

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Sagade, Satchit; Scharnowski, Stefan; Theissen, Erik; Westheide, Christian

## Working Paper A tale of two cities: Inter-market latency and fast-trader competition

SAFE Working Paper, No. 430

**Provided in Cooperation with:** Leibniz Institute for Financial Research SAFE

*Suggested Citation:* Sagade, Satchit; Scharnowski, Stefan; Theissen, Erik; Westheide, Christian (2024) : A tale of two cities: Inter-market latency and fast-trader competition, SAFE Working Paper, No. 430, Leibniz Institute for Financial Research SAFE, Frankfurt a. M., https://doi.org/10.2139/ssrn.3475442

This Version is available at: https://hdl.handle.net/10419/303051

#### 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.



## WWW.ECONSTOR.EU



Satchit Sagade | Stefan Scharnowski | Erik Theissen | Christian Westheide

# A Tale of Two Cities – Inter-Market Latency and Fast-Trader Competition

SAFE Working Paper No. 430 | July 2024

## Leibniz Institute for Financial Research SAFE Sustainable Architecture for Finance in Europe

info@safe-frankfurt.de | www.safe-frankfurt.de Electronic copy available at: https://ssrn.com/abstract=3475442

## A Tale of Two Cities – Inter-Market Latency and Fast-Trader Competition<sup>\*</sup>

# Satchit Sagade<sup>1,2</sup>, Stefan Scharnowski<sup>3</sup>, Erik Theissen<sup>3,4</sup>, Christian Westheide<sup>2,5</sup>

<sup>1</sup>Nasdaq
<sup>2</sup>Leibniz Institute for Financial Research SAFE
<sup>3</sup>University of Mannheim
<sup>4</sup>Centre for Financial Research, Cologne
<sup>5</sup>University of Vienna

This version: July 10, 2024

#### Abstract

We examine the impact of increasing competition among the fastest traders by analyzing a new low-latency microwave network connecting exchanges trading the same stocks. Using a difference-in-differences approach comparing German stocks with similar French stocks, we find improved market integration, faster incorporation of stock-specific information, and an increased contribution to price discovery by the smaller exchange. Liquidity worsens for large caps due to increased sniping but improves for mid caps due to fast liquidity provision. Trading volume on the smaller exchange declines across all stocks. We thus uncover nuanced effects of fast trader participation that depend on their prior involvement.

**Keywords:** Latency, Market Fragmentation, Arbitrage, Liquidity, Price Efficiency, High-Frequency Trading

JEL classification: G10, G14, G15

<sup>\*</sup>We thank Sergei Glebkin (discussant), Richard Payne (discussant), Michael Schneider (discussant), Andriy Shkilko (discussant), and conference and seminar seminar participants at the 1st Future of Financial Information conference, the 4th Sydney Market Microstructure Meeting, the Paris December Finance Meeting, Annual Meeting of the German Finance Association, the University of Mannheim, the University of Frankfurt, and the University of Vienna for helpful comments and discussions. All errors remain our own. The authors gratefully acknowledge financial support from the Research Center SAFE, funded by the State of Hessen research initiative LOEWE. The authors thank Equipex PLADIFES ANR-21-ESRE-0036 (France 2030) for providing the EUROFIDAI BEDOFIH database. An earlier version of this paper appeared as chapter 1 of Stefan Scharnowski's doctoral dissertation at the University of Mannheim. The analysis, interpretation and conclusion of this study reflect solely the views of the authors. *Email addresses:* satchit.sagade@nasdaq.com (Satchit Sagade), scharnowski@uni-mannheim.de (Stefan Scharnowski), theissen@uni-mannheim.de (Erik Theissen), christian.westheide@univie.ac.at (Christian Westheide)

## 1 Introduction

Trading in most of the world's largest equity markets is fragmented across multiple exchanges which often have their matching engines situated in different geographical locations. This geographical dispersion induces high frequency traders (HFTs) to invest in technology allowing the rapid transmission of market data between trading venues. This, in turn, makes inter-market latency a critical dimension of the competition among HFTs. This competition is distinct from intra-exchange competition and is characterized not only by speed differences. Rather, when traders use the fastest, wireless networks, they face capacity constraints inherent in this technology. Such constraints are particularly important when a large cross-section of instruments is traded, as is commonly the case in equity markets. While there is substantial research on the broader implications of market fragmentation and on speed competition within exchanges, the empirical consequences of inter-market latency competition, particularly in scenarios with a diverse range of traded instruments, have attracted much less attention.

This paper investigates the effects of varying inter-market latency among market participants by exploiting the deployment of the first commercially available microwave link between London and Frankfurt. Prior to this event, the vast majority of trading firms relied on comparatively slower fiber optic cables for data transmission between these two financial hubs. We refer to these traders as *slow* high frequency traders. However, some traders, referred to as *fast* high frequency traders, apparently operated privately-owned faster microwave connections and used them for market-making in the largest stocks. The introduction of the new link significantly broadened the set of market participants with access to the fastest mode of information transmission, thereby increasing the number of fast, relative to the number of slow, high frequency traders and thus altering the competitive landscape.

Employing difference-in-differences methodology, our analysis focuses on the effects of the newly established microwave link on traders' order submission behavior, trading activity, various aspects of market quality, and price discovery. Our treatment group consists of German stocks that are traded both on the Frankfurt-based Deutsche Börse's Xetra system (henceforth Xetra) and on Chi-X, located near London. The control group comprises French stocks traded on Euronext and Chi-X, which were, at the time, located near London.<sup>1</sup> French stocks are a particularly well-suited control group for several reasons. First, and most importantly, the communication speed between Euronext and Chi-X is not affected by the microwave link between Frankfurt and London. Additionally, the French and German equity markets are roughly equally-sized, operate similar trading systems, and stocks from both countries are actively traded on Chi-X. Our sample consists of a broad cross-section of large-cap and mid-cap firms from both countries. The large spread in the size of our sample firms allows us to shed light on the potentially heterogeneous impact of cross-market communication latencies on stocks of different size.

There is no existing theoretical model that directly parallels the setting of our empirical analysis. However, the model of Baldauf and Mollner (2020) comes closest. It features a multi-market setting with liquidity traders, information traders, high-frequency market

<sup>&</sup>lt;sup>1</sup>In November 2010, Euronext relocated its exchange matching engine from Aubervilliers, near Paris, to Basildon, near London.

makers and snipers. This model serves as the foundation from which we derive hypotheses on the effect of the introduction of the microwave network on liquidity. Besides liquidity we consider several other important outcome variables, but do not derive formal hypotheses from the Baldauf and Mollner (2020) model due to the differences between the model setup and the institutional setting we analyze.

We show that, following the introduction of the microwave link, integration of pricing between the two exchanges increased, as evidenced by a reduction in arbitrage opportunities. Specifically, we find, for the whole cross-section of stocks, a decline in the frequency, profitability, and persistence of arbitrage opportunities. These findings demonstrate that the new network enabled faster information transmission between the trading venues and increased the speed of adjustment of prices on one venue to information generated on the other.

In the second part of our analyses, we focus on trading and quoting activity. We find a negative effect of the introduction of the microwave network on trading volume on Chi-X, especially for large-cap stocks. A possible explanation is that traders without access to the microwave network cannot adjust their quotes on Chi-X fast enough to information generated in Frankfurt and therefore relocate orders from Chi-X to Xetra.<sup>2</sup> The stronger effect for large caps is consistent with the notion that users of the new microwave connection primarily engage in aggressive, opportunistic trading in these stocks,

<sup>&</sup>lt;sup>2</sup>New information on German stocks is predominantly produced in Germany. Consequently, a fast data connection between Frankfurt and London is clearly beneficial to adjust quotes in the London-based Chi-X to new information but is not needed to the same degree to adjust quotes in the Frankfurt-based Xetra system. As a result, the introduction of the microwave network puts slower market participants at a larger disadvantage on Chi-X than on Xetra.

whereas for mid caps, microwave-using liquidity providers become active for the first time.

We also find that the introduction of the microwave network affects the amount of order message traffic. For both mid and large caps, we find large negative effects on the number of messages affecting the best bid-offer (BBO) on Chi-X. These results are consistent with the hypothesized decrease in the number of competitive limit order traders that is attributable to the first-time entry of the fastest liquidity providers in the mid-cap segment, and increased fast aggressive activity among large caps, respectively. At the same time, there are positive, albeit smaller and statistically weaker, effects on the message traffic at the BBO on Xetra, possibly due to the migration of some participants away from Chi-X. For large caps, the negative effect on the number of messages also extends to orders away from the BBO, albeit to a lesser degree. This effect may be due to the increased exposure to snipers, as orders behind the BBO for both groups of stocks, again suggesting a shift of participants' activity towards the primary venue.

Enhanced capacity for fast information transmission between trading venues should also improve the responsiveness of prices to stock-specific information relative to market-wide information. We test this hypothesis by regressing price changes on Chi-X on contemporaneous and past price changes on Xetra and, for the market-wide information, price changes of futures traded on Eurex, which is also located in Frankfurt. The analysis reveals that, subsequent to the introduction of the microwave link, both mid-cap and large-cap stock prices exhibit a more pronounced reaction to stock-specific information and a relatively subdued response to market-wide information. Moreover, the time taken to react to stock-specific information decreases noticeably with the availability of the new network.

Besides the speed and nature of information transfer, we also consider its directionality. We employ high-resolution information shares (Hasbrouck, 2021) to the order book mid-point prices at a millisecond resolution. We find no significant changes for large caps, with Xetra maintaining an information share of about 65 percent. However, the apparent first-time entry of the fastest liquidity providers does appear to enhance the role of Chi-X in the price discovery process of mid-cap stocks. Although Xetra's information remains dominant and higher than that for large-cap stocks, we estimate a roughly 10 percentage point drop in its information share as the result of the new microwave link.

The effects on liquidity are in line with the predictions we derive from the Baldauf and Mollner (2020) model. For mid-cap stocks, we observe large reductions in quoted and effective spreads on Chi-X. Price impacts drop by about the same amount, such that there is no statistically significant effect on realized spreads. These results are consistent with the first-time entry of fast liquidity providers with access to a microwave network, which helps reduce adverse selection. The depth at the top of the order book decreases, presumably because one or a small set of fast liquidity providers replaces a mix of liquidity providers and other traders. Apart from depth, which appears unaffected, we find qualitatively similar but smaller effects on Xetra. These results suggest that liquidity providers are able to price in information from Chi-X faster than before, which contributes to a reduction in adverse selection.

For large-cap stocks, the results are consistent with an increased activity of snipers

on Chi-X. Though there is no statistically significant effect on spreads or depth, price impacts increase and, correspondingly, realized spreads decrease. The lack of an effect on spreads indicates that liquidity providers were earning rents (most likely caused by tick-size constraints) in the pre-event period. On Xetra, individually statistically insignificant effects on spreads and price impact add up to a significant reduction in realized spreads. This finding can also be explained by increasing sniping activity, even though the effect is less pronounced than on Chi-X due to a smaller amount of information originating in London as compared to Frankfurt.

Finally, we assess effects on volatility and price efficiency, the latter proxied by return autocorrelations and variance ratios. For mid-cap stocks, volatility is reduced and price efficiency improved on Chi-X, with no significant effects on Xetra. There is no significant effect on large-cap volatility, though the result for autocorrelation on Xetra implies a slightly decreased price efficiency.

Our interpretation of the results is based on three premises. First, the introduction of the new microwave link increased the number and activity of fast high-frequency traders. Second, due to the inherent capacity limits of microwave data transfer, the privately owned networks existing before October 2012 were predominantly used to trade large-cap stocks. Third, the additional fast high-frequency traders that became active after October 2012 focused on sniping in large caps and market making in mid caps. Although we cannot formally prove these claims, we believe they are plausible and supported by our empirical findings.

Regarding our first claim, it is conceivable that the new microwave network essentially

replaced the prior privately owned networks, thus not effectively increasing the number of very fast traders. However, in this case, we would not expect to see any first-order effect on the outcome variables under investigation with the introduction of the new network.

Our second claim is that before October 2012, the fastest traders focused on large stocks. The capacity of pre-existing networks was likely insufficient for a trader to be simultaneously active in a large number of stocks (such as the 60 stocks in our sample). Therefore, traders had to focus on small subsets of instruments and most likely chose actively traded stocks. Furthermore, the first faster trader who chooses to become active in a stock is likely to act as a liquidity provider. This choice is driven by the opportunity to earn the spread and avoid adverse selection when trading against slower traders. In contrast, sniping opportunities would be limited in a market where slower traders post wide spreads.

When the new network became available, subscribers needed to decide how to allocate their newly acquired data transfer capacity. Because market making in large caps was already dominated by incumbents with access to private networks, sniping in large caps was one option. Alternatively, traders could move towards less liquid mid-cap stocks with a focus on market making. The fact that the new network had 12 subscribers does not imply there were 12 competitors in each stock. Rather, due to capacity limits, each subscriber had to focus on a subset of stocks.

Our findings that (1) spreads for mid-caps decrease on Chi-X while (2) quoting activity at the best bid/ask decreases is consistent with the notion that a few fast traders with a cost advantage replaced a larger number of slower liquidity providers. Simultaneously, the increase of the price impact for large caps aligns with our conjecture of increased sniping activity in these stocks. While these results do not prove our premises, we see no alternative explanation consistent with our findings.

Our paper adds to the literature on trading speed and speed differences between traders. Most theoretical contributions to this literature (such as Aït-Sahalia and Sağlam, 2024, Bernales, 2019, Biais et al., 2015, Bongaerts and Van Achter, 2021, Budish et al., 2015, Foucault et al., 2016, Hoffmann, 2014 and Menkveld and Zoican, 2017) consider a single market setting. Only few papers consider multiple market settings. Some of them, such as Baldauf and Mollner (2021) and Pagnotta and Philippon (2018), focus on competition between trading venues, a topic that is beyond the scope of our paper. Others focus on the role and impact of cross-market arbitrage. While arbitrageurs are traditionally considered as contrarian traders (Gromb and Vayanos, 2002, 2010) who positively affect market outcomes by enforcing the law of one price and/or providing liquidity, there may be a dark side to arbitrage. Foucault et al. (2017) model the presence of toxic arbitrage opportunities that emerge when market makers adjust their quotes to new information with a delay. Such arbitrage opportunities arise when market makers compete with faster traders, and they harm liquidity because market makers price the resulting adverse selection costs into the bid-ask spread. Kozhan and Tham (2012) highlight the presence of execution risk in high-frequency arbitrage strategies due to a crowding out effect associated with competing arbitrageurs. Baldauf and Mollner (2020) analyze the interaction of two major HFT strategies, market making and sniping, in a multi-market setting. We describe their model in more detail in Section 3.

Most empirical studies that investigate the effects on market quality of increases in trading speed do so in a single-market setting. The speed increases under investigation are caused by either the introduction of and/or upgrade to colocation services, or by general speed improvements due to system upgrades. While the majority of papers, including Brogaard et al., 2015, Boehmer et al., 2021, Frino et al., 2014 and Hollifield et al., 2017, document improvements in liquidity, Ye et al. (2013) obtain the opposite result and Aquilina et al. (2022) quantify adverse selection costs due to snipers picking off stale orders.

The empirical literature on speed differences in a multiple-markets setting is scant.<sup>3</sup> Rzayev et al. (2023) study the effects of latency on market quality in a two-markets setting. They focus on toxic versus non-toxic arbitrage opportunities and find that (consistent with the predictions made by Foucault et al., 2017) inter-market latency reductions are associated with a deterioration of liquidity after the former and with liquidity improvements after the latter.

Shkilko and Sokolov (2020) investigate the impact on liquidity and transaction costs of precipitation-induced changes in communication latencies between index futures traded in Chicago and ETFs traded in New York. Precipitation disrupts microwave networks, forces high frequency traders (HFTs) to temporarily downgrade to slower fiber optic networks, and thus effectively eliminates the speed dispersion across HFT firms for short periods

<sup>&</sup>lt;sup>3</sup>An early predecessor of the literature on trading and communication speed is a paper by Garbade and Silber (1978). They study the effects of two technological innovations from the 19th century, the introduction of a domestic telegraph system in the US and that of the first Trans-Atlantic cable between London and New York. They observe increased integration (proxied by inter-market price differentials) between geographically dispersed markets (such as regional stock exchanges in the US) around the introduction of these two technologies.

of time. In addition, Shkilko and Sokolov (2020) analyze an event<sup>4</sup> that resulted in a permanent reduction of speed differences. In both settings the authors find, consistent with theoretical models of speed-induced adverse selection such as Foucault et al. (2017), improvements in liquidity.

Our empirical setting is distinct from that in Shkilko and Sokolov (2020) in important respects. First, we consider a setting in which HFTs operate in two parallel markets for the *same* securities whereas Shkilko and Sokolov (2020) consider index futures contracts and associated ETFs. Second, we analyze a broad cross-section of stocks spanning a wide range in liquidity rather than focusing on a set of highly liquid instruments. Third (as we show in Section 5) price discovery in the stocks we analyze is bi-directional whereas, as recently re-confirmed by Dobrev and Schaumburg (2023), stock index futures lead ETFs, making price discovery in the setting analyzed by Shkilko and Sokolov (2020) largely uni-directional. Fourth, we analyze the "switching on" of an additional microwave link that considerably increased the number of fast high frequency traders, an event that is inherently different from the precipitation-induced switching on and off of *all* fast high frequency trading analyzed by Shkilko and Sokolov (2020). Given our emphasis on the variation in the extent of microwave network utilization, exogenous shocks resulting from precipitation are unsuitable for our setting.<sup>5</sup>

These differences in the empirical setting allow us to contribute to the literature in

<sup>&</sup>lt;sup>4</sup>In December 2012 Quincy data started to sell information transmitted through a microwave network to subscribers (rather than selling bandwidth on the network), thereby making the high-speed information available to a large number of traders.

<sup>&</sup>lt;sup>5</sup>Moreover, we note that granular time-stamped data on precipitation for the area between Frankfurt and London is not available because the area covers three or four countries (depending on the exact route), not all of which provide usable data, as well as parts of the North Sea.

multiple ways. Analyzing identical securities traded in different markets enables us to study the implications of trading speed on market integration and the competitive dynamics between the two trading venues under investigation. Considering a broad cross-section of stocks is potentially important because there is empirical evidence suggesting that the beneficial effects of high frequency trading are more pronounced for large than for small firms (Boehmer et al., 2021). Finally, by studying a setting in which price discovery is bi-directional we can analyze whether the introduction of the microwave connection affects the contributions to price discovery of the trading venues we analyze.

Our findings collectively highlight the complex consequences of changing the inter-market latency distribution of trading firms in a fragmented and geographically dispersed market. The impact of these decisions – owing to the physical limitations that restrict the bandwidth of microwave connections – vary in the cross-section of financial instruments. Ultimately, the welfare implications of augmenting or restricting the capabilities and capacities of the fastest market participants hinges on a careful consideration of the trade-offs across multiple dimensions of outcomes. Our results are also relevant for trading venues when choosing the location of their exchange data centers and the different connectivity options in a geographically dispersed market. For example, CBOE and Euronext relocated their data center in 2012 and 2022, respectively.<sup>6</sup> Finally, our results also imply that the distribution of intermarket latencies across traders is a

<sup>&</sup>lt;sup>6</sup>CBOE in 2012 moved its main option exchange from Chicago closer to the stock exchanges in New Jersey with the goal of reducing latency. In contrast, market participants raised concerns about latency arbitrage due to Euronext's decision in 2022 to move its matching engine from Basildon near London (close to its competitors' data centers) to Bergamo, Italy. Finally, trading venues and HFT firms also actively choose the location of microwave towers in and around exchanges' data centers to optimize access speeds.

relevant aspect that trading venues should consider when defining their trading rules, fees, and other parameters.

## 2 Institutional Background

Two institutional factors – automation and regulation – have led to the evolution of HFT in modern equity markets. Advances in computational power and innovations in networking technologies combined with decreasing computing costs accelerated the shift towards electronic trading. Concurrently, regulatory initiatives such as the Regulation National Market System (RegNMS) in the US and the Markets in Financial Instruments Directive (MiFID) in the EU led to the emergence of new trading venues. By 2012, the US and EU markets had seen a significant increase in the number of trading venues, with nine public limit-order markets in the US and three in the EU competing with the primary listing exchanges for individual stock order flows.<sup>7</sup> These new venues, by investing in enhanced trading speeds and offering lower transaction fees, have become attractive to HFT firms. These firms, in turn, provide liquidity, enabling these venues to maintain competitiveness. HFT firms additionally ensure that prices in fragmented markets remain arbitrage-free by swiftly eliminating discrepancies arising from local demand and supply imbalances.

HFT and technology firms also invest in improving (both in terms of transmission speed and reliability) the connections between different financial centers. In Europe, this

<sup>&</sup>lt;sup>7</sup>In addition, several dozen dark pools, internalization platforms, and other off-exchange venues also compete with these exchanges.

involves connecting the cities hosting different exchanges such as London Stock Exchange (London, UK), Euronext (Basildon, UK)<sup>8</sup>, Xetra (Frankfurt, Germany), and Chi-X (Slough, UK) (see Figure 1) by setting up or improving existing fiber optic connections and, increasingly, installing microwave networks. Microwave networks are faster than fiber optic cables: transmission of data through the air happens at nearly the speed of light whereas transmission through cables, because the signal bounces off the walls of the cables, is only about two thirds as fast.

#### [Insert Figure 1 about here]

In October 2012, Perseus Telecom – a provider of low-latency networks – launched operations of the first *commercially available* microwave network between London and Frankfurt. Before that, there were a small number of preexisting microwave networks owned by HFT firms linking the two cities.<sup>9</sup> The new network enabled a significant reduction in round-trip latency, decreasing it by nearly half from 8.3 to "less than 4.6" milliseconds.<sup>10</sup> This performance compares to a theoretical round-trip latency of 4 milliseconds based on transmission at the speed of light. Perseus, in its press release and in a later interview with its CEO, announced that: (i) the initial capacity had been

<sup>&</sup>lt;sup>8</sup>This information pertains to our sample period. In 2022, Euronext moved to a data center in Bergamo, Italy.

<sup>&</sup>lt;sup>9</sup>Alexandre Laumonier, an anthropologist who studied the evolution of the usage of microwave networks by HFTs, reports at https://sniperinmahwah.wordpress.com/2014/10/02/hft-in-my-backyard-iii/: "Here is Vigilant Global, a Montreal-based prop trading firm. Vigilant is now a subsidiary of DRW. [...] Another [competitor] said: 'I actually admire Vigilant/DRW. They were first.[...]' Yes, Vigilant were the first to build a network in Europe. They probably initiated their network 2011, perhaps even in 2010.", and, referring to a planning application made in September 2012, "Once Jump [...] apologized: 'The two largest dishes proposed for World Class Wireless [Jump Trading's subsidiary making the planning application] were actually installed in advance of the submission of the planning applications.'

<sup>&</sup>lt;sup>10</sup>Shkilko and Sokolov (2020) report similar improvements in data transmission speeds resulting from upgrades to networks linking Chicago and New York.

quickly sold out to approximately twelve subscribers each utilizing identical capacities, and (ii) the network had been running at 99% uptime. This network was a component of a larger super-network owned by Perseus connecting Xetra and Eurex in Frankfurt to London Stock Exchange, BATS, Chi-X, Euronext, and Liffe in and around London.<sup>11</sup> In the interview, the CEO also referred to bandwidth constraints as a limitation of microwave as compared to fiber networks: "The capacity of of fiber networks is at 1000 times that of a wireless network. [...] So you have to think very carefully of what signals you are going to send between markets on your wireless network." He explained that "we can serve 12 customers with a 10 megabits per second slice each." This capacity corresponds to about one kilobyte per millisecond. Thus, the new network substantially increased the number of microwave users, though each trading firms' bandwidth was still contrained, such that they needed to make strategic choices as to how to use the newly available network.

The period around the introduction of this microwave network thus provides a unique opportunity to empirically evaluate the effects of enhanced speed available to more market participants on market outcomes in fragmented markets.

### **3** Theoretical Background and Hypotheses

To the best of our knowledge, there is no theoretical model that fully captures our empirical setting. The one that comes closest is Baldauf and Mollner (2020). In this

<sup>&</sup>lt;sup>11</sup>See https://intelligenttradingtechnology.com/hardware/blog/perseus-beats-rivals-live-london-frankfurt -microwave-service (saved to Dropbox) and https://prnewswire.com/news-releases/perseus-telecom-announ ces-fastest-available-microwave-service-in-european-financial-markets-193012631.html for the press release concerning Perseus' microwave network and see http://www.marketswiki.com/wiki/Jock\_Percy for the interview with the CEO at the time.

section we provide a brief sketch of their model and then "translate" it to our setting in order to derive testable implications. In the model, a single security is traded on several trading venues, organized as (identical) anonymous limit order books. The value of the security is either +1 or -1 and is initially unknown to all agents. A single trader who may either be a liquidity trader (with probability  $1 - \lambda$ ) or an information trader (with probability  $\lambda$ ) arrives and trades one share. A liquidity trader is equally likely to buy or sell and can only trade at her home exchange (determined randomly). An information investor can send orders to all exchanges. Liquidity and information traders can only use immediate-or-cancel (IOC) orders, i.e., limit orders that are executed on arrival if marketable but are deleted (rather than added to the book) if not marketable.

The model features two types of high frequency traders, high frequency market makers (HFT-MM) and an infinite number of high frequency snipers. The snipers can only submit IOC orders, whereas the HFT-MM can only submit post-only orders. These are limit orders that are canceled if they are marketable. Bertrand competition between the HFT-MMs ensures that they earn zero profit in expectation when at least two HFT-MMs are present.

In the equilibrium of the model, liquidity traders, information traders, HFT-MM and snipers use the following strategies. A liquidity trader simply buys or sells one share at her home exchange. An information trader does nothing if she does not learn the asset value. If she learns the asset value she sends an IOC order to buy [sell] one unit of the asset with a price limit of +1 [-1] to all exchanges. There is one "active" HFT-MM who places limit buy and sell orders in the book.<sup>12</sup> The equilibrium bid and ask prices are such that the HFT-MM is exposed to adverse selection risk.

If a trade occurs at one (or more) exchanges, the HFT-MM will immediately send out cancellation messages for her exposed orders, which are now at risk of being unprofitable. This action is predicated on the understanding that a second trade implies the activity of an information trader who has learned the asset value. If trades occur on *more* than one exchange, HFT-MM and snipers conclude that an information trader traded. They learn the asset value based on the direction of the trades. Snipers react immediately by submitting IOC orders to trade against the now stale HFT-MM orders while the HFT-MM sends cancellation messages.

Time in the model progresses in very small increments, denoted as  $\epsilon$ . The latency of an order (or a cancellation), defined as the time between submission and processing of the message, is either  $1\epsilon$  or  $3\epsilon$ . The latency durations are random and governed by exogenous probabilities, which may be different for traders and HFTs. In instances where two or more orders arrive simultaneously, ties are broken randomly. To illustrate how the model operates, let us consider three cases.

**Case 1.** A liquidity trader arrives and executes a trade for one unit of the asset at her home exchange. Depending on the latency, that trade executes at time  $1\epsilon$  or at time  $3\epsilon$ . Upon observing the trade, the active HFT-MM cancels all her outstanding quotes. The cancellations, subject to their own latency, are processed either at  $2\epsilon$  or  $4\epsilon$  if the initial

<sup>&</sup>lt;sup>12</sup>Note that, if one HFT-MM has standing quotes priced at the break-even bid and ask prices in the book, other HFT-MMs have no incentive to add depth because only information traders and snipers would be interested in trading a second unit, resulting in a loss to the HFT-MM.

trade was at  $1\epsilon$ , or at  $4\epsilon$  or  $6\epsilon$  if it was at  $3\epsilon$ . In this scenario, the active HFT-MM earns a profit in expectation (the half-spread) and nothing else happens.

**Case 2.** An information trader arrives, learns the asset value, sends orders to all exchanges, but only one of her trades executes at time  $1\epsilon$ . Upon observing the trade at time  $1\epsilon$  the HFT-MM cancels all her outstanding orders. Meanwhile, the snipers remain inactive, as at time  $1\epsilon$  cases 1 and 2 are indistinguishable to them. Some of the HFT-MM's order cancellations get processed at time  $2\epsilon$ . At time  $3\epsilon$  the remaining orders submitted by the information trader execute at those exchanges where the HFT-MM's cancellation messages were not processed at time  $2\epsilon$ . The snipers observe these trades, infer the information trader's information, and submit orders to those exchanges where the HFT-MM still has standing quotes. The snipers' orders reach the exchanges at either  $4\epsilon$  or  $6\epsilon$ . The orders arriving at time  $6\epsilon$  are canceled upon arrival because, at that time, the book is empty. However, those arriving at time  $4\epsilon$  arrive simultaneously with the HFT-MM's cancellation messages. Given that (a) ties are broken randomly and that (b) there is an infinite number of snipers, one sniper order is processed before the HFT-MM's cancellation message with probability 1. Thus, the sniper earns a profit while the HFT-MM incurs a loss.

**Case 3.** An information trader arrives, learns the asset value, sends orders to all exchanges and more than one of her trades executes at time  $1\epsilon$ . Upon observing these trades at time  $1\epsilon$ , the HFT-MM cancels all her outstanding orders, wheras the snipers submit orders to exploit the now-stale quotes. Again, because of the infinite number of snipers, there will be sniper orders at every exchange at time  $2\epsilon$ , and even when the HFT-MM's cancellation message also arrives at time  $2\epsilon$ , a sniper's order will be processed first with probability 1. The HFT-MM earns a profit in 1. In 2 she suffers a loss on a subset of exchanges while in 3 she will suffer a loss on all exchanges.

It is evident that the Baldauf and Mollner (2020) model does not perfectly align with our empirical setting. Unlike the model's assumption of numerous identical markets, our analysis specifically focuses on two distinct markets, Xetra and Chi-X. These markets differ significantly, with Xetra being the primary market for our sample stocks and likely the first to reflect firm-specific fundamental information in its prices. We expect, and later provide empirical evidence, that Xetra contributes more significantly to price discovery and attracts a higher proportion of captive liquidity traders, thus being more liquid compared to Chi-X. Even though we focus on trading on Xetra and Chi-X, traders receive information from other markets (such as markets for ETFs, index and single stock futures, and equity options) as well. Trades in these markets are potentially informative on the presence of information traders. The fact that these additional markets are located in Frankfurt supports the argument that new information tends to be impounded into prices on Xetra first.

Further, Baldauf and Mollner (2020) assume that there are at least two HFT-MMs and an infinite number of snipers. In our empirical setting there are "slow" traders, traditional HFTs (those without access to a microwave connection) and fast HFTs (those with access to a microwave connection). Note that this distinctions matters most to trading on Chi-X, whereas, due to the fact that most information originates in Frankfurt, the set of competitive HFTs on Xetra is larger even before the introduction of the new microwave connection. In the model of Baldauf and Mollner (2020) all markets are identical, meaning that the probability of observing a trade is the same on all markets. In contrast, in our setting trades are more likely to occur on Xetra because there are more liquidity traders in Xetra than on Chi-X, and because new information is predominantly generated in Frankfurt. Thus, the reduction in inter-market latency is likely to be more important for Chi-X and, consequently, our empirical predictions focus on effects on Chi-X. However, there may also be repercussions on Xetra because some information is likely to be generated in London as well. Traders using the microwave link to trade on that information will impose additional adverse selection risk on other traders on Xetra.

We focus on a change in the number of fast HFTs and implicitly lump together slow traders and traditional HFTs. As noted previously, there is evidence that there had been fast HFTs (using a privately owned microwave link) already before the commercial microwave link became available. Therefore, the introduction of the commercial microwave connection increases the number of fast HFTs from a small number to a larger number, albeit still far from infinite. We acknowledge that Baldauf and Mollner (2020) do not model a finite number of snipers, but we posit that the qualitative implications of having fewer snipers are intuitive, and we will use this sort of intuitive reasoning below.

Microwave links, constrained by their limited capacity, necessitate a strategic selection from their users regarding the information and messages they prioritize for transmission. Initially, when the number of fast HFTs is small, it is likely that these HFTs primarily focus on trading large stocks. This preference stems from the greater profit opportunities typically associated with actively traded instruments. However, as the market evolves with the entrance of additional fast HFTs, these new entrants need to choose which stocks to prioritize, and to determine their role as either market makers or snipers. We will analyze the implications of these choices, discussing the dynamics related to large and mid-cap stocks in detail to understand how these decisions affect market behavior and trading strategies.

In the market for large-cap stocks, the presence of incumbent fast HFTs likely leads to a scenario where the potential profits from market making are largely exploited. According to the model, the presence of merely two HFT-MMs is sufficient to reach the competitive level of spreads. Consequently, additional fast HFTs entering the market have little incentive to engage in market making, tending instead to assume the role of snipers. In doing so, they impose additional adverse selection costs on the HFT-MMs, resulting in larger price impacts. Typically, HFT-MMs would respond to this by widening their quoted spreads. However, in real-world markets, the existence of a discrete minimum tick size complicates this response. This tick size constraint allows HFT-MMs to maintain profitability even when spreads are at their competitive minimum, determined by the tick size. Therefore, the increased adverse selection costs imposed by additional snipers do not necessarily lead to wider quoted spreads but may lead to a reduction in realized spreads instead. Note that the discussions above refers to quoted spreads, though neither the model nor specifics of our institutional setting suggest distinct effects on effective as opposed to quoted spreads, such that we arrive at identical predictions for these two measures of market liquidity.

These considerations lead to the following hypotheses.

#### Hypothesis 1: Effects of the microwave link on large-cap liquidity

21

#### **a.** Price impacts increase on Chi-X.

c. Similar effects to those in a. and b. exist, in smaller magnitude, on Xetra.

Given the strong competitive pressure in the markets for large caps, some of the additional fast HFTs will find it worthwhile to specialize in mid-cap stocks. Before the introduction of the commercial microwave link, fast HFTs are unlikely to have been active in these markets on Chi-X, presenting new profit opportunities to fast HFT market makers (HFT-MMs). Owing to their speed advantage, these fast HFT-MMs will face less adverse selection compared to incumbent liquidity providers, allowing them to offer tighter spreads. This reduced adverse selection risk also suggests a decrease in price impacts. However, the presence of a discrete minimum tick size means that these lower adverse selection costs do not necessarily lead to reduced spreads but could actually increase realized spreads. While the entry of additional snipers may undo this effect to some extent, it is unlikely to completely offset it. The limited capacity of the commercial microwave link means that only a finite number of snipers will find it advantageous to target mid-cap stocks. Recall that, in the Baldauf and Mollner (2020) model, it takes just two HFT-MMs to reach competitive spread levels, but an infinite number of snipers are required to exploit all profit opportunities. With fewer snipers in the real world, HFT-MMs face less adverse selection risk and can consequently offer lower spreads and / or earn increased realized spreads) at unchanged quoted and effective spread levels. Because the bulk of information on German stocks originates in Germany, effects on liquidity in Xetra are expected to be

**b.** Quoted and effective spreads increase and/or the realized spreads decrease on Chi-X.

weaker but point in the same direction.

These considerations lead to the following hypotheses.

#### Hypothesis 2: Effects of the microwave link on mid-cap liquidity

a. Price impacts decrease on Chi-X.

**b.** Quoted and effective spreads decrease and/or the realized spreads increase on Chi-X.

c. Similar effects to those in a. and b. exist, in smaller magnitude, on Xetra.

### 4 Data, Descriptives, and Empirical Methodology

#### 4.1 Data and Sample Selection

Our sample comprises German stocks with a primary listing on Xetra, serving as our treatment group, and French stocks with a primary listing on Euronext as control group. The German sample contains the 30 constituents of the DAX index and the 30 largest stocks of the MDAX index, whereas the French sample contains all the constituents of the CAC 40 and CAC Next 20 indices. We use the index composition as of June 30, 2012. For each stock, we obtain data from the respective primary listing exchange – Deutsche Börse Xetra or Euronext Paris – and from Chi-X. We focus on the four months around the introduction of the microwave link, spanning August to December 2012. The precise date when the link became operational is not known, and it is unlikely that all twelve trading firms began using it simultaneously. Therefore, to account for a potential staggered adoption period, we exclude the month of October 2012 from our analysis.

We obtain data from Eurofidai's high-frequency database (BEDOFIH). For Euronext Paris and Chi-X, the database contains messages of both the orders entering the order book and the completed transactions. For Xetra, it contains orderbook snapshots at irregular intervals, updates to the orderbooks in between these snapshots, and transaction records. The data resolution is at the microsecond-level for Euronext Paris and at the millisecond-level for Chi-X. As for Xetra, the resolution is at the millisecond level until 23 November 2012 and at the microsecond-level thereafter.<sup>13</sup> To maintain consistency and minimize resolution-related errors, we aggregate all data recorded at the microsecond level to millisecond resolution. We supplement the BEDOFIH dataset with data on market capitalization and trading volume on other trading venues (BATS and Turquoise, both operating electronic open limit order books where all our sample stocks are traded) from Thomson Reuters Eikon.

For each exchange, we reconstruct the orderbooks using the quote updates, which reflect changes due to additions, cancellations, or executions of orders. This approach allows us to capture the precise state of the orderbook at any given moment during the trading session. We focus on trading between 09:05 CET and 17:25 CET, thus excluding the opening and closing auctions as well as the first and last five minutes of the continuous trading session. We also exclude the period associated with the intraday call auction on Xetra for all treatment and control firms. We drop stocks with missing data or errors in the message feed. Our final sample contains 111 stocks, among which 56 are German and 55 are French stocks.

<sup>&</sup>lt;sup>13</sup>This change in resolution does not affect the frequency of updates to the real time data feeds relied upon by trading firms.

#### 4.2 Variable Definitions

In this subsection, we define the specific metrics used for our analysis. For each stock and day, we compute these measures separately for the primary exchange (Xetra or Euronext) and Chi-X. We winsorize the variables at the 99% level.

We measure trading and quoting activity by computing several variables: trading volume, median trade size, number of trades, and message frequency. Trading volume and trade size are measured in log euros, the number of trades is the log count of all executions, and message frequency is defined as the log sum of all new limit orders and cancellations. We split the message frequency into messages that affect the best bid or ask and those that only affect the book at deeper levels.

As measures of quoting liquidity we compute time-weighted averages of relative quoted spreads and top-of-book depth.<sup>14</sup> Specifically, denote the best bid and ask prices and corresponding quantities as  $P^{bid}$ ,  $P^{ask}$ ,  $Q^{bid}$ , and  $Q^{ask}$  and let  $M = \frac{P^{ask} + P^{bid}}{2}$ . Our variables are then defined as:

$$Quoted = \frac{P^{ask} - P^{bid}}{M}$$
$$Depth_{top} = \ln \left( P^{bid} Q^{bid} + P^{ask} Q^{ask} \right)$$

We measure trading liquidity using the effective spread, realized spread, and price impact. The effective spread quantifies the actual transaction cost paid by the trader submitting a marketable order; the realized spread is the compensation earned by the

<sup>&</sup>lt;sup>14</sup>Our results remain qualitatively unchanged if we use depth within 50bps of the midquote instead.

trader whose limit order is executed, after adjusting for any adverse selection cost; and the price impact measures the information content of the trade. Specifically, for each trade at time t:

$$Effective = \frac{2D_t(P_t - M_t)}{M_t}$$
$$Realized_{\Delta} = \frac{2D_t(P_t - M_{t+\Delta})}{M_t}$$
$$Price Impact_{\Delta} = \frac{2D_t(M_{t+\Delta} - M_t)}{M_t}$$

where  $P_t$  is the transaction price,  $M_t$  is the midquote just before the trade,  $M_{t+\Delta}$  is the midquote  $\Delta$  seconds after the trade, and  $D_t$  is +1 (-1) if the trade is initiated by the buyer (seller). For each stock and day, we compute trade-size weighted averages of these values for  $\Delta = 10$  seconds.<sup>15</sup>

We define intraday volatility as the standard deviation of log quote midpoint returns measured at one second intervals. As measures of short-term price efficiency, we employ the absolute values of return autocorrelation and variance ratios, also based on quote midpoints. The price efficiency measures compute deviations of prices from a random walk benchmark. As in Boehmer and Kelley (2009), we take absolute values of these variables in order to capture deviations on either side of the random walk benchmark to ensure that our measures reflect the magnitude of inefficiencies irrespective of their direction. Specifically, first denote the log return over a time interval  $\Delta$  as  $r_t^{\Delta} = ln(M_t) - ln(M_{t-\Delta})$ .

<sup>&</sup>lt;sup>15</sup>Our results remain qualitatively unchanged if we use  $\Delta = 1$  second instead.

Our price efficiency measures are given by

Abs. 
$$AC_{\Delta} = \left| Corr(r_t^{\Delta}, r_{t-\Delta}^{\Delta}) \right|$$
  
Abs.  $VR_{n,m} = \left| 1 - \frac{n Var(r_t^m)}{m Var(r_t^n)} \right|$ 

We compute autocorrelations for  $\Delta = 10$  seconds, and variance ratio for (n, m) = (10, 30)seconds.<sup>16</sup>

Finally, we also measure cross-market arbitrage activity and information shares. The corresponding variables are defined in Section 5.1 and Section 5.4, respectively.

#### 4.3 Descriptive Statistics

We present descriptive statistics, based on average values for July 2012, in Table 1. Building on the reasoning presented in Section 3, we perform our analyses separately for large-cap and mid-cap firms. Within each country, we split our sample based on the firms' average market capitalization during July 2012.<sup>17</sup> The descriptive statistics for mid-caps are shown in Panel A, those for large caps in Panel B.

The French mid caps are larger than their German counterparts (average market capitalization  $\in 5.1$  billion as compared to  $\in 3.5$  billion), they have higher trading volume (average daily volume on the four major electronic limit order books  $\in 30.1$  million as compared to  $\in 17.4$  million) and higher message traffic (230,400 messages as compared to  $\approx 7,400$ ). French mid caps trade at a one-cent spread more frequently than their German

<sup>&</sup>lt;sup>16</sup>Again, our results remain qualitatively unchanged if we use different window lengths instead.

<sup>&</sup>lt;sup>17</sup>Our results remain qualitatively unchanged if we split the sample by index membership.

peers (10.5% versus 3.2% on the primary market and 7.6% versus 1.9% on Chi-X), meaning that the tick size-constraint is more binding for the French stocks. The level of market fragmentation (measured by the Herfindahl Hirschmann Index (HHI) calculated from the market shares of the primary exchange, Chi-X, BATS and Turquoise) is similar for French and German mid caps (36.7 as compared to 34.0), as are the market shares of the primary exchange (68.9% for French and 65.9% for German mid caps) and Chi-X (20.2% versus 22.5%). French large caps are also larger than their German peers (average market capitalization  $\in$  27.6 billion as compared to  $\in$  23.8 billion) and have higher message traffic (407,400 as compared to 351,100), but they have *lower* average daily trading volume ( $\notin$  108.6 million as compared to  $\notin$  147.9 million). The fraction of time a stock trades at a one-cent spread is similar for French and German large caps (14.8% versus 15.1% on the primary market and 13.3% versus 12.3% on Chi-X). Similarly, the level of market fragmentation as measured by the HHI and the market shares of the primary exchange and Chi-X are similar for French and German stocks (HHI 37.0 versus 35.5, market share primary market / Chi-X 69.0 / 21.0 versus 67.2 / 23.3).

#### [Insert Table 1 about here]

Figure 2 shows the market shares of the primary exchange, Chi-X, BATS and Turquoise for all four groups of stocks (French and German mid caps and large caps) during the sample period. The two markets under investigation, the primary exchange and Chi-X, account for approximately 90% of trading. This fraction is similar across the four groups and stable during the sample period.

#### [Insert Figure 2 about here]

The Appendix contains additional summary statistics for all dependent variables considered in this article separated by country and stock category.

#### 4.4 Empirical Approach

To assess the impact of the introduction of the microwave network between Frankfurt and London on various dimensions of market quality, we employ difference-in-differences (DiD) methodology. Given that the event might have differential effects on the primary market (PM) and Chi-X (CHIX), we first perform our estimation separately for each venue. Specifically, letting  $DEPVAR_{it}^k$  be the variable of interest for stock *i*, day *t*, and venue  $k \in \{PM, CHIX\}$ , we estimate the following regression:

$$DEPVAR_{it}^k = \alpha + \delta DE_i \times EVENT_t + \eta_i + \kappa_t + \varepsilon_{it}$$

where  $DE_i$  is an indicator variable equal to one if stock *i* is listed on Xetra and zero otherwise,  $EVENT_t$  is an indicator variable that takes the value one for all days after October 2012 and zero for days before. The interaction term  $DE_i \times EVENT_t$  captures the effect of the event on the treatment group in the post-event period.  $\varepsilon_{it}$  is the error term. We employ stock  $(\eta_i)$  and date  $(\kappa_t)$  fixed effects, and cluster standard errors by stock and date. Our coefficient of interest,  $\delta$ , measures the differential impact of the microwave network introduction on Xetra-listed stocks compared to the control group. We plot the development of all dependent variables used in these regressions in the Internet Appendix. German and French stocks generally closely co-move before the introduction of the microwave link, further supporting the choice of control group.

Furthermore, we employ a triple-differences (DiDiD) approach, allowing us to compare the DiD effects on Chi-X with those on Xetra. Our theoretical discussion suggests different expected effects on these two venues, and the DiDiD analysis enables us to empirically test these hypotheses. Thus, we estimate the following regression:

$$DEPVAR_{it}^{CHIX} - DEPVAR_{it}^{PM} = \alpha + \delta DE_i \times EVENT_t + \eta_i + \kappa_t + \varepsilon_{it}$$

We estimate the above two equations separately for large and mid-cap firms to study the effects of HFTs' differential trading strategies for stocks of different market capitalization and liquidity.

# 5 Traders' Cross Market Response and Trading Activity

In this section, we provide evidence for changes in traders' cross-market responses and trading activity following the introduction of the microwave link. We consider four aspects that are affected by the increase in the number of the fastest market participants. First, we consider the most direct consequences of reduced inter-market latency: its impact on the frequency, duration, and profitability of arbitrage opportunities. Second, we consider how general measures of trading activity in the markets under consideration are affected. Third, our analysis focuses on the speed at which stock-specific and market-wide information is assimilated into prices, given the enhanced aggregate capacity of fast connections. Fourth, we estimate cross-venue information shares in high resolution using the Hasbrouck (2021) VECM to analyze effects on the two venues' relative contributions to price discovery.

#### 5.1 Arbitrage Opportunities

We define arbitrage opportunities as instances where the market-wide highest bid exceeds the market-wide lowest ask, resulting in crossed order books. An arbitrage opportunity can be resolved in three different ways: (1) a trader observes the changes in the order books that create an arbitrage opportunity and acts by submitting orders (e.g. to exploit the opportunity) or canceling orders on one or both markets; (2) a trader submits or cancels orders on one market without being aware of the existence of an arbitrage opportunity; (3) a trader submits or cancels orders on both markets nearly simultaneously (possibly delaying order submission to the nearer market to account for latency differences and assure that order messages arrive on both markets at approximately the same time), thereby causing the appearance of an arbitrage opportunity that is, however, not exploitable. Importantly, only the first scenario will be affected by changes in cross-market communication latency. Because of the existence of the second and third scenarios, there will exist arbitrage opportunities lasting for periods shorter than the fastest possible communication time between the two markets. Although our data does not allow to differentiate between these three types of arbitrage opportunity resolutions, we expect that the event we analyze will impact primarily the first type, with no important effects on the other two. Consequently, the overall effect of the event should be driven by changes in the first type of resolution. Therefore, if traders can respond more quickly on one trading venue based on information from another, we anticipate a reduction in the average time required to eliminate arbitrage opportunities. We define the duration of an arbitrage opportunity as the interval from the initial crossing of the order book to the point where the inside spread is at least zero, regardless of any intervening order book changes.

We determine the theoretical gross arbitrage profit as the profit available if one were to trade against the standing orders so as to maximize one's gain. E.g., if the difference between the highest bid and the lowest ask is one tick, the theoretically available arbitrage profit is the size of a tick multiplied by the minimum of the number of shares available at the highest bid and the lowest ask. The total potential profit from an arbitrage opportunity is then the maximum such value obtained while the arbitrage opportunity persists, which is the maximum profit – before trading fees – that an arbitrageur could have achieved by instantaneously trading on both markets at any time during the arbitrage opportunity under consideration.

For the purpose of our analysis, we keep all arbitrage opportunities that last up to 10ms, which account for 84% of such opportunities in our sample. We exclude longer-lasting ones as their non-exploitation is likely not attributable to limitations in speed. Instead, these instances are more plausibly explained by traders perceiving a high risk associated with not being able to execute both legs of the trade successfully. It is important to note that all arbitrage opportunities involve inherent risks. These risks stem from the fact that latency is not zero and order books are constantly changing, creating a scenario where an opportunity identified may disappear before it can be exploited. Likely differences between the execution risks for different kinds of stocks also imply that differences in the level of arbitrage durations have no obvious interpretation. To account for execution fees, we consider only arbitrage opportunities with a theoretical gross profit of at least  $\in 1.^{18}$ The results are robust to using different windows and using other profitability cutoffs.

Table 2 presents the results of DiD analyses of variables related to arbitrage opportunities. The top panel shows results for average durations of arbitrage opportunities. Overall durations are shorter for French stocks compared to German stocks, which is to be expected because the geographical distance between the primary market and Chi-X is smaller for French stocks. Post-event, arbitrage durations for German stocks decrease significantly, whereas there is no notable change for French stocks. The DiD estimates of the reduction are -0.34ms for large caps and -0.58ms for mid caps, both of which are statistically significant. This larger decrease in arbitrage durations for mid caps as opposed to large caps is consistent with our assumption that, before the introduction of the Perseus network, very fast traders were already active in large-cap stocks but were less prevalent or absent in mid-cap stocks.

The second panel of Table 2 reveals a noteworthy decrease in the number of arbitrage opportunities post-introduction of the microwave link, suggesting enhanced market synchronization. The DiD analysis indicates that this reduced latency leads, on average, to a decrease of about 2.2 arbitrage opportunities per stock and day for mid caps and

<sup>&</sup>lt;sup>18</sup>On Xetra, the standard execution fees are determined as a percentage of the transaction value, subject to a minimum fee of  $\in 0.60$  per order. Reductions for designated market makers and Xetra liquidity providers apply, such that fees paid by likely arbitrageurs are lower. Fees on Euronext are of a similar magnitude, and those on Chi-X are lower. Thus, a  $\in 1$  gross profit would result in a positive net profit.

of about 4.3 such opportunities for large caps. Notably, while the absolute reduction in the number of opportunities is larger for large-cap stocks, the relative reduction is more pronounced for mid-cap stocks, indicating a differential impact based on market capitalization.

The third panel shows that the reduced number of arbitrage opportunities is also reflected in their total value per stock and day, with a reductions by  $\in 11.21$  for mid caps and by  $\in 19.08$  for large caps on average. These reductions correspond to over 40% for mid-caps and about 20% for large caps of their respective pre-event values. Thus, the absolute reduction is again larger for large-cap stocks than for mid-cap stocks, whereas the opposite is true for the relative change.<sup>19</sup>

[Insert Table 2 about here]

### 5.2 Trading Activity

In this subsection, we examine the effects of the microwave network's introduction on measures of trading and quoting activity. Table 3 shows averages per exchange, capitalization group, and period, as well as DiD estimates.

#### [Insert Table 3 about here]

The top panel shows that the log number of trades in German stocks on Chi-X was negatively affected for both groups of stocks, with DiD estimates of -0.115 and -0.147

<sup>&</sup>lt;sup>19</sup>Figure A1 in the appendix shows the distribution of the duration of potential arbitrage opportunities. For German stocks, there is generally a decrease in the number of potential arbitrage opportunities lasting 5 to 6 milliseconds and an increase in shorter-lived arbitrage opportunities. For French stocks, there are no strong shifts in the distribution of the duration of arbitrage opportunities.
for mid and large caps, respectively. This finding suggests that an increased participation by the fastest traders does not lead to incremental volume, presumably because other traders reduce their participation in turn. In contrast, we find no statistically significant effects on trading on the primary exchanges for either group of stocks.

We next observe, in the second panel, no significant effects on the log of median trade sizes on Chi-X, where trades are generally smaller than on the primary exchange. In contrast, the trade size on Xetra increases both in absolute terms and relative to Euronext, with DiD estimates of about 0.082 and 0.053 for mid and large caps, respectively.

The third panel shows that the effects on the log of the daily trading volume in euros for Chi-X are consistent with what we would expect based on those for the number of trades and for trade size. For the primary exchange, there is no significant effect on mid-cap trading volume. The estimate for large-cap stocks of 0.085, which is significant at the 10 percent level, indicates that Xetra experiences a relative gain in trading volume.

The first three panels of Table 3 can be interpreted jointly. It appears that relatively larger trades happen on Xetra whereas the number of trades on Chi-X drops after the introduction of the microwave network, which is consistent with the idea that investors reduce their splitting of orders across multiple exchanges and instead submit larger orders to Xetra. This might be explained by concerns that the microwave connection better allows fast traders to respond on one exchange to events on the respective other, such that order splitting across multiple exchanges is associated, when attempting to use marketable orders, with a higher risk of vanished liquidity on the second exchange by the time the investor's order reaches the market, and with increased adverse selection when using limit orders.

In the next two panels we consider the log number of order messages that affect and do not affect the best bid or offer (BBO) prices or volumes on the exchange under consideration, respectively. On Chi-X, the amount of order message traffic at the BBO for German stocks decreases, with DiD estimates of -0.357 and -0.517, respectively. These decreases exceed the reduction in the trading activity described earlier. The results for Xetra are in the opposite direction to those for Chi-X, though economically smaller and statistically weaker, consistent with the effects on trading volume. The effect on the amount of order messages behind the BBO on Chi-X is significant only for large caps, and at -0.256 smaller than the effect at the BBO but still substantial. In contrast, there are significantly positive effects of about 0.18 on the message volume behind the BBO at Xetra for both mid and large caps.

A potential explanation for the observed shifts in order message flow and sophisticated consumption of liquidity can be linked to the evolving competitive landscape among liquidity suppliers. The launch of the microwave network intensified the competition faced by 'slow' HFT market makers, particularly from their 'fast' counterparts. As new information predominantly originates in Frankfurt, fast HFT market makers gain a competitive edge on Chi-X due to their ability to rapidly adjust quotes in response to this information. This advantage, however, is less relevant on Xetra, where the proximity to the source of information negates the need for high-speed data connections for quote adjustments. Consequently, slow market makers may recalibrate their strategies in response to this increased competition. On Chi-X, where they are now at a relative disadvantage, we observe a reduction in quotation activity, leading to decreased message traffic and a reduction in the number of trades. Conversely, on Xetra, these market makers might increase their order submissions, contributing to higher message traffic on the primary market. The observed changes in sophisticated aggressive trading also align with this narrative. Slower and less sophisticated traders, finding themselves increasingly outmatched by the fastest traders on Chi-X, may choose to shift their trading activities away from this venue. This migration is reflective of a strategic realignment in response to the heightened competition and speed differentials brought about by the introduction of the microwave network.

# 5.3 Response to stock-specific and market-wide information

We have argued that prior to the introduction of the microwave link, there probably had been no fast HFTs trading mid caps. Consequently, the introduction of the microwave link may have implications for cross-market price discovery.

In the model of Baldauf and Mollner (2020), there is only one asset, such that there is no distinction between firm-specific and market-wide information. This contrasts with the real world, where stock prices are affected by both types of information. Initially, prior to the introduction of the commercial microwave link, only few HFTs had access to private microwave connections. Given the capacity constraints of these private links, it is plausible that these HFTs to a significant degree focused on trading based on market-wide information. However, the introduction of the commercial network expanded the overall capacity of microwave links, potentially leading some fast HFTs to shift their focus towards trading on firm-specific information. Consequently, we expect that, following the network introduction, firm-specific information – primarily generated in Frankfurt, as previously argued – will begin to affect prices on Chi-X more rapidly than before. Nonetheless, we do not anticipate a similar change in the speed of reaction to market-wide information.

The preceding discussion implies that returns on one market subsequent to returns observed on the other market result from traders responding to information observed in the latter. Similar to the reasoning in Section 5.1, correlated returns on the two markets can also result from a trader acting on both markets nearly simultaneously, ensuring a persistent positive relation between returns on the two venues at speeds faster than the theoretically possible communication latency. We do not anticipate this type of activity to be affected by the event. Therefore, changes in the speed at which market-wide and stock-specific information are incorporated into prices can be interpreted as the effects of faster information transmission between markets.

To test for such effects, we estimate the following regression model for each stock and date:

$$r_{i,t}^{Chi-X} = \alpha_i + \sum_{k=0}^{10} \beta_{i,k} r_{m,t-k}^{Eurex} + \sum_{k=0}^{10} \gamma_{i,k} r_{i,t-k}^{Xetra} + \varepsilon_{i,t}$$
(1)

where i indicates the stock, t the point in time at millisecond resolution, and m the market as proxied for by the near-term futures contract on the DAX30 large-cap index.<sup>20</sup> In Table 4, we report the means of the *beta* and *gamma* coefficients before and after the event. Furthermore, we test for differences by regressing these coefficients on a dummy

<sup>&</sup>lt;sup>20</sup>Note that the futures on the mid-cap index MDAX is too illiquid to be of use here. However, the large-cap index DAX is representative of and thus suitable as a proxy for the market.

variable taking on the value 1 in the period after the introduction of the microwave link. We cluster the standard errors by stock and date.

We then plot the values of the estimated coefficients in Figure 3 (a) and (b). Moreover, Figure 3 (c) and (d) show the cumulative sum of these coefficients within k milliseconds. For both large-cap and mid-cap stocks, it is apparent that subsequent to the introduction of the new microwave link the amount of stock-specific information priced in within 4 milliseconds, i.e., faster than would be possible without a microwave link, increases. In turn, the additional information priced in only after delays achievable by fiber connections decreases for large caps, resulting in a modest increase in the total sensitivity of their prices on Chi-X to changes of Xetra prices within the past 10 milliseconds. Such a decrease is not observable for mid caps, presumably because there was very little use of stock-specific information in these stocks before the event, when they were quite illiquid on Chi-X. With respect to the information from futures contracts, we observe a decrease for all of the 10 milliseconds under consideration. This finding is consistent with the notion that market-wide information is displaced by stock-specific information as increased data bandwidth at fast speed becomes available.

In conclusion, we find evidence that stock prices on Chi-X more quickly incorporate stock-specific information subsequent to the introduction of the new microwave connection, while the importance of market-wide information is reduced.

[Insert Table 4 about here]

[Insert Figure 3 about here]

39

# 5.4 Cross-Venue Information Shares

In this subsection, we shift our focus from the speed and type of information incorporated into prices on Chi-X to examining the impact of faster cross-market connectivity on the venue of information incorporation. Specifically, we investigate the relative contribution to price discovery of two trading venues and shifts thereof around the introduction of the microwave link. We rely on the methodology pioneered by Hasbrouck (1995, 2003). The price dynamics in this setting can be expressed by the VECM

$$\Delta P_t = \gamma \beta P_{t-1} + \sum_{i=1}^k \phi_i \Delta P_{t-i} + \varepsilon_t \tag{2}$$

where  $P_t = [p_t^{PM}, p_t^{CHIX}]$  is the vector of log quote midpoints,  $\beta P_{t-1}$  is the pre-specified, zero-mean cointegrating vector with  $\beta = (1, -1)'$ ,  $\gamma$  is a 2 × 1 error correction vector, the second term on the right side is the VECM's autoregressive component of order k, and  $\varepsilon_t$ is a zero-mean error vector with variance-covariance matrix  $\Omega$ . Our estimation is based on quote midpoints observed with a millisecond resolution to ensure that any effect associated with reducing round-trip latency between Frankfurt and London from 8.3ms to 4.6ms is picked up by the model. We therefore employ the high resolution version of the method developed by Hasbrouck (2021). We set k = 10000 to capture autoregressive effects up to 10 seconds in the past and impose constraints on the autoregressive coefficients. Specifically, we let  $\phi_1$  unrestricted and assume  $\phi_2 = \ldots = \phi_{10}, \phi_{11} = \ldots = \phi_{100}, \phi_{101} = \ldots = \phi_{1000},$ and  $\phi_{1001} = \ldots = \phi_{10000}$ .<sup>21</sup> We note that using midquotes rather than transaction prices

 $<sup>^{21}</sup>$ We rely on code made available by Joel Hasbrouck to estimate the VECM. The code along with its documentation is available at http://people.stern.nyu.edu/jhasbrou/Research/

(or combinations of both) avoids the biases described in Hagströmer and Menkveld (2023). The high data resolution of one millisecond allows us to obtain very tight upper and lower bounds for the information shares. This is because at low resolutions time aggregation of events increases the contemporaneous correlation in  $\Delta P_t$ , whereas at high resolutions the sequence of events is determined more accurately. At one millisecond resolution, the median difference between the upper and lower bounds across all stock-days is 1.3%.<sup>22</sup>

We estimate the primary market information share separately for each stock-day. To perform statistical inference, we analyze the distribution of these information shares across all stock-days in our sample. Table 5 contains the results. We report the mean and median values of the lower and upper bounds of the primary market information share as well as their average.

## [Insert Table 5 about here]

The information share in the primary market is larger for German than for French large-cap stocks, with a mean of 64.7% compared to 54.0%. The introduction of the microwave connection has no significant effect on the large-cap primary market information shares. However, for mid-caps, where the primary market information share is generally much higher, we obtain very different results. The primary market information shares decline sharply after the introduction of the network, with an effect size of nearly 10 percentage points.

HRVAR/HRVARindex.html.

<sup>&</sup>lt;sup>22</sup>We estimate the VECM at one second resolution and obtain wider (and, as a result, overlapping) information share bounds for the primary market and Chi-X.

As IS mainly measures which market prices in information first (Yan and Zivot, 2010), this result points to a large shift of information discovery to Chi-X, while Xetra altogether remains dominant. It appears the faster communication between the two venues allows information arising in Chi-X – possibly in response to price changes of other assets on exchanges in and near London – to travel to Xetra. Such effects are not apparent for large caps, where the magnitude of the dominance of Xetra remains unchanged.

Overall, these results complement the trading activity effects documented above. An increase in fast liquidity provision by HFTs for mid-cap firms leads to an increase in Chi-X's competitiveness relative to Xetra, further leading to an increase in its contribution to price discovery. For large-cap firms, the two markets are already highly integrated and equally competitive such that we do not observer a similar competition-induced shift in information shares towards Chi-X.

# 6 Market Quality Implications

In this section, we examine the impact of the microwave link on various measures of market quality. We begin by analyzing the link's effects on market liquidity, which we tie, to a large degree, to our discussion of theory and the resulting hypotheses of Section 3. Following this, we explore the link's implications for market volatility and price efficiency.

## 6.1 Liquidity and Transaction Costs

We begin our analysis by examining the effects on quoted and traded liquidity, considering Chi-X and Xetra separately. As elaborated in Section 3, we anticipate a differential impact not only across these two trading venues but also between large-cap and mid-cap stocks. This expectation stems largely from variations in the existing trader populations and their adaptive responses to the introduction of the commercial microwave link. Further, the findings in Section 5, which indicate Xetra's more significant contribution to price discovery compared to Chi-X, reinforce our expectation of more pronounced effects on Chi-X. The introduction of the microwave link, altering the speed differences, is more important in a context where most of the information originates from the other end of this connection.

All findings related to liquidity are detailed in Table 6. We first analyze the impact on large-cap liquidity, guided by predictions from Hypothesis 1, which are based on theory and the institutional context. On Chi-X, we observe a significant positive effect on price impact, aligning with our expectation of increased sniping activity. This effect, approximately 0.8 basis points, is economically substantial in comparison to the average of about 5 basis points. Conversely, on Xetra, the smaller volume of information flow in the reverse direction appears insufficient to produce a similar effect, with the observed impact being an order of magnitude smaller and not statistically significant. The positive effect on adverse selection on Chi-X is mirrored by a negative effect on realized spreads of about half a basis point. This suggests that liquidity providers previously earned rents, likely because the tick size constrained their price competition. These rents are either reduced or eliminated due to heightened adverse selection. The observation of slightly negative average realized spreads is a common phenomenon in modern equity markets (see, e.g., Aquilina et al., 2022) attributable to limit order trades by directional traders, who are not included in our model. For these traders, a marginally negative realized spread may still be preferable to the costs associated with marketable orders. Overall, we observe small and statistically insignificant effects on both effective and quoted spreads on Chi-X. In contrast, Xetra shows a minor negative impact on realized spreads, about 0.3 basis points, which exceeds the negligible effect on price impacts. This result could imply reduced inventory costs, possibly due to improved integration of trading activities by liquidity providers across both markets. However, we caution against over-interpreting this modest finding. Lastly, when assessing market depth, measured by the size of the best bid and ask for large caps, we find no significant effects.

Shifting our focus to mid-cap stocks, we anticipate differing outcomes as outlined in Hypothesis 2. Whereas these stocks previously had apparently not been traded using information from microwave connections by the fastest market participants, likely due to network capacity constraints, our findings indicate effects that align with the entry of fast liquidity providers. The relative speed of these providers, in comparison to the average liquidity consumer, appears to lead to a marked decrease in adverse selection. This is evident from the significant negative effect on price impact, which amounts to approximately 3.8 basis points. On Xetra, we notice a smaller, yet also negative, impact on price impact. This aligns with the expectation of reduced adverse selection in the rarer instances of information being transmitted to, rather than from, Frankfurt. We do not observe significant changes in realized spreads, which corroborates the supposition that tick sizes do not notably constrain competition among liquidity providers. The market for liquidity provision was competitive both pre- and post-event, with the primary difference being the exposure to adverse selection, a factor common to all liquidity providers. Consequently, the magnitude of effects on effective spreads mirrors those on price impacts. The impact on quoted spreads on Chi-X is even more pronounced. This suggests that prior to the introduction of the new microwave link, liquidity on Chi-X was often costly, leading traders to engage on this platform only when spreads were exceptionally narrow. Finally, we note a negative effect on depth on Chi-X. This decrease is likely attributable to the diminished competitiveness of slower traders who engage in large limit orders, especially following the entry of the fastest liquidity providers.

[Insert Table 6 about here]

# 6.2 Volatility and Price Efficiency

In this subsection, we consider effects of the increase in the population of traders with the fastest information transmission technology on stock return volatility and price efficiency, i.e., the deviations of prices from a random walk. Although the theoretical framework of Baldauf and Mollner (2020) does not explicitly provide testable hypotheses for these variables, we nevertheless extend our analysis to assess the impact of the change in speed dispersion also on these important dimensions of market quality. We present the results in Table 7.

## [Insert Table 7 about here]

For large caps, we observe no significant effect on our measure of volatility, the one-second standard deviation of quote midpoint returns, on either of the two exchanges. Neither is there evidence for a significant effect on volatility for mid caps on Xetra. However, we do find evidence for a large negative effect of more than 2 bps for mid caps on Chi-X, which amounts to more than half of the pre-event level. The latter result is consistent with the effects on liquidity observed in Section 6.1 and aligns with the idea that smaller spreads coincide with more precise quote midpoint prices and smaller short-term changes to those prices.

In our examination of price efficiency, defined as the deviation of quote midpoint prices from a random walk, we analyze two metrics: the absolute autocorrelation of 10-second quote midpoint returns and the variance ratio of 10- and 30-second midpoint returns. For enhanced legibility, we have scaled these measures up by a factor of 100. For large caps, we find no significant effects on Chi-X. However, on Xetra, we find a slight adverse effect on price efficiency, with effect sizes of about 0.6 and 0.7 for autocorrelation and variance ratio, respectively, though only the former is statistically significant. For mid caps, the results for price efficiency mirror those for volatility. There is no evidence for a significant effect on Chi-X of about 7.7 and 11.5 units for autocorrelation and variance ratio, respectively. These results indicate that after the introduction of the commercial microwave link, the more liquid order book not only exhibits reduced volatility but also decreased predictability of price changes. In conclusion, our findings offer compelling evidence that the first-time market entry of the fastest traders, who apparently choose to engage in liquidity provision, which we infer is the case for mid caps in our setting, has beneficial effects on various measures of market quality. By contrast, as observed for large-cap stocks, if there is already competitive fast liquidity provision, the incremental entry of fast snipers can have adverse effects on liquidity.

# 7 Conclusion

Trading in today's fragmented financial markets is characterized by speed differences both in terms of the capability to quickly respond to information on an individual trading venue and in terms of access to the fastest ways of transmitting information and order messages between venues trading the same or related securities. The fastest, wireless, networks connecting geographically dispersed trading centers have capacity constraints due to physical limits. These constraints are particularly relevant when transmitting data through air, which nowadays is the common approach as it minimizes latency. The limited bandwidth of the networks is a particularly severe restriction when, as is the case in today's fragmented equity markets, a large number of liquid financial instruments is traded in parallel markets, and when the valuation of these instruments depends on idiosyncratic information.

We analyze the introduction of the first commercially available microwave link between Frankfurt and London, which increased the speed of communication between Frankfurt-based Xetra and London-based Chi-X for the subscribers to the link. We find that this event affects trading activity and market quality of both large-cap and mid-cap stocks. We establish that the increased communication speed of a subset of market participants enhances the synchronization of prices across the two venues, as evidenced by a reduction in the number, duration, and value of arbitrage opportunities. We further find that the event reduces trading volume and quoting activity on the smaller venue Chi-X, whereas the larger (and home) venue Xetra is largely unaffected in this regard. The new microwave link results in faster incorporation of stock-specific information originating in the home market into prices, thereby reducing the relative importance of market-wide information transmitted from Frankfurt to London. Information shares for large-cap stocks are unaffected by the event whereas for mid caps we find an increase in information shares for Chi-X and a corresponding decrease for Xetra. This result may be due to the fact that proprietary microwave networks were apparently not used to communicate stock-specific information for mid caps.

We document changes in market liquidity that are consistent with predictions that we derive from the model in Baldauf and Mollner (2020). The sniping activities of fast market participants in large caps appear to increase adverse selection on Chi-X and reduce rents earned by liquidity providers which presumably result from minimum tick size-constraints. At the same time, liquidity for mid-cap stocks on Chi-X improves, most likely because some market participants use the microwave link for market making activities in those stocks. The reduced spreads on Chi-X go hand in hand with a decrease in volatility and a higher degree of price efficiency.

Altogether, our results are nuanced and do not suggest that additional fast connections between trading venues are unambiguously good or bad. They rather suggest that having a small number of fast competitive liquidity providers with access to the fastest inter-market connection is beneficial to liquidity and price efficiency. However, a larger number of such market participants may be harmful to liquidity because of increased sniping activity. Our findings further imply that the improved liquidity upon the entry of fast liquidity suppliers may not necessarily benefit the venue operator because there is no overall increase in trading volume.

# References

- Aït-Sahalia, Y. and M. Sağlam (2024). High frequency market making: The role of speed. Journal of Econometrics 239(2), 105421.
- Aquilina, M., E. Budish, and P. O'neill (2022). Quantifying the high-frequency trading "arms race". *The Quarterly Journal of Economics* 137(1), 493–564.
- Baldauf, M. and J. Mollner (2020). High-frequency trading and market performance. The Journal of Finance 75(3), 1495–1526.
- Baldauf, M. and J. Mollner (2021). Trading in Fragmented Markets. Journal of Financial and Quantitative Analysis 56(1), 93–121.
- Baldauf, M. and J. Mollner (2022). Fast Traders Make a Quick Buck: The Role of Speed in Liquidity. *Journal of Financial Markets* 58, 100621.
- Bernales, A. (2019). Make-take decisions under high-frequency trading competition. Journal of Financial Markets 45, 1–18.
- Biais, B., T. Foucault, and S. Moinas (2015). Equilibrium Fast Trading. Journal of Financial Economics 116(2), 292–313.
- Boehmer, E., K. Y. L. Fong, and J. Wu (2021). Algorithmic Trading and Market Quality: International Evidence. Journal of Financial and Quantitative Analysis 56, 2659–2688.
- Boehmer, E. and E. K. Kelley (2009). Institutional Investors and the Informational Efficiency of Prices. *Review of Financial Studies* 22(9), 3563–3594.

- Bongaerts, D. and M. Van Achter (2021). Competition among liquidity providers with access to high-frequency trading technology. *Journal of Financial Economics* 140(1), 220–249.
- Brogaard, J., B. Hagströmer, L. Nordén, and R. Riordan (2015). Trading fast and slow: Colocation and liquidity. *The Review of Financial Studies* 28(12), 3407–3443.
- Budish, E., P. Cramton, and J. Shim (2015). The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *Quarterly Journal* of Economics 130(4), 1547–1621.
- Dobrev, D. and E. Schaumburg (2023). High-Frequency Cross-Market Trading: Model Free Measurement and Testable Ipplications. *Working Paper*.
- Foucault, T., J. Hombert, and I. Roşu (2016). News Trading and Speed. The Journal of Finance 71(1), 335–382.
- Foucault, T., R. Kozhan, and W. W. Tham (2017). Toxic Arbitrage. The Review of Financial Studies 30(4), 1053–1094.
- Frino, A., V. Mollica, and R. I. Webb (2014). The Impact of Co-Location of Securities Exchanges' and Traders' Computer Servers on Market Liquidity: The Impact of Co-Location on Market Liquidity. *Journal of Futures Markets* 34(1), 20–33.
- Garbade, K. D. and W. L. Silber (1978). Technology, Communication and the Performance of Financial Markets: 1840-1975. *Journal of Finance* 33(3), 819–832.

- Gromb, D. and D. Vayanos (2002). Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs. *Journal of Financial Economics* 66(2-3), 361–407.
- Gromb, D. and D. Vayanos (2010). Limits of Arbitrage: The State of the Theory. *NBER* Working Paper (15821).
- Hagströmer, B. and A. J. Menkveld (2023). Trades, quotes, and information shares. Available at SSRN 4356262.
- Hasbrouck, J. (1995). One Security, Many Markets: Determining the Contributions to Price Discovery. *Journal of Finance* 50(4), 1175–1199.
- Hasbrouck, J. (2003). Intraday Price Formation in U.S. Equity Index Markets. Journal of Finance 58(6), 2375–2399.
- Hasbrouck, J. (2021). Price discovery in high resolution. Journal of Financial Econometrics 19(3), 395–430.
- Hoffmann, P. (2014). A Dynamic Limit Order Market with Fast and Slow traders. Journal of Financial Economics 113(1), 156–169.
- Hollifield, B., P. Sandas, and A. Todd (2017). Latency Arbitrage When Markets Become Faster. *Working Ppaper*.
- Kozhan, R. and W. W. Tham (2012). Execution Risk in High-Frequency Arbitrage. Management Science 58(11), 2131–2149.

52

- Menkveld, A. J. and M. A. Zoican (2017). Need for Speed? Exchange Latency and Liquidity. The Review of Financial Studies 30(4), 1188–1228.
- Pagnotta, E. S. and T. Philippon (2018). Competing on Speed. Econometrica 86(3), 1067–1115.
- Rzayev, G. Ibikunle, and T. Steffen (2023, November). The Market Quality Implications of Speed in Cross-Platform Trading: Evidence from Frankfurt-London Microwave Networks. *Journal of Financial Markets* 66, 100853.
- Shkilko, A. and K. Sokolov (2020). Every Cloud Has a Silver Lining: Fast Trading, Microwave Connectivity and Trading Costs. *Journal of Finance* 75(6), 2899–2927.
- Yan, B. and E. Zivot (2010). A structural analysis of price discovery measures. Journal of Financial Markets 13(1), 1–19.
- Ye, M., C. Yao, and J. Gai (2013). The Externalities of High Frequency Trading. *Working Paper*.

#### Table 1: Summary Statistics for Sample Stocks

This table shows summary statistics for the stocks in the sample based on the pre-treatment period of August and September 2012. Market Cap. is the market capitalization in 1bn EUR. Trading volume is the daily trading volume at the primary listing exchange and Chi-X, Bats, and Turquoise, in 1mn EUR. # Messages is the message traffic as measured by the number of added or canceled limit orders on the primary listing exchange and Chi-X, in 1000. Tick Constrained shows the fraction of the day the quoted spread is at one tick, in percent.  $HHI_{Exchanges}$  is the normalized Herfindahl index for trading volume at the four exchange, in percent. Market Share shows the relative market share of the primary listing exchange / of Chi-X relative to all four exchanges, in percent. Stocks are sorted into mid and large caps by the medians of market capitalization for the two main listing exchanges during July 2012.

		Germa	n stock	s	French stocks						
	Mean	P5	P50	P95	Mean	P5	P50	P95			
Panel A: Mid caps											
Market Cap.	3.5	1.8	3.1	5.7	5.1	2.5	4.8	8.5			
Trading Volume	17.4	2.0	12.5	57.4	30.1	7.0	26.7	69.0			
# Messages	87.4	28.1	71.9	213.4	230.4	40.4	223.5	618.9			
Tick Constrained <sub>PM</sub>	3.2	0.7	2.1	9.3	10.5	2.8	6.9	26.0			
Tick Constrained <sub>Chi-X</sub>	1.9	0.2	1.1	5.1	7.6	0.9	5.0	23.2			
$\mathrm{HHI}_{\mathrm{Exchanges}}$	34.0	19.4	31.2	61.5	36.7	24.9	36.7	48.2			
Market $Share_{PM}$	65.9	53.1	65.1	83.8	68.9	59.9	68.9	76.3			
Market $\text{Share}_{\text{Chi-X}}$	22.5	10.5	23.1	33.0	20.2	15.7	20.3	25.2			
Panel B: Large caps											
Market Cap.	23.8	7.1	18.1	60.3	27.6	9.5	22.5	81.7			
Trading Volume	147.9	44.4	109.0	310.1	108.6	20.3	86.2	277.8			
# Messages	351.1	142.1	265.1	707.0	407.4	115.1	346.2	890.6			
Tick Constrained <sub>PM</sub>	15.1	1.5	9.9	40.1	14.8	5.4	10.6	37.2			
Tick Constrained <sub>Chi-X</sub>	12.3	1.5	7.5	35.4	13.3	3.0	8.7	32.0			
$\mathrm{HHI}_{\mathrm{Exchanges}}$	35.5	25.2	34.4	50.6	37.0	28.1	38.5	44.6			
Market $Share_{PM}$	67.2	58.3	66.8	77.6	69.0	63.2	70.3	74.2			
Market $Share_{Chi-X}$	23.3	15.9	23.9	31.2	21.0	16.9	20.6	25.7			

#### Table 2: Arbitrage Opportunities

This table shows summary statistics and difference-in-difference results for arbitrage opportunities between the primary listing exchanges and Chi-X.  $Arb._{Duration}$  is the average duration of an arbitrage opportunity in milliseconds.  $Arb._{Number}$  is the daily number of arbitrage opportunity per stock.  $Arb._{Value}$  is the daily theoretical maximum gross arbitrage profit per stock in EUR. Only arbitrage opportunities lasting up to 10ms and with a profitability of at least 1 EUR are included. The column denoted DiD show the difference-in-difference coefficients. The before and after periods are August–September 2012 and November–December 2012, respectively. Stocks are sorted into mid and large caps by the medians of market capitalization for the two main listing exchanges during July 2012. Cluster robust t-statistics are given in parentheses. \*\*\*, \*\*, \* denotes significance at the 1%, 5%, 10%–level, respectively.

		German	stocks	French s	stocks	
		Before	After	Before	After	DiD
Arb. Duration	Mid caps	4.06	3.62	2.08	2.21	$-0.576^{***}$ (-5.30)
	Large caps	4.09	3.82	2.15	2.22	$-0.339^{***}$ (-3.10)
	Mid caps	5.92	4.02	2.48	2.76	$-2.183^{***}$ (-2.95)
ArD. Number	Large caps	20.19	15.55	6.80	6.49	$-4.328^{***}$ (-2.86)
	Mid caps	25.89	16.27	9.03	10.62	$-11.210^{***} \\ (-2.66)$
Arb. Value	Large caps	97.27	73.02	28.98	23.82	$-19.080^{**}$ (-2.57)

#### Table 3: Trading and Quoting Activity

This table shows summary statistics and difference-in-difference results for trading and quoting activity. # Trades is the log number of transactions. Trade Size<sub>p50</sub> is the median trade size during a day in InEUR. Tr. Volume is the trading volume in InEUR. # Messages BBO (notBBO) is the message traffic (not) affecting the best bid or offer, measured by the log number of new and canceled quotes. The columns denoted DiD show the difference-in-difference coefficients. In the estimations for  $DiD_{Chi-X}$ , the treatment group are German stocks on Chi-X and the control group French stocks on Chi-X. Similarly, for  $DiD_{PM}$  we take data from the two primary listing exchanges Xetra and Euronext.  $DiD_{Diff}$  gives the difference-in-difference estimates on the differences between the primary listing exchange and Chi-X. The before and after periods are August–September 2012 and November–December 2012, respectively. Stocks are sorted into mid and large caps by the medians of market capitalization for the two main listing exchanges during July 2012. Cluster robust t-statistics are given in parentheses. \*\*\*, \*\*, \* denotes significance at the 1%, 5%, 10%–level, respectively.

		Chi-X <sub>0</sub>	$\operatorname{hi-X}_{\operatorname{German}}$ $\operatorname{Chi-X}_{\operatorname{French}}$			$\mathrm{PM}_{\mathrm{German}}$		$\mathbf{PM}_{\mathbf{French}}$				
		Before	After	Before	After	${\rm DiD}_{{\rm Chi}-{\rm X}}$	Before	After	Before	After	${ m DiD_{PM}}$	${ m DiD}_{ m Diff}$
	Mid caps	6.67	6.67	7.26	7.37	$-0.115^{*}$ (-1.72)	6.89	6.77	7.89	7.80	$-0.041 \\ (-0.59)$	-0.074 (-1.54)
#1rades	Large caps	8.32	8.25	8.15	8.22	$-0.147^{***}$ (-2.61)	8.16	8.06	8.54	8.41	$\begin{array}{c} 0.029 \\ ( \ 0.57) \end{array}$	$-0.176^{***}$ (-3.73)
Trade Size <sub>P50</sub>	Mid caps	7.63	7.65	7.66	7.70	$-0.026 \\ (-0.62)$	8.09	8.21	8.01	8.06	$\begin{array}{c} 0.082^{**} \\ ( 2.15) \end{array}$	$-0.107^{***}$ (-2.61)
	Large caps	8.25	8.28	7.99	8.03	$-0.015 \\ (-0.33)$	9.07	9.11	8.54	8.53	$\begin{array}{c} 0.053 \\ ( \ 1.63) \end{array}$	$-0.068 \\ (-1.62)$
The Volume	Mid caps	14.52	14.56	15.19	15.35	$-0.118^{*}$ (-1.68)	15.44	15.44	16.23	16.22	$\begin{array}{c} 0.005 \\ ( \ 0.06) \end{array}$	$-0.123^{**}$ (-2.44)
Ir. volume	Large caps	16.94	16.86	16.47	16.56	$-0.178^{***}$ (-3.49)	17.72	17.68	17.42	17.29	$0.085^{*}$ ( 1.85)	$\begin{array}{c} -0.263^{***} \\ (-5.82) \end{array}$
//Magga mag BBO	Mid caps	9.19	9.04	9.50	9.71	$-0.357^{***}$ (-3.54)	9.25	9.04	10.57	10.27	$0.089 \\ (0.88)$	$-0.446^{***}$ (-4.97)
#Messages DDO	Large caps	10.54	10.52	10.26	10.76	$-0.517^{***}$ (-4.57)	10.55	10.45	11.18	10.94	$\begin{array}{c} 0.147^{*} \\ ( 1.75) \end{array}$	$-0.664^{***}$ (-5.87)
	Mid caps	9.77	9.70	10.35	10.31	$-0.034 \\ (-0.36)$	10.20	10.09	11.27	10.98	$\begin{array}{c} 0.184^{**} \\ ( 2.27) \end{array}$	$-0.219^{***}$ (-3.05)
#inessages notBBO	Large caps	11.44	11.25	11.22	11.29	$-0.256^{***}$ (-2.63)	11.72	11.56	11.80	11.46	$\begin{array}{c} 0.181^{**} \\ ( \ 2.01 ) \end{array}$	$-0.437^{***}$ (-5.63)

#### Table 4: Market Information

This table shows coefficients from a regression of returns of German stocks at Chi-X on lagged returns of the same stock at Xetra and returns of DAX index futures traded at Eurex. Lags are given in milliseconds. The columns *Diff* and *t-stat* show the results for a t-test of differences between the coefficients before and after the introduction of the microwave link. The before and after periods are August–September 2012 and November–December 2012, respectively. Stocks are sorted into mid and large caps by the median of market capitalization at the primary listing exchange during July 2012.

			Μ	lid Caps			$\mathbf{L}$	arge Cap	s
Market	Lag	Before	After	Diff	t-stat	Before	After	Diff	t-stat
	0	0.0030	0.0025	-0.0005	-0.36	0.0043	0.0045	0.0002	0.37
	1	0.0044	0.0047	0.0003	0.27	0.0041	0.0059	0.0018	$3.79^{***}$
	2	0.0074	0.0076	0.0002	0.12	0.0077	0.0091	0.0015	$2.66^{***}$
	3	0.0159	0.0198	0.0039	1.39	0.0199	0.0260	0.0061	$5.59^{***}$
	4	0.0224	0.0510	0.0286	$7.95^{***}$	0.0275	0.0556	0.0281	$11.17^{***}$
Xetra	5	0.0408	0.0550	0.0142	$2.02^{**}$	0.0767	0.0707	-0.0061	$-2.53^{**}$
	6	0.0121	0.0275	0.0154	$2.29^{**}$	0.0495	0.0358	-0.0137	$-10.54^{***}$
	7	0.0147	0.0108	-0.0038	-1.43	0.0212	0.0168	-0.0044	$-6.20^{***}$
	8	0.0088	0.0065	-0.0022	-1.16	0.0144	0.0119	-0.0024	$-3.68^{***}$
	9	0.0044	0.0048	0.0004	0.22	0.0095	0.0088	-0.0007	-0.82
	10	0.0020	0.0021	0.0000	0.00	0.0067	0.0057	-0.0011	$-2.58^{***}$
	0	0.0011	0.0004	-0.0007	$-2.05^{**}$	0.0036	0.0026	-0.0010	$-4.60^{***}$
	1	0.0016	0.0006	-0.0010	$-3.93^{***}$	0.0050	0.0041	-0.0009	$-3.77^{***}$
	2	0.0016	0.0009	-0.0006	$-1.77^{*}$	0.0068	0.0060	-0.0008	$-2.52^{**}$
	3	0.0016	0.0013	-0.0002	-0.71	0.0092	0.0084	-0.0008	$-2.15^{**}$
	4	0.0039	0.0029	-0.0010	$-1.75^{*}$	0.0214	0.0164	-0.0050	$-4.41^{***}$
Futures	5	0.0044	0.0031	-0.0013	$-1.97^{**}$	0.0249	0.0186	-0.0064	$-5.05^{***}$
	6	0.0035	0.0027	-0.0007	-1.34	0.0190	0.0154	-0.0036	$-4.03^{***}$
	7	0.0030	0.0023	-0.0007	-1.58	0.0152	0.0119	-0.0033	$-4.58^{***}$
	8	0.0027	0.0019	-0.0008	$-2.30^{**}$	0.0119	0.0098	-0.0021	$-4.08^{***}$
	9	0.0024	0.0016	-0.0008	$-3.46^{***}$	0.0090	0.0077	-0.0013	$-3.05^{***}$
	10	0.0021	0.0014	-0.0007	$-2.88^{***}$	0.0071	0.0059	-0.0012	$-4.04^{***}$

57

#### Table 5: Information Shares

This table shows summary statistics and difference-in-difference results for information shares. Info Share  $_{PM}$  is the average of the lower and upper bound of the high-resolution information share of the primary market in percent. Info Share  $_{PM}$  <sup>Lower</sup> and Info Share  $_{PM}$  <sup>Upper</sup> are the respective bounds. The column denoted DiD show the difference-in-difference coefficients. The before and after periods are August–September 2012 and November–December 2012, respectively. Stocks are sorted into mid and large caps by the medians of market capitalization for the two main listing exchanges during July 2012. Cluster robust t-statistics are given in parentheses. \*\*\*, \*\*, \* denotes significance at the 1%, 5%, 10%–level, respectively.

		German stocks		French s	stocks	
		Before	After	Before	After	$\operatorname{DiD}$
${\rm InfoShare}_{\rm PM}$	Mid caps	86.12	73.16	63.07	60.04	$-9.933^{***}$ (-4.82)
	Large caps	64.67	64.70	55.87	53.98	$1.929 \\ (1.18)$
Lefe Classe Lower	Mid caps	85.99	72.95	60.59	57.30	$-9.755^{***}$ (-4.61)
$\mathrm{Info}\mathrm{Share}_{\mathrm{PM}}$	Large caps	64.08	64.34	52.65	50.50	$2.410 \\ (1.36)$
I C Cl Upper	Mid caps	86.26	73.31	65.52	62.72	$-10.150^{***}$ (-5.01)
Into Snare PM <sup>oppor</sup>	Large caps	64.79	64.99	59.02	57.30	$\begin{array}{c} 1.923 \\ ( \ 1.17) \end{array}$

#### Table 6: Liquidity

This table shows summary statistics and difference-in-difference results for various liquidity measures .  $Pr. Impact_t$  is the volume-weighted price impact at horizon t, i.e. twice the signed difference between transaction prices and the quote midpoint after time t. Realized<sub>t</sub> is the volume-weighted difference between effective spreads and the price impact at time horizon t. Effective is the volume-weighted effective spread, i.e. twice the signed difference between transaction prices and quote midpoints immediately before the transactions. Effective spreads, price impact, and realized spreads are presented in basis points of the transaction price. Quoted is the time-weighted quoted spread in basis points of the quote midpoint.  $Depth_{top}$  is the total quoted depth at the best bid and ask in lnEUR. The columns denoted DiD show the difference-in-difference coefficients. In the estimations for  $DiD_{Chi-X}$ , the treatment group are German stocks on Chi-X and the control group French stocks on Chi-X. Similarly, for  $DiD_{PM}$  we take data from the two primary listing exchanges Xetra and Euronext.  $DiD_{Diff}$  gives the difference-in-difference estimates on the differences between the primary listing exchange and Chi-X. The before and after periods are August–September 2012 and November–December 2012, respectively. Stocks are sorted into mid and large caps by the medians of market capitalization for the two main listing exchanges during July 2012. Cluster robust t-statistics are given in parentheses. \*\*\*, \*\*, \* denotes significance at the 1%, 5%, 10%–level, respectively.

		Chi-X <sub>0</sub>	-X <sub>German</sub> Chi-Y		$hi-X_{French}$		$\mathrm{PM}_{\mathrm{German}}$		$\mathrm{PM}_{\mathrm{French}}$			
		Before	After	Before	After	${ m DiD}_{ m Chi-X}$	Before	After	Before	After	${ m DiD_{PM}}$	${ m DiD}_{ m Diff}$
Dr. Lucrast	Mid caps	13.71	9.24	7.51	6.84	$-3.804^{**}$ (-2.34)	9.94	8.74	6.84	7.14	$-1.508^{***}$ (-4.07)	$-2.296 \\ (-1.52)$
11. Impact <sub>10s</sub>	Large caps	4.90	4.56	5.61	4.50	$\begin{array}{c} 0.769^{***} \\ ( \ 4.11 ) \end{array}$	5.21	4.79	5.43	4.92	$\begin{array}{c} 0.089 \\ ( \ 0.49) \end{array}$	$\begin{array}{c} 0.680^{***} \\ (5.70) \end{array}$
Mid caps Realized <sub>10s</sub> Large cap	Mid caps	-0.01	0.04	-1.42	-1.20	$-0.165 \ (-0.63)$	-0.04	-0.08	-0.72	-1.07	$\begin{array}{c} 0.309 \\ ( \ 1.12 ) \end{array}$	$-0.475^{***}$ (-3.18)
	Large caps	-0.58	-0.70	-1.09	-0.68	$-0.528^{***}$ (-4.40)	-0.36	-0.50	-0.70	-0.55	$-0.289^{**}$ (-2.13)	$-0.240^{***}$ (-2.77)
Effective	Mid caps	13.79	9.34	6.10	5.66	$-4.011^{**}$ (-2.39)	9.93	8.63	6.15	6.08	$-1.238^{***}$ (-3.51)	$-2.773^{*}$ (-1.77)
Ellective	Large caps	4.31	3.86	4.52	3.82	$\begin{array}{c} 0.245 \\ ( \ 1.63) \end{array}$	4.85	4.29	4.74	4.37	-0.194 (-1.15)	$\begin{array}{c} 0.439^{***} \\ ( \ 4.69 ) \end{array}$
Queted	Mid caps	24.85	14.39	8.64	7.91	$-9.740^{***}$ (-2.63)	11.79	9.97	7.56	7.29	$-1.541^{***}$ (-4.74)	$-8.199^{**}$ (-2.31)
Quotea	Large caps	5.88	5.15	6.11	5.13	$\begin{array}{c} 0.252 \\ ( \ 1.13) \end{array}$	5.42	4.66	5.69	5.19	$-0.264 \\ (-1.35)$	$\begin{array}{c} 0.517^{***} \\ ( \ 4.30 ) \end{array}$
Darth	Mid caps	9.66	9.63	9.67	9.82	$-0.193^{***}$ (-2.59)	10.11	10.20	11.06	11.14	$\begin{array}{c} 0.015 \\ ( \ 0.31) \end{array}$	$-0.207^{***}$ (-3.25)
$Deptn_{top}$	Large caps	10.79	10.90	10.35	10.50	$-0.039 \ (-0.55)$	11.23	11.28	11.66	11.70	$\begin{array}{c} 0.005 \\ ( \ 0.10) \end{array}$	$-0.045 \\ (-0.83)$

#### Table 7: Price Efficiency and Volatility

This table shows summary statistics and difference-in-difference results for volatility and price efficiency.  $Volatility_t$  is the standard deviation of quote midpoint returns, where the returns are calculated at time horizon t, in basis points. Abs.  $AC_t$  is the absolute first order autocorrelation of quote midpoint returns, where the returns are calculated at time horizon t, in basis points. Abs.  $AC_t$  is the absolute first order autocorrelation of quote midpoint returns, where the returns are calculated at time horizon t, in basis points. Abs.  $VR_{10/30sec}$  is the absolute deviation from unity of the appropriately scaled variance ratio using 10sec and 30sec quote midpoint returns, in percent. The columns denoted DiD show the difference-in-difference coefficients. In the estimations for  $DiD_{Chi-X}$ , the treatment group are German stocks on Chi-X and the control group French stocks on Chi-X. Similarly, for  $DiD_{PM}$  we take data from the two primary listing exchanges Xetra and Euronext.  $DiD_{Diff}$  gives the difference-in-difference estimates on the differences between the primary listing exchange and Chi-X. The before and after periods are August–September 2012 and November–December 2012, respectively. Stocks are sorted into mid and large caps by the medians of market capitalization for the two main listing exchanges during July 2012. Cluster robust t-statistics are given in parentheses. \*\*\*, \*\*, \* denotes significance at the 1%, 5%, 10%–level, respectively.

		<b>Chi-X</b> <sub>0</sub> Before	German After	<b>Chi-X</b> Before	French After	${ m DiD}_{ m Chi-X}$	PM <sub>G</sub> Before	erman After	$\mathbf{PM}_{\mathbf{F}}$ Before	rench After	${ m DiD_{PM}}$	${ m DiD}_{ m Diff}$
Volatility	Mid caps	3.87	1.56	1.30	1.08	$-2.091^{***}$ (-4.34)	1.06	0.88	1.11	0.97	$-0.041 \\ (-0.90)$	$-2.049^{***}$ (-4.34)
Volatility <sub>1sec</sub>	Large caps	1.13	0.93	1.11	0.90	$\begin{array}{c} 0.011 \\ ( \ 0.21) \end{array}$	0.98	0.84	1.02	0.84	$\begin{array}{c} 0.048 \\ (1.29) \end{array}$	$-0.036 \\ (-0.97)$
	Mid caps	15.55	7.65	5.01	4.82	$-7.698^{***}$ (-7.14)	5.08	5.07	5.11	5.51	$-0.402 \\ (-1.00)$	$-7.296^{***}$ (-6.08)
ADS. $AC_{10sec}$	Large caps	4.94	4.57	4.19	4.25	$-0.431 \\ (-1.11)$	4.31	4.56	4.43	4.09	$\begin{array}{c} 0.588^{**} \\ ( \ 2.07) \end{array}$	$-1.018^{***}$ (-2.58)
AL- VD	Mid caps	23.38	12.17	7.18	7.48	$-11.500^{***}$ (-7.14)	7.73	8.33	7.55	8.06	$\begin{array}{c} 0.092 \\ ( \ 0.16) \end{array}$	$-11.592^{***}$ (-6.62)
ADS. V $K_{10/30sec}$	Large caps	7.47	7.45	6.33	6.57	$-0.253 \\ (-0.45)$	6.49	7.06	6.66	6.56	$\begin{array}{c} 0.673 \\ ( \ 1.56 ) \end{array}$	$-0.926^{*} \\ (-1.90)$

Figure 1: Geography of Key European Equity Markets

This graph shows the locations of four major European exchanges: London Stock Exchange (London), BATS and Chi-X (Slough), Euronext (Basildon), and Xetra (Frankfurt).





These graphs show the market shares of daily total trading volume of different exchanges, separate for French and German stocks. Stocks are sorted into mid and large caps by the medians of market capitalization for the two main listing exchanges during July 2012.



62



#### Figure 3: Stock-specific vs. Market-wide Information

These graphs show coefficients from a regression of returns of German stocks at Chi-X on lagged returns

of the same stock at Xetra and returns of DAX index futures traded at Eurex. Lags are given in milliseconds. Panels a and b show the coefficient values while panels c and d show the cumulative sum of the coefficients up to a given lag. Stocks are sorted into mid and large caps by the medians of market capitalization for the two main listing exchanges during July 2012. (a) Mid Caps
(b) Large Caps

63

# Appendix

In this appendix we report additional summary statistics. We then show the distribution of the duration of arbitrage opportunities.

We first present summary statistics for the variables used in the analyses in Section 5 and Section 6 separately for the primary markets (Xetra and Euronext) and for Chi-X. Table A1 and Table A2 contain the summary statistics for mid-cap and large-cap stocks, respectively. Panel A (B) of the tables reports the statistics for the trading activity, liquidity, and price efficiency variables for Chi-X (the primary markets), and Panel C reports the statistics for the arbitrage variables and information shares.

Variables measuring trading and quoting activity such as #Trades, Tr.Volume, and #MessagesBBO are largely comparable across the primary markets for mid and large-cap stocks. The same holds for quoting and trading liquidity (for example,  $Pr.Impact_{10s}$ , Effective, and Quoted) for large cap stocks. However, liquidity for French mid-cap stocks is notably higher compared to German stocks on both the primary market and on Chi-X. This is evident from the lower quoted and effective spreads as well as lower price impacts. With the exception of mid-cap stocks on Chi-X, price efficiency on the primary markets and Chi-X is largely similar for French and German stocks. For mid cap stocks, prices on Chi-X more closely resemble a random walk for French stocks compared to German stocks. Cross-market arbitrage strategies between the primary market and Chi-X seem more profitable for German mid cap and large-cap stocks compared to similarly-sized French stocks. This is evident from the higher frequency of arbitrage opportunities (Arb.Number)

and the gross arbitrage profits  $(Arb._{Value})$ . Consistent with the geographical location of Xetra, Euronext, and Chi-X, the duration of arbitrage opportunities is shorter for French stocks compared to German stocks. Finally, the higher information share of Xetra compared to Euronext for mid and large cap stocks suggests that Chi-X's contribution to price discovery is higher for French stocks compared to German stocks.

Figure A1 shows the distribution of the duration of potential arbitrage opportunities. For German stocks, there is generally a decrease in the number of potential arbitrage opportunities lasting 5 to 6 milliseconds and an increase in shorter-lived arbitrage opportunities. For French stocks, there are no strong shifts in the distribution of the duration of arbitrage opportunities.

	(	ermar	ı stock	s	French stocks					
	Mean	P5	P50	P95	Mean	P5	P50	P95		
Panel A: Chi-X										
#Trades	6.7	4.6	6.8	8.1	7.3	6.1	7.3	8.4		
Trade $Size_{P50}$	7.6	7.2	7.6	8.1	7.7	7.0	7.7	8.1		
Tr. Volume	14.5	12.2	14.6	16.3	15.2	13.9	15.2	16.5		
#Messages BBO	9.2	8.1	9.2	10.3	9.5	8.3	9.5	10.6		
#Messages notBBO	9.8	8.2	9.8	11.3	10.3	8.8	10.5	11.9		
% Trade Full Level	63.6	50.4	64.3	74.0	66.3	52.3	67.5	76.8		
Pr. Impact <sub>10s</sub>	13.7	4.7	9.0	50.1	7.5	3.6	7.0	13.1		
Realized <sub>10s</sub>	-0.0	-6.9	-0.5	8.2	-1.4	-4.4	-1.4	1.4		
Effective	13.8	4.3	8.5	52.2	6.1	3.3	5.5	11.6		
Quoted	24.9	6.1	13.5	109.6	8.6	4.6	8.0	15.6		
$Depth_{top}$	9.7	9.1	9.6	10.4	9.7	9.1	9.7	10.3		
Volatility <sub>1soc</sub>	3.9	0.8	2.8	11.7	1.3	0.5	1.0	2.9		
Abs. AC10sec	15.5	0.6	15.1	37.9	5.0	0.3	3.8	13.2		
Abs VB <sub>10</sub> (ac	23.4	1.0	20.6	57.0	7.2	0.0	49	20.1		
The second second		110	2010			0.1	110			
Panel B: Primary N	1arket									
#Trades	6.9	5.5	6.9	8.1	7.9	6.7	7.9	9.0		
Trade Size <sub>P50</sub>	8.1	7.5	8.1	8.7	8.0	7.4	8.0	8.5		
Tr. Volume	15.4	13.8	15.4	17.2	16.2	14.8	16.3	17.6		
#Messages BBO	9.2	8.2	9.3	10.3	10.6	8.5	10.8	11.8		
#Messages notBBO	10.2	9.0	10.1	11.7	11.3	9.7	11.4	12.6		
% Trade Full Level	58.9	46.9	59.4	69.2	42.8	28.5	43.7	53.4		
Pr. Impact <sub>10s</sub>	9.9	5.0	8.8	19.3	6.8	3.8	6.4	11.3		
Realized <sub>10s</sub>	-0.0	-4.1	-0.4	5.6	-0.7	-3.3	-0.8	2.4		
Effective	9.9	4.6	8.7	21.0	6.1	3.4	5.5	12.0		
Quoted	11.8	5.5	10.5	21.4	7.6	4.2	6.8	13.7		
$Depth_{top}$	10.1	9.5	10.1	10.8	11.1	10.2	11.1	11.7		
Volatility <sub>1sec</sub>	1.1	0.5	0.9	2.1	1.1	0.4	0.9	2.4		
Abs. AC10sec	5.1	0.4	4.1	13.0	5.1	0.4	4.1	13.0		
Abs. VR <sub>10/30sec</sub>	7.7	0.5	5.6	22.0	7.6	0.5	5.6	20.5		
Panel C: Cross mar	ket meas	ures								
A 1	05.0	0.0	10.0	00.1	0.0	0.0	0.4	40.0		
Arp. Value	25.9	0.0	10.8	88.1	9.0	0.0	2.4	42.0		
Arb. Number	5.9	0.0	3.0	19.0	2.5	0.0	1.0	10.0		
() rb	4.1	2.0	4.0	6.0	2.1	1.0	2.0	4.0		
AID. Duration										
Info Share PM	86.1	58.5	91.1	98.9	63.1	44.6	62.8	83.0		
Info Share $_{PM}$ Info Share $_{PM}$ <sup>Lower</sup>	$     86.1 \\     86.0 $	$58.5 \\ 58.3$	$91.1 \\ 91.0$	$98.9 \\98.8$	$63.1 \\ 60.6$	$\begin{array}{c} 44.6\\ 41.0 \end{array}$	$\begin{array}{c} 62.8\\ 60.4 \end{array}$	$83.0 \\ 81.5$		

Table A1: Summary Statistics for Mid-Cap Stocks

This table shows summary statistics for the variables used in the analyses for the mid-cap stocks in the sample based on the pre-treatment period of August and September 2012. The variables are defined as in the main paper.

	6	Fermar	ı stock	s	French stocks					
	Mean	P5	P50	P95	Mean	P5	P50	P95		
Panel A: Chi-X										
#Trades	8.3	7.4	8.3	9.5	8.1	6.8	8.1	9.5		
Trade Size <sub>P50</sub>	8.3	7.8	8.2	8.8	8.0	7.5	8.0	8.6		
Tr. Volume	16.9	15.7	16.9	18.2	16.5	14.8	16.5	17.9		
#Messages BBO	10.5	9.7	10.5	11.5	10.3	9.0	10.3	11.3		
#Messages notBBO	11.4	10.4	11.4	12.5	11.2	9.5	11.3	12.6		
% Trade Full Level	62.9	47.5	64.0	74.7	64.5	49.8	65.8	76.3		
Pr. Impact <sub>10s</sub>	4.9	2.9	4.5	8.3	5.6	3.1	5.4	8.9		
$Realized_{10s}$	-0.6	-1.8	-0.6	0.6	-1.1	-2.8	-1.0	0.5		
Effective	4.3	2.4	4.1	7.2	4.5	2.4	4.4	7.0		
Quoted	5.9	2.9	5.6	9.8	6.1	3.0	6.1	9.7		
$\mathrm{Depth_{top}}$	10.8	9.8	10.7	11.9	10.3	9.6	10.2	11.4		
Volatility <sub>1sec</sub>	1.1	0.6	0.9	2.3	1.1	0.6	0.9	2.2		
Abs. AC10000	4.9	0.3	3.4	14.1	4.2	0.2	3.2	11.4		
Abs. $VR_{10/20coo}$	7.5	0.4	5.8	21.0	6.3	0.4	4.8	17.0		
Panel B: Primary M	Iarket									
//Thur days	0.0	7.0	0.1	0.0	0 5	7.0	0.0	0.7		
# Trades	0.2	1.2	0.1	9.2	0.0	1.2	0.0	9.7		
Trade Size <sub>P50</sub>	9.1 17.7	0.0 16 4	9.1	9.0	0.0	0.0	0.0	100		
II. voluille	105	10.4	10.5	19.0	11.4	10.7	11.0	10.0		
#Messages DDO	10.5	9.7	10.5	11.4	11.2	10.2	11.0	12.1		
# Messages Hot DDO	61.0	10.7	62.8	12.0	11.0	21 5	11.9	52.0		
70 Hade Full Level	01.0	44.7	02.0	11.0	40.2	31.5	43.0	00.0		
$Pr. Impact_{10s}$	5.2	3.0	4.9	8.9	5.4	3.0	5.3	8.3		
$Realized_{10s}$	-0.4	-1.7	-0.4	1.1	-0.7	-2.4	-0.7	1.1		
Effective	4.8	2.8	4.5	7.8	4.7	2.6	4.5	7.4		
Quoted	5.4	3.0	5.3	8.0	5.7	3.1	5.6	8.8		
$Depth_{top}$	11.2	10.3	11.2	12.3	11.7	11.0	11.6	12.5		
$Volatility_{1sec}$	1.0	0.5	0.8	2.0	1.0	0.5	0.9	2.0		
Abs. $AC_{10sec}$	4.3	0.3	3.5	11.3	4.4	0.3	3.5	11.5		
Abs. $VR_{10/30sec}$	6.5	0.5	5.1	18.3	6.7	0.4	5.1	17.9		
Panel C: Cross mar	ket meas	ures								
Arb. Value	97.3	4.3	67.3	325.1	29.0	0.0	11.3	117.1		
Arb. Number	20.2	1.0	13.0	68.0	6.8	0.0	3.0	29.0		
Arb. Duration	4.1	3.0	4.1	5.3	2.2	1.0	2.0	4.0		
Info Chano	647	16.6	69 7	86.0	EE O	20 5	E 4 4	76.9		
Into Share PM	04.1	40.0	03.1	80.9	55.9 59.7	39.5	54.4	10.3		
Info Share PM	04.1	40.0	03.4	80.0	52.7	30.3 49.0	01.1 57.0	14.4		
Into Share PM Opper	64.8	46.9	64.1	87.1	59.0	43.2	57.8	78.2		

## Table A2: Summary Statistics for Large-Cap Stocks

This table shows summary statistics for the variables used in the analyses for the large-cap stocks in the sample based on the pre-treatment period of August and September 2012. The variables are defined as in the main paper.



These graphs show the distribution of the duration of potential arbitrage opportunities between the primary market and Chi-X for German or French stocks, respectively. The duration is given in milliseconds. Only arbitrage opportunities lasting up to 10ms and with a profitability of at least 1 EUR are included. All panels show the before and after periods separately, where the before and after periods are August-September 2012 and November-December 2012, respectively. The before and after periods are August-September 2012 and November-December 2012, respectively. Stocks are sorted into mid and large caps by the medians of market capitalization for the two main listing exchanges during July 2012.



# A Tale of Two Cities – Inter-Market Latency and Fast-Trader Competition

# **Internet Appendix**

Figure I1 to Figure I13 show the time series plots of the most important trading activity, liquidity, and price efficiency variables of interest. We separately plot the time series evolution for large-cap versus mid-cap stocks on Chi-X and the primary market.

## Figure I1: Number of Trades

These graphs show the development of Number of Trades as defined in the main paper.


### Figure I2: Trade Size<sub>P50</sub>

These graphs show the development of Trade  $\mathrm{Size}_{\mathrm{P50}}$  as defined in the main paper.



### Figure I3: Trading Volume

These graphs show the development of Trading Volume as defined in the main paper.



#### Figure I4: #Messages BBO

These graphs show the development of #Messages BBO as defined in the main paper.



#### Figure I5: #Messages notBBO

These graphs show the development of #Messages not BBO as defined in the main paper.



### Figure I6: Price $Impact_{10s}$

These graphs show the development of Price  $Impact_{10s}$  as defined in the main paper.



### Figure I7: Realized Spreads<sub>10s</sub>

These graphs show the development of Realized  $\rm Spreads_{10s}$  as defined in the main paper.





These graphs show the development of Effective Spread as defined in the main paper.



### Figure I9: Quoted Spread

These graphs show the development of Quoted Spread as defined in the main paper.



## Figure I10: Depth<sub>top</sub>



These graphs show the development of  $\mathrm{Depth}_{\mathrm{top}}$  as defined in the main paper.

### Figure I11: Volatility<sub>1sec</sub>

These graphs show the development of Volatility  $_{\rm 1sec}$  as defined in the main paper.





These graphs show the development of Abs. Autocorrelation  $_{10sec}$  as defined in the main paper.



# Figure I13: Abs. Variance $Ratio_{10/30sec}$

These graphs show the development of Abs. Variance  $\text{Ratio}_{10/30\text{sec}}$  as defined in the main paper.





# **Recent Issues**

No. 429	Christine Laudenbach, Ulrike Malmendier, Alexandra Niessen- Ruenzi	The Long-Lasting Effects of Experiencing Communism on Attitudes Towards Financial Markets
No. 428	Christine Laudenbach, Stephan Siegel	Personal Communication in an Automated World: Evidence from Loan Repayments
No. 427	Gül Demirtaş, Christian Strenger, Tobias Tröger	Sustainability Skills Disclosure for Boards: An Essential Prerequisite for Assessing Sustainability Competence
No. 426	Peter Andre, Ingar Haaland, Christopher Roth, Mirko Wiederholt, Johannes Wohlfart	Narratives about the Macroeconomy
No. 425	Leo Kaas, Georgi Kocharkov, Nicolas Syrichas	Understanding Spatial House Price Dynamics in a Housing Boom
No. 424	Matteo Bagnara	The Economic Value of Cross-Predictability: A Performance-Based Measure
No. 423	Tobias Berg, Florian Heider	Leverage and Risk-Taking in a Dynamic Model
No. 422	Julian Detemple	Thoughts about the Dictator and Trust Game
No. 421	Claes Bäckman, Patrick Moran, Peter van Santen	Mortgage Design, Repayment Schedules, and Household Borrowing
No. 420	Konstantin Bräuer, Andreas Hackethal, Tobin Hanspal, Samuel M. Hartzmark	Educating Investors about Dividends
No. 419	Andreas Hackethal, Philip Schnorpfeil, Michael Weber	Inflation and Trading
No. 418	Ioannis Asimakopoulos, Tobias Tröger	Reform of the CMDI Framework – Driving Off with the Breaks On
No. 417	Thiago Fauvrelle, Max Riedel, Mathias Skrutkowski	Collateral Pledgeability and Asset Manager Portfolio Choices During Redemption Waves
No. 416	Olga Balakina, Claes Bäckman, Anastasiia Parakhoniak	Beyond Connectivity: Stock Market Participation in a Network
No. 415	Monica Billio, Aoife Fitzpatrick, Carmelo Latino, Loriana Pelizzon	Unpacking the ESG Ratings: Does One Size Fit All?

Leibniz Institute for Financial Research SAFE | www.safe-frankfurt.de | info@safe-frankfurt.de