does not preclude having them together in a single regression specification. Panel B presents the results of 52 regressions, one for each stock, for the model specified in equation (1). For each independent variable, we provide the average coefficient value, a *t*-statistic computed from the cross-sectional distribution of the regression coefficients, and several summary statistics (e.g., the number of positive and negative coefficients).

We observe a rather striking relationship: the loadings on the HFT similarity measures that represent competition between HFT firms in the three underlying common strategies are all negative. A closer look suggests that for both MSG and LMT, greater similarity between strategies of HFT firms that follow the second principal component (*MktSimPC2*) seems to drive a significant portion of the relationship: its mean coefficient is the most negative, the *t*-statistic computed from the cross-sectional distribution of the coefficients points to a highly significant relationship, and there are 41 stocks (out of 52) for which this variable has negative loadings in the individual stock regressions.<sup>38</sup> This principal component represents the underlying common strategy that we associate with market making based on the

---

<sup>37</sup> The decrease in HFT activity in 2011 and 2012 was attributed to a decline in volume and volatility in the market. See, for example “High-speed trading no longer hurtling forward” by Nathaniel Popper, *The New York Times*, October 14, 2012, and “High frequency trading loses its luster” by Ivy Schmerken, *Wall Street and Technology*, April 1, 2013.

<sup>38</sup> Sign tests and Wilcoxon tests reject the hypotheses that the mean and median of *MktSimPC2* are equal to zero (with *p*-values <0.001). We also reject the hypothesis that the coefficients on *MktSimPC2* are equal to zero with a joint F-test from a Seemingly Unrelated Regression Estimation (SURE).



component scores regressions in Section 3. The magnitude of the effect is also economically meaningful. An increase in *MktSimPC2* from 0 to 0.5, for example, would result in a decrease of 16.1% (MSG) to 34.4% (TRD) in 1-second volatility.<sup>39</sup> As expected, the volatility of individual stocks increases with the volatility of the market. In contrast, the loadings on the HFT similarity measures are predominantly negative, suggesting that HFT competition does not serve as a proxy for market volatility, but rather represents a separate and distinct effect.<sup>40</sup> Our results, therefore, contrast with Jarrow and Protter's (2012) prediction that competition between HFT firms that manifest as more highly correlated HFT activity would increase volatility.

Why would greater competition between HFT firms that pursue market-making strategies decrease rather than increase return volatility? At short horizons (e.g., 1-second intervals), volatility is driven predominantly by the price impact of trades (as opposed to the public release of fundamental news, which is relatively rare for a stock), and competition between HFT firms could affect this price impact. For example, Jovanovic and Menkveld (2016) posit that HFT intermediaries incorporate "hard information," such as price changes in same-industry stocks or the market index, into their algorithms. Greater competition to turn these cross-stock "private" signals into public information implies lower adverse selection costs and hence a lower permanent price impact of trades. The negative loadings could also be driven by the effects of HFT competition on the temporary, rather than the permanent, price impact of trades if they reflect reduced rents from market-making activity or better inventory control. If returns and order flows of various stocks are correlated, efficient market making would necessitate algorithms that consider multiple stocks (e.g., Ho and Stoll 1983). Competition between such market-making algorithms would decrease the temporary price impact of trades (or realized spread), lowering short-horizon volatility. To investigate the channel through which HFT competition affects volatility, we therefore need to consider the price impact of trades.<sup>41</sup>

---

<sup>39</sup> While an increase in *MktSimPC2* from 0 to 0.5 is a rather large change, even a one-standard-deviation increase in *MktSimPC2* would decrease short-horizon volatility between 4.3% (MSG) and 17.8% (TRD).

<sup>40</sup> Out of the 52 stocks, 37 are cross-listed in the U.S. A possible concern is that the model in equation (1) is misspecified for these stocks because it excludes market return, volatility, and volume from the U.S. We therefore use the SPY exchange-traded fund that tracks the S&P 500 Index to construct variables at the one-second frequency for the U.S. market return, volatility, and volume (where we use the SPY volume as a proxy for overall U.S. market volume). The correlations between the U.S. and Canadian variables is modest, which facilitate adding the three U.S. variables to the regressions in equation (1). The addition of the U.S. market variables does not materially affect our results. In particular, most stocks have a negative and statistically significant loading on *MktSimPC2*.

<sup>41</sup> We thank a referee for suggesting this analysis.

Let  $p_t$  be the execution price of a trade,  $m_t$  be the prevailing midquote at the time of the trade execution, and  $m_{t+\tau}$  be the midquote  $\tau$  seconds after the trade. The total price impact of a trade has often been decomposed in the market microstructure literature into permanent and temporary components. The permanent price impact is the difference between the midquote at  $t+\tau$  and the prevailing midquote ( $m_{t+\tau} - m_t$ ), which is attributed to information asymmetry (or adverse selection). The temporary price impact is the difference between the trade price and the midquote at  $t+\tau$  ( $p_t - m_{t+\tau}$ ), and is often thought of as the reversal needed to compensate liquidity providers (also called the realized spread). The length of time used for the price adjustment ( $\tau$ ) has varied historically. SEC rule 11Ac1-5, which became effective in 2001, prescribed computing the realized spread for reporting purposes using  $\tau=5$  minutes. However, with the advent of HFT algorithms that respond to the market in milliseconds, researchers have been using intervals as short as 5 seconds (e.g., O’Hara, Saar, and Zhong (2017)). The permanent or temporary price impact is computed for each trade, and then averaged over all trades (often using volume weighting) to obtain a measure of these price impacts for an aggregation interval,  $T$ , such that  $T > \tau$ . Our paper focuses on very short intervals with lengths of 1, 10, and 60 seconds. To compute measures of permanent or temporary price impacts for a 1-second interval, we need to assume that price adjustments for all trades are completed within the same second (i.e., that  $\tau$  is much smaller than 1 second), which can be unrealistic.<sup>42</sup> We therefore use  $\tau=10$  seconds for the computation of the price impact of each trade and  $T=60$  seconds as the aggregation interval to obtain measures of permanent and temporary price impacts.<sup>43</sup>

To investigate whether the influence of HFT competition on volatility is driven by the manner in which it affects the permanent price impact of trades, we first regress for each stock the return volatility measure on the volume-weighted permanent price impact measure using all 60-second intervals in the sample periods:

$$|r_{it}| = b_{0i} + b_{1i}PPI_{it} + error_{it}$$

We then take the fitted values from this regression,  $|\hat{r}|_{PPI_{it}}$ , which describe the portion of short-horizon volatility that is explained by the permanent price impact, and use them as the dependent variable in the regression specified in equation (1). Similarly, to investigate whether the negative impact of HFT

---

<sup>42</sup> One cannot specify  $\tau > T$  such that a long period is used to obtain the permanent or temporary price impacts for each trade and then average these price impacts for a short aggregation interval because most of the price adjustment would occur outside the aggregation interval.

<sup>43</sup> In the empirical implementation we use relative price impact measures, dividing the \$CAD price impact by the prevailing midquote, to be consistent with the common usage of these measures in the literature.

competition on volatility is driven by a reduction in the temporary price impact of trades, we regress the volatility measure on the volume-weighted temporary price impact measure:

$$|r_{it}| = b_{0i} + b_{1i}TPI_{it} + error_{it}$$

and use the fitted values ( $|\hat{r}|_{TPI_{it}}$ ) as a dependent variable for the regression in equation (1).

Panel A of Table 6 shows the results using the permanent price impact fitted values.<sup>44</sup> We observe that the only underlying common strategy that exhibits a consistent negative relationship using all three representations of HFT activity (MSG, TRD, and LMT) is the one captured by the second principal component, which we associate with market making. In particular, 45 out of 52 stocks have negative loadings on *MktSimPC2* in these stock-by-stock regressions using MSG, and the mean coefficient is highly statistically significant. The loadings on *MktSimPC3* do not have a predominant direction, and those on *MktSimPC1* show a direction only for TRD, and it is positive rather than negative. Panel B of Table 6 shows the results of the regressions using the temporary price impact fitted values as the dependent variable. We observe a similar albeit weaker picture: the mean coefficient on *MktSimPC2* is negative for MSG, TRD, but TRD is not statistically significant.

How do these tests help us interpret the volatility results? Table 5 shows that short-horizon volatility loads negatively on competition between HFT firms that pursue the underlying common strategy that we associate with market making. Market makers affect price formation through their liquidity provision, which often determines the price impact of trades. We find that the portions of volatility explained by both the permanent and temporary price impacts appear to drive the negative relationship with HFT competition. The strong result with respect to the permanent price impact can be viewed in the context of the intermediation model in Jovanovic and Menkveld (2016): greater competition means hard information signals are revealed more quickly, and there is therefore less adverse selection in the trading environment. The temporary price impact result we obtain suggests that lower volatility could reflect competition between market makers reducing the compensation they earn (the realized spread).

To summarize, we find that the short-horizon volatility of most stocks loads negatively on a measure of similarity in the strategies of HFT firms that follow the underlying common strategy we associate with market making. This is contrary to concerns that correlated HFT activity may increase market fragility. Therefore, competition between HFT firms, especially those that pursue market making,

---

<sup>44</sup> The regressions presented in Table 6 use 60-second intervals for both the dependent and independent variables.

could benefit the market in more than one way: not just through the possible reduction in economic rents but also through a reduction in short-horizon volatility. It is important to stress, though, that this result is observed during “normal” times, and may not hold during a market breakdown episode like the U.S. Flash crash in May of 2010.

#### *4.2 HFT Competition and Market Concentration*

Equity trading markets in the U.S., Canada, and many other countries are characterized by multiple trading venues on which stocks can be traded. Such trading fragmentation may create negative externalities in the form of lower price integrity and higher costs as liquidity is scattered across trading venues. Against these negative externalities, proponents of this structure argue that the competition it induces between trading venues results in lower fees and greater innovation in terms of trading technology and services.

HFT is pertinent to the debate on market fragmentation. Brokers route customer orders to trading venues using routing tables that take into account multiple parameters, including past execution quality as well as rebates and fees. Under the best execution requirement, however, brokers need to make reasonable efforts to ensure that customer orders receive the best prices. Therefore, a key criterion in the routing of each order necessitates determining which trading venues display the best price when the order arrives, and brokers are required to take into account information from all venues on which a security is traded. As such, trading venues that display the best prices or the smallest spreads attract more customer order flow from brokers. Given the central role HFT firms play in setting prices and spreads, a trading venue’s market share depends on these firms’ activity. In particular, HFT liquidity provision on new trading venues can be essential to fostering competition between venues.<sup>45</sup> Of course, fragmentation of the trading environment among multiple trading venues creates profit opportunities for HFT firms in moving liquidity from one venue to another and ensuring that prices are the same across the venues. Hence, trading venue proliferation and HFT activity are interrelated. What is less well understood, however, is how competition between HFT firms, which is reflected in the pursuit of similar strategies by multiple firms, affects the competitiveness of trading venues and the concentration of trading in the market. This is the question we explore in this section.

---

<sup>45</sup> Menkveld (2013) notes that a new trading venue in Europe became viable only when a large HFT firm began trading on it.

To examine market concentration of trading volume in stock  $i$ , we compute the Herfindahl-Hirschman Index ( $HHI_i$ ) of market share for the five trading venues: Alpha ATS Limited Partnership (ALF), Chi-X Canada ATS (CHX), Omega ATS (OMG), Pure Trading (PTX), and the Toronto Stock Exchange (TSX). The lower the  $HHI_i$ , the less concentrated the market.<sup>46</sup> To construct a measure of similarity in HFT strategies for each stock, we compute the correlation coefficient between the activities of pairs of HFT firms over all time intervals, and average the correlations for all pairs of HFT firms in each product category. If  $MSG_{ik}$  is the message activity of HFT firm  $k$  in stock  $i$  and interval  $t$ , then  $Cor(MSG_k, MSG_l)_i$  is the correlation between the activities of HFT firms  $k$  and  $l$  in stock  $i$  computed over all intervals. There are  $0.5 * 8 * (8 - 1) = 28$  such correlations for a group of eight HFT firms that follow the underlying common strategy represented by the first principal component, and our stock-specific similarity measure,  $StockSimPC1_i$ , is computed as the average of these correlations for the 28 pairs. We compute  $StockSimPC2_i$  and  $StockSimPC3_i$  in a similar fashion for the groups of firms that follow the underlying common strategies we associate with the second and third principal components, respectively.

It is important to note that the relationship between these two variables—market concentration and HFT competition—is best viewed as being jointly determined in equilibrium rather than as one causing the other. While competition between HFT firms could enhance liquidity provision and hence the viability of smaller trading venues, the multiplicity of trading venues could also increase the profit opportunities from arbitrage activity and market making and lead to intensified competition between HFT firms. Therefore, we examine the cross-sectional association between  $HHI_i$  and the measures of similarity in HFT strategies without assuming causality. Still, we would like to control for several variables that could influence the level of concentration in a stock. Specifically, we use three variables—market capitalization, price level, and the standard deviation of 30-minute returns over the sample period—to control for heterogeneity in fundamental attributes across stocks. We employ two additional variables, time-weighted average spread and market-wide bid-and-offer depth, to control for the liquidity environment of a stock. Lastly, since our goal is to study the effect of similarity in HFT strategies rather than the magnitude of HFT activity per se, we add the magnitude of HFT activity in stock  $i$  as a control

---

<sup>46</sup> The HHI is computed as the sum of squared market shares of the trading venues, and with five trading venues it is always between 0.2 (if volume is equally divided among the five trading venues) and 1 (if volume concentrates on one trading venue). When we use 1-second intervals to compute  $HHI$  for each stock, the average value for the sample is 0.879, and it declines to 0.778 and 0.629 when we use 10-second and 60-second intervals, respectively.

variable.<sup>47</sup>

Table 7 presents the partial correlations between *HHI* and the HFT similarity measure controlling for market capitalization, price level, volatility, spread, depth, and the magnitude of HFT activity.<sup>48</sup> The results are presented side-by-side for three HFT activity variables (MSG, TRD, and LMT) and three interval lengths (I1, I2, and I3). We observe very intuitive results: the partial correlations between *HHI* and *StockSimPC1* are negative and highly statistically significant in all permutations, and there is a similar strong negative relationship between *HHI* and *StockSimPC2*. If one is willing to accept our observation in Section 3 that the first principal component is related to cross-venue arbitrage, then we would expect competition in this underlying common strategy to be negatively related to market concentration because it improves the quality of prices displayed on the smaller trading venues. Competition between HFT firms that follow the second principal component, which we associate with market making, should exhibit a strong negative relationship with market concentration if competition in liquidity provision on smaller venues is important to attracting order flow from brokers. Competition between HFT firms and competition between trading venues appear tightly linked in a “competition begets competition” manner.

We would like to study the channel through which competition between HFT firms could affect the market share of trading venues. As we mention above, best execution practices result in brokers routing customer orders to trading venues that post the best prices. If HFT competition gives rise to better prices on smaller trading venues, it could in principle help these venues increase their market share and thereby reduce concentration in the market. To examine this channel, we investigate whether greater similarity in HFT strategies on each of the five trading venues increases the percentage of time it displays the best prices or narrowest spreads.

Panel A of Table 8 presents summary statistics on the market share of the five trading venues and two measures of their competitiveness (or viability): (i) the percentage of time that the trading venue posts either the best bid or the best ask in the market (where the market is defined as the aggregation of all five trading venues), and (ii) the percentage of time that the bid-ask spread on the trading venue is the

---

<sup>47</sup> For example, if *StockSimPC1<sub>i</sub>* is computed using MSG, we add as a control variable the aggregate number of submissions, cancellations, and executions in stock *i* that the eight HFT firms send to the market.

<sup>48</sup> Unlike the market-wide similarity measures used in Section 4.1, the stock-specific similarity measures for the first two principal components are highly correlated (e.g., 0.71 using MSG). We therefore analyze each one separately rather than having all three measures in a single model.

narrowest spread in the market.<sup>49</sup> We observe that there is a dominant trading venue in Canada with a market share of 69.26% that displays the best price 92% of the time (averaged across all stocks in our sample). Similarly, the largest trading venue has the narrowest spread 76.8% of the time (compared with 54.2%, 36.1%, 26.9%, and 9.4% for the other four trading venues).

To study how HFT competition affects the viability of trading venues, we compute the HFT similarity measures for each venue by considering activity on that venue only. However, many of the HFT firms are active on only some of the trading venues. To ensure that our results could be interpreted in the same manner for all trading venues, we restrict attention in this test to the HFT firms that send at least 10,000 messages to each of the trading venues during our sample period.<sup>50</sup> We construct the measure of similarity in the strategies of HFT firms in each stock as before: computing the correlation coefficient between the activities of pairs of HFT firms over all time intervals, and averaging the correlations for the pairs of HFT firms. The two differences from before are that we compute a separate similarity measure for each trading venue using activity on that venue only, and we focus on a set of eight firms that send messages to all trading venues. We use this similarity measure,  $StockSim_{iv}$ , to represent competition between HFT firms in stock  $i$  on trading venue  $v$ . For each trading venue, we run the following cross-sectional regression:

$$C_{iv} = a_0 + a_1 StockSim_{iv} + a_2 Activity_{iv} + a_3 Spread_{iv} + a_4 BBOdepth_{iv} + a_5 MktCap_i + a_6 Price_i + a_7 Volatility_i + error_{iv}, \quad (2)$$

where  $C_{iv}$  stands for one of the two competitiveness measures for stock  $i$  on trading venue  $v$ .<sup>51</sup> As control variables we include the magnitude of HFT activity on the specific trading venue ( $Activity_{iv}$ ), spreads and depth (now computed from the best prices on a single trading venue), market capitalization, price level, and the standard deviation of 30-minute returns over the sample period.

To present results for the five trading venues and the three interval lengths side-by-side in Panel B of Table 8, we report only the coefficient on  $StockSim_{iv}$  from each regression (with heteroscedasticity-consistent  $t$ -statistics). We identify an interesting pattern: the coefficients for the smaller trading venues

---

<sup>49</sup> We note that two or more trading venues could potentially be at the best bid or ask at the same time (or have the same narrowest spread). We also recognize that activity on one trading venue could affect activity on another trading venue as in the multi-market inventory model of Lescourret and Moinas (2015).

<sup>50</sup> There are eight such HFT firms, four of which follow the underlying common strategy identified by the first principal component and four of which follow the underlying common strategy identified by the second principal component.

<sup>51</sup> We present the results for MSG (total messages sent by the HFT firm) to economize on the size of the table. Results using TRD and LMT are similar in nature.

are positive, although not all of them are statistically significant. For the largest trading venue in terms of market share, however, we find the opposite result: the coefficient is negative and statistically significant in almost all regressions. Therefore, the effect of similarity in HFT strategies on the viability or competitiveness of a trading venue depends on the nature of that venue: it benefits smaller venues that introduce competition into the market and detracts from the dominant position of the largest trading venue.

The disparity in results between the smaller trading venues and the dominant venue suggests that reverse causality is less likely to be a concern in these regressions. We also believe that the economics of trading on these venues is such that the percentage of time a trading venue posts the best prices is determined by the strategies of the HFT firms, rather than vice versa, because HFT firms are the dominant players in terms of liquidity provision. Our results therefore suggest that one driver of the negative relationship we document between market concentration and HFT competition is that HFT competition enhances the viability of smaller trading venues (in terms of displaying better prices and smaller spreads), which increases their market share.

## **5. Conclusions**

In this paper, we examine product market competition in the HFT industry. The difficulty in pursuing such a study has to do with the idiosyncrasies of an industry in which the product is often not publicly defined or easily observed by market participants. While regulators can collect data on the actions of HFT firms, attempting to understand the product attributes (i.e., the essence of the strategy that generates the actions) is often difficult if not outright impossible. Against this backdrop, the idea behind our study is to define the main product categories in this industry by looking at whether multiple HFT firms follow related strategies. The empirical tool we use, principal component analysis, enables us to employ detailed data on the actions of individual HFT firms to identify underlying strategies that are common to multiple firms. In and of itself, our finding that there are at least three main product categories is important because it shows that thinking about HFT as a single entity in terms of its impact on the market may afford limited insights. We show that these three underlying common strategies differ from one another not just in terms of types of HFT activity (e.g., cross-venue messages or passive liquidity provision), but also in terms of the market and limit-order book conditions associated with them. These strategies serve different functions



and should be regarded separately by both regulators and market participants.

While much of the literature in economics regards competition as a good attribute that lowers economic rents and facilitates innovation, there are cases in which such competition can have negative byproducts.<sup>52</sup> Since competition between HFT firms manifests itself in more highly correlated actions across stocks, it is quite natural to be concerned that this competition will result in higher short-horizon volatility. We find the opposite result: the short-horizon volatility of stocks loads negatively on competition between HFT firms that follow the underlying common strategy we associate with market making. This negative relationship appears to be driven by the manner in which HFT competition in this particular product category affects both the permanent and temporary components of the price impact of trades.

Lastly, the market structure for equity trading is currently fragmented among many trading venues. When brokers search for the best prices in the market to deliver best execution to their customers, the dominant role HFT firms play in liquidity provision and price arbitrage in the fragmented marketplace means that they directly influence the market share of trading venues. We find that greater HFT competition on smaller trading venues is helpful for the competitive position of these venues, but the opposite result is observed for the dominant trading venue in Canada. This is likely one of the drivers behind the negative relationship we document between market concentration and HFT competition.

While our findings contribute several important insights to the study of the industrial organization of the HFT industry, much remains unknown. In particular, there is no evidence pertaining to the process of product development and testing in this industry, which has implications for both barriers to entry (or the amount of funding required to develop profitable algorithms) and the stability of the markets (e.g., runaway algorithms that can wreak havoc and cause market disruptions). The importance of specific strategies for the well-functioning of our markets is also not completely understood. We view these as open questions, and hope that future studies will be able to deepen our knowledge of this fascinating industry.

---

<sup>52</sup> For example, product market competition can make it more difficult to infer a manager's action given a firm's output (Golan, Parlour, and Rajan 2015), and in general could exacerbate the problem of incentivizing managers (e.g., Scharfstein 1988).

## References

- Ait-Sahalia, Yacine and Mehmet Saglam, 2016, High Frequency Market Making, Working paper, Princeton University.
- Aldridge, Irene, 2013, *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems*, Second Edition, John Wiley & Sons, Hoboken, New Jersey.
- Anand, Amber, and Kumar Venkataraman, 2016, Market conditions, fragility, and the economics of market making, *Journal of Financial Economics* 121, 327-349.
- Baron, Matthew, Jonathan Brogaard, and Andrei Kirilenko, 2012, The trading profits of high frequency traders, working paper, Princeton University.
- Benos, Evangelos, James Brugler, Erik Hjalmarsson, and Filip Zikes, 2015, Interactions among high-frequency traders, Working Paper No. 523, Bank of England.
- Benos, Evangelos, and Satchit Sagade, 2016, Price discovery and the cross-section of high-frequency trading, forthcoming in the *Journal of Financial Markets*,.
- Besanko, David, David Dranove, Mark Shanley, and Scott Schaefer, 2010, *Economics of Strategy*, Fifth Edition, John Wiley & Sons, Hoboken, New Jersey.
- Biais, Bruno, Thierry Foucault, and Sophie Moinas, 2015, Equilibrium fast trading, *Journal of Financial Economics* 116, 292-313.
- Boehmer, Ekkehart, Kingsley Y. L. Fong, and Julie J. Wu, 2015, International evidence on algorithmic trading, Working paper, Singapore Management University.
- Breckenfelder, H. Johannes, 2013, Competition between high-frequency traders, and market quality, Working paper, Stockholm School of Economics, SHoF & SIFR.
- Brogaard, Jonathan A., and Corey Garriott, 2016, High-frequency trading competition, Working paper, University of Washington and Bank of Canada.
- Brogaard, Jonathan A., Terrence Hendershott, and Ryan Riordan, 2014, High frequency trading and price discovery, *Review of Financial Studies* 27, 2267-2306.
- Brogaard, Jonathan A., Terrence Hendershott, and Ryan Riordan, 2016, Price discovery without trading: Evidence from limit orders, Working paper, University of Washington.
- Brogaard, Jonathan A., Allen Carrion, Tjobait Moyaert, Ryan Riordan, Andriy Shkilko, and Konstantin Sokolov, 2017, High-frequency trading and extreme price movements, forthcoming in the *Journal of Financial Economics*.
- Brown, Nerissa C., Kelsey D. Wei, Russ Wermers, 2014, Analyst recommendations, mutual fund herding, and overreaction in stock prices, *Management Science* 60, 1-20.

- 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.
- Chaboud, Alain P., Benjamin Chiquoine, Erik Hjalmarsson, and Clara Vega, 2014, Rise of the machines: Algorithmic trading in the foreign exchange market, *Journal of Finance*, LXIX, 5, October, 2045-2084.
- Choi, Nicole, and Richard W. Sias, 2009, Institutional industry herding, *Journal of Financial Economics* 94, 469-491.
- Clark-Joseph, Adam D., 2014, Exploratory trading, Working paper, University of Illinois.
- Comerton-Forde, Carole, Katya Malinova, and Andreas Park, 2016, Regulating dark trading: order flow segmentation and market quality, Working paper, University of Toronto and FIRN.
- Easley, David, Marcos M. Lopez de Prado, and Maureen O'Hara, 2012, The volume clock: Insights into the high-frequency paradigm, *Journal of Portfolio Management* 39, pp. 19-29.
- Foucault, Thierry, Johan Hombert, and Ioanid Rosu, 2016, News trading and speed, *Journal of Finance* 71, 335-382.
- Golan, Limor, Christine A. Parlour, and Uday Rajan, 2015, Competition, Managerial slack, and corporate governance, *Review of Corporate Finance Studies* 4, 43-68.
- Goldstein, Michael A., Pavitra Kumar, and Frank C. Graves, 2014, Computerized and high-frequency trading, forthcoming, *The Financial Review*, 49, 177-202.
- Hagstromer, Bjorn, and Lars Norden, 2013, The diversity of high-frequency traders, *Journal of Financial Markets* 16, 741-770.
- Hagstromer, Bjorn, Lars Norden, and Dong Zhang, 2014, How aggressive are high-frequency traders? *The Financial Review*, 49, 395-419.
- Han, Jungsuk, Mariana Khapko, and Albert S. Kyle, 2014, Liquidity with high-frequency market making, Working paper, Swedish House of Finance.
- Hasbrouck, Joel, and Gideon Saar, 2013, Low-latency trading, *Journal of Financial Markets* 16, 646-679.
- Hendershott, Terrence, Charles M. Jones, and Albert J. Menkveld, 2011, Does algorithmic trading improve liquidity? *Journal of Finance* 66, 1-33.
- Hirschey, Nicholas H., 2016, Do high-frequency traders anticipate buying and selling pressure? Working paper, London Business School.
- Ho, Thomas S.Y., and Hans R. Stoll, 1983, The dynamics of dealer markets under competition, *Journal of Finance* 38, 4, September, 1053-1074.

- Hoberg, Gerard, and Gordon Phillips, 2016, Text-based network industries and endogenous product differentiation, forthcoming in the *Journal of Political Economy*.
- Hoffmann, Peter, 2014, A dynamic limit order market with fast and slow traders, *Journal of Financial Economics* 113, 156-169.
- Hotelling, Harold, 1929, Stability in competition, *Economic Journal* 39, 41-57.
- IIROC, 2012, The HOT study: Phase I and II of IIROC's study of high frequency trading activity on Canadian equity marketplaces, IIROC report.
- IIROC, 2014, Identifying trading groups: methodology and results, IIROC report.
- Jarnecic, Elvis, and Mark Snape, 2014, The provision of liquidity by high-frequency participants, *The Financial Review*, 49, 371–394.
- Jarrow, Robert A., and Philip Protter, 2012, A dysfunctional role of high frequency trading in electronic markets, *International Journal of Theoretical and Applied Finance*, 15, 3, 1-15.
- Jones, Charles M., 2013, What do we know about high-frequency trading? Working paper, Columbia Business School.
- Jovanovic, Boyan, and Albert J. Menkveld, 2016, Middlemen in Limit Order Markets, Working paper, New York University.
- Khandani, Amir E., and Andrew W. Lo, 2007, What happened to the quants in August 2007? *Journal of Investment Management* 5, 29-78.
- Khandani, Amir E., and Andrew W. Lo, 2011, What happened to the quants in August 2007? Evidence from factors and transactions data, *Journal of Financial Markets*, 14, 1-46.
- Kirilenko, Andrei, Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun, 2016, The flash crash: The impact of high frequency trading on an electronic market, forthcoming in the *Journal of Finance*.
- Korajczyk, Robert A., and Dermot Murphy, 2016, High frequency market making to large institutional trades, Working paper, Northwestern University.
- Lance, Charles E., 1988, Residual Centering, Exploratory and Confirmatory Moderator Analysis, and Decomposition of Effects in Path Models Containing Interactions, *Applied Psychological Measurement* 12, 163-175.
- Lescourret, Laurence, and Sophie Moinas, 2015, Liquidity supply across multiple trading venues. Working paper. ESSEC Business School.
- Malinova, Katya, Andreas Park, and Ryan Riordan, 2014, Do retail traders suffer from high frequency traders? Working paper, University of Toronto.
- Menkveld, Albert J., 2013, High frequency trading and the new-market makers, *Journal of Financial Markets*, 16, 712–740.

- Menkveld, Albert J., and Marius A. Zoican, 2016, Need for speed? Exchange latency and liquidity, Working paper, Duisenberg School of Finance - Tinbergen Institute Discussion Paper.
- O'Hara, Maureen, Gideon Saar, and Zhuo Zhong, 2017, Relative tick size and the trading environment, Working paper, Cornell University.
- Pedersen, Lasse Heje, 2009, When everyone runs for the exit, *International Journal of Central Banking* 5, 177-199.
- Rosu, Ioanid, 2016, Fast and slow informed trading, Working paper, HEC.
- Salop, S., 1979, Monopolistic competition with outside goods, *Bell Journal of Economics* 10, 141-156.
- Scharfstein, David, 1988, Product market competition and managerial slack, *Rand Journal of Economics* 19, 147-55.
- Securities and Exchange Commission (SEC), 2014, Concept Release on Equity Market Structure, 34-61358.
- Shy, Oz, 1995, *Industrial Organization: Theory and Applications*, MIT Press, Cambridge, Massachusetts.
- Subrahmanyam, Avanidhar, and Hui Zheng, 2016, Limit order placement by high-frequency traders, Working paper, UCLA.
- Tirole, Jean, 1988, *The Theory of Industrial Organization*, MIT Press, Cambridge, Massachusetts.
- Wermers, Russ, 1999, Mutual fund herding and the impact on stock prices, *Journal of Finance*, LIV, 2, April, 581-622.

**Table 1**  
**Summary Statistics**

Our sample consists of 52 stocks from the S&P/TSX60 Index. We rank the daily returns of the S&P/TSX Composite Index from June 2010 through March 2011, and select the 10 worst days (down days), the 10 best days (up days), and the 10 days closest to and centered on zero return (flat days) for a 30-day sample period. Panel A presents summary statistics for the sample stocks: market capitalization, price, the standard deviation of 30-minute returns, daily volume (in Canadian dollars), and daily return. Panel B presents summary statistics for the 31 high-frequency trading (HFT) firms that we identify using data from the Investment Industry Regulatory Organization of Canada (IIROC). These 31 firms are further categorized into three subgroups according to market share of volume: MS1 (market share of at least 4%; 4 firms), MS2 (market share of between 1% and 4%; 6 firms), and MS3 (the rest of the HFT firms). Market share (in terms of volume or trades) is computed by dividing the trading undertaken by each HFT firm by total trading in the market. Trades consist of executions of both marketable orders and non-marketable orders. Sub/Canc is the number of (non-marketable) limit-order submissions and cancellations (i.e., all non-execution messages) that a firm sends to the market, where a modification of an order counts as a cancellation and a resubmission. Messages are all the orders (both marketable and non-marketable) and cancellations that a firm sends to the market. Orders/Trades Ratio is defined for each HFT firm as Messages divided by Trades. The mean order-to-trade ratio is the average of the order-to-trade ratios of the individual HFT firms, and therefore due to its non-linear nature and the heterogeneity of the firms is not equal to the cross-sectional mean number of messages divided by the cross-sectional mean number of trades. CrossEndInventory is the number of times per day that an HFT firm's intraday inventory position crosses its end-of-day inventory position. CrsVenueActivity is the number of daily 10-millisecond time-stamps in which an HFT firm is both (1) sending messages (orders or cancellations) to TSX, and (2) sending messages to another trading venue in the same or adjacent time-stamp. We compute the measures for each HFT firm using all days in our sample period, and then provide in the table cross-firm means and medians for all HFT firms as well as for subgroups by market share.

Panel A: Sample Stocks

Days	Stocks		MktCap (Million CAD)	Price (CAD)	StdRet (30min Ret)	CADVVolume (1000; daily)	Return (daily)
All	S&P60	Mean	19,412	39.1	0.40%	78,156	-0.04%
		Median	11,401	36.9	0.40%	51,120	-0.04%
Down Days	S&P60	Mean	19,203	38.3	0.43%	81,195	-1.72%
		Median	11,299	36.4	0.43%	55,892	-1.56%
Flat Days	S&P60	Mean	19,499	39.7	0.34%	70,614	-0.08%
		Median	11,840	37.4	0.34%	50,078	-0.07%
Up Days	S&P60	Mean	19,535	39.5	0.39%	82,903	1.66%
		Median	11,574	37.0	0.37%	54,097	1.46%

Panel B: HFT Firms

HFT Firms		MktShr (Volume)	MktShr (Trades)	Trades (daily)	Sub/Can (daily)	Messages (daily)	Ord/Trd Ratio	CrossEnd Inventory	CrsVenue Activity
All	Mean	1.50%	1.70%	19,445	1,056,838	1,063,974	92.7	14.9	42,775
	Med.	0.34%	0.44%	5,035	97,146	97,288	40.4	2.4	619
MS1	Mean	6.58%	8.95%	102,035	5,466,547	5,496,423	49.5	73.0	216,133
	Med.	7.23%	8.60%	98,048	3,272,863	3,312,989	42.2	70.1	177,952
MS2	Mean	2.67%	2.08%	23,730	982,137	995,868	38.3	13.4	38,307
	Med.	2.79%	2.11%	24,003	147,090	162,729	6.9	4.3	1,821
MS3	Mean	0.19%	0.22%	2,489	238,236	239,157	116.5	4.3	9,667
	Med.	0.05%	0.04%	714	40,129	40,268	71.4	1.8	168

**Table 2****Principal Components Analysis: Loadings**

This table presents the loadings from a principal component analysis of high-frequency trading (HFT) strategies. The representation of HFT activity that we use to characterize the strategies is MSG, which comprises all messages an HFT firm actively sends to the market in an interval (submission of non-marketable limit orders, cancellation of non-marketable limit orders, and marketable limit orders that result in trade executions). We conduct the analysis using one-second intervals. For the purpose of the analysis, we view the 31 HFT firms as the “variables,” while the observations are all intervals during the 30-day sample period for all sample stocks. The principal component analysis uses the varimax orthogonal rotation, and the first three principal components are retained for further analysis. The loading of an HFT firm on each of the principal components signifies the extent to which the firm’s activity corresponds to the underlying common strategy represented by that principal component; it is equivalent to the bivariate correlation between the firm’s measure and the principal component, and is therefore between  $-1$  and  $1$ . For each principal component, the loadings are sorted from the most positive to the most negative. An asterisk indicates all principal components that are  $\geq 0.35$ .

HFT	PC1 Loading	HFT	PC2 Loading	HFT	PC3 Loading
F14	0.76 *	F27	0.71 *	F17	0.51 *
F16	0.67 *	F08	0.52 *	F23	0.48 *
F04	0.57 *	F31	0.41 *	F19	0.47 *
F24	0.54 *	F05	0.40 *	F20	0.43 *
F31	0.48 *	F02	0.39 *	F26	0.33
F28	0.46 *	F28	0.38 *	F12	0.22
F20	0.40 *	F30	0.34	F08	0.20
F17	0.38 *	F06	0.27	F18	0.18
F29	0.33	F12	0.24	F14	0.15
F01	0.32	F26	0.22	F29	0.13
F21	0.26	F10	0.17	F21	0.12
F27	0.26	F01	0.16	F03	0.12
F06	0.26	F20	0.13	F10	0.08
F07	0.23	F04	0.12	F30	0.05
F26	0.23	F24	0.11	F05	0.03
F08	0.22	F09	0.11	F02	0.03
F11	0.15	F23	0.10	F06	0.03
F12	0.08	F11	0.07	F16	0.02
F19	0.07	F18	0.06	F25	0.01
F05	0.04	F14	0.05	F13	0.01
F22	0.03	F03	0.04	F27	0.00
F10	0.01	F15	0.03	F09	0.00
F15	0.01	F25	0.02	F22	-0.01
F09	0.00	F22	0.01	F15	-0.01
F25	0.00	F13	0.01	F04	-0.02
F03	0.00	F29	0.00	F28	-0.07
F13	0.00	F16	-0.03	F31	-0.09
F02	-0.04	F21	-0.07	F24	-0.13
F30	-0.06	F07	-0.12	F07	-0.13
F23	-0.07	F19	-0.12	F11	-0.13
F18	-0.08	F17	-0.16	F01	-0.30

**Table 3**  
**Regressions using Principal Component Scores**

This table presents regressions of principal component scores on variables that represent the market environment. In the PCA of the high-frequency trading (HFT) activity, the 31 HFT firms are the “variables,” while the observations are all 1-second intervals during the 30-day sample period for all sample stocks. The PCA uses the varimax orthogonal rotation, and the first three principal components are retained for further analysis. There are 13 right-hand-side variables in each regression that describe attributes of the HFT strategies, as well as the market environment. We identify only the HFT firms that simultaneously (i.e., in the same interval) submit messages to multiple trading venues, and HFTcrossmsg is the aggregate number of these messages for the 31 HFT firms. HFTliqsupply is the aggregate number of passive (i.e., liquidity supplying) trades for the HFT firms, and |HFTinventory| is the absolute value of the aggregate inventory position of all HFT firms in Canadian dollars (cumulative from the beginning of the day and assuming that all of them start the day with zero inventory). The next two variables represent the degree of integration of the Canadian market. Specifically, PriceAlign is the percentage of time that the three trading venues with the highest market share post the same bid or ask prices, while SpreadAlign is the percentage of time that the three trading venues have the same bid–ask spread. The next three variables represent the state of liquidity in the market (aggregated across all trading venues): total depth up to 10 cents from the market-wide best bid or offer (MWBBO), depth imbalance at the MWBBO (defined as the absolute value of the difference between the number of shares at the bid and at the ask), and percentage MWBBO spread. The last five variables represent market conditions in the interval: return (computed from the last transaction price in each interval), volatility (computed as the absolute value of return), the average time between trades in the interval, return on an ETF that represents the market, and ETF volatility. All right-hand-side variables are standardized to have zero mean and unit standard deviation to make the coefficients comparable across variables. We report the regression coefficients together with *t*-statistics computed using double-clustered (interval and stock) standard errors.

Variable	PC1		PC2		PC3	
	Coef.	<i>t</i> -stat.	Coef.	<i>t</i> -stat.	Coef.	<i>t</i> -stat.
HFTcrossmsg	0.6717	8.25	0.1548	1.55	0.3330	3.74
HFTliqsupply	-0.0285	-1.07	0.3556	10.61	-0.0973	-4.12
HFTinventory	0.0051	1.68	-0.0047	-1.40	0.0013	0.90
PriceAlign	-0.0050	-0.41	0.0299	2.60	-0.0386	-4.45
SpreadAlign	-0.0261	-5.00	0.0159	2.53	0.0013	0.33
Depth10	-0.0428	-4.84	0.0737	4.76	-0.0278	-2.43
TopDepthImb	0.0057	2.55	-0.0193	-2.74	0.0078	2.10
%Spread	-0.0077	-2.07	0.0030	1.43	0.0068	2.01
Return	0.0015	2.04	0.0007	0.91	0.0007	0.97
Return	0.0262	2.54	-0.0023	-0.39	0.0256	3.45
Time-bet-Trades	-0.0850	-5.54	-0.1051	-6.53	0.0151	1.73
ETF Return	-0.0004	-1.65	0.0002	0.62	0.0004	0.71
ETF Return	0.0006	1.39	0.0000	0.00	0.0032	4.24
Intercept	0.0000	0.00	0.0000	0.00	0.0000	0.00
R <sup>2</sup>	49.61%		27.40%		9.84%	



**Table 4**  
**Publicly Available Measures**

This table presents correlations between variables that we use to represent high-frequency trading (HFT) activity and measures that can be constructed from publicly available order-level data. The data we obtain are unique in terms of our ability to identify HFT firms, and we use these data to create three variables that describe their activity: MSG (all messages HFT firms send to the market), TRD (all their trades), and LMT (submissions and cancellations of non-marketable limit orders). We also use PCA to identify three underlying common strategies and use the principal component scores to represent their activity (PC1, PC2, and PC3). We consider four measures that can be constructed from publicly available order-level data:

- (i) Ord2Vol: Number of Limit Order Submissions/(1+Value of Trades (in CAD\$ 100))
- (ii) Canc2Trd: Number of Limit Order Cancellations/(1+Number of Trades)
- (iii) CancelRatio: Number of Limit Order Cancellations/(1+Number of Limit Order Submissions and Cancellations)
- (iv) RunsInProgress: Time-Weighted Number of Strategic Runs in an Interval

(where strategic runs are linked submissions, cancellations, and executions that are likely to be parts of a dynamic algorithmic strategy and that can be identified with publicly available order-level data). We use our order-level data, without taking advantage of the special user IDs or various other flags, to compute these four publicly available measures for the 52 sample stocks and all 1-second intervals in our sample period. Panel A presents the Pearson and Spearman correlations among the four publicly available measures. Panel B presents the correlations of the publicly available measures with the HFT activity variables we use in the paper (MSG, LMT, and TRD) aggregated across the 31 HFT firms and the component scores representing the three underlying common strategies. In Panel C we present the Spearman correlations of the publicly available measures with the HFT variables for the 1-second, 10-second, and 60-second intervals (I1, I2, and I3, respectively).

Panel A: Correlations among Publicly Available Measures

		Ord2Vol	Canc2Trd	CancelRatio	RunsInProgress
Pearson	Ord2Vol	1	0.943	0.238	0.135
	Canc2Trd		1	0.295	0.156
	CancelRatio			1	0.320
	RunsInProgress				1
Spearman	Ord2Vol	1	0.842	0.623	0.383
	Canc2Trd		1	0.850	0.417
	CancelRatio			1	0.353
	RunsInProgress				1

Panel B: Correlations of Publicly Available Measures with HFT Variables

		Ord2Vol	Canc2Trd	CancelRatio	RunsInProgress
Pearson	MSG	0.606	0.765	0.259	0.178
	LMT	0.611	0.769	0.259	0.178
	TRD	-0.048	0.012	0.127	0.104
	PC1	0.385	0.526	0.237	0.138
	PC2	0.075	0.130	0.148	0.176
	PC3	0.274	0.325	0.069	-0.054
Spearman	MSG	0.741	0.906	0.728	0.424
	LMT	0.743	0.907	0.728	0.426
	TRD	-0.113	0.238	0.247	0.138
	PC1	0.568	0.691	0.535	0.332
	PC2	0.037	0.140	0.124	0.210
	PC3	0.237	0.267	0.237	0.022

Panel C: Spearman Correlations of Publicly Available Measures with HFT Variables by Interval Length

		Ord2Vol	Canc2Trd	CancelRatio	RunsInProgress
I1	MSG	0.741	0.906	0.728	0.424
	LMT	0.743	0.907	0.728	0.426
	TRD	-0.113	0.238	0.247	0.138
	PC1	0.568	0.691	0.535	0.332
	PC2	0.037	0.140	0.124	0.210
	PC3	0.237	0.267	0.237	0.022
I2	MSG	0.115	0.731	0.411	0.568
	LMT	0.120	0.734	0.412	0.569
	TRD	-0.544	0.045	0.122	0.344
	PC1	0.077	0.519	0.254	0.469
	PC2	0.163	0.569	0.330	0.202
	PC3	-0.164	0.111	0.074	0.329
I3	MSG	0.192	0.549	0.293	0.623
	LMT	0.196	0.553	0.295	0.623
	TRD	-0.341	-0.053	-0.060	0.555
	PC1	0.256	0.549	0.368	0.220
	PC2	0.181	0.437	0.207	0.519
	PC3	-0.211	-0.046	-0.095	0.432

**Table 5**  
**Short-Horizon Stock Volatility and HFT Competition**

This table presents the results of regressions of short-horizon stock volatility on time-varying, market-wide measures that capture similarity in high-frequency trading (HFT) strategies across stocks in each of the product categories (or underlying common strategies). For each stock, we estimate the following regression over all 1-second intervals in the sample period:

$$|r_{it}| = a_{0i} + a_{1i}MktSimPC1_t + a_{2i}MktSimPC2_t + a_{3i}MktSimPC3_t + a_{4i}MktPC1_t + a_{5i}MktPC2_t + a_{6i}MktPC3_t + a_{7i}MktRet_t + a_{8i}|MktRet_t| + a_{9i}MktVol_t + error_{it}$$

where  $|r_{it}|$  is the absolute value of the return on stock  $i$  in interval  $t$ , which is our measure of interval return volatility,  $MktRet_t$  is the value-weighted return of all stocks in our sample,  $|MktRet_t|$  is the absolute value of the market return, and  $MktVol_t$  is the aggregate volume in all stocks. We compute the similarity measures that serve as proxies for HFT competition in the three product categories as follows. For each interval  $t$ ,  $MktSimPC1_t$  is the average over the cross-stocks correlations of all pairs of HFT firms from among the eight firms that load significantly on the first principal component, and  $MktSimPC2_t$  ( $MktSimPC3_t$ ) is computed in an analogous fashion for the firms that load significantly on the second (third) principal component. We add the magnitude of HFT activity of the firms that load on the three principal components as control variables:  $MktPC1_t$ ,  $MktPC2_t$ , and  $MktPC3_t$ . We compute the similarity measures for three variables that represent HFT strategies: MSG (all messages HFT firms send to the market), TRD (all their trades), and LMT (submissions and cancellations of non-marketable limit orders). Panel A shows the Pearson correlations between the measures of similarity in HFT strategies and all control variables in the regressions. Panel B presents for each independent variable the average coefficient across the 52 stocks,  $t$ -statistic computed from the cross-sectional distribution of the regression coefficients, the number of negative coefficients, the number of negative coefficients that are significant at the 5% level (from a two-sided test), the number of positive coefficients, the number of significant positive coefficients, and the average  $R^2$ .

Panel A: Correlations of HFT Similarity Measures with all other Independent Variables

	MSG			TRD			LMT		
	<i>MktSimPC1</i>	<i>MktSimPC3</i>	<i>MktSimPC3</i>	<i>MktSimPC1</i>	<i>MktSimPC3</i>	<i>MktSimPC3</i>	<i>MktSimPC1</i>	<i>MktSimPC3</i>	<i>MktSimPC3</i>
<i>MktSimPC1</i>	1	0.282	0.194	1	0.278	0.194	1	0.308	0.137
<i>MktSimPC2</i>	0.282	1	-0.012	0.278	1	-0.012	0.308	1	0.114
<i>MktSimPC3</i>	0.194	-0.012	1	0.194	-0.012	1	0.137	0.114	1
<i>MktPC1</i>	-0.073	0.013	-0.036	-0.076	0.012	-0.035	0.001	-0.031	-0.151
<i>MktPC2</i>	0.018	0.130	-0.064	0.016	0.129	-0.064	-0.031	0.020	-0.170
<i>MktPC3</i>	-0.147	-0.007	-0.021	-0.149	-0.009	-0.021	-0.034	-0.062	-0.131
<i>MktRet</i>	0.001	-0.001	0.000	0.001	-0.001	0.000	-0.005	-0.003	0.013
$ MktRet $	-0.025	0.013	-0.031	-0.027	0.010	-0.031	-0.113	-0.108	-0.215
<i>MktVol</i>	0.013	0.009	-0.004	0.013	0.008	-0.004	-0.007	-0.008	-0.132

Panel B: Regressions of Stock Short-Horizon Volatility on HFT Competition

		Avg. Coef	Avg. t-stat	# Coef < 0	# t-stat < -1.96	# Coef > 0	# t-stat > 1.96	Avg. R <sup>2</sup>
MSG	<i>Intercept</i>	9.39E-06	3.99	8	7	44	44	6.13%
	<i>MktSimPC1</i>	-2.52E-06	-0.81	33	30	19	13	
	<i>MktSimPC2</i>	-8.90E-06	-4.29	41	39	11	10	
	<i>MktSimPC3</i>	-6.84E-07	-0.68	37	29	15	9	
	<i>MktPC1</i>	-1.20E-09	-0.63	34	33	18	14	
	<i>MktPC2</i>	7.96E-09	3.87	10	7	42	39	
	<i>MktPC3</i>	9.25E-09	5.76	8	7	44	42	
	<i>MktRet</i>	-6.66E-03	-1.58	33	8	19	3	
	<i> MktRet </i>	8.56E-01	10.81	0	0	52	52	
	<i>MktVol</i>	2.02E-12	10.90	4	2	48	45	
TRD	<i>Intercept</i>	4.37E-05	7.27	1	0	51	50	12.56%
	<i>MktSimPC1</i>	-1.01E-05	-4.84	42	27	10	3	
	<i>MktSimPC2</i>	-1.91E-05	-9.78	48	43	4	0	
	<i>MktSimPC3</i>	-2.96E-05	-3.80	48	41	4	3	
	<i>MktPC1</i>	-2.35E-08	-0.71	34	15	18	9	
	<i>MktPC2</i>	8.13E-07	0.27	25	11	27	12	
	<i>MktPC3</i>	4.03E-07	4.95	11	2	41	28	
	<i>MktRet</i>	1.51E-03	0.15	23	4	29	4	
	<i> MktRet </i>	9.87E-01	20.81	0	0	52	52	
	<i>MktVol</i>	-2.20E-13	-0.41	31	10	21	11	
LMT	<i>Intercept</i>	9.42E-06	4.00	8	7	44	44	6.12%
	<i>MktSimPC1</i>	-2.54E-06	-0.82	33	31	19	13	
	<i>MktSimPC2</i>	-8.65E-06	-4.24	41	39	11	10	
	<i>MktSimPC3</i>	-7.39E-07	-0.73	37	29	15	9	
	<i>MktPC1</i>	-7.74E-10	-0.42	34	34	18	15	
	<i>MktPC2</i>	6.98E-09	3.52	12	7	40	39	
	<i>MktPC3</i>	8.96E-09	5.76	7	7	45	42	
	<i>MktRet</i>	-6.64E-03	-1.58	33	8	19	3	
	<i> MktRet </i>	8.58E-01	10.88	0	0	52	52	
	<i>MktVol</i>	2.05E-12	10.86	4	2	48	45	

**Table 6****Price Impact Portions of Short-Horizon Volatility and HFT Competition**

This table presents results on the channel through which high-frequency trading (HFT) competition affects volatility by isolating the portions of short-horizon volatility explained by the permanent and temporary price impacts of trades. Let  $p_t$  be the execution price of a trade,  $m_t$  be the prevailing midquote at the time of the trade execution, and  $m_{t+\tau}$  be the midquote  $\tau$  seconds after the trade. The permanent price impact is the difference between the midquote at  $t+\tau$  and the prevailing midquote ( $m_{t+\tau} - m_t$ ), which is attributed to information asymmetry (or adverse selection). The temporary price impact is the difference between the subsequent midquote and the trade price ( $p_t - m_{t+\tau}$ ), and is often thought of as the reversal needed to compensate liquidity providers (also called the realized spread). We compute the permanent or temporary price impacts for each trade using  $\tau=10$  seconds and then average them (using volume weighting) to obtain a measure of these price impacts for an aggregation interval of  $T=60$  seconds. To investigate whether the influence of HFT competition on volatility is driven by the manner in which it affects the permanent price impact of trades, we first regress for each stock the return volatility measure on the volume-weighted permanent price impact for all 60-second intervals in the sample periods:

$$|r_{it}| = b_{0i} + b_{1i}PPI_{it} + error_{it}$$

We then use the fitted values from this regression,  $|\hat{r}|_{PPI_{it}}$ , as the dependent variable in the following regression:

$$|\hat{r}|_{PPI_{it}} = a_{0i} + a_{1i}MktSimPC1_t + a_{2i}MktSimPC2_t + a_{3i}MktSimPC3_t + a_{4i}MktPC1_t + a_{5i}MktPC2_t + a_{6i}MktPC3_t + a_{7i}MktRet_t + a_{8i}|MktRet_t| + a_{9i}MktVol_t + error_{it}$$

where the independent variables are defined as in Table 5. Similarly, to investigate whether the negative impact of HFT competition on volatility is driven by a reduction in the temporary price impact of trades, we regress the volatility measure on the volume-weighted temporary price impact:

$$|r_{it}| = b_{0i} + b_{1i}TPI_{it} + error_{it}$$

and use the fitted values ( $|\hat{r}|_{TPI_{it}}$ ) as the dependent variable. We present the regression analysis for the permanent price impact fitted values (in Panel A) and the temporary price impact fitted values (in Panel B) using three variables that represent HFT strategies: MSG (all messages HFT firms send to the market), TRD (all their trades), and LMT (submissions and cancellations of non-marketable limit orders). The table presents for each independent variable the average coefficient across the 52 stocks,  $t$ -statistics computed from the cross-sectional distribution of the regression coefficients, the number of negative coefficients, the number of negative coefficients that are significant at the 5% level (from a two-sided test), the number of positive coefficients, the number of significant positive coefficients, and the average  $R^2$ .

Panel A: Permanent Price Impact Portion of Volatility and HFT Competition

		Avg. Coef	CS t-stat	# Coef < 0	# t-stat < -1.96	# Coef > 0	# t-stat > 1.96	Avg. R <sup>2</sup>
MSG	<i>Intercept</i>	4.97E-04	22.67	0	0	52	52	4.70%
	<i>MktSimPC1</i>	9.09E-06	0.61	29	12	23	17	
	<i>MktSimPC2</i>	-8.27E-05	-7.51	45	33	7	1	
	<i>MktSimPC3</i>	3.22E-06	0.57	22	9	30	11	
	<i>MktPC1</i>	1.35E-09	7.00	8	3	44	35	
	<i>MktPC2</i>	9.67E-12	0.05	25	9	27	9	
	<i>MktPC3</i>	-1.30E-09	-4.95	39	30	13	5	
	<i>MktRet</i>	2.97E-04	0.23	25	1	27	0	
	<i> MktRet </i>	9.22E-02	16.07	0	0	52	51	
<i>MktVol</i>	9.00E-14	1.14	20	6	32	2		
TRD	<i>Intercept</i>	5.09E-04	21.77	0	0	52	52	4.20%
	<i>MktSimPC1</i>	3.56E-05	2.86	21	3	31	18	
	<i>MktSimPC2</i>	-4.74E-05	-3.89	39	21	13	3	
	<i>MktSimPC3</i>	-2.88E-06	-0.48	22	6	30	8	
	<i>MktPC1</i>	9.01E-08	7.04	4	0	48	29	
	<i>MktPC2</i>	1.88E-08	2.86	11	5	41	20	
	<i>MktPC3</i>	-1.43E-07	-6.66	46	30	6	2	
	<i>MktRet</i>	-1.47E-03	-0.89	25	2	27	0	
	<i> MktRet </i>	9.38E-02	13.13	1	0	51	48	
<i>MktVol</i>	-1.22E-13	-1.15	33	9	19	1		
LMT	<i>Intercept</i>	4.97E-04	22.87	0	0	52	52	4.74%
	<i>MktSimPC1</i>	1.35E-05	0.88	29	11	23	18	
	<i>MktSimPC2</i>	-8.85E-05	-7.71	44	34	8	2	
	<i>MktSimPC3</i>	3.08E-06	0.55	22	9	30	11	
	<i>MktPC1</i>	1.30E-09	6.95	8	3	44	34	
	<i>MktPC2</i>	7.28E-11	0.36	25	7	27	9	
	<i>MktPC3</i>	-1.25E-09	-4.87	38	27	14	5	
	<i>MktRet</i>	1.77E-04	0.14	25	1	27	0	
	<i> MktRet </i>	9.19E-02	16.01	0	0	52	51	
<i>MktVol</i>	6.85E-14	0.89	23	6	29	2		

Panel B: Temporary Price Impact Portion of Volatility and HFT Competition

		Avg. Coef	CS t-stat	# Coef < 0	# t-stat < -1.96	# Coef > 0	# t-stat > 1.96	Avg. R <sup>2</sup>
MSG	<i>Intercept</i>	5.34E-04	22.83	0	0	52	52	1.10%
	<i>MktSimPC1</i>	4.30E-06	0.75	20	6	32	9	
	<i>MktSimPC2</i>	-1.21E-05	-3.14	38	16	14	1	
	<i>MktSimPC3</i>	-1.85E-06	-0.62	29	9	23	4	
	<i>MktPC1</i>	7.61E-11	1.76	16	3	36	17	
	<i>MktPC2</i>	2.08E-10	2.80	23	1	29	9	
	<i>MktPC3</i>	-3.51E-11	-0.66	36	16	16	5	
	<i>MktRet</i>	1.33E-04	0.20	26	11	26	3	
	<i> MktRet </i>	2.32E-02	4.11	15	0	37	21	
<i>MktVol</i>	-2.97E-14	-0.70	37	7	15	2		
TRD	<i>Intercept</i>	5.35E-04	22.76	0	0	52	52	1.01%
	<i>MktSimPC1</i>	1.07E-05	2.40	13	1	39	16	
	<i>MktSimPC2</i>	-2.73E-06	-1.04	29	10	23	1	
	<i>MktSimPC3</i>	-1.41E-06	-0.63	22	4	30	7	
	<i>MktPC1</i>	7.59E-09	1.64	12	2	40	13	
	<i>MktPC2</i>	-7.93E-10	-0.38	26	6	26	3	
	<i>MktPC3</i>	-1.43E-09	-0.22	36	14	16	5	
	<i>MktRet</i>	-1.68E-04	-0.24	27	9	25	4	
	<i> MktRet </i>	2.51E-02	4.19	15	0	37	22	
<i>MktVol</i>	-8.35E-14	-1.50	34	4	18	0		
LMT	<i>Intercept</i>	5.34E-04	22.86	0	0	52	52	1.10%
	<i>MktSimPC1</i>	4.26E-06	0.77	18	6	34	9	
	<i>MktSimPC2</i>	-1.28E-05	-3.21	38	16	14	1	
	<i>MktSimPC3</i>	-1.77E-06	-0.59	27	8	25	3	
	<i>MktPC1</i>	7.00E-11	1.64	17	3	35	16	
	<i>MktPC2</i>	2.19E-10	2.93	23	0	29	9	
	<i>MktPC3</i>	-2.89E-11	-0.55	35	14	17	5	
	<i>MktRet</i>	1.16E-04	0.18	26	11	26	3	
	<i> MktRet </i>	2.31E-02	4.11	15	0	37	22	
<i>MktVol</i>	-3.20E-14	-0.77	37	7	15	2		

**Table 7****Market Concentration and HFT Competition**

This table presents the results of an analysis that relates concentration of trading across the trading venues to competition between high-frequency trading (HFT) firms in three product categories (or underlying common strategies). The five trading venues we investigate are organized as electronic limit-order books and together execute approximately 97.7% of the trading volume in Canada during our sample period. Our sample consists of 52 stocks from the S&P/TSX 60 (S&P). To examine market concentration, we compute the Herfindahl-Hirschman Index ( $HHI_i$ ) of market share of the five trading venues for each stock. The  $HHI$  is computed as the sum of squared market shares of the trading venues, and the lower the  $HHI$  the less concentrated the market. We use data from the Investment Industry Regulatory Organization of Canada (IIROC) to identify 31 high-frequency trading (HFT) firms. To represent HFT competition we construct a measure of similarity in strategies of the HFT firms in each of the product categories. To construct the measure for stock  $i$  for the first principal component ( $StockSimPC1$ ), we compute the correlation coefficient between the activities of pairs of HFT firms in stock  $i$  over all time intervals, and average across all pairs of the eight firms. We construct the measures for the firms that load significantly on the second and third principal components ( $StockSimPC2$  and  $StockSimPC3$ , respectively) in an analogous fashion. We examine three representations of HFT activity: (i) the number of “messages” (MSG) HFT firms send to the market, where messages are defined as submissions and cancellations of nonmarketable limit orders as well as executions of marketable limit orders, (ii) trades (TRD), and (iii) submissions/cancellations of nonmarketable limit orders (LMT). All measures are computed separately for three interval lengths: I1 (1-second intervals), I2 (10-second intervals), and I3 (60-second intervals). The table shows the partial correlations between  $HHI$  and the three similarity measures when we control for market capitalization, price level, volatility, spread, depth, and the magnitude of activity of the same HFT firms that are included in the similarity measure. We report in parentheses  $p$ -values for the (two-sided) hypothesis that the partial correlation is equal to zero.

	MSG			TRD			LMT		
	I1	I2	I3	I1	I2	I3	I1	I2	I3
<i>StockSimPC1</i>	-0.539 (<.001)	-0.587 (<.001)	-0.689 (<.001)	-0.375 (0.010)	-0.424 (0.003)	-0.535 (<.001)	-0.539 (<.001)	-0.588 (<.001)	-0.686 (<.001)
<i>StockSimPC2</i>	-0.564 (<.001)	-0.483 (0.001)	-0.490 (0.001)	-0.502 (<.001)	-0.422 (0.004)	-0.500 (<.001)	-0.554 (<.001)	-0.500 (<.001)	-0.536 (<.001)
<i>StockSimPC3</i>	0.031 (0.840)	-0.086 (0.571)	-0.150 (0.319)	-0.075 (0.622)	-0.111 (0.462)	-0.038 (0.802)	0.022 (0.885)	-0.083 (0.581)	-0.152 (0.313)



## Table 8

### Trading Venue Competitiveness and the Similarity in HFT Strategies

This table presents summary statistics for measures of the competitiveness (or viability) of trading venues and regression results that relate them to competition between high-frequency trading (HFT) firms in three product categories (or underlying common strategies). The five trading venues we investigate are all organized as electronic limit order books and together execute 97.7% of the volume in Canada during our sample period. We denote the five trading venues in the table by the letters A through E. Panel A provides cross-sectional summary statistics for market share in terms of volume, as well as for the two measures of trading venue viability or competitiveness: (i) %TimeBestPrices, defined as the percentage of time that the trading venue posts either the best bid or the best ask in the market (where the market is defined as the aggregation of all five trading venues), and (ii) %TimeSmallSpread, defined as the percentage of time that the bid–ask spread on the trading venue is the narrowest spread in the market. In Panel B we examine whether correlated activity of HFT firms on a particular trading venue is helpful for the competitive position of the trading venue by manifesting in better prices and spreads. To ensure that our results could be interpreted in the same manner for all trading venues, we focus on the HFT firms that engage in substantial activity on all five trading venues, which we define as sending at least 10,000 messages to each of the venues. There are eight HFT firms that satisfy this criterion, and we compute an HFT competition measure,  $StockSim_{iv}$ , that provides information on the similarity in strategies of these HFT firms for a given stock on a particular trading venue. We use MSG (the number of messages HFT firms send to the market) to represent the strategies. For each trading venue  $v$ , we run the following cross-sectional regression:

$$C_{iv} = a_0 + a_1 StockSim_{iv} + a_2 Activity_{iv} + a_3 Spread_{iv} + a_4 BBOdepth_{iv} + a_5 MktCap_i + a_6 Price_i + a_7 Volatility_i + error_{iv}$$

where  $C_{iv}$  is one of the two viability measures, and as control variables we include the magnitude of HFT activity on the specific trading venue ( $Activity_{iv}$ ), spreads and depth (now computed from the best prices on a single trading venue), market capitalization, price level, and the standard deviation of 30-minute returns over the sample period. We report in Panel B only the coefficients on  $StockSim_{iv}$  from each regression (with heteroscedasticity-consistent  $t$ -statistics in parentheses) for the five trading venues and the three time intervals (I1, I2, and I3).

Panel A: Summary Statistics for Market Share and Competitiveness Measures

		A	B	C	D	E
%Market Share of Volume	Mean	14.3%	11.8%	0.4%	1.9%	69.3%
	Std. Dev.	6.7%	3.4%	0.6%	1.3%	8.0%
	Min	4.8%	3.7%	0.0%	0.4%	45.1%
	25 <sup>th</sup> Perc	9.5%	9.8%	0.1%	1.2%	64.3%
	Median	13.5%	12.2%	0.1%	1.6%	70.4%
	75 <sup>th</sup> Perc	16.7%	14.0%	0.5%	2.3%	75.9%
	Max	34.8%	17.6%	2.7%	7.2%	82.8%
	<i>N</i>	52	52	52	52	52
%TimeBestPrices	Mean	65.2%	83.2%	32.7%	66.0%	92.0%
	Std. Dev.	22.6%	13.7%	17.4%	33.0%	6.1%
	Min	15.0%	39.6%	8.1%	0.1%	71.1%
	25 <sup>th</sup> Perc	50.9%	80.6%	20.8%	37.0%	89.8%
	Median	71.9%	87.0%	28.0%	79.6%	92.5%
	75 <sup>th</sup> Perc	84.9%	91.0%	40.7%	96.3%	96.7%
	Max	92.2%	98.0%	82.8%	99.3%	98.9%
	<i>N</i>	52	52	52	52	52
%TimeSmallSpread	Mean	36.1%	54.2%	9.4%	26.9%	76.8%
	Std. Dev.	24.0%	15.8%	14.3%	19.9%	8.0%
	Min	1.8%	13.2%	0.1%	0.0%	50.4%
	25 <sup>th</sup> Perc	13.4%	44.2%	1.9%	6.4%	72.8%
	Median	37.0%	58.1%	3.3%	27.3%	78.5%
	75 <sup>th</sup> Perc	57.1%	65.3%	10.7%	44.9%	82.5%
	Max	83.5%	83.8%	69.3%	67.6%	89.8%
	<i>N</i>	52	52	52	52	52


Panel B: Coefficients on  $StockSim_{iv}$  from Regressions on Competitiveness Measures

	Interval	Venue A	Venue B	Venue C	Venue D	Venue E
%TimeBestPrices	I1	0.658 (2.15)	0.220 (0.92)	0.970 (1.27)	4.281 (3.20)	-0.340 (-3.64)
	I2	0.704 (2.85)	0.325 (1.46)	1.474 (2.17)	2.091 (1.11)	-0.199 (-2.15)
	I3	0.661 (2.70)	0.348 (1.76)	1.427 (3.31)	0.781 (0.53)	-0.096 (-1.26)
%TimeSmallSpread	I1	1.218 (2.64)	0.891 (2.61)	0.735 (1.17)	2.472 (2.77)	-0.875 (-4.63)
	I2	1.188 (2.77)	0.891 (2.73)	0.966 (1.79)	1.375 (1.39)	-0.660 (-3.37)
	I3	1.042 (2.60)	0.794 (2.68)	0.818 (2.48)	0.539 (0.72)	-0.438 (-2.72)

## Figure 1

### HFT Firms: Number of User IDs

In this figure, we plot the number of active user IDs for each of the 31 high-frequency trading (HFT) firms (sorted from the largest to the smallest). Our data come from the Investment Industry Regulatory Organization of Canada (IIROC). The record of each message contains a user ID, and we have a mapping from user IDs to trading firms. We identify 31 HFT firms. Some of these firms have direct market access (DMA) arrangements with dealers while others are dealers engaged in proprietary trading. We choose to work at the HFT-firm level (i.e., to aggregate all user IDs belonging to the same firm) to make our analysis robust to whatever gaming could be occurring in terms of a firm's discretionary assignment of user IDs to its orders. The 31 HFT firms we identify send orders via 572 user IDs during our sample period.




## Figure 2


### Similarity in HFT Strategies: Total versus Directional

In this figure, we compare measures of similarity in the strategies of high-frequency trading (HFT) firms for total versus directional activity. We use data from the Investment Industry Regulatory Organization of Canada (IIROC) to identify 31 HFT firms. These firms are further categorized into three subgroups according to market share—MS1 (market share > 4%), MS2 (market share of between 1% and 4%), and MS3 (the rest). We compare the magnitudes of the market-wide similarity measures in Panel A, and the stock-specific similarity measures in Panel B for three variables that describe total HFT activity (defined as buy plus sell orders) and directional (“Net”) HFT activity (defined as buy minus sell orders): MSG (all messages HFT firms send to the market), TRD (all their trades), and LMT (submissions and cancellations of non-marketable limit orders). The market-wide similarity measures indicate whether the strategies of HFT firms are correlated across stocks in a particular time interval. For each one-second time interval, we compute the correlation coefficient between the activities of two HFT firms across the stocks in the sample, and average the correlations for all pairs of firms. The stock-specific similarity measures indicate whether the strategies of HFT firms are correlated over time for a particular stock. For each stock, we compute the correlation coefficient between HFT activities of any two HFT firms, and average these across all pairs of firms.

Panel A: Market-wide Correlations



Panel B: Stock-specific Correlations




### Figure 3

#### Similarity in HFT Strategies and Market Conditions

This figure compares measures of similarity in the strategies of high-frequency trading (HFT) firms for varying market conditions. We use data from the Investment Industry Regulatory Organization of Canada (IIROC) to identify 31 HFT firms. These firms are further categorized into three subgroups according to market share—MS1 (market share > 4%), MS2 (market share of between 1% and 4%), and MS3 (the rest). We rank the daily returns of the S&P/TSX Composite Index from June 2010 through March 2011, and select the 10 worst days (down days), the 10 best days (up days), and the 10 days closest to and centered on zero return (flat days), for a total of 30 days (sample period). We compare the magnitudes of the market-wide similarity measures in Panel A, and the stock-specific similarity measures in Panel B for down days, flat days, and up days. For each period we examine the similarity in HFT strategies in terms of the number of “messages” (MSG) they send to the market, where messages are defined as submissions and cancellations of non-marketable limit orders as well as execution of marketable limit orders. The market-wide similarity measures indicate whether the strategies of HFT firms are correlated across stocks in a particular time interval. For each one-second time interval, we compute the correlation coefficient between the activities of two HFT firms across the stocks in the sample, and average the correlations for all pairs of firms. The stock-specific similarity measures indicate whether the strategies of HFT firms are correlated over time for a particular stock. For each stock, we compute the correlation coefficient between HFT activities of any two HFT firms, and average these across all pairs of firms.

Panel A: Market-wide Correlations



Panel B: Stock-specific Correlations

