

Chen, Marie; Garriott, Corey

Working Paper

High-frequency trading and institutional trading costs

Bank of Canada Staff Working Paper, No. 2018-8

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Chen, Marie; Garriott, Corey (2018) : High-frequency trading and institutional trading costs, Bank of Canada Staff Working Paper, No. 2018-8, Bank of Canada, Ottawa,
<https://doi.org/10.34989/swp-2018-8>

This Version is available at:

<https://hdl.handle.net/10419/197861>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Staff Working Paper/Document de travail du personnel 2018-8

High-Frequency Trading and Institutional Trading Costs



by Marie Chen and Corey Garriott

Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Bank of Canada Staff Working Paper 2018-8

February 2018

High-Frequency Trading and Institutional Trading Costs

by

Marie Chen¹ and Corey Garriott²

¹ London School of Economics
y.chen145@lse.ac.uk

² Financial Markets Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
cgarriott@bankofcanada.ca

Acknowledgements

We are grateful to Jason Allen, David Cimon, Jean-Sébastien Fontaine, Darcey McVanel, Ryan Riordan, Harri Vikstedt and Jonathan Witmer for their advice and suggestions. We are also grateful to the Montréal Exchange for providing data and interviews.

Abstract

Using bond futures data, we test whether high-frequency trading (HFT) is engaging in back running, a trading strategy that can create costs for financial institutions. We reject the hypothesis of back running and find instead that HFT mildly improves trading costs for institutions. After a rapid increase in the number of HFTs, trading costs as measured by implementation shortfall decrease by 27 basis points for smaller-sized positions (\$2–\$10 million notional). For larger-sized positions there is no significant effect. We explain the improvement as being the consequence of HFT reducing effective spreads and per-trade price impacts.

Bank topics: Financial markets; Market structure and pricing; Financial system regulation and policies

JEL codes: G20, G14, L10

Résumé

Au moyen de données relatives aux contrats à terme sur obligations, nous vérifions si les opérateurs qui pratiquent la négociation à haute fréquence ont recours au parasitisme (*back running*), une stratégie de négociation qui consiste à répliquer les transactions d'opérateurs institutionnels et qui est susceptible d'imposer des coûts aux institutions financières. Nous rejetons l'hypothèse du parasitisme et concluons plutôt que la négociation à haute fréquence réduit légèrement les coûts de transaction pour les institutions. À la suite d'une hausse rapide du nombre d'opérations de négociation à haute fréquence, les coûts de transaction (mesurés selon le décalage lié à l'exécution) ont diminué de 27 points de base pour les transactions de petite taille (notionnel de 2 à 10 millions de dollars) et n'ont pratiquement pas changé pour les transactions de grande taille. Selon nous, cette baisse s'explique par le fait que la négociation à haute fréquence réduit les écarts effectifs et l'incidence sur les prix par transaction.

Sujets : Marchés financiers, Structure de marché et fixation des prix, Réglementation et politiques relatives au système financier

Codes JEL : G20, G14, L10

Non-technical summary

High-frequency trading (HFT) is a computerized trading strategy that derives revenue from trading a security for a short period of time, often holding a long or short position in a security for mere moments before exiting the position. This new style of trading has been a source of controversy because it is unclear how its short-termism can be of any value to society. A large literature has grown to address the question: see “The Economics of High-Frequency Trading: Taking Stock” (Menkveld 2016, *Annual Review of Financial Economics*). The general result is that HFT is automating (for good or for ill) a set of trading activities that have historically been performed by human beings. This research paper supports the literature using the case of Canadian bond futures.

One cause for concern is that HFT might use its technology to detect when large financial institutions are trading and then trade to move prices against them. This strategy would be profitable because it would exploit a need at a financial institution to trade, either because the institution must fill an order from a client or because it must act on new information. HFT could exploit this need by “back running”: buying when the institution buys, causing the price to rise more than it would have, and then selling at the aggravated high price. This would create costs for the institution as it pays a higher price. The cost matters to society, as institutions manage most of our savings. While the strategy is nothing new, HFT is arguably going to be better at it than humans were.

In this paper, we investigate whether HFT is indeed loading costs on institutions in this manner. We find no evidence that it is trading in the same direction as institutions when they are building large positions in Canadian bond futures. In fact, HFT trades in the other direction, absorbing the institutional trading. This is consistent with HFT’s usual role in automating the services provided by a human financial intermediary. To drive home the point, we show that trading costs for institutions improve after more HFTs begin trading the bond futures. The reason for this is that HFT, in competing to trade with institutions, diminishes effective spreads and price impacts—a result that has been found in a variety of studies.

High-frequency trading (HFT) is criticized by institutional traders, who argue it increases their execution costs.¹ They argue that HFT uses predatory strategies such as back running, which is trading in the same direction as an institution until the price moves and then divesting the position (Yang and Zhu 2016). Predation can load costs on institutions such as mutual funds and insurance companies, which manage most of the savings in the economy on behalf of individual and corporate savers. There is mounting evidence that HFT can trade in the same direction as institutions (Hirschey 2017), and the question is whether HFT is raising institutional costs.

In this paper, we test for HFT back running and ask whether increased HFT presence affects trading costs for institutions. We use data on the future for the 10-year Canadian bond, the basic hedging tool in Canadian fixed income and the single most-traded exchange contract in Canada. To test for back running, we track HFT order flow as institutions build large positions using a series of trades. To assess the cost of building a position, we compute the implementation shortfall, the cumulative price impact of the series of trades. For statistical control, we compare the outcomes in bond futures with those in Canadian interest-rate futures, which trade on the same exchange and share interest-rate risks, but which lack HFT presence because they have poor features for HFT, namely a crowded order book and low volatility.

Using the data, we reject the hypothesis that HFT engages in back running. HFT trades against the direction of institutional positions and not with them, which is the opposite of how a back runner should behave (Yang and Zhu 2016; Brunnermeier and

¹ See, for example, section IV in BIS Markets Committee (2016); or the 2015 Liquidnet survey of institutional clients, which found 76% of institutional traders are concerned about HFT.

Pedersen 2005). HFT does not back run regardless of whether institutions are working large or small positions, using long or short amounts of time, or creating permanent or transitory price impacts.

Instead, we find that greater HFT presence decreases trading costs for institutions, though only for the smallest positions that institutions construct. This reduction in trading costs can be seen in Figure 1, which graphs implementation shortfall and the number of HFTs present on bond futures. After a permanent spike in the number of HFTs, the shortfall for small positions (less than \$10 million) decreases by around 30 basis points.

FIGURE 1 ABOUT HERE

To control for endogeneity, we estimate the treatment effect using a difference-in-differences (DiD) event study in which interest-rate futures are the control group. The estimated treatment effect is 26.9 basis points and is robust to controls for volume, volatility and lags, as well as to variation in the study-window size. However, HFT has no statistical effect on trading costs when institutions construct positions larger than \$10 million.

It could be puzzling that HFT improves trading costs for institutions, because HFT is thought to lack the capital commitment that institutions require from a counterparty (BIS Markets Committee 2016). We therefore seek to explain how HFT in the sample could be improving implementation shortfall. As a preliminary, we first verify that HFTs do engage with institutional volumes. Indeed, the HFT group is counterparty to 30–40% of this segment of volume. Then we ask if HFT could be improving outcomes by committing

capital. The answer is no. The HFT group bears a small proportion of the positions constructed, on the order of \$100,000 to \$500,000 compared with the average position of \$12 million. The puzzle is that HFT is interacting with institutional volumes and improving outcomes, yet it is not dedicating capital.

Our explanation for how HFT decreases costs for institutions is that it creates a limited improvement in price stability. We run the DiD event study on three measures of the per-trade price impact experienced by institutions as they are building positions. HFT improves per-trade price impacts measured at the 5-second and 30-second horizons but has no statistical effect when per-trade price impacts are measured at the 5-minute horizon. Such an ephemeral contribution helps explain why HFT does not improve shortfall for institutions building larger positions, as these positions can be worked for an hour or more. The finding is consistent with Brogaard, Hendershott and Riordan (2014), in which HFT limits short-term price movements by trading against price changes.

Some of our findings are confirmed by other work on HFT and institutions. We find HFT's volume share as a counterparty to institutions declines as institutions build positions, as in Korajczyk and Murphy (2016) and Tong (2015). However, in our case the decline is not dramatic, and we do not find larger declines when institutions build larger positions. As institutions build positions, HFT decreases its trading with them by at most seven percentage points (out of a 30–40% share).

Some of our findings contrast with other work. In contrast to Van Kervel and Menkveld (2016), we reject the hypothesis of back running. On average, the HFT group in our sample never holds an aggregate position in the same direction as institutions: neither when institutions appear to be trading for liquidity (as modelled by Brunnermeier and

Pedersen 2005), nor when they appear to be trading on information (as modelled by Yang and Zhu 2016). There are two likely explanations for this difference in results. First, it is costlier for a back runner to manipulate a futures price because futures markets have high trading volumes and a professional clientele, decreasing the price impact achievable by a back runner. Second, the large HFT positions that Van Kervel and Menkveld (2016) observe on the Swedish NASDAQ are probably offset in other markets (alternative, foreign or derivative) because HFTs are reluctant to bear large positions with material risk (Baron et al. 2016; Menkveld 2013). The HFT positions we observe in Canadian bond futures are likely representative because there is only one market for the futures, and there is no electronic access to the natural hedge in the underlying bond market.

A distinguishing feature of our study is that we use a control group and perform DiD. Korajczyk and Murphy (2016) also study a market-structure event using a single-difference event study. In contrast to our work, they find HFT supplies less liquidity after their event. It is natural that our results would differ because the events in our papers differ; ours is an increase in HFT presence and theirs is an intensification of latency asymmetries. Latency asymmetries can have poor consequences for market quality (Shkilko and Sokolov 2016; Chakrabarty et al. 2014), whereas greater HFT presence can improve it (Brogaard and Garriott 2017).

1. Data

The data were granted by the Montréal Exchange (MX), founded in 1874 as the Montréal Stock Exchange and today operating as Canada's derivatives exchange. The data contain one report on every quotation, update, fill and cancel for the CGB contract, which is the

future on the Government of Canada 10-year bond, as well as for the BAX contract, which is the future on the Canadian Banker's Acceptance rate (the rate used for the CDOR index), during the years 2009–2014. We refer to the CGB contract as the bond future and the BAX contract as the rates future. The CGB was the only liquid bond future in Canada during the sample period. Each report contains fields such as price, quantity, millisecond timestamp, side (buy or sell) and initiation (buyer or seller). The CGB contract size is for par value \$100,000; the BAX is for par value of \$1 million.

Using the data, we compute a time series of several statistics: the average *price*, the contract's midquote price in dollars; the total *number of trades* (using single counting); the total *trading volume*, the total electronic number of contracts traded during normal trading hours (using single counting); the average *bid-ask spread*, the difference between ask and bid prices in cents;² the average *inside depth*, the sum of the quantities of all outstanding visible limit orders at the bid and ask prices; and the *volatility*, the two-week rolling standard deviation of the daily close price log return, annualized.

Figure 2 plots the trading volume of bond and interest-rate futures over the sample period. Table 1 gives additional summary statistics.

FIGURE 2 AND TABLE 1 ABOUT HERE

² In fixed-income futures it is not convention to compute *relative* or percentage bid-ask spreads, because the denominator of the ratio would par value, which is always the same for a future, e.g., \$100,000 for the bond future. Dividing by a constant would not change the relative levels of the metric.

In the figure, bond-future volumes grow through the sample period, whereas volumes for rates futures are level after 2011. The two series begin to diverge after 2012, the year after the number of HFTs on the bond future spikes (see Figure 1).

Table 1 shows that the bond and rates futures have different liquidity. Although the two futures have comparable bid-ask spreads, particularly later in the sample, rates futures have far deeper queues at the best prices (by two orders of magnitude) and are much less volatile (by one order of magnitude). It makes sense that HFTs would prefer the bond future, as HFTs prefer more volatile markets with shorter queues (Baron et al. 2016). Accordingly, we do find HFT is far more active in the bond future. There is a level difference in certain liquidity metrics for bond and rates futures, making the comparison less satisfying in a DiD setup. However, though the level of liquidity is sometimes different, the trends are correlated, enabling DiD. As a preview, the correlation of the implementation shortfall metrics will be 0.659, and the correlation in per-trade price impacts will be 0.645.

From the order-book data, we drop early-hour and after-hour activity, defined as activity prior to the 8:20 a.m. market open and after the 4:00 p.m. market close. For bond futures, only one contract has significant volume at a time, so we keep only data on the front contract. For rates futures, there is ample volume through the curve. To choose a good comparison contract, we keep whichever of the first three rates futures (ordered by expiry) has the highest daily volume.

a. Identification of HFT on MX

To identify market participants as HFTs, we draw from the literature three standard criteria that the participants' behaviour must satisfy (Kirilenko et al. 2017): a minimum

trading threshold, a maximum overnight inventory threshold, and a buy-sell switching rate threshold.

The *minimum trading threshold* requires the participant to generate at least an average of 100 contracts (\$10 million) of daily volume on the product. This criterion is meant to ensure the participant is frequently trading. The *overnight inventory threshold* requires the participant to carry an overnight inventory of zero during four days out of five through the sample. This criterion is meant to ensure the participant is reluctant to take risk overnight, a defining feature of HFT (Baron et al. 2016; Menkveld 2013). The *buy-sell switching rate threshold* requires the participant to follow buy trades with sells or sell trades with buys at least a fourth of the time. This criterion is meant to ensure the participant does not trade directionally intraday, another defining feature of HFT (versus, e.g., algorithmic order execution). Our set of identified HFTs is robust to the relaxation of any one of the thresholds. If we were to remove either the trading, inventory, or switching criteria, we would identify the same set of participants.

Figure 3 plots the volume share of HFT in the bond and rates futures.

FIGURE 3 ABOUT HERE

HFT has a significant presence in the bond future but is hardly present in rates futures. For the bond future, HFT's volume share increases from around 5% in 2009 to around 30% in 2014. However, for rates futures, HFT's volume share is zero in the beginning and peaks at around 3% in 2012. HFT presence in the bond future probably explains the continued growth in its volume relative to rates futures. HFT's volume share in

the bond future is correlated with the increases in bond-future volumes in Figure 2. Rates futures, which have no HFT, do not see the same volume growth, suggesting that HFT is important to explain the difference.

Table 2 gives summary statistics on the behaviour of participants split in two groups: HFT and nonHFT. The statistics are averaged over members of the two groups. The statistics are the previously defined *number of trades*, *switch rate*, and *overnight inventory*.

TABLE 2 ABOUT HERE

The average HFT is becoming more important to the bond future. While all MX participants are trading more on average, HFTs grew their trading more than nonHFTs did. The average member of the HFT group quadruples its trades from 2009 to 2014, whereas the average member of the nonHFT group only doubles its trades. Compared with nonHFT, the HFT group has double the switch rate and carries two orders of magnitude less inventory overnight.

b. Identification of “trade strings”

To study how HFT behaves when institutions work large orders, we identify events in which institutions are building large positions using the proxy of “trade strings.” The idea is to locate series of trades nearby in time and in which one of the sides is building a large position. Quantitatively, a trade string is defined as at least four consecutive unidirectional (all buy or all sell) trades in which one side builds a net position at least \$2 million in size

and with no more than a 20-minute gap between the side's trades.³ We have tried the analysis for gap sizes of 10 minutes and 30 minutes, and it does not change the results. Furthermore, we allow for attempts at misdirection by institutions, often called "tape-colouring," by allowing trade strings that have 5% of their volume in violation of unidirectionality. We have tried the analysis for a 0% violation of unidirectionality, and it does not change the results.

c. Computation of implementation shortfall

Using the trade strings, we compute the *implementation shortfall*. Implementation shortfall is a common performance benchmark for institutions (Harris 1998). The idea is to benchmark the all-in price of buying or selling a security versus some notion of the fair price of the security, which is usually proxied by the price when the trading began. We compute the implementation shortfall of a trade string by subtracting the midquote contemporaneous to the first trade from the signed, volume-weighted average price of the trades in the string. We compute the shortfalls for both the bond future and the interest-rate futures and average them both weekly and monthly.

Table 3 gives average characteristics of the trade strings by year and by four volume bins: \$2-\$10 million, \$10-\$25 million, \$25-\$100 million, and \$100+ million. The statistics are: *implementation shortfall*, as defined above; *number*, the raw number of strings in the category; *volume per string*, the par value in millions of the number of contracts purchased or sold in the string; *trades per string*, the number of trades in the string; *market share per*

³ This means, for example, a participant that purchases \$25 million from 9:00 a.m. to 9:10 a.m. and then another \$25 million from 10:00 a.m. to 10:10 a.m. conducts two separate strings of \$25 million and not one of \$50 million.

string, the volume of the string in percentage of daily volume (single counted); average *duration*, two times the volume-weighted execution time from start to finish of the string in minutes (multiplied by two since raw duration is a half-life); and average *aggressiveness*, the percentage of the volume in the string executed via marketable orders.

TABLE 3 ABOUT HERE

A few trends are worth noting. As time passes, institutions distribute their trading in smaller and smaller pieces. For example, to build a position sized \$10–25 million, 34 trades would be used in 2009 versus 54 trades in 2014. Moreover, at least for positions smaller than \$100 million, institutions shift to using a more and more aggressive mix of trades. For example, to build a position sized \$10–25 million, the mix of orders was 39.5% marketable in 2009 versus 45.1% marketable in 2014. Last, just for the smallest positions (\$2–10 million), it took longer and longer every year: from 14.1 minutes in 2009 to 21.6 minutes in 2014. It is possible the move to smaller trades, more aggressive trades, and a longer duration of trading could be (in part) a response to the increasing presence of HFT in derivatives trade.

d. Computation of other liquidity metrics for trade strings

In addition to the implementation shortfall, we compute weekly and monthly averages of two per-trade liquidity metrics solely for the trades in the trade strings (and not for all the trades on the market): the per-trade *effective spread*, the signed difference between the trade price and contemporaneous midquote; and the per-trade *price impact*, the signed

difference between the contemporaneous midquote and the midquote some time interval after the trade. We compute these metrics only for trades in the trade strings to narrow the focus of the liquidity metrics to outcomes experienced by institutions when negotiating large positions. It could be that HFTs improve the quality of the market on average but worsen it specifically when they detect the presence of an institution, so we focus on the institutional case, as many papers have covered the average case. The price-impact metric is computed for intervals of 5 seconds, 30 seconds and 5 minutes to understand whether a reduction in price impact is permanent or is merely a delay in the impact.

2. Methodology

To find a statistical instrument to control for endogeneity, we search for a shock to the presence of HFT, as in Brogaard and Garriott (2017). We identify November 2011 to February 2012 as a period of shock because the number of HFTs on the bond future increases significantly and at least one market participant identified as an HFT exits (see Figure 1). We identify a treatment effect by comparing outcomes for the bond future around this shock to those of the rates future, which lacks HFT presence.

The model we use to identify the treatment effect of increased HFT presence is the DiD event study. Since the treatment and control observations are clearly heteroskedastic, we write the DiD model in first differences, which is more conservative than attempting to model the heteroscedasticity:

$$(metric_{CGB,t} - metric_{BAX,t}) = c + \beta treatment_t + \varepsilon_t,$$

where *metric* is the dependent variable; the index *CGB* or *BAX* indicates the treatment or control group; the index *t* is the week, the interval of observation of the study; the variable *treatment* is a dummy equalling one after November 2011; the estimated coefficient *c* is the first difference of the DiD; and the estimated coefficient β is the difference in differences (the coefficient of interest). The DiD model appears to lack fixed effects for the treatment and control groups because they have been differenced out. Since this is a single time series, there is no need to cluster standard errors by date or by security.

The event window of the study is the three months before and three months after November 2011. We drop the week of the Christmas holiday 2011. The event study is run on weekly averages of: *implementation shortfall* for various subsamples of the trade strings; the *effective spread* of the trades in the trade strings; and various per-trade *price impact* metrics for the trades in the trade strings. *Implementation shortfall* is the difference between the volume-weighted average price of a string and the midquote contemporaneous to the first trade. *Effective spread* is the signed difference between the trade price and contemporaneous midquote for trades in the strings; *price impact* is the signed difference between the contemporaneous midquote and the midquote some time interval after the trade for trades in the strings.

The interval of observation is weekly rather than daily because there are not always trade strings sized larger than \$100 million on any given day and because the IS metrics at a daily frequency tend to be volatile. If the price should rise during the trading day as the result of macroeconomic news, it will exogenously raise the IS metrics of buyers and lower the IS metrics of sellers, which are not always balanced in size. Using a weekly average rather than a daily average smooths these outliers.

In our robustness checks, we vary the size of the event window from two months to nine months and find it has no impact on the results. Also in the robustness checks, we check whether the treatment effect can be explained using control variates,

$$(metric_{CGB,t} - metric_{BAX,t}) = c + \beta treatment_t + \delta(X_{CGB,t} - X_{BAX,t}) + \varepsilon_t,$$

where all indices, variables and coefficients are the same as before with the addition of X , a vector of three control variates—the *trading volume* for the relevant contract, the rolling two-week close-price *volatility* of the relevant contract, and the *lagged value* of the metric for the relevant contract—and the estimated coefficient δ , the effect of the control variates.

3. HFT improves implementation shortfall on smaller institutional trade strings

We begin by looking at long-term evidence on implementation shortfall in the period 2009–2014 and correlate it with the growth of HFT.

FIGURE 4 ABOUT HERE

Figure 4 plots parallel trends of the shortfall metric for both the bond future and rates future. The vertical black bar is for November 2011, the beginning of the event used in the event study. Each panel of the figure shows the shortfall for a different size of position (\$2–10 million, \$10–25 million, \$25–100 million, and \$100 million+), the format for most of our figures. Although the shortfall series have different levels, there are correlations in the trends (that we have made easier to view by using two y-axis scales). The correlation

between the monthly shortfall metric of the average bond-future trade string and the average rates-future trade string is 0.659.

In Panel A, the shortfall for positions sized \$2–10 million exhibits a clear break on the event date, an improvement of perhaps 30 basis points, which is permanent through the subsequent three years of the sample. There is no break for the control group, which has slightly worse shortfall in the post-period. The other panels do not evidence an impact clearly attributable to the event. Positions sized \$10–25 million see improvements later in 2012, which could conceivably be due to the development of HFT presence on the exchange, but we lack a clear instrument for this period and can say nothing. As for positions sized \$25 million and greater, there is no obvious change in the metric's behaviour before and after the event.

This is the result one should expect from HFT. The business strategy of HFT is to make quick profits using its comparative advantages in quantitative analysis and ultra-low response times (Menkveld 2013; Carrion 2013; Hendershott and Riordan 2013). HFT trades in a way that stabilizes prices (Brogaard et al. 2016, Riordan and Storkenmaier 2012) but is ultimately done on an agency basis (Baron et al. 2016), as HFT does not have a comparative advantage against traditional financial intermediaries in the supply of risk capital. Thus, it would be surprising to find that HFT has any effect on the cost to build larger positions, which benefit much more by increased risk capital than by short-term price stability. It is a little surprising even to find that HFT lowers costs for building small positions. Nevertheless, it is the smaller positions that can benefit materially from a tighter effective spread and a milder short-term price impact, the usual outcomes of increased HFT

(as in, e.g., Chaboud et al. 2014; Carrion 2013; Jones 2013; Hendershott, Jones and Menkveld 2011).

It is interesting to note that some of the shortfall metrics hover around zero. Zero is an achievable shortfall if a trader uses an intelligent balance of passive and aggressive trades, crossing the spread only when necessary to obtain immediacy (Griffiths et al. 2000). As Table 3 shows, larger positions are built a little more aggressively than smaller positions, perhaps reflecting a greater need for immediacy. As the sample ages, positions smaller than \$100 million are built more aggressively, indicative of a lower cost of immediacy, perhaps due to HFT.

To confirm the graphical evidence, we use the DiD model to search for a treatment effect in a window surrounding the event date. Table 4 gives the results as well as the robustness check.

TABLE 4 ABOUT HERE

Panel A identifies a treatment effect of 27.9 basis points for positions sized \$2–10 million, and the effect is significant at the 5% level. There are no statistically identifiable effects for other sizes of institutional positions. Panel B shows the identified effect for \$2–10 million is robust to control variates and to the size of the event window. In fact, the treatment effect is larger after adding the control variables.

In a second table, Table 5, we use the DiD model to search for a treatment effect in the trade strings now binned by duration: strings that take 0–10 minutes, 10–60 minutes, and 60 minutes or more.

TABLE 5 ABOUT HERE

The regression identifies a treatment effect of 12 basis points for positions requiring 0–10 minutes to construct. As before, there are no statistically identifiable effects for other durations of institutional positions, though for positions requiring 10–60 minutes, the effect just misses significance.

4. HFT is a frequent institutional counterparty but takes small exposures

In seeking to understand why HFT is having this effect, we first check whether HFT is indeed transacting with the institutions when they are building positions and whether HFT remains in the market while positions are worked. If HFT does not, it cannot be the cause of the improvement. Figure 5 shows the HFT group's volume share as a counterparty to institutions when they are building positions, by position size and during volume quartiles of the lifetime of the position.

FIGURE 5 ABOUT HERE

HFT is a consistent counterparty to all positions built by institutions in all bins and during all quartiles of their lifetimes. HFTs are counterparty to 35% of the volume of the smallest trade strings, 30% of the next largest, 27% of the next largest, and 25% of the largest. There is evidence that certain HFTs withdraw during the trade strings, but in contrast to Korajczyk and Murphy (2016), we find this is more so for the smaller orders.

The largest withdrawal is seven percentage points (out of 35) for the \$2–10 million category.

5. There is no evidence of HFT back running

Now we search for HFT dynamics that are consistent with theoretical predictions on predation. In Brunnermeier and Pedersen (2005), a predatory trader infers the presence of a large investor who has a liquidity need and exploits this need in two stages. First, the predator trades in the *same* direction as the large investor, amplifying the investor’s price impact. Second, after the price is impacted, the predator profits by unwinding its position at the new price. Similarly, in Yang and Zhu (2016), a predatory back runner infers the presence of a large investor that is trading on information and, as before, trades first in the *same* direction as the investor and then in the opposite.

If the predator’s position in the asset is graphed alongside the victim’s position, the necessary observable outcome is a “hump” in the predator’s net order flow during and after the period of the investor’s activity.⁴ We engage the theory by creating graphs of HFT and nonHFT net order flow while institutions build positions and search for the “hump” shape predicted by theory. First, we search for the hump shape in trade strings sorted by volume bin. Second, we search for the hump shape in trade strings sorted by informedness proxies.

⁴ Van Kervel and Menkveld (2016) provide a set of intuitive graphs showing an idealized back runner’s position in time.

a. There is no evidence of back running in trade strings binned by volume

In Figure 6, we graph average HFT and nonHFT net order flow as institutions build positions via trade strings segregated in bins by size: \$2–10 million, \$10–25 million, \$25–100 million, and more than \$100 million. The HFT position is the middle line (red). The top line (blue) is the position of institutions building the position; the bottom line (green) is the position of all other nonHFTs on the exchange. The graphs plot the trade strings by deciles of time in the lifetime of the string.

FIGURE 6 ABOUT HERE

In every bin, the top line (institutions building the positions) and bottom line (other nonHFT) are almost mirror images, whereas the middle line (HFT) is persistently just underneath zero, consistent with a role as a minor liquidity supplier. The ultimate bearers of the risk to a trade string are the other nonHFTs. HFTs hold positions on the order of one to five contracts (\$100,000 to \$500,000) throughout the life of the string whatever its size, a tiny amount compared with positions of size \$25 million or \$100 million. Accordingly, the red (HFT) line is hard to distinguish from zero for the largest positions. There is no hump shape in the red line, which is inconsistent with the hypothesis that HFT is, on aggregate, trading to amplify the price impact of institutional trading. Since we observed HFTs are transacting with these institutions in Figure 5, it must be that HFTs are trading with institutions but laying off any exposures quite soon afterward.

b. There is no evidence of back running in trade strings binned by informedness

We now search for HFT predation during trade strings when the strings are binned by their long-term or “permanent” price impact. We do this to engage more directly with the theoretical literature. It is not clear whether our data provide a test of every model on predation or back running, since models accord motives to institutional investors that might not be present in our data. Specifically, Brunnermeier and Pedersen (2005) assume investors who are trading for liquidity reasons, whereas Yang and Zhu (2016) assume investors who are trading on information. Arguably, futures markets should by nature contain both types of investor because futures markets are conventionally understood to be markets both for risk transfer and for price discovery (Newman 1998). Nevertheless, we check that both types are present in our data to ensure that we search for predatory trading as conceived by both frameworks. We check that the data contain groups of trade strings that appear to have information and groups that appear to lack information (i.e., that are motivated by liquidity). Then we check for HFT predation during strings when segregated by the apparent investment motive.

Since we do not observe *ex ante* the motives of any market participant, we lack an *ex ante* criterion to group strings by informedness. As a substitute, we sort the strings by two different *ex post* criteria and check for back running under both criteria. First, we bin trade strings by their size and duration. Large institutional positions that are worked slowly are more likely to be informed because informed traders trade more slowly than uninformed traders to hide in their flow (Kyle 1985). Indeed, in the data, the permanent price impacts are higher for positions worked more slowly, particularly if the position is large (more than \$25 million). Second, we also bin trade strings by initial price impact, measured as the

signed change of the midquote in basis points from the start to end of the string. Trade strings that have larger initial price impact arguably have more impact to contribute in the long term, which we also find is reflected in the data.

We find the long-term or “permanent” price impacts for trade strings can be large. To illustrate, Figure 7 graphs the mean price impact and its 95% confidence interval for the trade strings when binned by duration and volume. There are two duration bins: strings worked in less than one hour and strings worked in greater than one hour. The volume bins are the same as before, though to economize we combine the two smallest volume bins (\$2–10 million and \$10–25 million). In the figure, the price impacts are graphed for various time lags: at string finish, at end of day, and at five 24-hour increments after finish.

FIGURE 7 ABOUT HERE

For the trade strings that are worked slowly, long-term price impacts range from one to three cents and are statistically significant five days after the string. A positive price impact of two cents, multiplied by the sample average five-day daily trading value of \$135 million, means an increase in costs to buyers (or in revenue to sellers) of \$2.7 million per string.

Similarly, Figure 8 graphs the price impact for trade strings in bins by initial price impacts in four ranges: -1–1 cent, 1–3 cents, 3–6 cents, and 6+ cents.

FIGURE 8 ABOUT HERE

For the trade strings that have positive initial-price impact, permanent price impacts range from 1 to 10 cents and are again statistically significant five days after the string. A positive price impact of 4 cents, using the same math, means an increase in costs to buyers (or in revenue to sellers) of \$5.4 million per string.

For none of the bins, whether sorted by duration or by initial price impact, does HFT act in a way consistent with back running or predation. Figure 9 graphs HFT and nonHFT net order flow as institutions build positions via trade strings segregated in the duration bins. As before, the HFT position is the middle line (red). The top line (blue) is the position of institutions building the position; the bottom line (green) is the position of all other nonHFTs on the exchange. The graphs plot the trade strings by deciles of time in the lifetime of the string.

FIGURE 9 ABOUT HERE

As in section 5.a., no matter the bin, the top line (institutions building the position) and bottom line (other nonHFT) are almost mirror images, whereas the middle line (HFT) is flat and located just underneath zero, consistent with a role as a minor liquidity supplier. There is no hump shape in the red line, which is inconsistent with the hypothesis that HFTs are, on aggregate, trading to amplify the price impact of institutions. Figure 9 graphs HFT and nonHFT net order flow as institutions build positions via trade strings segregated in the duration bins. Again, the HFT line is always flat and located just underneath zero.

FIGURE 10 ABOUT HERE

Since HFT is never observed to be predated or back running in the data, even on trade strings with large and permanent price impact, its behaviour is inconsistent with the trading strategies formalized both by Brunnermeier and Pedersen (2005) and by Yang and Zhu (2016).

6. HFT's reduction of implementation shortfall can be explained by a reduction in pre-trade price impact

In this section, we explain how HFT improves shortfall, which is a puzzling improvement because HFT is not supplying risk capital: As section 4 shows, HFT offloads the risk of its trades with institutions even before they have finished building a position. We explain the improvement by showing that HFT has its typical effect of improving short-term price impact but *not* improving long-term price impact. In other words, HFT delays but does not eliminate price impact, which is helpful to institutions building smaller positions but marginal to those building larger ones.

Institutions building large positions would not benefit from a short-lived delay in price impact because they stay in the market much longer than the delay. Large positions in the sample are worked over periods four times as long as small positions and have implementation shortfall that is multiples of the bid-ask spread. However, institutions working smaller positions can benefit by the typical improvements in liquidity attributable to HFT as these institutions never moved the price far in the first place. If HFT causes the price to remain constant for 30 seconds longer than it would have otherwise, institutions can trade a little more volume at a mildly tighter effective spread.

Figure 11 plots parallel trends of effective spreads and price impacts for the bond future and rates future through the sample.

FIGURE 11 ABOUT HERE

There are parallel trends in price impacts but not in effective spreads. The rates future is a poor control for effective spreads because the rates future was tick-constrained by the minimum spread during the sample period. However, the price impacts are better correlated, with correlations of 0.600, 0.645 and 0.648 for price impacts measured at the 5-second, 30-second and 5-minute lag.

In Panel A, there is a visible improvement in effective spreads after the event date. We do not plot the control group's effective spreads after September 2014 as the contract decimalized on this date. In Panels B–D, there is a long-term improving trend in price impact measured at 5-second, 30-second and 5-minute lags, and it is matched in part by improvements in the rates future. Here, the DiD model has value in distinguishing treatment from trend. The improvements in Panels B and C will be robust to differencing, but the “improvement” in Panel D is not.

Table 6 reports the results of the DiD event study on effective spreads and price impacts. Again, these metrics are computed using only the trades that are part of the trade strings, meaning we focus the analysis on those trades that matter to institutions.

TABLE 6 ABOUT HERE

The event has statistically significant effects on effective spreads, 5-second price impacts, and 30-second price impacts with magnitudes of 14 basis points, 4 basis points, and 5 basis points. However, there is no statistically significant effect on the 5-minute price impact. We report a robustness check on the 5-second price impact showing the effect is the same if we add control variates and vary the event window; in unreported results, it is the same for the effective spread and 30-second price impact, whereas the 5-minute price impact is insignificant. Our conclusion is that HFT is tightening spreads and creating more price stability but only in the short run. The price impact of institutional trades eventually makes its way into prices but at a greater lag.

7. Conclusions

We use a DiD event study to test whether HFT improves implementation shortfall for institutional-sized orders, defined as unidirectional strings of trades by one of the sides and of value at least \$2 million. HFT does improve shortfalls but only for positions sized \$2–10 million, the smallest category we study. We explain the improvement as the result of HFT's conventional role in tightening effective spreads and eliminating short-lived price discrepancies. The same event-study methodology shows that HFT improves price impacts at 5- and 30-second horizons but not at 5-minute horizons. This is the sort of improvement that would benefit institutions building smaller positions, which are worked quickly, but not larger positions, which are worked for periods of time much longer than five minutes.

The results in the paper confirm some of the existing literature. As in Korajczyk and Murphy (2016), the HFTs in our sample reduce their exposure to institutional positions as institutions work the positions. However, in contrast with Van Kervel and Menkveld

(2016), we find no evidence of back running. Instead, using our results on price impact, we argue HFT can be thought of as a “shock absorber” for institutions. HFT is counterparty to a significant portion of institutional volumes and holds them for a short time before returning the position to the market opportunistically. This decreases the temporary price impact of these trades, though not the longer-term price impact, so the contribution of HFT is limited.

Bibliography

- Baron, Matthew D., Jonathan Brogaard, Björn Hagströmer, and Andrei A. Kirilenko. (2016). "Risk and return in high-frequency trading."
- BIS, Markets Committee. (2016). "Electronic trading in fixed income markets." Bank for International Settlements Markets Committee 7.
- Brogaard, Jonathan, Allen Carrion, Thibaut Moyaert, Ryan Riordan, Andriy Shkilko, and Konstantin Sokolov. (2016). "High-frequency trading and extreme price movements."
- Brogaard, Jonathan, and Corey Garriott. (2017). "High-frequency trading competition."
- Brogaard, Jonathan, Terrence Hendershott, and Ryan Riordan. (2014). "High-frequency trading and price discovery." *Review of Financial Studies*.
- Brunnermeier, Markus K., and Lasse Heje Pedersen. (2005). "Predatory trading." *The Journal of Finance* 60, no. 4: 1825-1863.
- Carrion, Allen. (2013). "Very fast money: High-frequency trading on the NASDAQ." *Journal of Financial Markets* 16, no. 4: 680-711.
- Chaboud, Alain P., Benjamin Chiquoine, Erik Hjalmarsson, and Clara Vega. (2014). "Rise of the machines: Algorithmic trading in the foreign exchange market." *The Journal of Finance* 69, no. 5: 2045-2084.
- Chakrabarty, Bidisha, Pankaj K. Jain, Andriy Shkilko, and Konstantin Sokolov. (2014). "Speed of market access and market quality: Evidence from the SEC naked access ban."
- Griffiths, Mark D., Brian F. Smith, D. Alasdair S. Turnbull, and Robert W. White. (2000). "The costs and determinants of order aggressiveness." *Journal of Financial Economics* 56, no. 1: 65-88.
- Harris, Lawrence. (1998). "Optimal dynamic order submission strategies in some stylized trading problems." *Financial Markets, Institutions & Instruments* 7, no. 2: 1-76.
- Hendershott, Terrence, Charles M. Jones, and Albert J. Menkveld. (2011). "Does algorithmic trading improve liquidity?" *The Journal of Finance* 66, no. 1: 1-33.
- Hendershott, Terrence, and Ryan Riordan. (2013). "Algorithmic trading and the market for liquidity." *Journal of Financial and Quantitative Analysis* 48, no. 4: 1001-1024.
- Hirschey, Nicholas. (2017). "Do high-frequency traders anticipate buying and selling pressure?"
- Jones, Charles M. (2013). "What do we know about high-frequency trading?"
- Kirilenko, Andrei, Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun. (2017). "The flash crash: High-Frequency Trading in an electronic market." *The Journal of Finance* 72, no. 3: 967-998.

- Korajczyk, Robert A., and Dermot Murphy. (2016). "High frequency market making to large institutional trades."
- Kyle, Albert S. (1985). "Continuous auctions and insider trading." *Econometrica: Journal of the Econometric Society*: 1315-1335. 27, no. 8: 2267-2306.
- Menkveld, Albert J. (2013). "High frequency trading and the new market makers." *Journal of Financial Markets* 16, no. 4: 712-740.
- Menkveld, Albert J. (2016) "The economics of high-frequency trading: Taking stock." *Annual Review of Financial Economics* 8: 1-24.
- Newman, Peter, ed. (1998). *The New Palgrave Dictionary of Economics and the Law*. Springer.
- Riordan, Ryan, and Andreas Storkenmaier. (2012). "Latency, liquidity and price discovery." *Journal of Financial Markets* 15, no. 4: 416-437.
- Shkilko, Andriy, and Konstantin Sokolov. (2016). "Every cloud has a silver lining: Fast trading, microwave connectivity and trading costs."
- Tong, Lin. (2015). "A blessing or a curse? The impact of high frequency trading on institutional investors."
- Van Kervel, Vincent, and Albert J. Menkveld. (2016). "High-frequency trading around large institutional orders."
- Yang, Liyan, and Haoxiang Zhu. (2016). "Back-running: Seeking and hiding fundamental information in order flows."

Figure 1: Implementation shortfall for institutional positions of sizes \$2–10M and the number of active HFTs in inverted scale

This figure plots the average monthly number of HFTs in inverted scale (zero on top) and the average monthly implementation shortfall for trade strings sized \$2–\$10 million for the CGB contract, the future on the Government of Canada 10-year bond. A *trade string* is a sequence of trades by one of the sides, 95% in the same direction (buy or sell), with no more than 20 minutes between any two trades.

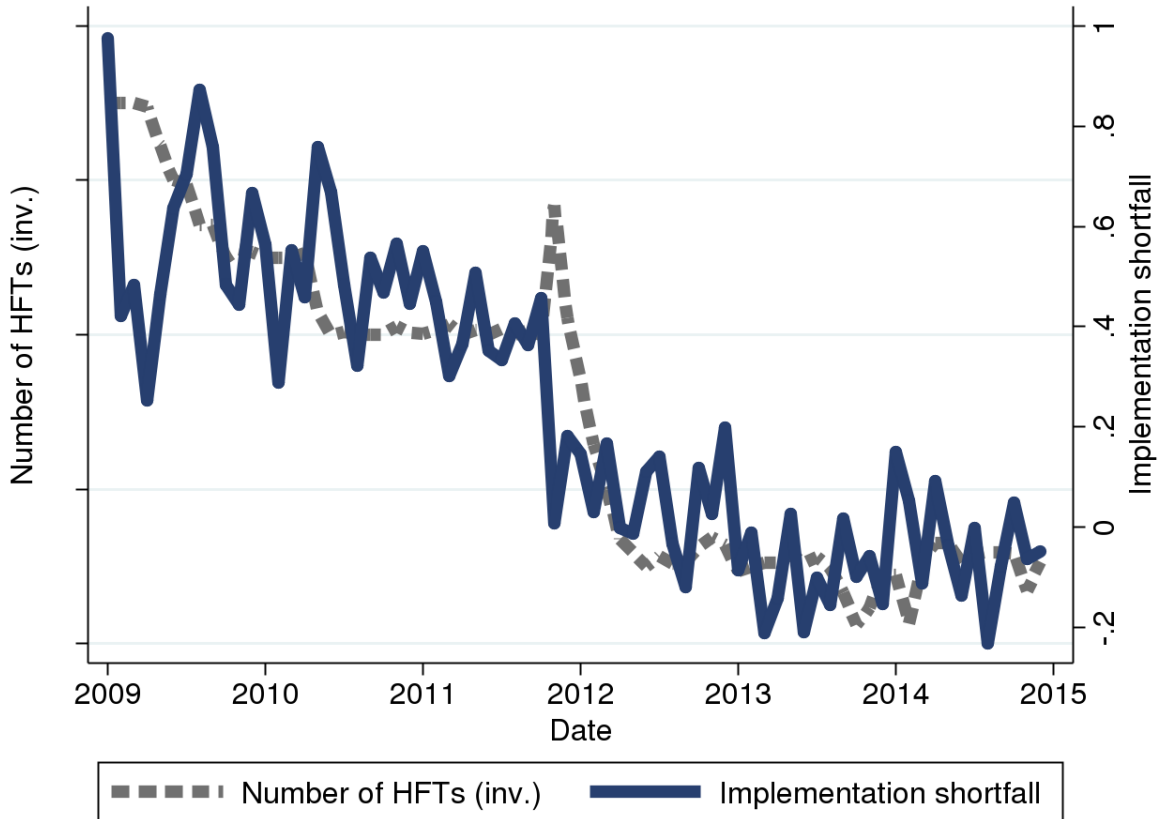


Figure 2: Trading volume of the CGB and BAX contracts

This figure gives average monthly trading volume (in number of contracts) for the CGB (the future on the Government of Canada 10-year bond) and for the front three contracts of the BAX (the future on the Canadian Banker's Acceptance rate). *Trading volume* is the total electronic number of contracts traded during normal trading hours.

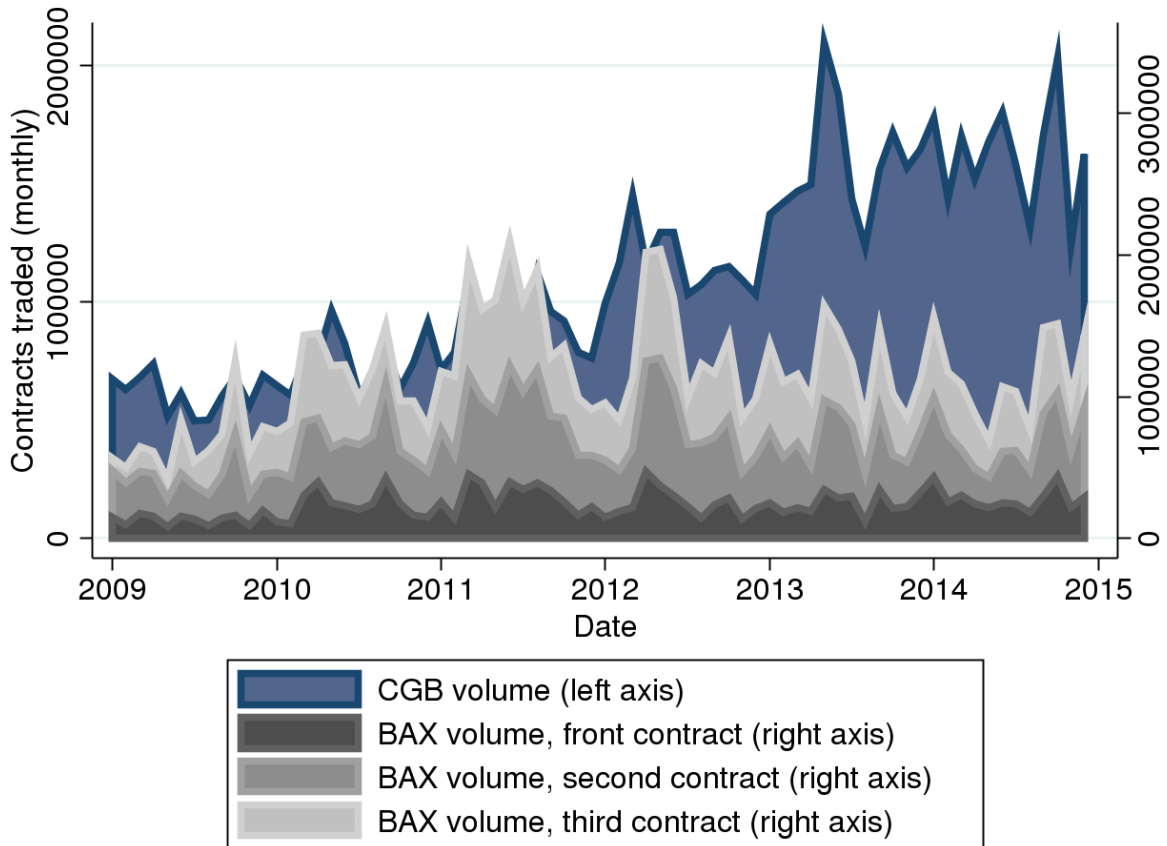


Figure 3: HFT volume share of the CGB and BAX contracts

This figure gives average monthly percent HFT volume share for the CGB (the future on the Government of Canada 10-year bond) and for the front three contracts of the BAX (the future on the Canadian Banker's Acceptance rate).

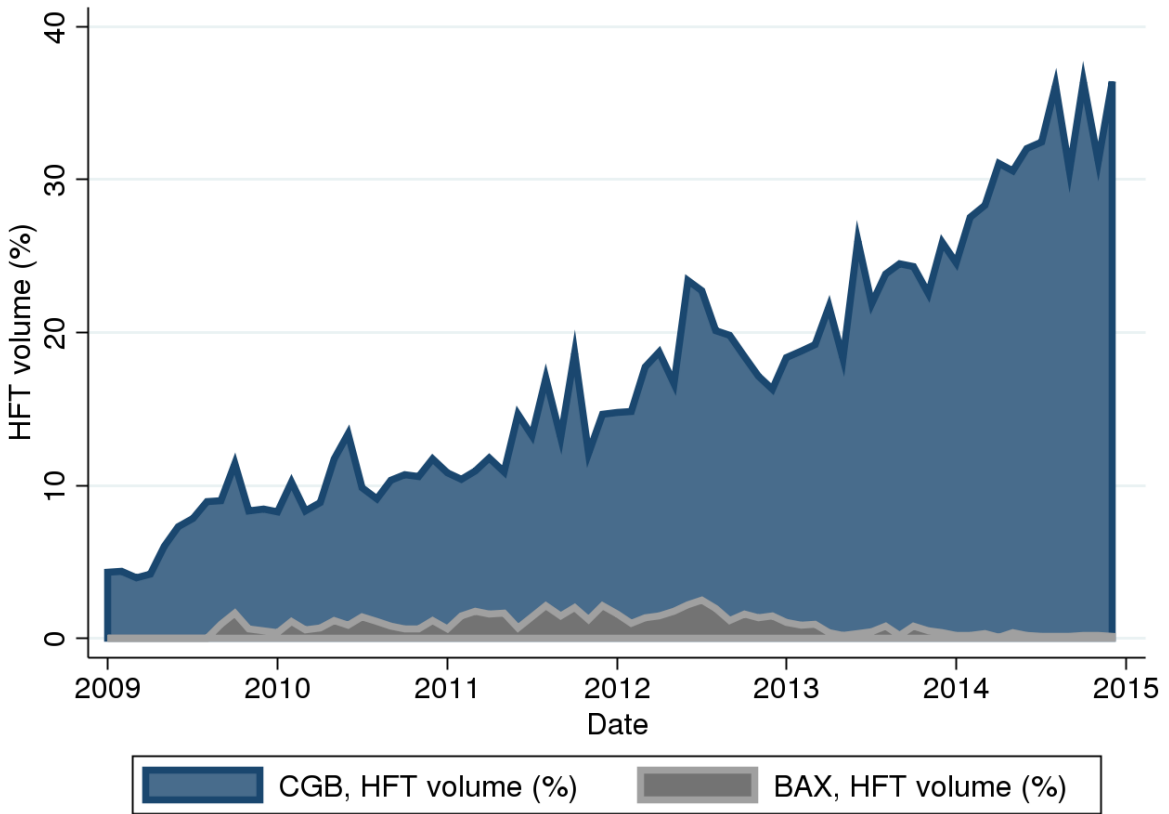
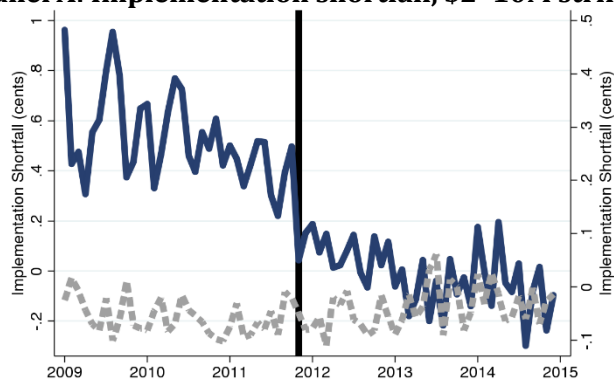



Figure 4: CGB implementation shortfall compared with BAX implementation shortfall, by volume bin

This figure shows the weekly average implementation shortfall of trade strings for the CGB contract (in solid blue), the future on the Government of Canada 10-year bond, compared with that for the BAX contract (in dashed gray), the future on the Canadian Banker's Acceptance rate, 2009–2014. A *trade string* is a sequence of trades by one of the sides, 95% in the same direction (buy or sell), with no more than 20 minutes between any two trades. *Shortfall* is the difference between the volume-weighted average price of a string and the midquote contemporaneous to the first trade. The vertical black bar is the event month.

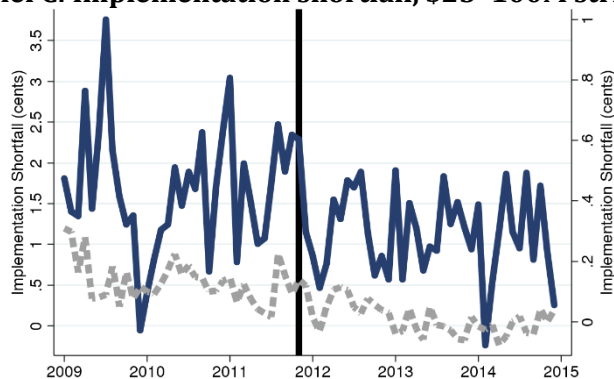
Panel A: Implementation shortfall, \$2–10M strings



Panel B: Implementation shortfall, \$10–25M strings



Panel C: Implementation shortfall, \$25–100M strings



Panel D: Implementation shortfall, \$100M+ strings

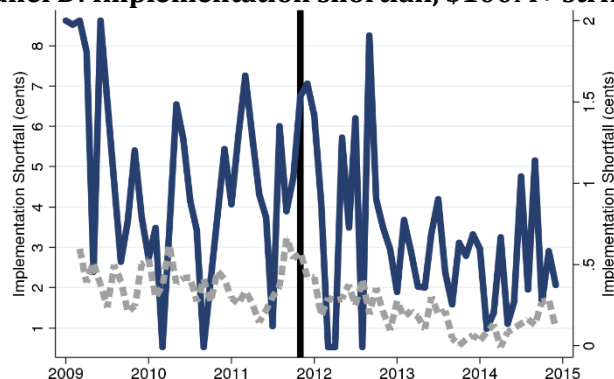




Figure 5: HFT volume share as counterparty of institutional trades, by volume bins

This figure gives the percent of volume institutions transacted with HFT during a trade string for the CGB contract. A *trade string* is a sequence of trades by one of the sides, 95% in the same direction (buy or sell), with no more than 20 minutes between any two trades. Each panel is a subsample of trade strings by volume. The graph in each panel divides trade strings into four quartiles at cutoffs of 25%, 50% and 75% of the volume.


Panel A: HFT volume share, \$2–10M strings



Panel B: HFT volume share, \$10–25M strings



Panel C: HFT volume share, \$25–100M strings



Panel D: HFT volume share, \$100M+ strings

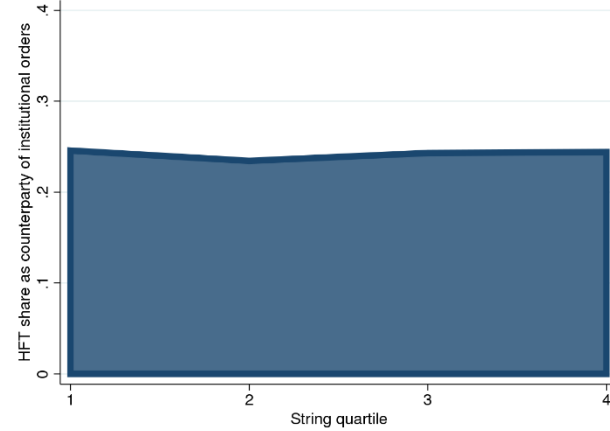




Figure 6: Net HFT and nonHFT inventory positions with respect to trade strings, by volume bin

This figure gives the net inventory position held by institutions building a position (blue), the net position of HFTs in the market (red), and the net position of all other participants in the market (green), during 10 time deciles during the life of a trade string for the CGB contract, the future on the Government of Canada 10-year bond. The series sum to zero by definition.


Panel A: HFT and nonHFT positions, \$2–10M strings



Panel B: HFT and nonHFT positions, \$10–25M strings



Panel C: HFT and nonHFT positions, \$25–100M strings



Panel D: HFT and nonHFT positions, \$100M+ strings





Figure 7: Permanent price impact of institutional trade strings, in bins by size and duration

This figure gives the price impact of trade strings at the time lags: immediately after the string's finish, at end of day, and at five 24-hour increments after finish. In black is the mean price impact, and in gray is the 95% confidence interval. A *trade string* is a sequence of trades by one of the sides, 95% in the same direction (buy or sell), with no more than 20 minutes between any two trades. The panels subsample the trade strings in six bins by volume and duration. *Volume* is the size of the total net position acquired during the trade string; *duration* is two times the volume-weighted execution time from start to finish of the string in minutes.

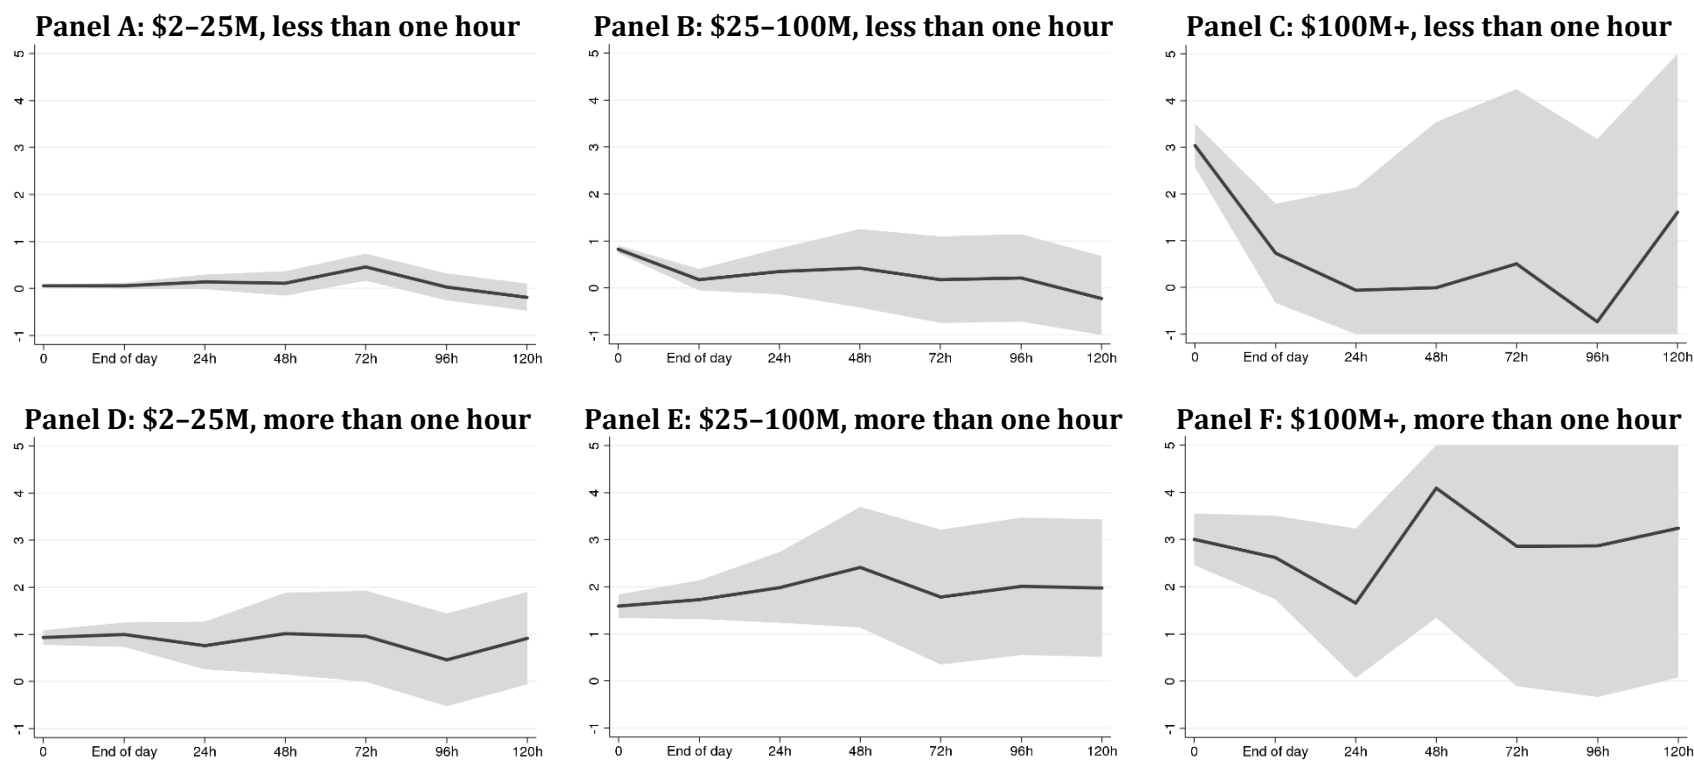


Figure 8: Permanent price impact of institutional trade strings, in bins by initial price impact

This figure gives the price impact of trade strings at the time lags: immediately after the string's finish, at end of day, and at five 24-hour increments after finish. In black is the mean price impact, and in gray is the 95% confidence interval. A *trade string* is a sequence of trades by one of the sides, 95% in the same direction (buy or sell), with no more than 20 minutes between any two trades. The panels subsample the trade strings in four bins by immediate price-impact. Immediate *price impact* is the signed difference between the midquote at the start of the trade string and the midquote immediately after the trade string finishes.

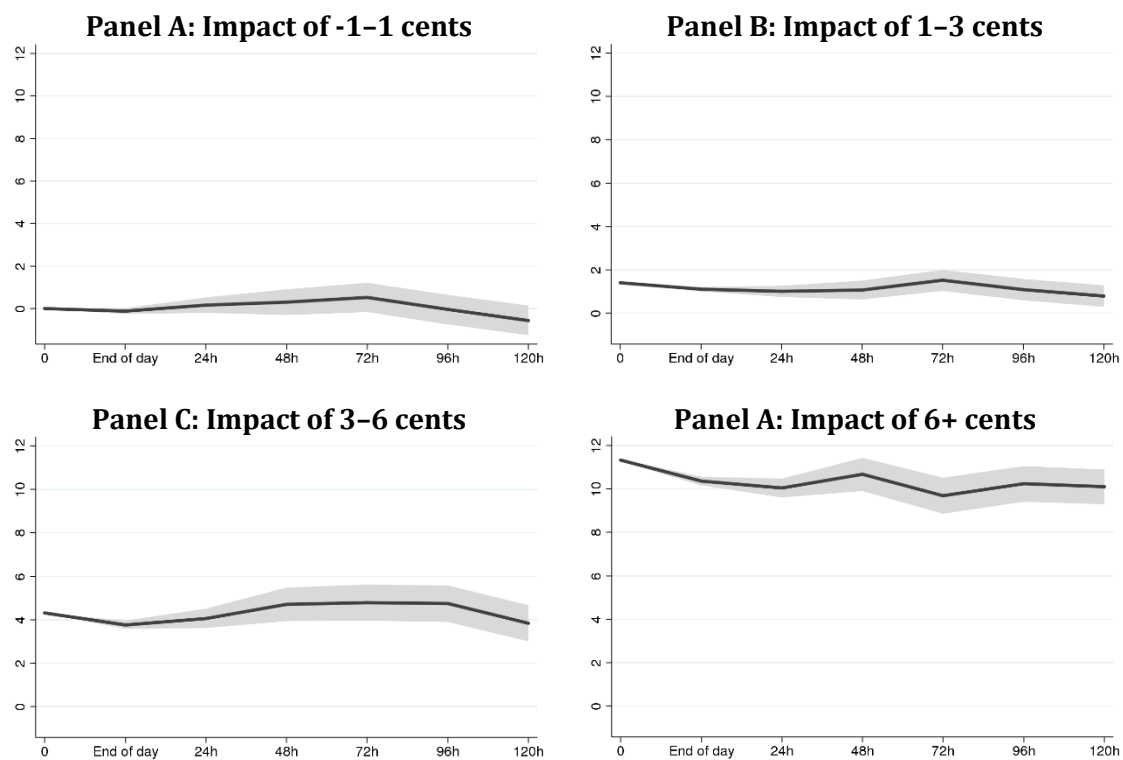


Figure 9: Net HFT and nonHFT inventory positions with respect to trade strings, by duration and volume bin

This figure gives the net inventory position held by institutions building a position (blue), the net position of HFTs in the market (red), and the net position of all other participants in the market (green), during 10 time deciles during the life of a trade string for the CGB contract, the future on the Government of Canada 10-year bond. The series sum to zero by definition. The panels subsample the trade strings by volume and duration bins. *Volume* is the size of the total net position acquired during the trade string; *duration* is two times the volume-weighted execution time from start to finish of the string in minutes.

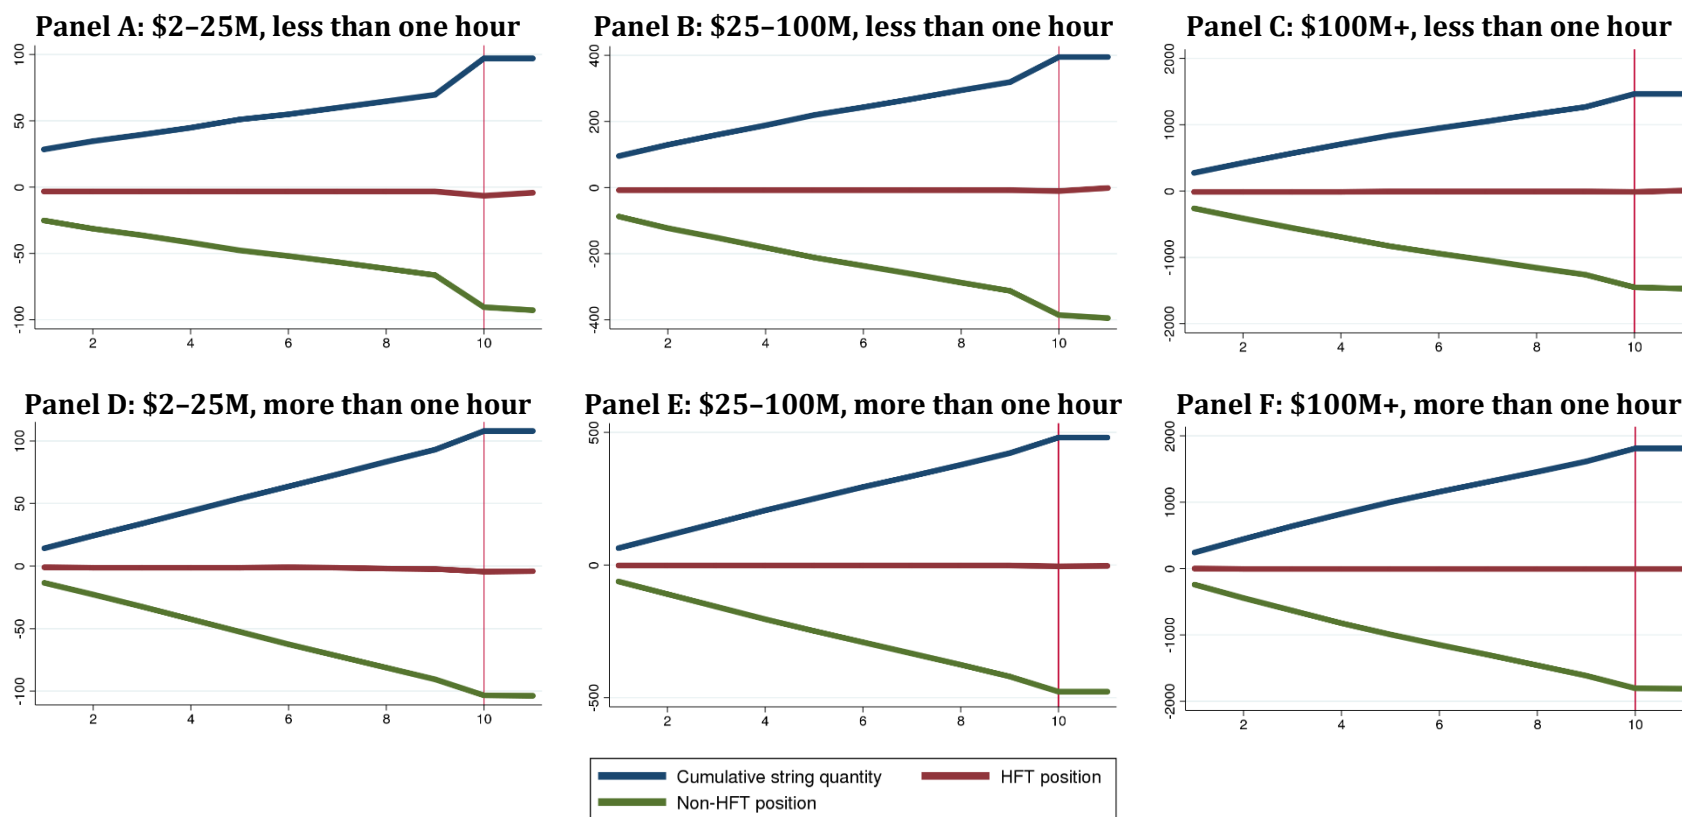


Figure 10: Permanent price impact of institutional trade strings, in bins by initial price impact

This figure gives the net inventory position held by institutions building a position (blue), the net position of HFT in the market (red), and the net position of all other participants in the market (green), during 10 time deciles during the life of a trade string for the CGB contract, the future on the Government of Canada 10-year bond. The series sum to zero by definition. The panels subsample the trade strings in four bins by immediate price-impact. Immediate *price impact* is the signed difference between the midquote at the start of the trade string and the midquote immediately after the trade string finishes.

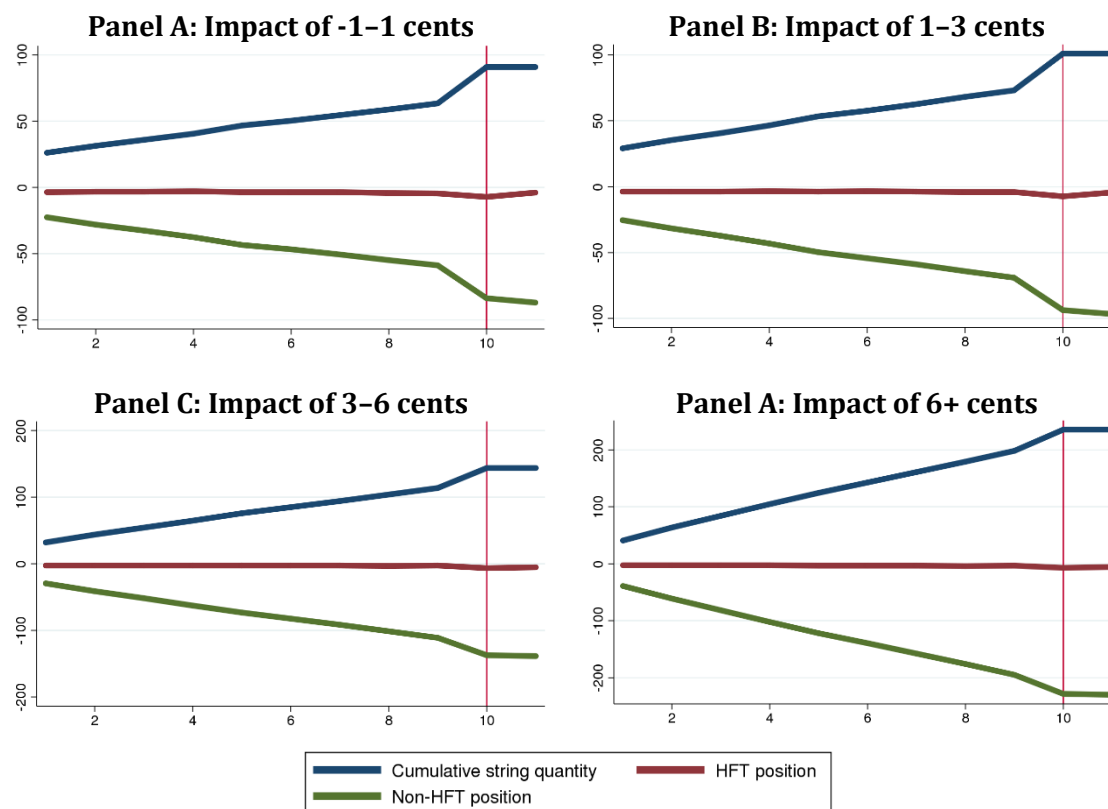
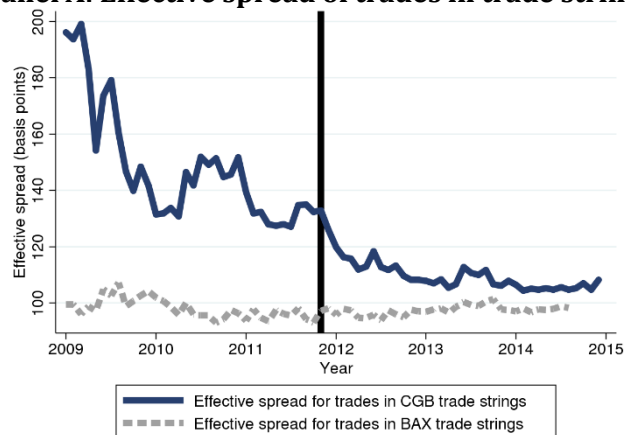


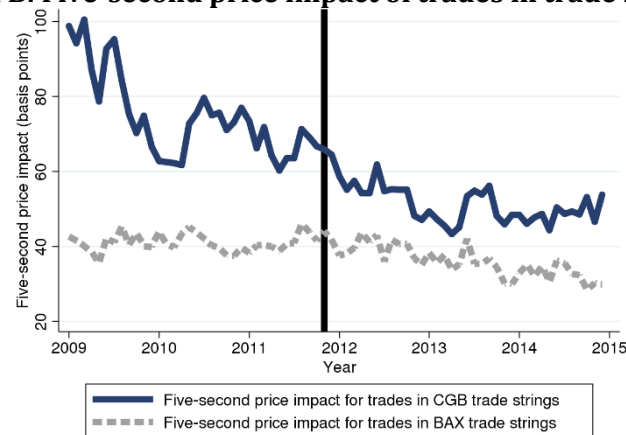
Figure 11: CGB liquidity metrics compared with BAX liquidity metrics

This figure gives liquidity metrics for trades in trade strings for the CGB contract (the future on the Government of Canada 10-year bond) compared with those for the BAX contract (the future on the Canadian Banker's Acceptance rate). *Effective spread* is the signed difference between the trade price and contemporaneous midquote; *price impact* is the signed difference between the contemporaneous midquote and the midquote some time interval after the trade.


Panel A: Effective spread of trades in trade strings



Panel B: Five-second price impact of trades in trade strings



Panel C: Thirty-second price impact of trades in trade strings



Panel D: Five-minute price impact of trades in trade strings




Table 1: Summary statistics for the bond and rates futures, 2009–2014

This table gives yearly averages and standard deviations of the daily observations of various statistics for the CGB (the future on the Government of Canada 10-year bond) and for the BAX (the future on the Canadian Banker's Acceptance interest rate). *Price* is the contract's midquote price in dollars. *Number of trades* is the number of trades (single counting). *Trading volume* is the total electronic number of contracts traded during normal trading hours (using single-counting). *Bid-ask spread* is the difference between ask and bid prices in cents. *Inside depth* is the sum of the quantities of all outstanding visible limit orders at the bid and ask prices. *Volatility* is the two-week rolling standard deviation of the daily close price log return, annualized.

Panel A: Statistics for the bond future						
	Years					
	2009	2010	2011	2012	2013	2014
Price	122.3 (2.568)	121.7 (3.155)	126.3 (5.227)	135.4 (2.602)	132.0 (2.889)	134.1 (3.341)
Number of trades	4800.6 (1411.9)	5912.7 (1809.5)	8396.7 (2169.5)	11736.8 (2773.8)	15421.2 (4009.1)	16529.2 (4294.0)
Trading volume	15361.3 (4567.1)	18440.3 (6016.1)	22350.7 (5598.5)	28260.4 (6925.6)	38664.0 (10669.6)	39899.6 (10112.8)
Bid-ask spread	2.172 (0.512)	1.642 (0.184)	1.488 (0.124)	1.217 (0.0751)	1.171 (0.0671)	1.094 (0.0397)
Inside depth	28.43 (9.404)	30.14 (9.307)	30.30 (7.016)	52.51 (14.30)	79.04 (22.88)	88.25 (17.17)
Volatility	7.103 (2.264)	6.132 (1.367)	6.467 (1.246)	5.163 (1.110)	5.024 (1.445)	5.007 (1.350)

Panel B: Statistics for rates futures						
	2009	2010	2011	2012	2013	2014
Price	99.39 (0.159)	98.64 (0.317)	98.60 (0.336)	98.71 (0.107)	98.68 (0.0693)	98.69 (0.0713)
Number of trades	378.3 (196.2)	628.1 (313.9)	789.4 (360.4)	656.8 (337.4)	517.8 (338.1)	407.9 (265.7)
Trading volume	9307.6 (5662.8)	13335.1 (6712.5)	17199.4 (7922.8)	15272.8 (7713.3)	16116.6 (7697.0)	14922.0 (7082.9)
Bid-ask spread	1.016 (0.126)	1.026 (0.0901)	1.017 (0.0835)	1.002 (0.0819)	0.990 (0.0863)	0.837 (0.238)
Inside depth	1952.5 (1917.3)	1724.8 (1282.4)	2056.7 (1354.8)	3143.1 (1685.2)	10770.0 (8337.4)	11790.0 (9444.8)
Volatility	0.819 (0.340)	1.038 (0.362)	1.115 (0.357)	0.832 (0.313)	0.526 (0.265)	0.437 (0.274)

Table 2: Summary statistics on participant groups trading the CGB, 2009–2014

This table gives yearly averages and standard deviations of the daily observations of trading statistics for market participants trading the CGB contract, the future for the Government of Canada 10-year bond. The statistics are averaged over the members of the HFT and nonHFT groups. *Number of trades* is the average daily number of trades conducted by members of the participant group (with double counting). *Switch rate* is the average percent of trades in which a buy follows a sell or a sell follows a buy for members of the participant group. *Overnight inventory* is the end-of-day inventory position in number of contracts for members of the participant group.

Panel A: Statistics for HFTs						
	Years					
	2009	2010	2011	2012	2013	2014
Number of trades	648.8 (354.7)	781.2 (555.2)	1266.4 (837.7)	1412.2 (1098.4)	1710.3 (1740.4)	2670.8 (2820.9)
Switch rate	30.01 (9.474)	28.54 (10.17)	21.79 (8.078)	27.26 (11.19)	27.64 (10.73)	27.93 (12.07)
Overnight inventory	0.362 (1.604)	0.681 (3.491)	1.278 (3.767)	1.545 (3.905)	2.820 (7.623)	5.378 (12.87)

Panel B: Statistics for nonHFTs						
	2009	2010	2011	2012	2013	2014
Number of trades	400.7 (664.1)	437.3 (662.2)	598.7 (856.4)	808.5 (1279.6)	1060.0 (1794.3)	980.2 (1576.5)
Switch rate	12.85 (13.07)	12.39 (11.48)	11.44 (12.23)	9.232 (10.10)	8.653 (8.397)	8.768 (8.457)
Overnight inventory	410.6 (580.2)	430.2 (575.9)	468.7 (649.3)	541.6 (751.8)	689.9 (948.0)	682.3 (947.3)

Table 3: Summary statistics for bond-future trade strings by volume bin

This table gives yearly averages of statistics for trade strings on the CGB, the future on the Government of Canada 10-year bond. A *trade string* is a sequence of trades by one of the sides, 95% in the same direction (buy or sell), with no more than 20 minutes between any two trades. *Implementation shortfall* is the difference between the volume-weighted average price of a string and the midquote contemporaneous to the first trade. *Number* is the raw number of strings in the category. *Volume per string* is the par value in millions of the number of contracts purchased or sold in the string. *Trades per string* is the number of trades in the string. *Market share per string* is the volume of the string in percentage of daily volume (single counted). *Duration* is two times the volume-weighted execution time from start to finish of the string in minutes (multiplied by two since raw duration is a half-life). *Aggressiveness* is the percentage of the volume in the string executed via marketable orders.

	Years					
	2009	2010	2011	2012	2013	2014
Positions of \$2–10M						
Implementation shortfall	0.580¢	0.544¢	0.355¢	0.073¢	-0.089¢	-0.037¢
Number	14,977	17,789	20,254	23,687	26,222	26,513
Volume per string	\$4.7M	\$4.7M	\$4.7M	\$4.7M	\$4.7M	\$4.8M
Trades per string	15.5	15.8	18.4	19.8	19.2	20.0
Market share per string	0.32%	0.27%	0.22%	0.17%	0.13%	0.12%
Duration	14.1m	14.9m	16.7m	16.1m	16.9m	21.6m
Aggressiveness	40.1%	41.2%	42.3%	44.4%	42.0%	42.1%
Positions of \$10–25M						
Implementation shortfall	0.763¢	0.852¢	0.722¢	0.563¢	0.297¢	0.310¢
Number	5,693	6,651	7,814	10,009	12,267	11,730
Volume per string	\$14.7M	\$14.9M	\$14.7M	\$14.7M	\$15.0M	\$14.9M
Trades per string	34.1	37.7	46.2	53.3	50.9	53.6
Market share per string	1.0%	0.83%	0.67%	0.53%	0.40%	0.38%
Duration	24.6m	30.8m	29.4m	27.6m	24.4m	29.5m
Aggressiveness	39.5%	43.5%	42.0%	47.0%	45.3%	45.1%
Positions of \$25–100M						
Implementation shortfall	1.76¢	1.46¢	1.80¢	1.12¢	1.12¢	1.08¢
Number	2,389	2,767	2,850	4,008	5,241	4,918
Volume per string	\$42.0M	\$42.7M	\$41.9M	\$42.5M	\$42.5M	\$42.7M
Trades per string	98.4	106.6	126.7	157.1	147.6	154.0
Market share per string	2.7%	2.3%	1.8%	1.5%	1.1%	1.1%
Duration	58.8m	62.0m	60.0m	65.5m	53.4m	60.1m
Aggressiveness	39.5%	42.8%	40.7%	48.3%	45.8%	45.2%
Positions of \$100M+						
Implementation shortfall	6.88¢	4.14¢	4.90¢	3.26¢	2.41¢	2.09¢
Number	212	209	240	321	537	536
Volume per string	\$179.1M	\$149.5M	\$164.0M	\$154.2M	\$175.2M	\$172.5M
Trades per string	653.0	390.8	580.6	523.2	663.4	773.5
Market share per string	10.2%	7.3%	6.7%	5.2%	4.3%	4.1%
Duration	135.7m	121.3m	120.3m	91.2m	118.4m	152.1m
Aggressiveness	53.2%	45.9%	46.2%	43.4%	41.0%	48.5%

Table 4: Difference-in-differences event study on implementation shortfall, by volume bin

This table reports regression coefficients from a DiD event study using weekly average data from the CGB (the future on the Government of Canada 10-year bond) and the BAX (the future on the Canadian Banker’s Acceptance rate). The treatment group is the CGB, and the control group is the BAX. For the study on this table, the event window is the three months before and three months after the first month that there was an increase in the number of HFT trading the CGB, November 2011. The study reports treatment effects for trade strings grouped in four bins by volume. Panel B reports treatment effects for trade strings in the \$2–10 million bin for various sizes of event window.

Implementation shortfall is the difference between the volume-weighted average price of a string and the midquote contemporaneous to the first trade, in basis points.

Asterisks represent statistical significance at the 5%, 1% and 0.1% level.

Panel A: Effect on implementation shortfall, by volume bin

	Volume bin of trade string			
	\$2–10M	\$10–25M	\$25–100M	\$100M+
DiD	-26.86* (-2.10)	31.82 (1.24)	-72.55 (-1.39)	38.08 (0.21)
Constant	41.63*** (4.04)	66.68** (3.49)	176.40*** (4.31)	515.60*** (4.47)
Observations	27	27	27	26
R-squared	0.153	0.059	0.073	0.002

Panel B: Effect on shortfall for orders sized \$2–10M, by event window

	Size of event window				
	2 mo.	3 mo.	4 mo.	6 mo.	9 mo.
DiD	-51.04* (-2.72)	-54.96* (-2.61)	-51.19* (-2.65)	-58.91*** (-4.99)	-53.54*** (-5.47)
Volume	0.00 (0.64)	0.00 (1.07)	0.00 (1.92)	0.00** (3.20)	0.00* (2.52)
Volatility	-3.62 (-0.78)	-5.46 (-1.21)	-3.94 (-1.20)	-3.77** (-2.82)	-2.84 (-1.71)
Lagged level	-0.40 (-2.09)	-0.20 (-1.18)	-0.22 (-1.50)	-0.27* (-2.10)	-0.30** (-2.87)
Constant	78.23 (1.87)	76.70 (1.81)	64.71* (2.55)	70.62*** (5.97)	67.42*** (6.32)
Observations	16	24	32	50	76
R-squared	0.458	0.293	0.274	0.415	0.412

Table 5: Difference-in-differences event study on implementation shortfall, by duration bin

This table reports regression coefficients from a DiD event study using weekly average data from the CGB (the future on the Government of Canada 10-year bond) and the BAX (the future on the Canadian Banker's Acceptance rate). The treatment group is the CGB, and the control group is the BAX. For the study on this table, the event window is the three months before and three months after the first month that there was an increase in the number of HFT trading the CGB, November 2011. Panel A reports treatment effects for trade strings grouped in three bins by string duration. Panel B reports treatment effects for trade strings in the 0–10 minute bin for various sizes of event window.

Implementation shortfall is the difference between the volume-weighted average price of a string and the midquote contemporaneous to the first trade, in basis points.

Asterisks represent statistical significance at the 5%, 1% and 0.1% level.

Panel A: Effect on shortfall for trade strings, by duration of string

	Duration bin of trade string		
	0–10 min.	10–60 min.	60 min.+
DiD	-12.33* (-2.08)	-44.29 (-1.97)	52.55 (0.42)
Constant	50.66*** (15.02)	49.34** (3.14)	182.15 (1.69)
Observations	27	27	27
R-squared	0.144	0.134	0.007

Panel B: Effect on shortfall for strings of duration 0–10 min., by event window

	Size of event window				
	2 mo.	3 mo.	4 mo.	6 mo.	9 mo.
DiD	-15.19* (-2.63)	-15.40 (-1.91)	-15.62* (-2.06)	-17.70** (-3.18)	-14.32** (-3.25)
Volume	-0.00 (-0.61)	-0.00 (-0.57)	-0.00 (-0.20)	0.00 (1.06)	0.00 (0.73)
Volatility	-4.44* (-3.09)	-4.07 (-1.93)	-3.16 (-1.72)	-1.03 (-1.03)	-0.15 (-0.20)
Lagged level	0.38 (1.62)	0.18 (0.82)	0.21 (1.14)	0.17 (1.10)	0.16 (1.37)
Constant	70.65** (3.16)	73.65** (3.43)	61.10*** (4.23)	46.50*** (5.43)	41.04*** (6.48)
Observations	16	24	32	50	76
R-squared	0.399	0.280	0.309	0.347	0.276

Table 6: Difference-in-differences event study on effective spreads and price impacts

This table reports regression coefficients from five DiD event studies that each have different event windows. The event windows are the indicated number of months before and after the first month that there was an increase in the number of HFT trading the CGB, November 2011. As in Table 4, the studies use weekly average data from the CGB (the future on the Government of Canada 10-year bond) and the BAX (the future on the Canadian Banker's Acceptance rate). The treatment group is the CGB, and the control group is the BAX.

Effective spread is the signed difference between the trade price and contemporaneous midquote for trades in the strings, in basis points; *price impact* is the signed difference between the contemporaneous midquote and the midquote some time-interval after the trade for trades in the strings, in basis points.

Asterisks represent statistical significance at the 5%, 1% and 0.1% level.

Panel A: Effect on effective spreads and price impacts				
	Liquidity metric			
	Effective spread	5-second price impact	30-second price impact	5-minute price impact
DiD	-15.05*** (-7.09)	-4.83** (-3.39)	-5.85** (-3.05)	-4.78 (-1.52)
Constant	39.85*** (23.17)	22.21*** (24.81)	23.66*** (15.71)	21.67*** (9.26)
Observations	27	27	27	27
R-squared	0.672	0.312	0.273	0.085

Panel B: Effect on five-second price impact, by size of event window					
	Size of event window				
	2 mo.	3 mo.	4 mo.	6 mo.	9 mo.
DiD	-7.95** (-3.69)	-6.10** (-2.99)	-4.26* (-2.16)	-4.79** (-2.97)	-6.80** (-2.69)
Volume	0.00 (0.73)	-0.00 (-1.38)	-0.00 (-0.44)	0.00 (0.15)	0.00 (0.95)
Volatility	-0.56 (-1.44)	-0.45 (-0.85)	0.10 (0.17)	0.80* (2.35)	-0.22 (-0.37)
Lagged level	-0.71** (-3.33)	-0.27 (-1.24)	-0.07 (-0.41)	0.11 (1.09)	0.22 (1.96)
Constant	41.23*** (6.12)	33.11*** (4.53)	22.65*** (4.01)	14.18*** (5.56)	18.38** (3.36)
Observations	16	24	32	50	76
R-squared	0.541	0.399	0.283	0.473	0.328