

Insider Trading Under the Microscope *

Andriy Shkilko[†]

December 6, 2019

Abstract

In financial markets, informed agents play a central role in price formation. Theory models offer a variety of predictions on the behavior of such agents; from aggressive and therefore quickly revealing, to stealthy and largely undetectable. I examine these predictions using a comprehensive intraday dataset that contains all orders and trades of a prominent group of privately informed agents – company insiders. When trading on price-relevant information, insiders usually submit large liquidity-demanding orders, and prices adjust quickly. Consistent with theory, insider aggressiveness is attributable to competition, trading urgency, and the value of information. Back-running by other market participants further increases the speed of price adjustment.

Keywords: price discovery, private information, insider trading, informed agent behavior

*I thank Markus Baldauf, Tarun Chordia, David Cicero, Matti Keloharju, Albert Menkveld, Emiliano Pagnotta, Andreas Park, Gideon Saar, Kumar Venkataraman, Ingrid Werner, and seminar/conference participants at Aalto University, Chapman University, Goethe University Frankfurt, University of Melbourne, Central Bank Conference on the Microstructure of Financial Markets, European Finance Association, and Financial Intermediation Research Society.

[†]Wilfrid Laurier University, ashkilko@wlu.ca.

I. Introduction

Trading behavior of privately informed agents and their role in the price discovery process is central to market structure research. Theory models of such trading generally assume that informed agents aim to avoid detection by other market participants before having had a chance to complete their trades. In some models, the informed do so less than seamlessly (e.g., Kyle, 1985; Holden and Subrahmanyam, 1992; Baruch, Panayides, and Venkataraman, 2017). As a result, they leave traces in order flow, market makers infer their presence, and prices adjust. Other models offer a different view arguing that strategic timing may make informed agents difficult to detect (e.g., Foster and Viswanathan, 1996; Back, Cao, and Willard, 2000; Collin-Dufresne and Fos, 2016). Consequently, market reactions to informed activity are weak, price discovery is slow, and traditional metrics of information asymmetry misfire.

In this paper, I examine trading by an important group of privately informed agents – company insiders. To do so, I use an intraday audit trail dataset that identifies every insider order and trade submitted to the Toronto Stock Exchange (TSX) over a 2.25-year period. The intraday dimension of insider trading has never been explored in the literature due to the lack of data, yet it is crucial for testing theory predictions. The above-mentioned theory models posit that insiders attempt to stay under the radar to prevent market makers from learning about their presence. Market maker learning is generally an intraday phenomenon; when trading was largely manual, quotes adjusted to order flow within minutes (Bessembinder, 2003a), while in modern electronic markets such adjustments take mere sub-seconds (Conrad and Wahal, 2019). As such, the intraday dimension is important for understanding interactions between the informed and other market participants as well as the process of price formation.

Overall, the data show that when insiders trade on price-relevant information, detection avoidance is not their primary concern. First, they choose to demand liquidity, using marketable orders rather than limit orders. Second, they trade large quantities relatively quickly, creating sizable order imbalances. Third, the way in which insider parent orders are split suggests a degree of urgency; child orders submitted early in the day are significantly larger than the ones submitted later. Finally, insiders do not appear to time periods of uninformed trading, greater volume, or better liquidity, during which they may be better able to avoid detection. Overall, the trading patterns observed in the data are more consistent with the theory models, in which insiders choose urgency over patience.

The TSX audit trail data identify every insider order and trade, but do not contain information on insider identities. To take a deeper look into theory predictions, I supplement these data with (i) a dataset containing insider identities and job titles (e.g., CEO, CFO, board member) and (ii) a corporate news dataset. Holden and Subrahmanyam (1992) suggest that when multiple insiders trade simultaneously, they will compete to incorporate private information into prices and therefore will choose to execute aggressively. The data confirm this prediction showing that aggressiveness of insider trades, as well as their traceability, increases when multiple insiders trade on the same day.

Kaniel and Liu (2006) and Baruch, Panayides, and Venkataraman (2017) model insiders, who choose aggressiveness levels based on the magnitude of non-execution costs. They posit that if the value of private information is high, the cost of non-execution will also be high, and insiders will build positions more aggressively. The data support this explanation; insider trade aggressiveness is positively related to long-term returns that follow insider trades.¹

¹A somewhat different cost-related argument comes through the models of Carré, Collin-Dufresne,

Chau and Vayanos (2008) and Caldentey and Stacchetti (2010) propose that when there is urgency or uncertainty in timing of information incorporation into prices (incorporation may occur at a random time T), aggressiveness of insider trades will increase. I examine this possibility focusing on periods of price discovery that follow corporate announcements. Huddart, Ke, and Shi (2007) show that insiders often profitably trade soon after such announcements, as it takes time for the market to interpret the more subtle aspects of corporate filings. During such periods, insiders generally do not know how long the market will take to fully process the information and as such may trade with greater urgency. The data confirm this notion; insiders trade more aggressively right after corporate announcements compared to similar non-announcement periods.

Theory models generally assume that insiders have sufficient trading skill and understand how the market functions. In an empirical setting that focuses on illegal insider trading, Kacperczyk and Pagnotta (2019) suggest that this may not always be the case; some insiders may lack investment and trading skill and therefore not engage in strategic timing. To examine this possibility, I sub-divide insiders into (i) CFOs and financial industry employees and (ii) the rest. The first group may be more cognizant of the trading process either due to education or professional exposure and therefore may trade without leaving as much of a trace. The results are consistent with this possibility; trades by the financial industry insiders and CFOs are less traceable than those of their non-finance counterparts.

Most theory models discussed so far envision the market that consists of insiders, market makers, and noise traders. Huddart, Hughes, and Levine (2001) and Yang and Zhu (2019) innovate by adding third-party traders, who respond to price pressures gen-

and Gabriel (2019) and Kacperczyk and Pagnotta (2019), in which the cost of legal penalties associated with illegal insider trading is weighted against the value of inside information. In my sample, this cost may be relatively minor, as the vast majority of insider trades are executed within legal limits.

erated by insiders and mimic their positions in the process called *back-running*. In my setting, back-running is present, and price pressures associated with it represent a sizable share of total pressures, accelerating the adjustment of prices in the direction of insider trades. Importantly, third-party traders appear to learn about insider presence by simply observing order flow rather than being tipped off by insider brokers as suggested by Geczy and Yan (2006) and Li, Mukherjee, and Sen (2018).

Up to this point, the discussion has largely abstracted from the role of brokerages that execute insider trades. Given their fiduciary duties and financial market expertise, the brokerages should guide insiders towards strategic timing and play an important role in order splitting and aggressiveness decisions. In interviews, the brokers however suggest that insiders rarely, if ever, consult with them about execution timing. Rather, insiders tend to place orders the day they wish to execute, thereby constraining the brokers' ability to maneuver. With this in mind, in the remainder of the paper I assume that brokers do not usually affect insiders' choices of trading days and have limited influence on trade timing within the day. Consequently, empirical analyses in the subsequent sections examine two dimensions: daily – the realm of insiders' sole decision making, and intraday – the dimension where brokers have some, if limited, influence.

To further examine the role of brokerages, I focus on their affiliation. Researchers often divide Canadian brokerages into two groups, (i) those affiliated with large banks and (ii) relatively small unaffiliated firms. The former category has greater resources and therefore more sophisticated trading systems such as smart order routers and transaction cost analysis units (McNally, Shkilko, and Smith, 2017). It is therefore conceivable that the brokerages from the former category are able to offer better execution quality to their insider clients. The results are consistent with this possibility; insider trades executed

through the bank-affiliated brokerages are less visible to the rest of the market compared to those executed through the small brokerages.

Institutional and literature background. By the virtue of their employment, insiders regularly have access to material non-public information. To maintain market integrity and fairness, many jurisdictions prohibit trading on such information. Securities regulators, law-enforcement agencies, and industry self-regulatory bodies such as the Investment Industry Regulatory Organization of Canada (IIROC) and the U.S. Financial Industry Regulatory Authority (FINRA) routinely monitor markets for suspicious activity and also rely on investigators and whistle-blowers to detect illegal insider trading. The firms also restrict insider trading by imposing blackout periods prior to material corporate announcements. Insiders are prohibited from trading during such periods (Bettis, Coles, and Lemmon, 2000).

Despite these restrictions, many studies find that insider trades are on average informed (e.g., Lakonishok and Lee, 2001; Cohen, Malloy, and Pomorski, 2012), while Piotroski and Roulstone (2005) and Huddart, Ke, and Shi (2007) suggest that at least some of this informedness comes from insider's superior ability to interpret firm-specific information and detect misvaluation. An important exception to the overall insider trade informedness applies to routine trades conducted for the purposes of diversification, personal liquidity, or voting rights (e.g., Chan and Lakonishok, 1995; Cohen, Malloy, and Pomorski, 2012). In empirical tests, I separate the uninformed insider trades from the informed and focus on the latter.

Until recently, academic research into insider trading has mainly focused on the relation between insider activity and subsequent returns. Put differently, the literature has asked: *When* do insiders usually trade? While the question of *when* has been studied

extensively, less attention has been paid to the question of *how*. How do insiders trade? Do they use aggressive market and marketable orders, or do they prefer cautious limit order executions? Do they split parent orders into child orders trying to avoid detection? Do they time liquidity? Do they attempt to hide among the uninformed? How do other market participants, and specifically market makers, react to insider presence?

Kyle (1985) is first to ask some of these questions in the context of a theory model. He suggests that insiders will break their parent orders into a series of child orders and execute gradually in an effort to prevent prices from reacting too quickly. The child orders will create small but persistent order imbalances, and market makers will gradually learn about insider presence and adjust their quotes (Figure 1a).

[Figure 1]

Holden and Subrahmanyam (1992) expand the Kyle framework to that with multiple insiders that compete to impound information into prices and as such trade impatiently. In their model, information is revealed rapidly, and prices adjust quickly (Figure 1b).² Prices also adjust quickly when there is a level of urgency and when there is uncertainty as to when information will become public (i.e., Chau and Vayanos, 2008; Caldentey and Stacchetti, 2010) and when the non-execution costs are high (i.e., Kaniel and Liu, 2006; Baruch, Panayides, and Venkataraman, 2017).

Based on Kyle's logic, the literature has developed a number of metrics aimed at inferring the presence of informed traders. One such metric (and perhaps the most frequently used one), the *price impact*, examines how liquidity provider quotes react to trades. Trades initiated by buyers that lead to positive changes in quoted prices are

²Models by Foster and Viswanathan (1996) and Back, Cao, and Willard (2000) show that if multiple insiders simultaneously trade on heterogeneous signals, aggressive behavior predicted by Holden and Subrahmanyam (1992) does not materialize.

deemed coming from informed traders. So are trades initiated by sellers that lead to negative price changes. The price impact metric is mainly used by the studies examining intraday data. For studies using daily data, the literature has developed a number of proxies that follow the logic of the price impact metric but are less computationally demanding (e.g., Amihud, 2002; Goyenko, Holden, and Trzcinka, 2009).

The Kyle framework has been the cornerstone of thinking about trading by the informed for a number of years. Given that conventional datasets do not usually flag the trades of the informed, until not long ago empirical studies have been unable to directly test Kyle's predictions. The status quo changed recently, driven by the emergence of new datasets that identify activity of various trader types. For instance, Collin-Dufresne and Fos (2015) examine price changes on days when activist investors build positions and find that conventional information asymmetry metrics do not reflect this activity. Collin-Dufresne and Fos (2016) follow up on this finding with a theory model that extends Kyle (1985) by suggesting that an insider (or any informed trader) could time their executions to periods of high noise trading to avoid detection. In the limit, such timing makes insiders invisible to the rest of the market, and prices only adjust when information is revealed through a public announcement (Figure 1c).

Is the ability to leave no footprints unique to the activist investors, or does it apply to other informed market participants such as company insiders? There are two reasons to believe that the former may be true. First, activist investing is the domain of rather sophisticated institutions (Brav, Jiang, Partnoy, and Thomas, 2008). It is conceivable that these firms' sophistication extends into trading, even if only through outsourcing executions to the most capable brokerages (Goldstein, Irvine, Kandel, and Wiener, 2009). The activists may also have some, if limited, leeway when it comes to trade timing.

In the meantime, not all informed investors may have the luxury of timing their transactions or access to the best brokerages. For example, company insiders often trade when prices do not reflect all relevant information. If such information has a short or uncertain lifespan, is too valuable, or if other agents also have it, liquidity timing may not be the insiders' primary concern. It is also possible that some insiders trade through less sophisticated brokerages. Recognizing the need to better understand how such trading occurs and how other market participants react to it, two recent studies examine a prominent subsample of insider trades – those prosecuted by the U.S. Securities and Exchange Commission (SEC) as illegal.

Kacperczyk and Pagnotta (2018) find that insiders who trade illegally are rather successful in liquidity timing, and consequently conventional metrics of information asymmetry, including the price impact, do not detect their presence. Using the same SEC sample, Ahern (2018) shows that an adjustment for trade urgency improves the metrics performance slightly, yet generally confirms Kacperczyk and Pagnotta's results. Ahern further points out that academic researchers and regulators often use the same metrics to detect informed activity. Illegal insiders, and especially financial industry insiders that represent a notable share of the SEC sample, are likely to be aware of these detection techniques and have strong incentives to adjust their trading accordingly. It is therefore possible that the SEC sample sheds light on a (rather sophisticated) segment, but not the entire universe of insider trades.

Akey, Gregoire, and Martineau (2019) examine the behavior of traders, who illegally obtain soon-to-be-made-public earnings information by hacking newswire services. Their data suggest that the hackers trade on such illegally obtained information aggressively, and prices adjust in the direction of their trades quickly. This result echoes Ahern's

inasmuch as the hackers' trading needs are urgent; they only have several hours before earnings information becomes public.

The findings of Ahern (2018), Kacperczyk and Pagnotta (2018), and Akey, Gregoire, and Martineau (2019) are important in that they empirically outline the spectrum of informed agents' trading strategies. Perhaps fortunately, the episodes of illegal insider trading and hacking are relatively rare, and much of private information flows into prices through insider trading activity that is legal. As such, one of the contributions of this study is to generalize the results from prior research using a sample that is representative of trades by all insiders, and point to the part of the theoretically predicted spectrum of strategies that represents the behavior of an average insider.

The rest of the paper is organized as follows. Section II describes the data, sample, and primary metrics. Section III reports the results of empirical analyses, with daily granularity, while Section IV zooms in to intraday granularity. Section V concludes.

II. Data, sample, and metrics

The main dataset used in this study is from the Toronto Stock Exchange (TSX) and includes 27 months between October 2004 and December 2006. Similar data for the same time period are used by Malinova and Park (2015), Anand and Venkataraman (2016), and McNally, Shkilko, and Smith (2017). The data represent an audit trail of all orders and trades that are processed by the TSX matching engine, time-stamped to a centisecond (one hundredth of a second), for a total of over 4.5 billion messages.

The data are uniquely suited for the purposes of this study, as they identify all insider orders and trades. In Canada, insiders are required by law to disclose their status to the brokerage when they open an account. Also, holders of existing accounts must

notify the brokerage immediately when their insider status changes. Brokerage order processing systems automatically mark each order submitted by an insider with a flag before passing it on to the exchange. The insider flag is invisible to market participants outside the brokerage and is only used for surveillance and compliance purposes by the exchange and the regulators. In addition to identifying insider orders and trades, the data also allow me to (i) see whether insiders provide or demand liquidity and (ii) infer the behavior of retail traders, other uninformed traders, market makers, and insider brokers.

To my knowledge, this is the only dataset available to academics that identifies insider orders and trades on the intraday level. As such, it is uniquely suitable for the task of examining insider trading strategies. This said, I acknowledge that because it covers the period between 2004 and 2006, the TSX dataset is not fully representative of today's markets dominated by ultra-low-latency algorithms. This said, using similar data Anand and Venkataraman (2016) show that liquidity in Canadian equities in 2004-2006 is supplied mainly by algorithmic endogenous liquidity providers (ELPs). While these ELPs may be slower than modern algorithms on a sub-second level, their behavior and order flow recognition abilities are likely comparable, especially given the size of order imbalances generated by insider trades. Put differently, if price pressures from insider trades are identifiable by other market participants during my sample period, they should also be identifiable today.

I supplement the TSX audit trail data with two additional datasets. One contains insider identities and job titles, and the other contains all news announcements associated with the sample firms. The first dataset comes from SEDI (the System for Electronic Disclosure by Insiders), the Canadian equivalent of the EDGAR database administered

by the SEC in the U.S. During my sample period, Canadian insiders must report their trades to SEDI within ten days of execution. Similar to the Form 4 filed by the U.S. insiders, these reports contain insider names, firm affiliation, job titles, and trade information such as volume traded on a particular day and execution price. The match rate between the TSX and SEDI datasets is high, at 98%. The second dataset is collected from the Globe Investor database maintained by The Globe and Mail, Inc., Canada's preeminent financial publisher. It includes announcements related to earnings, product and/or service, corporate governance changes, financial updates (e.g., dividends or share repurchases), among others.

An average Canadian public company is smaller and trades less frequently than its U.S. counterpart. Since I would like to examine the possibility that insiders trade strategically to conceal their activity among that of other market participants, I focus on relatively liquid stocks, in which such concealment is possible. Specifically, I require that a stock trades at least 10 times a day, and its price does not fall below \$5 during the sample period. I also require that a stock survives during the entire sample period and has at least one insider trade. The resulting sample includes 111 firms, and Table I reports the security characteristics for these firms.

[Table I]

The sample is populated by relatively large companies, with the mean market capitalization of CAD 9.4 billion. In the meantime, an average U.S. firm in the CRSP database during the same period is 2.8 times smaller after an exchange rate adjustment. Further comparison shows that the average sample firm is similar in size to an average NYSE-listed firm.³ In addition, the average stock price is CAD 36.07, similar to the average

³Due to relatively strict listing requirements, the NYSE firms have historically been larger than the average U.S. firm.

price of a U.S. stock. Finally, the average sample stock trades 547,746 shares a day, with an average trade size of 616 shares.

To measure trading costs and trade informedness, I use techniques conventional to market microstructure research, including the quoted bid-ask spreads, effective spreads, and price impacts of trades. During the sample period, Canadian equity trading is almost entirely concentrated on the TSX. As such, the TSX best bid and offer (the BBO) is a suitable proxy for the Canadian BBO. With this in mind, I compute the quoted bid-ask spread as the difference at time t for stock i between the best offer and the best bid and scale it by the midquote as follows: $(ask_{it} - bid_{it})/mid_{it}$, where the midquote is computed as $mid_{it} = (ask_{it} + bid_{it})/2$. When computing daily quoted spread averages, I weight the quoted spreads by the amount of time they are outstanding.

The quoted spreads represent prices at which liquidity providers are willing to trade. They do not necessarily reflect the prices at which trades occur. For instance, an average trade may occur during periods of relatively high information asymmetry and wider bid-ask spreads. As such, liquidity demanders may pay more for liquidity than implied by the average quoted spread. To measure the actual cost of liquidity demand, I compute effective spreads as twice the signed difference between the traded price and the corresponding midquote, scaled by the midquote: $2 \times I \times (price_{it} - mid_{it})/mid_{it}$, where the indicator I equals to 1 for buyer-initiated and -1 for seller-initiated trades.⁴

To measure trade informedness, I compute price impacts as twice the signed difference between the midquote 15 minutes after the trade and the midquote at the time of the trade, scaled by the latter: $2 \times I \times (mid_{it+15} - mid_{it})/mid_{it}$. The 15-minute horizon is conventional for price impact estimation during the sample period, but the results are

⁴The data identify the liquidity-demanding side of the trade, and therefore I do not need to use classification algorithms such as that of Lee and Ready (1991).

also robust to estimation horizons that range from 1 to 30 minutes. When computing aggregate daily statistics, I volume-weight effective spreads and price impacts. Finally, I compute intraday volatility as the difference between the high and low prices observed during the day, scaled by the high price and multiplied by 100. Panel B of Table I shows that an average stock has the bid-ask spread of 0.19%, the effective spread of 0.21%, the price impact of 0.13%, and intraday price volatility of 2.27%.⁵

Having described the sample and overall trading activity, I next introduce insider activity statistics in Table II. The average day when insiders are active sees 28.35 insider trades, for a total of 5,921 insider trades per firm during the sample period. Insider trades are almost three times larger than average and represent more than 8% of daily trading volume. Unconditionally, an 8% increase in volume should alert the market to the possibility of informed trading, yet Collin-Dufresne and Fos (2016) suggest that insider trading is likely to be conditional on increases in the overall volume. As such, it is not necessary that the incremental volume arising from insider transactions is sufficient for detection. In fact, Kacperczyk and Pagnotta (2019) show that the insiders, who trade illegally, represent nearly 10% of total daily volume, but strategic timing makes their activity largely undetectable.

[Table II]

Another notable statistic in Table II suggests that 55.47% of insider volume originates from limit orders and as such supplies liquidity. Kaniel and Liu (2006) suggest that insiders will provide liquidity when they trade on information that is long-lived. My ability to separate such trades from those occurring when insiders demand liquidity may shed new light on the aspect of opportunistic timing related to the horizon of information

⁵Chakrabarty, Moulton, and Shkilko (2012) report statistics of similar magnitudes for a 2005 sample of actively traded U.S. securities.

revelation. Further, insider purchases and sales are almost evenly split, with purchases amounting to 47.28% of insider-generated volume. In the subsequent sections, I check if insider purchases and sales affect the results differently by examining them separately.

The SEDI and the Globe Investor databases provide several additional statistics pertinent to the analysis. First, insiders trade in groups of two on average, although almost 70% of insider trading days have only one active insider. Second, in compliance with restrictions imposed on them by the law and the firm-specific blackout periods, insiders avoid trading in the week prior to earnings and other corporate announcements. Instead, a sizable portion of insider activity is concentrated in the three days following such announcements. Given the announcement frequency, unconditionally about 5% of insider trades should occur in the post-announcement periods. In Table II, this figure is considerably greater, consistent with the possibility that insiders often trade soon after corporate announcements attempting to benefit from superior understanding of the information in the announcement (Huddart, Ke, and Shi, 2007).

III. Empirical findings: daily granularity

A. Return timing: When do insiders trade?

I begin with a bird's-eye view of insider trading patterns. The results reported in this section examine prices around insider trades at daily granularity, conforming to the setup used in recent studies of illegal insider trading and thereby providing a base for the more granular intraday analyses that follow.

Figure 2 reports cumulative market-adjusted returns in the [-30; +30]-day window around insider purchases and sales. The results show that an average insider trade is contrarian; insiders buy after prices have been declining and sell after price increases.

The contrarian pattern is more pronounced for insider sales, but is clearly present for both sales and purchases. Further, an average insider trade appears to be informed; prices change in the direction of the trade (increase after insider purchases and decline after sales) in the subsequent weeks. These price changes are similar in magnitude for both purchases and sales. Given the similarities, I combine the results for purchases and sales in the subsequent figures and tables. To do so, I multiply the returns for insider sales by -1 as represented by the black solid line in Figure 2.

[Figure 2]

A number of early theory models assume that liquidity demand dominates the trading strategies of the informed (e.g., Kyle, 1985; Glosten and Milgrom, 1985; Glosten, 1994). More recently, Kaniel and Liu (2006) and Goettler, Parlour, and Rajan (2009) innovate by endowing informed traders with a choice between market and limit orders. They suggest that when information is more price-relevant or the trading need is urgent, the informed are expected to preference market and marketable orders. Figure 3 sheds some essential light on insider order choices showing a strong relation between an insider's decision to demand or supply liquidity and future price movements. More specifically, when insiders choose to demand liquidity, prices subsequently adjust in the direction of insider trades (increase after purchases and decline after sales), consistent with the possibility that these trades are information-driven. In the meantime, when insiders choose to supply liquidity, prices continue to move in the direction opposite to the insider trade, although the slope of the movement becomes less pronounced.

[Figure 3]

B. Price impacts, uninformed flow timing, and trading costs

Given the informed nature of insider trades, and particularly the trades originating from liquidity-demanding orders, in this section I examine market reactions to such trades using the price impact metric. As mentioned in an earlier section, price impacts, or intraday changes in midquotes subsequent to trades, are a conventional way of inferring whether market makers update their beliefs about asset values in response to liquidity-demanding order flow.

Table III reports event study statistics for abnormal price impacts, computed as the difference between the event-window value and the sample mean, scaled by the sample mean. Results for purchases and sales are rather similar, and I combine them. Three results stand out. First, insiders do not time periods when price impacts are low (uninformed trading is likely to be high during such periods) either when they supply or demand liquidity. In fact, price impacts are 5% greater than average in periods preceding liquidity-supplying insider trades. Second, price impacts are 9% greater than average on days when insiders demand liquidity, but not on days when they supply liquidity. This finding is consistent with the notion that the price impact metric is able to capture the presence of the informed.

[Table III]

Collin-Dufresne and Fos (2016) suggest that insiders may time periods of high uninformed volume and periods when uninformed volume is highly volatile to better conceal their own activity. The TSX dataset identifies trader accounts and therefore enables me to shed light on the possibility of such timing. To this end, I follow Malinova and Park (2015) and Korajczyk and Murphy (2018) and compute two proxies: (i) for retail investor activity and (ii) for all non-retail short-term uninformed activity.

To qualify a trading account as specializing in retail orders, I use two criteria. First, since retail investors often use odd lots, I require that more than 1% of the account's transactions are odd lots. I acknowledge that O'Hara, Yao, and Ye (2014) show that odd lots are often used by algorithmic and high-frequency traders. To account for this possibility, the second criterion requires that less than 10% of an account's transactions are short sales. This requirement is based on an established result that retail traders are considerably less likely to open short positions compared to institutional traders (Boehmer, Jones, and Zhang, 2008).

Next, to qualify a trading account as non-retail short-term uninformed, I compute its net position at the end of each trading day.⁶ For long (short) positions, I then compute return as (the negative of) the percentage difference between the closing price five days later and the volume-weighted average price of the position. The five-day window is conventional in the literature and has been used by Chan and Lakonishok (1995), Korajczyk and Murphy (2018), among others. It represents a compromise between a shorter window, which is more likely to be affected by transitory price effects, and a longer window, which is more likely to be affected by noise. I then compute the average return for each account over the sample period and rank the accounts. I consider accounts in the lowest return tercile to have a greater probability of being uninformed.

The results in Table III show that the share of total volume generated by retail and uninformed accounts is somewhat smaller prior to insider trades, but only when insiders provide liquidity (are uninformed). More specifically, retail and short-term uninformed activity is 4% and 3% lower than normal. In the meantime, the volatility of retail and uninformed volume is somewhat higher than normal, respectively, 4% and 2%. Notably, insiders do not appear to time periods of retail and uninformed activity when submitting

⁶To avoid double-counting, I exclude retail accounts identified previously.

liquidity-demanding (informed) orders.

In Table IV, I further examine insider trade determinants in a multivariate setting, paying particular attention to pre-trade return, price impact, uninformed trading and its volatility, as well as non-insider volume and liquidity costs as follows:

$$\begin{aligned}
 Prob(INSIDER = 1)_{it} = & \alpha_0 + \beta_1 RET_{it} + \beta_2 PRIMP_{it} + \beta_3 RETAIL_{it} \\
 & + \beta_4 \sigma RETAIL_{it} + \beta_5 VOL_{it} + \beta_6 SPREAD_{it} + \varepsilon_{it},
 \end{aligned} \tag{1}$$

where all explanatory variables are computed during a 30-day period preceding insider trade day t in stock i , and the variables are defined as follows: RET_{it} is the cumulative market-adjusted return multiplied by -1 in case of insider sales, $PRIMP_{it}$ is the average price impact, $RETAIL_{it}$ is the average share of retail volume in total volume, $\sigma RETAIL_{it}$ is the volatility of retail volume, VOL_{it} is non-insider trading volume, and $SPREAD_{it}$ is the effective spread. Untabulated analyses suggest that many of the above-mentioned control variables are correlated. To avoid concerns with the power of the tests related to multicollinearity, I examine each variable's relation to the incidence of insider trades in a setting where past return, RET , is the only other explanatory variable.

Earlier results suggest that insiders' decisions to demand or provide liquidity may be driven by different information sets. To examine this possibility, I estimate eq. 1 separately for liquidity demanding and supplying insider trades (Panel A) and for insider trades with high and low information content (Panel B). To define high and low information content trades, I split the sample along the median 30-day market-adjusted return that follows an insider trade. To save space, I only report the coefficients of interest from the respective regression models. Note that even though informed insiders tend to demand liquidity on average, not all informed insider trades are liquidity demanding. As

such, one should not expect the results in Panels A and B to match perfectly.

[Table IV]

This said, the results in the two panels are quite similar and show that none of the independent variables explain the timing of liquidity-demanding and informed insider trades. This finding is consistent with the view that when their information is price-relevant, insiders treat liquidity timing as an afterthought. In the meantime, liquidity-supplying trades and trades with low information content are contrarian and are more likely to occur after periods of lower retail activity and greater retail volume volatility.⁷ Specifically, the marginal effect estimate for *RET* suggests that a one standard deviation price decline (increase) in the previous 30 days increases the probability of a liquidity-providing insider purchase (sale) by 2.0%.⁸ Given that the unconditional probability of an insider trade is about 7% (Table II), this effect is economically non-trivial. Finally, it appears that insiders do not base their trading decisions on prior volume and spreads even when their trades are not motivated by information.

Having examined the daily-level determinants of insider activity, I switch focus to price impacts and trading costs on days when insiders trade via liquidity demanding orders. Table V examines these variables in the following regression setting:

$$\begin{aligned}
 DEPVAR_{it} = & \alpha_0 + \beta_1 INSIDER_{it} + \beta_2 PRIMP_{it} + \beta_3 SPREAD_{it} \\
 & + \beta_4 EFF.SPREAD_{it} + \beta_5 ABS.RET_{it} + \beta_6 RETAIL_{it} \\
 & + \beta_7 \sigma RETAIL_{it} + \beta_8 VOL_{it} + \varepsilon_{it},
 \end{aligned} \tag{2}$$

where $DEPVAR_{it}$ is the price impact, the quoted spread, or the effective spread in

⁷In untabulated results, similar relations hold for short-term uninformed volume and its volatility.

⁸Recall that to simplify exposition I multiply returns adjacent to insider sales by -1.

stock i on day t , $INSIDER_{it}$ is the dummy variable that equals to one on days when insiders trade via market or marketable orders, $ABS.RET_{it}$ is the absolute value of the cumulative return over a preceding 30-day period, and all other explanatory variables are as previously defined.

[Table V]

The table confirms the univariate results reported earlier in that the price impacts are greater on days when insiders take liquidity. Further, consistent with the notion that the price impact metric captures the adverse selection cost of market making, quoted and effective spreads are also greater on such days.

IV. Empirical findings: intraday granularity

A. *Intraday timing of insider transactions*

I begin the analysis in this section by asking whether insider trades are affected by returns, information asymmetries, and the state of liquidity on the intraday level. Since one of my goals is to shed light on the intraday process of price adjustment to trading by informed agents, I focus on the liquidity demanding insider trades that are shown in the previous sections to have relatively high information content.

Recall that according to the brokers, the daily-level trading decisions discussed in the previous section are mainly the insiders' prerogative. In the meantime, on the intraday level the brokerages may play a greater role in execution decisions, even though constrained by insiders' immediacy and quantity demands. When viewed at the daily granularity, liquidity-demanding insider trades do not seem to time past returns or liquidity. Given potentially greater brokerage involvement, does this result hold on the

intraday level? To begin answering this question, I estimate the following regression in the spirit of Barclay, Hendershott, and McCormick (2003):

$$\begin{aligned}
 Prob(INSIDER = 1)_{it} = & \alpha_0 + \beta_1 MOM_{it} + \beta_2 PRIMP_{it} + \beta_3 SPREAD_{it} \\
 & + \beta_4 VOL_{it} + \varepsilon_{it},
 \end{aligned}
 \tag{3}$$

where all explanatory variables are computed during a 15-minute period preceding an insider trade in stock i at time t . MOM_{it} is a 15-minute return multiplied by 1 for insider purchases and -1 for sales, $PRIMP_{it}$ is the average price impact, $SPREAD_{it}$ is the average effective spread, and VOL_{it} is the average traded volume. All explanatory variables are standardized at the stock level, and the regressions control for the intraday, day, and stock fixed effects. Controlling for the intraday fixed effects accounts for the possibility that insider trades may cluster during certain intraday periods that are accompanied by wider (early in the day) or narrower (later in the day) spreads. In turn, controlling for day fixed effects adjusts for the overall wider spreads and greater information asymmetry observed on days when insiders trade as discussed in the previous section.

The positive MOM coefficient in specification 1 of Table VI suggests that an average insider trade is momentum-chasing. Insiders purchase after periods of price increases and sell after periods of price declines. This finding may appear unexpected given that the daily results show no indication of insiders' being momentum traders. The following may explain this phenomenon. Since insiders trade in multi-trade sequences (Table II), the market may react to the first several trades in a sequence so quickly that an average insider trade occurs *after* such a reaction. Rapid price adjustments of this kind are consistent with predictions of Holden and Subrahmanyam (1992), Kaniel and Liu

(2006), Chau and Vayanos (2008), and Caldentey and Stacchetti (2010). To examine this possibility, specifications 2 through 5 estimate eq. 3 for the first insider trade of the day, the first five insider trades of the day, trades 5 through 10, and trades beyond the 10th.

[Table VI]

The results in specifications 2 through 5 support the notion that the market is quick to react to the first few insider trades in a sequence. While the first insider trade of the day is clearly contrarian, the contrarian aspect quickly gives way to momentum-chasing as trades beyond the fifth trade occur in the direction of past returns.

The data also shed initial light on intraday liquidity timing and information incorporation into prices. Insider transactions that happen early in the sequence follow periods of relatively low information asymmetry and spreads, consistent with intraday liquidity timing (specification 2). Insofar as the brokers influence insiders' trade patterns on the intraday level, such timing may be of their doing. This said, liquidity timing quickly fades for the subsequent insider trades; they follow periods of relatively high information asymmetry and spreads, likely reflective of the process of information incorporation started by the early trades in the sequence (specifications 3 and 4). The period of information incorporation is however relatively short-lived, and insider transactions at the end of the sequence occur after a decline in asymmetry (specification 5).

So far, the data suggest that prices react to insider trades quickly. To examine this issue in more detail, and particularly through the prism of market maker responses to the arrival of new information, I estimate the following regression:

$$PRIMP_{it} = \alpha_0 + \beta_1 INSIDER_{it} + \beta_2 MOM_{it} + \beta_3 VOLAT_{it} + \beta_4 VOL_{it} + \varepsilon_{it}, \quad (4)$$

where $PRIMP_{it}$ is the price impact of a trade in stock i at time t estimated at the 15-minute horizon, $INSIDER_{it}$ is a dummy variable equal to 1 when a trade is submitted by an insider and zero otherwise, $VOLAT_{it}$ is the difference between the high and low prices in the 15 minutes prior to the trade scaled by the high price, MOM_{it} is the return 15 minutes before the trade signed according to the trade direction (multiplied by 1 if the trade is buyer-initiated and by -1 if it is seller-initiated), and VOL_{it} is volume traded in the 15 minutes prior to the trade. All models control for stock, day, and intraday fixed effects.

Table VII (Panel A, specification 1) suggests that insider price impacts are greater than those of regular trades. Although this result is consistent with the notion of aggressive trading, it may also be due to poor trade timing. For instance, if insider trades tend to execute at times of high volume or volatility, market makers may react to them in a relatively strong manner (Bessembinder, 2003b). Further, given that insider trades are momentum-chasing on average, they may exhaust market maker inventory capacity leading to greater price impacts (Hendershott and Menkveld, 2014). Specification 2 examines these issues and shows that controlling for the above-mentioned variables does not change the fact that price impacts of insider trades are greater than those of an average trade.

[Table VII]

The discussion so far suggests that insider trades that occur early in intraday sequences have greater price impacts compared to those occurring later. To shed more light on this issue, Panel B of Table VII splits insider trades into three groups according to their position in a sequence. Group 1 contains trades 1 through 5, group 2 – trades 6 through 10, and the remaining trades are in group 3. The results corroborate the

implications of the previous table in that only the early trades in an insider sequence – mainly those in group 1 – have above-average price impacts. Trades in group 2 have price impacts that are also above average but much smaller in magnitude, and the price impacts of trades in the last group are about average. These results find further support in Figure 4 that reports individual coefficient estimates for trades 1 through 10 and group estimates for trades 11 through 15, 16 through 20, and beyond the 21st trade.

[Figure 4]

The solid line in Figure 4 captures the cumulative price impact from all trades in the insider sequence. The observed price adjustment pattern closely corresponds to that in Figure 1b, consistent with the notion of aggressive trading by the informed and rapid information incorporation into prices. It should however be noted that this pattern may also be attributed to, at least in part, back-running and information sharing. According to the former channel, order imbalances created by insiders may be observed by third-party traders. The latter channel in turn implies that insider brokers may share information about ongoing insider trades with their preferred clients. In both cases, third-party traders may open positions mimicking insider trades. In what follows, I attempt to differentiate between these channels.

B. Order imbalances and broker activity

The flags that allow me to identify insider orders and trades are only available in historical audit trail data, and no market participant, aside from the insider’s broker, sees them in real time. How does the market detect insider trades? To begin answering this question, I revisit the findings of McNally, Shkilko, and Smith (2017), who show that some insider brokers tend to share information about insider trades with third-

party clients. The clients subsequently open positions that mimic those being built by insiders, thereby temporarily increasing the broker's overall volume share. If such sharing persists in my sample, the mimicking trades may have large price impacts of their own, contributing to the pattern reported in Figure 4.

Relatedly, similar to the argument of Geczy and Yan (2006) and Li, Mukherjee, and Sen (2018), it is possible that information about ongoing insider sales is shared *within* the brokerage, particularly with the brokerage's market making unit. If this channel exists, the market makers tipped off by the insider's broker will rapidly adjust their quotes either before or soon after an insider sequence begins, and this adjustment will be limited to the market makers affiliated with the broker.

The data suggest that even if the information sharing channel exists, its effects are relatively small. First, Table VIII shows that insider trades create sizable order imbalances, which are further exacerbated by the mimicking non-insider trades. Overall order imbalances reach 11.9% in the 15 minutes after the first insider trade of the day, and it is not surprising that prices adjust accordingly. Second, volume shares of insider brokers do not appear to increase, inconsistent with information sharing. More specifically, an average insider broker executes 9.5% of all orders prior to the first insider trade and retains a similar share afterwards. This broker's execution share only slightly increases after the tenth trade.

[Table VIII]

It should be noted that although the volume shares of insider brokers remain stable, their third-party clients do contribute to the above-mentioned order imbalances. They however contribute commensurately with the clients of other brokerages. I acknowledge that using the same data McNally, Shkilko, and Smith (2017) find that third-party

order flow through insider brokerages increases, consistent with the possibility of broker information sharing. Importantly, these authors examine the entire universe of Canadian securities, the vast majority of which are very illiquid. It is for such illiquid securities that their information sharing results arise and as such do not replicate in my sample that focuses on the liquid stocks.

Next, to examine the possibility that insider brokers share information about upcoming or ongoing insider transactions with the affiliated market making units, I examine the quoting behavior of such units. I follow Anand and Venkataraman (2016) and identify user accounts flagged in the TSX database as a *specialist trader*, ST. Such accounts usually perform market making functions on the Toronto Stock Exchange, maintain two-sided quotes, and mostly trade via limit orders. I then define *affiliated* market makers as ST accounts that trade through the same brokerage as the insider.

Figure 5 reports insider price impacts computed from the quotes of affiliated market makers and quotes of all other liquidity providers. If the affiliated market makers receive tips from insider brokers, I expect their quotes to react faster than those of others. The data however do not support this expectation. In fact, there appears to be no discernible difference between the patterns of midquote adjustments by the two groups of liquidity providers.⁹

[Figure 5]

C. Insider trade patterns

The results discussed thus far show that insider executions tend to create sizable order imbalances. In the meantime, Table II suggests that insiders may split their parent

⁹The time and magnitude scales of price impacts in Figures 4 and 5 are different, and as such the two figures are not directly comparable.

orders into child orders, presumably to avoid creating these very imbalances. Table IX attempts to shed light on this matter by sub-dividing insider trading days when only one insider is present according to the number of executed child orders.

The data show that in 11.8% of cases insider parent orders are not split at all. In the remaining cases, parent orders are split, with 13.3% executing via two child orders, 36.9% executing via three to ten child orders, and 38% executing via 11 or more. Although it is surprising to see a large portion of parent orders executing via just one or two transactions, it may be that these orders are small and do not require splitting. Column 2 of Table IX does not support this possibility, as trades that execute in one or two transactions are, respectively, 8.6 or 3.4 times larger than an average trade in the surrounding 30-minute interval. Parent orders split into more than two children are also quite large; child orders in the range of 3 to 10 are more than 2.5 times larger than an average trade.

[Table IX]

Finally, column 3 shows that even when insiders use child orders, they do so in a way that suggests impatience. The column measures the percentage by which the first child order deviates from the size that would have been achieved through an even split of the parent. For instance, a parent order for 1,000 shares that is evenly split into four children should result in four orders of 250 shares. The deviation metric shows that the first child order is almost always larger than an evenly split benchmark. For instance, in the two-child sequences, the first order is 24% greater than the benchmark. This figure declines to 15.2% in the three-child sequences and averages 14.5% for two- to ten-child sequences. As such, impatience appears to be an important factor for the relatively large price impacts of insider trades.

In summary, insiders tend to demand a lot of liquidity quickly, and the market swiftly and efficiently reacts by adjusting prices. Even though insider brokers do not act in a nefarious manner, they seem unable to eliminate the effects of insider impatience. While such impatience may appear sub-optimal, theory literature proposes that it may, in fact, be rational and attributable to opportunistic considerations other than liquidity timing. I examine these considerations in what follows.

D. Price impact determinants

Theory literature suggests that insider aggressiveness may be affected by several opportunistic considerations. First, in Holden and Subrahmanyam (1992) when several insiders trade simultaneously, executions are more aggressive as insiders compete with each other to incorporate information into prices. Second, in Chau and Vayanos (2008) and Caldentey and Stacchetti (2010) when the information revelation horizon is short or uncertain, i.e., information may be incorporated into prices at a random time T , insiders act more aggressively. Third, in Kaniel and Liu (2006) and Baruch, Panayides, and Venkataraman (2017) if private information is more valuable, the cost of non-execution increases, and insiders increase aggressiveness.

To test these theory predictions, I use the SEDI data to split insider trading days into two groups: when only one insider trades and when two or more insiders are present. Further, I use the Globe Investor data to identify periods that follow corporate announcements. Huddart, Ke, and Shi (2007) show that during such periods complete information incorporation into prices has an uncertain horizon. Finally, I use post-trade price changes to identify inside information that is more valuable by splitting the 30-day post-insider returns into those above and below the median.

In addition to the above-mentioned considerations, the data allow me to examine the trading skill hypothesis of Kacperczyk and Pagnotta (2019), whereby some insiders may not have sufficient market knowledge or skill to engage in strategic trade timing. I assume that insiders affiliated with financial firms, as well as the CFOs of non-financial firms, may be more sophisticated in the matters of trading and may therefore appreciate the value of less aggressive executions. Finally, to examine possible effects of brokerage sophistication on insider trading patterns, I split the bank-affiliated and unaffiliated brokerages. Brokerages in the former group may provide higher quality executions due to their reliance on sophisticated order routing and transaction cost analysis techniques (McNally, Shkilko, and Smith, 2017).

I test the above-mentioned splits in the following regression setting:

$$PRIMP_{it} = \alpha_0 + \beta_1\delta_{it} + \beta_2\tilde{\delta}_{it} + \beta_3MOM_{it} + \beta_4VOLAT_{it} + \beta_5VOL_{it} + \varepsilon_{it}, \quad (5)$$

where δ_{it} is a dummy variable equal to 1 if (i) two or more insiders trade on the same day (MULT), (ii) the trading day is within three days of a corporate announcement, when complete information incorporation horizon is uncertain (POST-ANN), (iii) information is valuable (insider trades are followed by the 30-day return that is greater than the median), (VALU), (iv) none of the previous criteria hold (i.e., the trade is by a single insider, occurs outside of the post-announcement window, and information has relatively low value) making insider behavior likely non-opportunistic (NONOPP), (v) insider holds a CFO position, (vi) insider works for a financial firm (FIN), and (vii) insider trades through a sophisticated (bank-affiliated) brokerage (SOPH). Meanwhile, $\tilde{\delta}_{it}$ is a dummy variable equal to 1 for the remaining insider trades. Regressions for the CFO status exclude days when CFOs trade alongside non-CFOs. All the remaining variables

are as previously defined.

The results in Table X are broadly consistent with expectations. Trade aggressiveness increases when more than one insider is present, when the horizon of information incorporation is uncertain, and when information is more valuable. Notably, the remaining insider trades are still observable and are followed by abnormally large price impacts. The only group of insider trades that does not trigger abnormal price impacts is the non-opportunistic group (specification 4) that includes trades by insiders, who trade as singles, outside the post-announcement window, and on information that has relatively low value. Further, CFOs and financial industry insiders exhibit more cautious trading patterns, and sophisticated brokerages alleviate some of the price pressures caused by insider activity.

[Table X]

V. Conclusions

The theory literature makes a variety of predictions as to how economic agents may trade on private information. In some models, such trading is aggressive, and information is incorporated into prices quickly, while in others the informed trade carefully, and prices adjust slowly, if at all. I examine these predictions empirically, using a unique audit trail dataset that identifies all orders and trades submitted by company insiders to the Toronto Stock Exchange over a 2.25-year period.

The results suggest that when insiders trade on price-relevant information, they tend to focus on return timing rather than staying under the radar. As such, they usually demand liquidity and trade impatiently, executing large trades, and creating sizable order imbalances. Consequently, their trades are quickly identified by other market partici-

pants, and prices react accordingly.

The data further suggest that insider behavior is consistent with predictions of theory models that allow for opportunistic trade timing. Specifically, insider aggressiveness increases when they trade in groups of two or more, competing to incorporate information into prices; when information is more valuable; and when the information revelation horizon is short or uncertain. In addition, the data show that a background in finance and assistance from sophisticated brokerages moderately mitigate the above-mentioned aggressive patterns.

REFERENCES

- Ahern, Kenneth R., 2018, Do proxies for informed trading measure informed trading? evidence from illegal insider trades, working paper University of Southern California.
- Akey, Pat, Vincent Gregoire, and Charles Martineau, 2019, Retail insider trading and market price efficiency: Evidence from hacked earnings news, *Working paper*.
- Amihud, Yakov, 2002, Illiquidity and stock returns: cross-section and time-series effects, *Journal of Financial Markets* 5, 31–56.
- Anand, Amber, and Kumar Venkataraman, 2016, Market conditions, fragility, and the economics of market making, *Journal of Financial Economics* 121, 327–349.
- Back, Kerry, C Henry Cao, and Gregory A Willard, 2000, Imperfect competition among informed traders, *The journal of finance* 55, 2117–2155.
- Barclay, Michael J, Terrence Hendershott, and D Timothy McCormick, 2003, Competition among trading venues: Information and trading on electronic communications networks, *The Journal of Finance* 58, 2637–2665.
- Baruch, Shmuel, Marios Panayides, and Kumar Venkataraman, 2017, Informed trading and price discovery before corporate events, *Journal of Financial Economics* 125, 561–588.
- Bessembinder, Hendrik, 2003a, Issues in assessing trade execution costs, *Journal of Financial Markets* 6, 233–257.
- , 2003b, Selection biases and cross-market trading cost comparisons, *Working paper*.
- Bettis, J Carr, Jeffrey L Coles, and Michael L Lemmon, 2000, Corporate policies restricting trading by insiders, *Journal of Financial Economics* 57, 191–220.
- Boehmer, Ekkehart, Charles M Jones, and Xiaoyan Zhang, 2008, Which shorts are informed?, *The Journal of Finance* 63, 491–527.
- Brav, Alon, Wei Jiang, Frank Partnoy, and Randall Thomas, 2008, Hedge fund activism, corporate governance, and firm performance, *The Journal of Finance* 63, 1729–1775.
- Caldentey, René, and Ennio Stacchetti, 2010, Insider trading with a random deadline, *Econometrica* 78, 245–283.

- Carré, Sylvain, Pierre Collin-Dufresne, and Franck Gabriel, 2019, Insider trading with penalties, working paper Swiss Finance Institute and Ecole Polytechnique Fédérale de Lausanne.
- Chakrabarty, Bidisha, Pamela C Moulton, and Andriy Shkilko, 2012, Short sales, long sales, and the lee-ready trade classification algorithm revisited, *Journal of Financial Markets* 15, 467–491.
- Chan, Louis KC, and Josef Lakonishok, 1995, The behavior of stock prices around institutional trades, *Journal of Finance* 50, 1147–1174.
- Chau, Minh, and Dimitri Vayanos, 2008, Strong-form efficiency with monopolistic insiders, *Review of Financial Studies* 21, 2275–2306.
- Cohen, Lauren, Christopher Malloy, and Lukasz Pomorski, 2012, Decoding inside information, *Journal of Finance* 67, 1009–1043.
- Collin-Dufresne, Pierre, and Vyacheslav Fos, 2015, Do prices reveal the presence of informed trading?, *The Journal of Finance* 70, 1555–1582.
- , 2016, Insider trading, stochastic liquidity, and equilibrium prices, *Econometrica* 84, 1441–1475.
- Conrad, Jennifer, and Sunil Wahal, 2019, The term structure of liquidity provision, *Journal of Financial Economics*, forthcoming.
- Foster, F Douglas, and S Viswanathan, 1996, Strategic trading when agents forecast the forecasts of others, *Journal of Finance* 51, 1437–1478.
- Geczy, Christopher, and Jinghua Yan, 2006, Who are the beneficiaries when insiders trade? an examination of piggybacking in the brokerage industry, working paper University of Pennsylvania.
- Glosten, Lawrence R., 1994, Is the electronic open limit order book inevitable?, *Journal of Finance* 49, 1127–1161.
- , and Paul R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogenously informed traders, *Journal of Financial Economics* 14, 71–100.
- Goettler, Ronald, Christine Parlour, and Uday Rajan, 2009, Informed traders and limit order markets, *Journal of Financial Economics* 93, 67–87.

- Goldstein, Michael A, Paul Irvine, Eugene Kandel, and Zvi Wiener, 2009, Brokerage commissions and institutional trading patterns, *The Review of Financial Studies* 22, 5175–5212.
- Goyenko, Ruslan Y, Craig W Holden, and Charles A Trzcinka, 2009, Do liquidity measures measure liquidity?, *Journal of financial Economics* 92, 153–181.
- Hendershott, Terrence, and Albert J Menkveld, 2014, Price pressures, *Journal of Financial Economics* 114, 405–423.
- Holden, Craig, and Avanidhar Subrahmanyam, 1992, Long-lived private information and imperfect competition, *Journal of Finance* 47, 247–270.
- Huddart, Steven, John S Hughes, and Carolyn B Levine, 2001, Public disclosure and dissimulation of insider trades, *Econometrica* 69, 665–681.
- Huddart, Steven, Bin Ke, and Charles Shi, 2007, Jeopardy, non-public information, and insider trading around sec 10-k and 10-q filings, *Journal of Accounting and Economics* 43, 3–36.
- Kacperczyk, Marcin T., and Emiliano Pagnotta, 2018, Chasing private information, working paper Imperial College.
- , 2019, Becker meets kyle: Inside insider trading, working paper Imperial College.
- Kaniel, Ron, and Hong Liu, 2006, So what orders do informed traders use?, *Journal of Business* 79, 1867–1913.
- Korajczyk, Robert, and Dermot Murphy, 2018, High-frequency market making to large institutional trades, *Review of Financial Studies*, forthcoming.
- Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1336.
- Lakonishok, Josef, and Inmoo Lee, 2001, Are insider trades informative?, *Review of Financial Studies* 14, 79–111.
- Lee, Charles MC, and Mark J Ready, 1991, Inferring trade direction from intraday data, *The Journal of Finance* 46, 733–746.
- Li, Weikai, Abhiroop Mukherjee, and Rik Sen, 2018, Inside brokers, working paper Singapore Management University.

- Malinova, Katya, and Andreas Park, 2015, Subsidizing liquidity: The impact of make/take fees on market quality, *Journal of Finance* 70, 509–536.
- McNally, William J., Andriy Shkilko, and Brian F. Smith, 2017, Do brokers of insiders tip other clients?, *Management Science* 63, 317–332.
- O’Hara, Maureen, Chen Yao, and Mao Ye, 2014, What’s not there: Odd lots and market data, *The Journal of Finance* 69, 2199–2236.
- Piotroski, Joseph D, and Darren T Roulstone, 2005, Do insider trades reflect both contrarian beliefs and superior knowledge about future cash flow realizations?, *Journal of Accounting and Economics* 39, 55–81.
- Wang, Jiang, 1994, A model of competitive stock trading volume, *Journal of political Economy* 102, 127–168.
- Yang, Liyan, and Haoxiang Zhu, 2019, Back-running: Seeking and hiding fundamental information in order flows, *Review of Financial Studies*, forthcoming.

Table I
Sample characteristics

The table summarizes the characteristics of the sample stocks, including market capitalization (in Canadian dollars), stock price, daily volume (in # shares), trade size (in # shares), % bid-ask spread computed as the difference between the best ask and bid quotes scaled by the average of these quotes (the midpoint), % effective spread computed as twice the difference between the trade price and the corresponding midpoint scaled by the midpoint, % price impact computed as twice the difference between the midpoint at the time of a trade and the midpoint 15 minutes after the trade scaled by the former midpoint, and intraday volatility computed as the difference between the highest and the lowest prices of the day scaled by the highest price and multiplied by 100.

	mean	st. dev.	25%	median	75%
market capitalization, CAD	9.4B	12.7B	1.2B	2.7B	12.0B
price, CAD	36.07	21.72	17.50	31.65	50.05
daily volume, # shares	547,746	583,958	159,408	342,720	676,613
trade size, # shares	616	317	418	523	709
% bid-ask spread	0.19	0.14	0.08	0.15	0.26
% effective spread	0.21	0.15	0.09	0.18	0.29
% price impact	0.13	0.07	0.07	0.12	0.17
intraday midquote volatility, %	2.27	0.81	1.59	2.23	2.93

Table II
Insider trading statistics

The table contains insider trading statistics such as: (i) the number of insider trades per stock during the sample period, (ii) the number of insider trades on days when insider trading occurs, (iii) insider trade size relative to the average trade size, (iv) the share of volume represented by insider trades on days when such trades occur, (v) the share of insider trading volume executed via liquidity supplying orders (limit orders), (vi) the share of insider purchases in all insider transactions, (vii) the number of insiders trading on the same day, and (viii-xi) the percentage share of insider trades that occur in the three days prior to or in the three days after earnings and other corporate announcements.

	mean	st. dev.	25%	med	75%
# insider trades per stock	5,921	12,327	284	1,029	3,854
# insider trades per stock-day	28.35	36.50	7.14	14.06	27.99
insider trade size to average trade size	2.90	3.32	1.11	1.78	3.69
insider share of total volume, %	8.12	7.17	2.79	6.61	11.08
share of liquidity supply, %	55.47	18.93	45.09	58.61	69.41
share of purchases, %	47.28	29.31	21.03	45.45	73.56
insiders per day, #	2.14	4.44	1.00	1.00	2.00
insider trades prior to EAs, %	0.42	1.81	0.00	0.38	1.24
insider trades prior to non-EAs, %	0.10	0.11	0.00	0.08	0.15
insider trades after EAs, %	12.92	18.97	7.35	12.26	25.65
insider trades after non-EAs, %	9.31	14.78	3.86	8.74	22.21

Table III
Price impacts and uninformed activity

The table reports abnormal price impacts and two proxies for retail and uninformed investor activity and its volatility. The statistics are estimated (i) during the [-30; -1]-day window that precedes insider trades and (ii) on the event day (day 0). Price impacts are computed as twice the signed difference between the midquote 15 minutes after the trade and the midquote at the time of the trade, scaled by the latter. Retail volume originates from accounts generating substantial odd-lot activity and relatively low short-sale activity. Uninformed volume originates from accounts, whose positions are the least correlated with short-term future returns. For both of these metrics, I use the share of respective volume in total volume. Volatility of retail and uninformed activity, σ , is the standard deviation during the pre-event window. Abnormal values are computed as the difference between the event-window values and the sample mean, scaled by the sample mean. Asterisks *** and ** indicate whether the estimates differ from one using 1% and 5% significance thresholds.

	all [1]	liq. demand [2]	liq. supply [3]
price impact [-30; -1]	0.02**	0.00	0.05***
price impact [0]	0.06***	0.09***	0.00
retail volume [-30; -1]	-0.03***	0.01	-0.04***
retail volume [0]	-0.01	-0.01	0.00
uninformed volume [-30; -1]	-0.01**	0.01	-0.03**
uninformed volume [0]	0.00	0.00	0.00
σ retail volume [-30; -1]	0.02***	0.00	0.04***
σ uninformed volume [-30; -1]	0.01**	0.00	0.02**

Table IV
Insider trade determinants

The table reports coefficient estimates from the probit regression of the following form:

$$Prob(INSIDER = 1)_{it} = \alpha_0 + \beta_1 RET_{it} + \beta_2 PRIMP_{it} + \beta_3 RETAIL_{it} + \beta_4 \sigma RETAIL_{it} + \beta_5 VOL_{it} + \beta_6 SPREAD_{it} + \varepsilon_{it},$$

where all explanatory variables are computed during a 30-day period preceding insider trade day t in stock i , and the variables are defined as follows: RET_{it} is the cumulative market-adjusted return multiplied by -1 in case of insider sales, $PRIMP_{it}$ is the average price impact, $RETAIL_{it}$ is the average share of retail volume, $\sigma RETAIL_{it}$ is the volatility of retail volume, VOL_{it} is non-insider trading volume, and $SPREAD_{it}$ is the effective spread. Many of the above-mentioned control variables are correlated. To avoid concerns with low power of the tests related to multicollinearity, I examine each variable's relation to the probability of insider trades in a setting where past return, RET , is the only other explanatory variable. To save on space, the table reports only the coefficients of interest from the respective regression models. Panel A separates insider trading days into those when insiders demand liquidity and those when insiders supply liquidity. Panel B separates insider trading days into those followed by significant price adjustments in the direction of the trade (high informedness) and those followed by insignificant adjustments (low informedness). The separation is performed by splitting post-insider trade 30-day cumulative market-adjusted returns into two groups: above the median and below the median. All non-dummy variables are standardized at the stock level, and regressions control for stock fixed effects. p -Values are in parentheses, and the marginal effects are in square brackets. Asterisks *** and ** indicate whether the coefficient estimates differ from zero using 1% and 5% significance thresholds.

	RET	PRIMP	RETAIL	σ RETAIL	VOL	SPREAD
Panel A: liquidity demand and supply						
demand	0.005 (0.43) [0.000]	-0.006 (0.77) [-0.000]	-0.027 (0.30) [-0.001]	0.035 (0.14) [0.001]	0.008 (0.50) [0.001]	-0.025 (0.08) [-0.003]
supply	-0.063*** (0.00) [-0.002]	0.159*** (0.00) [0.006]	-0.041*** (0.01) [-0.002]	0.019*** (0.01) [0.002]	-0.057 (0.21) [-0.002]	-0.097 (0.09) [-0.004]
Panel B: insider informedness						
high	0.020 (0.10) [0.002]	0.027 (0.14) [0.003]	-0.003 (0.82) [0.000]	0.010 (0.21) [0.000]	0.008 (0.50) [0.001]	-0.025 (0.08) [-0.003]
low	-0.091*** (0.00) [-0.011]	0.043*** (0.00) [0.005]	-0.031*** (0.00) [-0.003]	0.018** (0.02) [0.002]	-0.021 (0.09) [-0.003]	-0.022 (0.11) [-0.003]

Table V
Price impacts, spreads and trading costs on insider trade days

The table reports coefficient estimates from the regression of the following form:

$$DEPVAR_{it} = \alpha_0 + \beta_1 INSIDER_{it} + \beta_2 PRIMP_{it} + \beta_3 SPREAD_{it} + \beta_4 EFF.SPREAD_{it} + \beta_5 ABS.RET_{it} + \beta_6 RETAIL_{it} + \beta_7 \sigma RETAIL_{it} + \beta_8 VOL_{it} + \varepsilon_{it},$$

where $DEPVAR_{it}$ is the price impact, or the inside spread, or the effective spread in stock i on day t , $INSIDER_{it}$ is the dummy variable that equals to one on days when insiders execute trades via market or marketable orders, $ABS.RET_{it}$ is absolute return, and other explanatory variables are computed as previously. All non-dummy explanatory variables are computed during a 30-day period preceding insider trade day and are standardized at the stock level. As such, the regressions control for stock fixed effects, and p -values are in parentheses. Asterisks *** and ** indicate whether the coefficient estimates differ from zero using 1% and 5% significance thresholds.

	PRIMP		SPREAD		EFF. SPREAD	
	[1]	[2]	[3]	[4]	[5]	[6]
INSIDER	0.041*** (0.00)	0.042*** (0.00)	0.019*** (0.00)	0.017*** (0.00)	0.022*** (0.00)	0.023*** (0.00)
PRIMP		0.083*** (0.00)				
SPREAD				0.435*** (0.00)		
EFF. SPREAD						0.427*** (0.00)
ABS. RET		0.024*** (0.00)		0.011*** (0.00)		0.012*** (0.00)
RETAIL		-0.007*** (0.00)		-0.004*** (0.10)		-0.004*** (0.42)
σ RETAIL		-0.001 (0.40)		-0.003 (0.37)		-0.005 (0.39)
VOL		-0.012*** (0.00)		-0.045*** (0.00)		-0.046*** (0.00)
INTERCEPT	-0.005 (0.15)	-0.005 (0.13)	0.001 (0.25)	0.001 (0.25)	0.001 (0.13)	0.001 (0.14)
Adj. R^2	0.00	0.02	0.00	0.22	0.00	0.19

Table VI
Intraday insider trade determinants

The table reports coefficient estimates from the regression of the following form:

$$Prob(INSIDER = 1)_{it} = \alpha_0 + \beta_1 MOM_{it} + \beta_2 PRIMP_{it} + \beta_3 SPREAD_{it} + \beta_4 VOL_{it} + \varepsilon_{it},$$

where all explanatory variables are computed during a 15-minute period preceding an insider trade at time t in stock i . MOM_{it} is a 15-minute return multiplied by the direction of the insider trade, $PRIMP_{it}$ is the average price impact, $SPREAD_{it}$ is the average spread, and VOL_{it} is the average volume. All explanatory variables are standardized at the stock level. The regressions control for stock, day, and intraday fixed effects. Specification [1] reports the results for an average insider trade. Specification [2] reports the results for the first insider trade of the day. Specifications [3], [4] and [5] report the results, respectively, for the first five, the second five and the remaining insider trades of the day. p -Values corresponding to heteroskedasticity-robust standard errors are in parentheses, and asterisks *** and ** indicate whether the coefficient estimates differ from zero using 1% and 5% significance thresholds.

	AVERAGE	FIRST	1-5	6-10	11+
	[1]	[2]	[3]	[4]	[5]
MOM	0.032*** (0.00)	-0.089*** (0.00)	0.002 (0.81)	0.096*** (0.00)	0.019** (0.00)
PRIMP	-0.012 (0.29)	-0.054*** (0.00)	0.028 (0.12)	0.076*** (0.04)	0.006 (0.23)
SPREAD	-0.009** (0.04)	-0.039*** (0.00)	0.002 (0.18)	0.033*** (0.00)	0.012 (0.11)
VOL	-0.083*** (0.00)	-0.172*** (0.00)	0.006 (0.28)	0.052*** (0.00)	0.005 (0.26)
INTERCEPT	-6.571*** (0.00)	-8.821*** (0.00)	-7.633*** (0.00)	-7.766*** (0.00)	-6.107*** (0.00)
pseudo- R^2	0.08	0.07	0.03	0.03	0.15

Table VII
Insider trade price impacts

The table contains coefficient estimates from the following regression model:

$$PRIMP_{it} = \alpha_0 + \beta_1 INSIDER_{it} + \beta_2 MOM_{it} + \beta_3 VOLAT_{it} + \beta_4 VOL_{it} + \varepsilon_{it},$$

where $PRIMP_{it}$ is the price impact in stock i at time t , $INSIDER_{it}$ is a dummy variable equal to 1 when a trade is submitted by an insider and zero otherwise, $VOLAT_{it}$ is the difference between the high and low prices in the 15 minutes prior to the trade scaled by the high price, MOM_{it} is the return 15 minutes before the trade signed according to the trade direction (e.g., a liquidity-demanding insider buy preceded by a positive return is consistent with momentum trading), and VOL_{it} is the volume traded in the 15 minutes prior to the trade. All non-dummy variables are standardized (demeaned and scaled by the standard deviation). Panel A reports coefficients from the base regressions that include all insider trades. Panel B divides insider trades into three groups according to the trade's place in an intraday sequence; respectively, for the first five, the second five and the remaining insider trades of the day. All models control for stock, day, and intraday fixed effects. p -Values corresponding to heteroskedasticity-robust standard errors are in parentheses, and asterisks *** and ** indicate statistical significance at the 1% and 5% levels.

Panel A: base regressions			
	[1]	[2]	[3]
INSIDER	0.034*** (0.00)	0.036*** (0.00)	
MOM		-0.037*** (0.00)	
VOL		0.007*** (0.00)	
VOLAT		0.024*** (0.00)	
INTERCEPT	-0.020*** (0.00)	-0.021*** (0.00)	
Adj. R^2	0.00	0.29	
Panel B: insider sequences			
	1-5	6-10	11+
INSIDER	0.054*** (0.00)	0.009** (0.04)	0.001 (0.29)

Table VIII
Order imbalances and insider broker share around insider trades

The table reports order imbalances and the share of volume executed by insider brokers around insider trades. For display purposes, I invert the numbers for insider sales to make them comparable to insider purchases. Order imbalances are computed as the difference between the buyer-initiated and the seller-initiated volume scaled by total volume. Insider broker share is computed as buyer-initiated (seller-initiated) volume originating from the brokerage that executes an insider purchase (sale) scaled by total volume. Imbalances are computed around the first trade in an insider sequence, whereas insider broker shares are computed around the first, fifth, and tenth trade. In the specification titled FIRST, event windows capture the 15-minute period prior to the first insider trade of the day and a 15-minute period subsequent to this trade. Specifications FIFTH and TENTH capture windows surrounding the fifth and the tenth insider trade. The post-window figures are computed with and without insider volume. Asterisks *** and ** indicate statistical significance at the 1% and 5% levels of the differences between the pre- and post- estimates. In specifications [2] through [4], the pre- estimate is the insider broker share prior to the first trade of the day.

	imbalance [1]	insider broker share		
		FIRST [2]	FIFTH [3]	TENTH [4]
pre-	-0.070	0.095		
post- without insider volume	0.015***	0.092	0.098	0.111**
post- with insider volume	0.119***	0.328***	0.224***	0.271***

Table IX
Child order distributions

The table contains statistics for insider child order sequences on days when only one insider is present, from one to more than eleven child orders. In addition to the share of all insider orders represented by a sequence length, the table reports the trade size ratio – the ratio of the trade size created by the child orders to the average trade size in the surrounding 30-minute interval. It also reports the percentage of the first child order’s deviation from a size that would have been achieved had the parent been evenly split.

sequence length	% share	trade size ratio	first trade deviation
1	0.118	8.568	-
2	0.133	3.418	0.240
3	0.096	2.494	0.152
4	0.069	3.376	0.160
5	0.046	3.268	0.150
6	0.042	2.839	-0.012
7	0.034	2.885	0.193
8	0.032	3.214	0.080
9	0.026	2.501	0.162
10	0.024	2.707	0.180
11 or more	0.380	1.216	-

Table X
Price impact determinants

The table examines the determinants of insider price impacts, among which are: simultaneous trading by several insiders, uncertain information incorporation horizon, trading on valuable information, a CFO status, financial industry affiliation, and trading through a sophisticated brokerage. The regressions for each determinant are estimated as follows:

$$PRIMP_{it} = \alpha_0 + \beta_1 \delta_{it} + \beta_2 \tilde{\delta}_{it} + \beta_3 MOM_{it} + \beta_4 VOLAT_{it} + \beta_5 VOL_{it} + \varepsilon_{it},$$

where δ_{it} is a dummy variable equal to 1 if (i) two or more insiders trade on the same day (MULT), (ii) the trading day is within three days after a corporate announcement, when complete information incorporation horizon is uncertain (POST-ANN), (iii) information is valuable (insider trades are followed by the 30-day return that is greater than the median), (VALU), (iv) none of the previous criteria hold (i.e., the trade is by a single insider, occurs outside of the post-announcement window, and information has relatively low value) making insider behavior likely non-opportunistic (NONOPP), (v) insider holds a CFO position, (vi) insider works for a financial firm (FIN), and (vii) insider trades through a sophisticated (bank-affiliated) brokerage (SOPH). Meanwhile, $\tilde{\delta}_{it}$ is a dummy variable equal to 1 for the remaining insider trades. Regressions for the CFO status exclude days when CFOs trade alongside non-CFOs. All the remaining variables are as previously defined. All non-dummy variables are standardized (demeaned and scaled by the standard deviation). All models control for stock, day, and intraday fixed effects with one exception – regressions for the financial industry affiliation do not include stock fixed effects. Heteroskedasticity-robust standard errors are in parentheses, and asterisks *** and ** indicate statistical significance at the 1% and 5% levels.

	MULT	POST-ANN	VALU	NONOPP	CFO	FIN	SOPH
	[1]	[2]	[3]	[4]	[5]	[6]	[7]
YES	0.071*** (0.00)	0.068*** (0.00)	0.085*** (0.00)	-0.002 (0.24)	0.026** (0.03)	0.033** (0.02)	0.029*** (0.00)
NO	0.022** (0.03)	0.029*** (0.00)	0.014** (0.04)	0.089*** (0.00)	0.052*** (0.00)	0.081*** (0.00)	0.058*** (0.00)




Figure 1

Price reactions to insider trades: theory

The figure outlines predictions of three groups of theory models that examine market reactions to insider trades. Figure (a) describes Kyle (1985), who suggests that market makers will learn about the presence of insider traders at a constant rate over time. Figure (b) depicts predictions of models that allow for an element of aggressiveness in insider trading, whether due to competition among several insiders as in Holden and Subrahmanyam (1992), or to information magnitude as in Kaniel and Liu (2006), or to uncertainty in timing as in Chau and Vayanos (2008) and Caldentey and Stacchetti (2010). Figure (c) describes a limit case of models that assume strategic timing by insiders, e.g., Foster and Viswanathan (1996), Wang (1994), Back, Cao, and Willard (2000), Collin-Dufresne and Fos (2016). In these models, insiders are successful concealing their trading activity, and private information may only be revealed at the time of the public announcement.

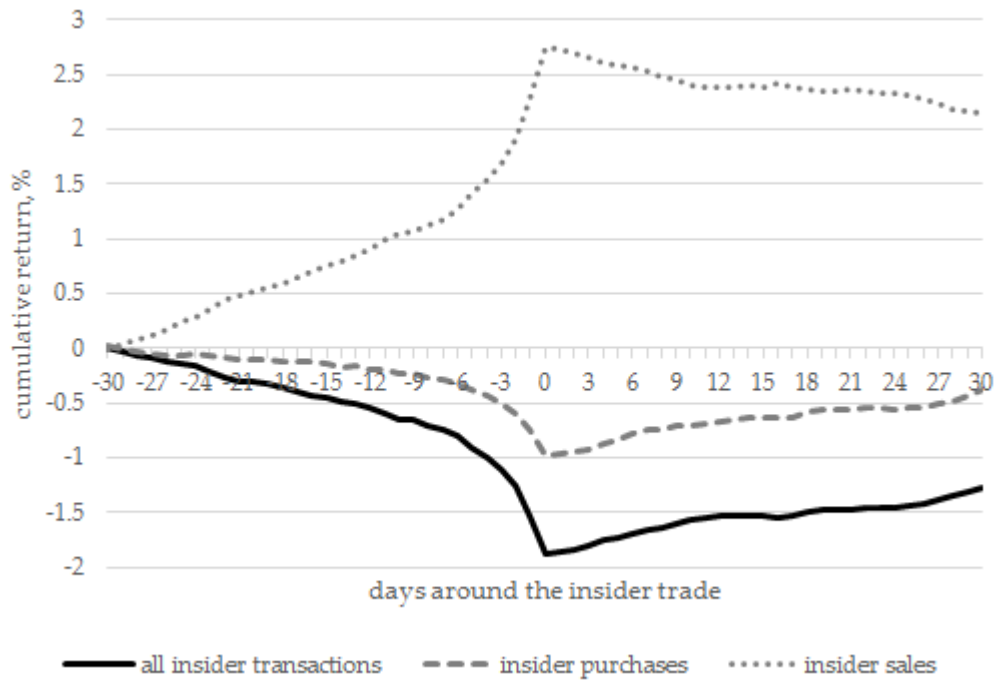


Figure 2
Cumulative returns around insider trades

The figure reports cumulative market-adjusted returns computed during the [-30; +30]-day event window surrounding insider trades. The dashed gray line represents returns around insider purchases. The dotted gray line represents returns around insider sales. The solid black line combines returns around purchases and sales by multiplying the returns around sales by -1.

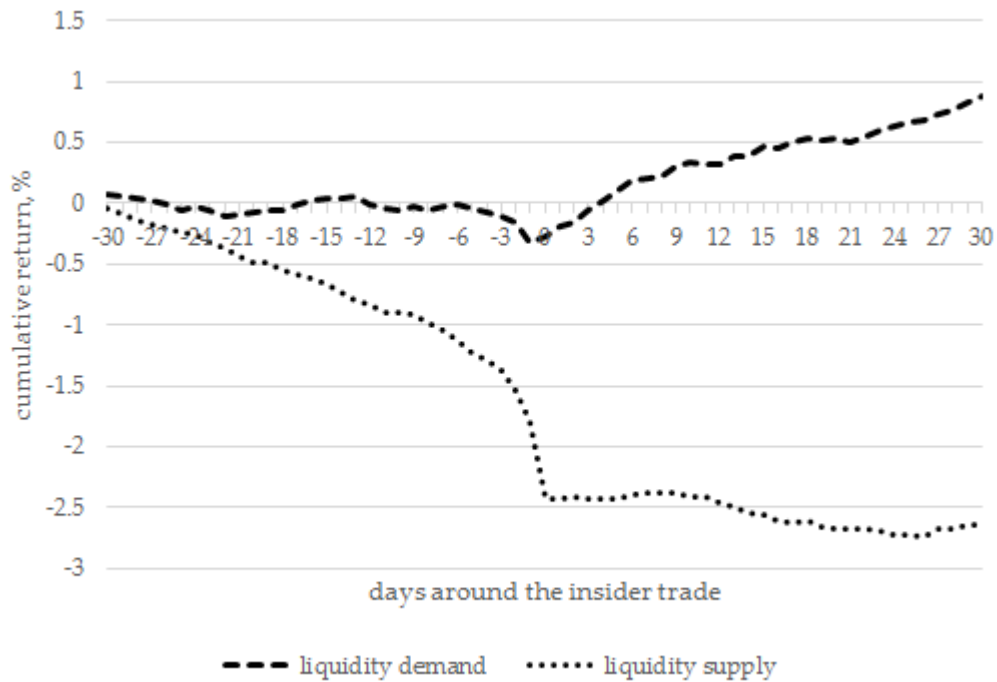


Figure 3

Cumulative returns when insiders provide/demand liquidity

The figure reports cumulative market-adjusted returns computed during the [-30; +30]-day event window surrounding insider trades. The dashed line captures returns around an average liquidity-demanding insider trade. The dotted line captures returns around an average liquidity-supplying insider trade.

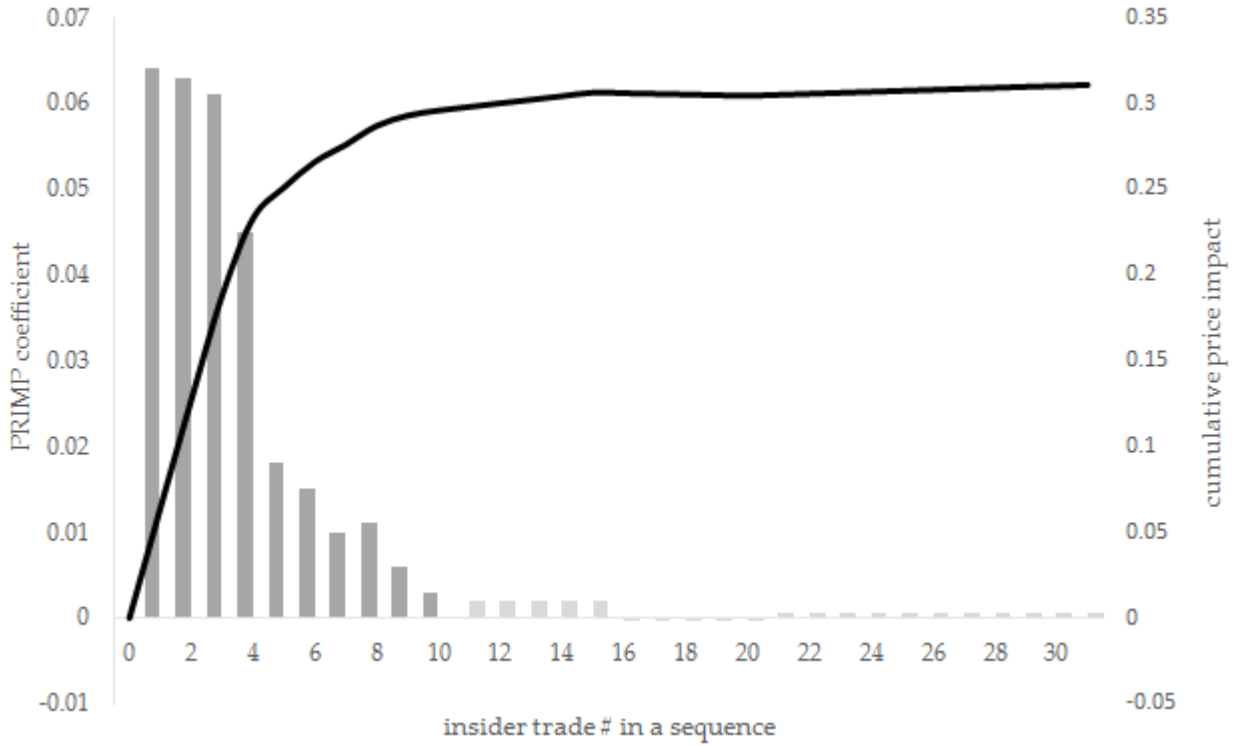


Figure 4
Price impacts in insider trade sequences

The figure reports coefficient estimates for the $INSIDER_{it}$ variable in eq. (4) estimated for each trade in multi-trade insider sequences. I estimate coefficients for the first ten trades individually (dark bars) and then proceed by estimating coefficients for trades 11-15, 16-20, and 21+ (lighter bars). The solid black line represents cumulative price adjustment from trade price impacts.




Figure 5

Price impacts for affiliated and unaffiliated market makers

The figure reports price impacts for the market making accounts affiliated with the brokerage that executes insider trades (solid line) and for market making accounts unaffiliated with the brokerage (dashed line).