Securing Passive Liquidity: the Impact of Europe's First Asymmetric Speed Bump on Market Quality

Caroline Le Moign[1]

Abstract

Speed bumps are an innovative market design response to high-frequency trading, aiming to safeguard market makers from the possible externalities of latency arbitrage. This study evaluates the impact of an asymmetric speed bump implemented on Eurex in 2019 on market quality. Utilizing a matched database of French equity option transactions, the liquidity changes are assessed through a difference-in-difference model, taking into account possible spillover effects. The delay had significant and positive impact on Eurex liquidity, confirming that the asymmetric speed bump effectively mitigated adverse selection against liquidity providers, but it also improved the liquidity of its competing platform.

Keywords: speed bump, options, HFT, latency arbitrage, market quality

JEL: D47, D82, G14, G15, G18

[1]University Paris 1 Panthéon-Sorbonne (Centre d'Économie de la Sorbonne). email: carolinelemoign@gmail.com. Declarations of interest: none. The views expressed in the article belong solely to the author. I am grateful to Gunther Capelle-Blancard for valuable guidance, and to Olena Havrylchyk, Thomas Renault and the CES seminary, as well as Laurence Daures-Lescourret and Guillaume Vuillemey for helpful discussions on this project.

1. Introduction

Evidence is mounting that current trading speeds on financial markets have reached the limits of increasing returns. The cost of latency arbitrage, i.e. the possibility for fast traders to react to stale quotes, appears as an additional cost paid by slower market participants (Aquilina et al. (2021)). Moreover, the possibility of being outrun may prompt liquidity providers to increase their spreads in order to minimise their losses, or to offer liquidity outside the order book (Biais et al. (2015)).

One way to reduce the externalities due to speed competition is to replace the continuous limit order book by different forms of periodic auctions (Budish et al. (2022),Budish et al. (2015), Farmer J. Doyne (2012)). Another potential market design response to high-frequency trading are speed bumps, which introduce a small delay between the moment at which an order is received by a trading platform and the moment at which it is processed by the platform (e.g., filled, cancelled, or added to the queue of limit orders). This delay can be symmetric (applied to all orders), asymmetric (only for aggressive market orders)[2] or deterministic/randomized. An asymmetric speed bump explicitly targets market participants who practice latency arbitrage, giving the liquidity provider a time advantage to review its price, while being protected by the delay imposed on incoming aggressive orders (Baldauf and Mollner (2020)).[3] Thus, these proposals of alternative market design expect to improve market quality, by encouraging liquidity providers to post more competitive prices on the order

[2] An order is "aggressive" when it triggers a trade. On the contrary, passive orders are orders that cannot be immediately executed.

[3] Note that latency arbitrage can also be prevented by other means. For example, the European platform Aquis does not permit aggressive non-client proprietary trading, i.e. does not allow HFT to take liquidity on their venue. Other options are also discussed in the literature, such as changing the structure of transaction fees to incentivise certain type of traders or transactions. These options include applying fees to cancellation orders, or introducing regular auctions to replace continuous trading. See for example Derchu et al. (2020).

book. However, assessing whether this holds true empirically requires market design changes.

In practice, since the creation of the IEX regulated market in the United States in June 2016, the desire to slow down the markets has also become popular with some trading platforms. In fact, as part of their strategy to attract market-makers, some platforms are willing to experiment with mechanisms that can diminish the comparative advantage of high-frequency traders, or are designed to protect the liquidity offered by market makers, such as speed bumps. Nevertheless, very few platforms have already implemented such a mechanism, and they typically do so on specific market segments, in order to evaluate their effectiveness.

The introduction of the first asymmetrical speed bump in Europe, on the Eurex platform in 2019, the German derivatives exchange, provides an opportunity to estimate if such mechanism can enhance market quality. The proposed mechanism is a small delay for aggressive orders only (3 milliseconds for French equity options), explicitly aiming at protecting market makers through the provision of "Passive Liquidity Protection" (PLP). In a market dominated by off-book volumes, Eurex expects a virtuous circle in which liquidity providers, no longer subject to latency arbitrage, would be encouraged to be more present on-book, thereby improving the overall market liquidity level. This improvement would ultimately enhance the attractiveness of the platform. Initially implemented in June 2019 on French and German equity options as a temporary experiment, Eurex has gradually extended the mechanism to other options in the following years. Since February 2022, it covers all equity and index options traded on the platform, a proof that Eurex was satisfied with the results.[4]

[4]Appendix A: Institutional set-up of the main speed bumps programmes provides further details on the PLP, but also on the current speed bump programmes in place or planned on other markets around the world.

This paper uses the natural experiment of the introduction of this asymmetric speed bump to assess its impact on market quality. It centers on French equity options, which possess the advantage of equal trading presence on both Eurex and its competing platform, Euronext. This paper evaluates the evolution of market quality on each platform and examines the potential spillover of market participants from one venue to another, since the explicit objective of the mechanism is to reduce the involvement of high-frequency traders (HFT). If no changes in the distribution of overall trading volumes between the two platforms are observed after the introduction of the speed bump, a simple proxy assessing the possible presence of HFT, shows an increase in HFT presence on Euronext and a decrease on Eurex after the speed bump, signaling a potential migration of those market participants from one platform to the other.

Using a matched database at transaction level for all the French equity options during the full year of the speed bump introduction, the changes in liquidity on the two platforms are assessed through two difference-in-difference models and five key liquidity indicators: quoted spread, effective spread, realized spread, price impact and market depth. The evolution of latency arbitrage and competition is assessed through traded volumes and a proxy of HFT presence. To assess for possible spillover effects between the two trading platforms, the main regression separates between the options that can be traded on both platforms and those that can only be traded on one. The main results show that implementing the asymmetrical speed bump led to a significant decrease in HFT presence on Eurex, and improved market liquidity on Eurex significantly, but also on the competing platform that did not implement the mechanism.

Those results are in line with the theoretical model of Baldauf and Mollner (2020) for a platform implementing an asymmetric speed bump: by providing an advantage to market-makers, the mechanism reduces adverse selection against the liquidity provider, which leads to their increased participation to the market and smaller spreads. The paper also confirms the empirical assessment of the introduction of the first symmetric speed bump on the IEX market, which improved market liquidity, with a significant decrease in spreads ranging from -1.8 to -2.9 basis points (Hu (2019)). They align with the fragmented market model proposed by Brolley and Cimon (2020), which anticipates enhancements in liquidity and trading volumes at the exchange with the delay, predicting informed investors to shift towards the 'conventional' exchange. However, our empirical findings differ as we also observe an improvement in liquidity at the competing exchange without a delay, contrary to their model's expectations. One plausible interpretation involves heightened trading competition on Euronext, potentially involving new market participants migrating from Eurex, encompassing both liquidity investors and speculators. Additionally, participants from both platforms that were operating off-book, may engage in competition over spreads. Another perspective, consistent with Brolley and Cimon (2020), suggests that the redistribution of delayed trading on each exchange might not be optimal due to an inadequately set delay length.

This paper contributes to the continuing debate on market structure design improvement, especially when taking into account speed (O'Hara (2015)), by showing that asymmetric speed bumps may be beneficial for market quality. The positive impact on market liquidity of the introduction of a small delay for aggressive orders confirms the expectations that *"asymmetric speed bumps change the rules of the game but in a small, surgical way"* (Baldauf and Mollner (2020)) and shows how minor changes in market design can potentially diminish the externalities of speed competition.

The main results are that in the short-term, the mechanism had a significant impact on the overall market liquidity of Eurex, with a significant quoted spread decrease of -15.9%, a decrease in effective spread (-0.6%), and in realized spread (-1.2%). A significant decrease of HFT presence is observed in the short term, as well as a decrease in traded volumes (-5.1%). Analysing the potential spillover effect between the two platforms with an extended regression on the matched sample, the biggest impact of the mechanism appears for options traded on Eurex that can also be traded on the competing platform, with an important decrease in their spreads (-8.5%), an increase of their depth (+17.3%) and a decrease of the presence of HFT on their transactions. For the options only traded on Eurex, the mechanism led to a significant decrease of HFT presence, with no other effect. On the other hand, the mechanism also had a positive and significant effect on the market liquidity of Euronext, with a significant decrease in both the spreads (-7.2%) and an increase of the proposed depth (+23.7%). Finally, a robustness check difference-in-difference comparing options traded only on one platform (Eurex and Euronext only) confirms that the mechanism decreased the presence of HFT on Eurex, even for contracts that were only available on the platform.

The paper proceeds as follows. Sections 2 and 3 describe the literature and the database overview. Sections 4 and 5 present the first impact analysis and the spillover analysis, including the matching process and regression set-up.

2. Literature review

Our article is related to the literature on latency delays, the main reference on asymmetric delays being Baldauf and Mollner (2020), which models the impact of an asymmetric speed bump similar to the one put in place by Eurex. Their model incorporates this market design change as an answer to the drawbacks of the presence of high-frequency trading. It reflects the trade-off observed in markets with increasing trading speed, arguing that the increase in HFT has reached a level where its benefits on improving the information available are diminishing: faster speed allows high-frequency traders to anticipate orders more effectively and reduce the price ranges offered (positive effect) but decreases the incentive to obtain and incorporate information into market prices, as informed traders cannot trade before high-frequency traders, reducing the information contained in prices (negative effect). Thus, the introduction of a speed bump improves market functioning. In their model, the asymmetric speed bump eliminates liquidity providers' expectations of the quote sniping mechanism, allowing them to keep an incentive to change their passive orders as soon as new information arises. Therefore, they can continue contributing to high-frequency market making, keeping the proposed price ranges tight, and the information contained in the prices at the same level. Brolley and Cimon (2020) extend latency delays analysis to a competitive market, assessing the interactions between exchanges with fixed/random delays and those without, providing a reference on the migration of traders between exchanges. In a model of fragmented markets, the delayed exchange sees improved liquidity and an increase in aggregate volumes, since informed speculators migrate to the main exchange, where widening bid–ask spreads are observed. Overall, the impact is a worsening of aggregate investor welfare, driven by the increase in delay costs borne by liquidity investors on the delayed exchange (i.e. costs incurred when they do not complete their orders in the expected time, coming from margin constraint or unmodeled risk aversion). A key component of their model is the adequacy of the length of the delay, by assuming that the probability of protection against sniping of a pre-existing order at the delayed venue is proportional to the length of the delay. Thus, delay calibration appears crucial.

As the speed bumps implemented so far have either been symmetrical or random, the empirical literature does not yet, to the best of our knowledge, offer any conclusions on this type of asymmetrical speed bump. However, Hu (2019) provides an analysis of the implementation of the symmetrical speed bump on the US IEX platform, showing that the mechanism improved market liquidity, by reducing adverse selection and tightening price ranges. By analysing two months of transactions and orders, a decrease in proposed and effective spreads by 1.8 to 2.9 basis points (-3%) is observed. Since then, the continued increase in IEX market share suggests that markets with this type of mechanism remain attractive to participants. This confirms that by handicapping all participants by the same delay, a symmetric deterministic speed bump does not change the relative likelihood of winning latency races for a trader and should therefore have only a limited effect on liquidity and investments in fast trading technologies. The other empirical reference on speed bumps analyses the random one put in place on the Canadian TSX Alpha platform in 2015, which consists of a random delay of 1 to 3 milliseconds for all market participants, except for traders who can pay an additional fee and respect a minimum order size to be immune to the speed bump.[5] Chen et al. (2017) observe that the random protection of liquidity providers leads to continued high-frequency market making, and to an increase in proposed and effective spreads of 0.6 basis points. The effects are most pronounced for stocks infrequently constrained by the minimum tick size, with higher prices, low depth and in which best bid and offer liquidity is frequently exhausted. However, the specificities and the possibilities to not be constrained by such mechanism make any comparison to their assessment difficult.

The implementation of an asymmetric speed bump assumes that the presence of high frequency traders is perceived to have a negative impact on liquidity.[6]

[5] See Appendix A for more details about the mechanism.

[6] However, there is also an important literature showing the positive effects of HFT, for instance showing the positive relationship between HFT and market quality (Hasbrouck and

The externalities of HFT have been widely documented, with the documentation of a link between higher speed and decreasing liquidity levels (Bongaerts et al. (2016)), highlighting their role in "flash crashes" (e.g. Kirilenko et al. (2017)), or their consumption of liquidity during periods of higher volatility (Megarbane et al. (2017)). However, their most important externality, which is the target of speed bumps, resides in "quote sniping", where market participants with a comparative speed advantage race to "overtake" passive orders by buying or selling a security at the stale price, before the liquidity provider is able to make a price adjustment. This type of arbitrage is sometimes referred to as "toxic", because it leads to an increase in the cost of liquidity provision, if liquidity providers include the possibility of being outrun on the bid and ask prices offered (Foucault et al. (2017), Budish et al. (2015)). It can also result in adverse selection for the slower traders, discouraging them from submitting orders, thereby reducing trading overall (Biais et al. (2015), Hoffmann (2014)). Breckenfelder (2019) confirms with an empirical analysis on Swedish data from 2009 to 2010, that competition between HFT increases high-frequency speculation activity by 11%, deteriorates liquidity and increases short-term volatility.[7] Aquilina et al. (2021) shows that these speculative "races" between HFT account on average for 22% of daily FTSE 100 trading volumes, and cost 0.42 basis points per trade. This cost is the additional cost of liquidity paid by market participants who do not practice latency arbitrage. They calculate that a market structure without these races, by a change in market design, would reduce the effective spread, i.e. the cost of liquidity for investors, by 17%.[8]

Saar (2013)), by reducing trade-related price discovery, thus improving liquidity (Hendershott et al. (2011)) and enhancing the informativeness of quotes (Brogaard et al. (2015)). Pagnotta and Philippon (2018) models how very high trading speeds allows the integration of information more quickly into prices and thus promotes market efficiency.

[7] Among the indicators studied, the average bid-ask spread increased by 5%, while the price impact, which measures the execution cost incurred by a transaction, increased by 23%.

[8] Furthermore, the winners and losers of these "races" are overwhelmingly the same HFT firms competing with each other - the gains and losses are therefore highly concentrated. By

However, the implementation of a speed bump does not eliminate the interest in investing in faster systems, since the first agent crossing the speed bump will be able to execute his order. The speed advantage is therefore not eliminated but it is reduced, as Khapko and Zoican (2021) show in their modelling: they estimate the decrease in technological investment to be 20% after the implementation of an asymmetric speed bump; and that increasing the magnitude of this speed bump by one standard deviation reduces low-latency investment further by 8%. However, they observe no change in the investment level in the case of a symmetric speed bump.

3. Database and HFT presence

The database comprises high-frequency data that covers all transactions in single name equity options under French jurisdiction[9] carried out on Eurex (Deutsche Börse group), the German derivatives market, and on MONEP in Paris (Euronext group)[10] and available on Refinitiv. The sample is created for the whole year of 2019, i.e. for 5 months before the implementation of the speed bump and 7 months afterwards. The database provides for each transaction and for the entire scope of the analysis: the price and quantities traded, as well as the prices and quantities available at the best limit. The constructed dataset encompasses the majority of the traded volumes on Eurex and nearly all volumes

separating the first winning and losing orders for each race, the top three firms make up 55% of the wins and 66% of the losses. For the top six companies, the figures were 82% and 87% respectively.

[9] "French equity options" in this analysis refers to equity options under French jurisdiction, i.e. whose most liquid market is located in France (as defined by Article 16 of Commission Delegated Regulation (EU) 2017/590, supplementing Regulation (EU) No 600/2014 with regard to regulatory technical standards for the reporting of transactions to competent authorities).

[10] Since only a very marginal number of these options are traded on other platforms (Euronext Amsterdam, the Italian derivatives market and ICE), these transactions are not taken into account.

traded on Euronext in 2019.[11]

This data underwent cross-referencing with the European Financial Instruments Reference Data System (FIRDS) to gather comprehensive information on these options, including their type, strike price, and expiration. Furthermore, the database was augmented with data on the underlying equities, sourced from the Euronext database provided by the French Market Authority. This supplementary information includes daily traded volumes, transaction counts, market capitalization, and closing prices of shares traded on Euronext. These additions facilitate the computation of intrinsic option value, which denotes the disparity between the option's strike price and the closing price of the underlying share each day, alongside its moneyness.

Table 1 presents the coverage and characteristics of the database, indicating that French equity options are traded[12] almost evenly on both platforms: 49% of volumes are traded on Euronext and 51% on Eurex, and this proportion remains relatively stable throughout the year.The options market is characterised by a vast array of instruments, with over 23,000 different options traded on Euronext and 26,000 on Eurex during the year. However, only a few of these instruments are actively traded on a daily basis, resulting in a low average number of daily transactions: 1,386 on Euronext and 1,588 on Eurex. For comparison, in 2019 on Euronext, French stocks were traded an average of over 323,000 times daily.

In 2019, options contracts traded covered a total of 105 French underlying equities. Among these, 63 underlying shares had option contracts offered on both Eurex and Euronext, 17 shares had options exclusively on Eurex, and

[11] We cannot account for the discrepancy in coverage between Refinitiv's data for the Eurex platform and Euronext. However, we anticipate this disparity in coverage to have no bearing on our results, given that we encompass the vast majority of volumes.

[12] A trade consists of the purchase or sale of one or more option lots: in most cases, one lot gives the right to buy or sell the equivalent of 100 equities. Thus,"volumes" are the number of lots traded and are not multiplied by 100 unless stated otherwise.

25 shares had options solely on Euronext. Figures 5 and 6 provide a visual representation depicting the number of options traded, the underlying shares, and the overlap between the two platforms in the matched sample.

Table 1: Overview of the database

Source	*Indicator*	*Euronext*	*Eurex*
Eurex,Euronext	Volumes traded in 2019 (quantity)	24,839,147	25,325,704
Database	Volumes traded in 2019 (qty)	24,749,541	21,358,436
Database	% volumes in the sample	99.5%	84.3%
Database	Total block volumes traded (qty)	16,553,248	12,642,197
Database	Total on-book traded volumes (qty)	8,196,293	8,716,239
Database	Median daily traded volumes (qty)	97,057	84,421
Database	Average daily traded volumes (qty)	92,033	83,265
Database	Average quantity by block transaction	1,797	1,454
Database	Average quantity by on-book transaction	24	23
Database	Transactions in 2019 (number)	351,218	401,085
Database	Average transactions/day (number)	1,386	1,588
Database	French options traded in 2019 (number)	23,130	25,967
Database	Average options traded per day (number)	569	653
Database	Underlying shares proposed (number)	87	80
Database	Underlying shares only on the platform (number)	25	17

Note: Traded volumes represent the number of French equity options contracts traded and are not multiplied by 100.

All indicators studied in the model solely account for on-book transactions. The options market features a large number of contracts traded outside of the order book, primarily as 'blocks' of traded lots.[13] At Euronext, 67% of option volumes traded in 2019 occurred off-book, within 'blocks' averaging 1,797 option contracts. On Eurex, 55% of volumes were traded off-book, within blocks averaging 1,454 contracts. These transactions remain excluded from the subsequent analysis for two reasons: i) they are unaffected by the speed bump, and ii) block trading differs from on-book trading, with most participants in the off-book options market focusing on dealing in larger quantities rather than prioritizing execution speed. Within the dataset, certain trades are excluded from the analysis. These exclusions encompass trades where the best buy/sell spread

[13]Blocks refer to substantial transactions (exceeding a threshold set by the platform), typically involving numerous contract lots, ranging from 50 to several thousand.

is incompletely reported, transactions with abnormal volumes (exceeding four times the proposed buy and sell depth), and options with a strike price of one cent.[14]

The main indicators on both markets during 2019 do not show any major change before and after the implementation of the speed bump on 3 June 2019, and no announcement effect is expected.[15] Daily traded volumes on both platforms decreased slightly during the second part of the year on both platforms (-8% daily volumes on Eurex and -12% on Euronext), as shown in figure 1, with no change in the distribution of aggregated volumes or open positions in the second half of the year between the two venues. Thus, the evolution of volumes does not seem to be impacted by the speed bump at first glance, with two possible explanations: either the mechanism did not attract any added volumes, or it did change the distribution of market participants and their traded volumes between the two platforms.

Thus, creating a proxy to assess the presence of HFT is crucial to understand the distribution of those market participants before and after the implementation of the speed bump. While numerous proxies are usable (**??**, **??**), our choices are limited by the accessible data required to substantiate them.[16] The proposed indirect proxy relies on the transaction timestamp within the database. We presume that trades occurring within ten microseconds of each other are indicative

[14]The exclusions due to unavailable bid or ask prices affected close to 117,900 transactions, accounting for 15% of the full sample of 795,000 transactions. Removing abnormal volumes and small strike prices impacted 1,375 transactions, representing 0.2% of the overall transaction sample.

[15]The possibility of putting in place such mechanism was announced by Eurex in October 2018, and then discussed through an extensive participant consultation. The final set-up information was announced on 8 March 2019, see Circular 027/19. However, no announcement effect is expected to impact our model, since the anticipations were built for more than 9 months - and the behavior of market participants is expected to switch once the program is activated, as is usually the case for latency arbitrageurs.

[16]Without access to the order book, measures such as quote intensity, order-to-quote, number of cancelled orders or message traffic cannot be developed to gauge HFT presence.

of HFT involvement.[17] Figure 2a shows the evolution of the volumes traded with this proxy as a share of the overall volumes traded on each platform. It shows a higher presence of HFT on Euronext (26% of daily volumes observed at the same microsecond on average before the speed bump) than on Eurex (17%). The fact that the speed bump is proposed by Eurex, which has a smaller share of HFT trading volumes, is consistent with the results from Brolley and Cimon (2020), which shows that only stand-alone exchange which have less revenues to lose from implementing a speed bump have the incentive to do so. After 3 June 2019, an increase in this share is observed on Euronext (35% of the volumes traded during the following 7 months of the year) and a decrease on average on Eurex (16%). This increase (resp. decrease) does not occur in the first weeks after the introduction of the mechanism, but in the following month and remains sustained throughout the end of the year. 21% of the contracts traded on Eurex and 29% of the contracts on Euronext are traded by HFT, and the volumes by instrument show an even distribution cross-section, with no skewed volumes for some contracts (see figure 2b).

If this proxy is limited in its possible assessment of the presence of HFT, it serves as a confirmation of the need to look at possible spillover effects from Eurex to Euronext. The lack of change in the overall traded volumes on-book, as well as of market share change between the two platforms, can reflect an increase of "slow traders" participation on Eurex, that were staying off-book before the speed bump, and a decrease of HFT participation; and on Euronext an important increase of HFT participation replacing "slow traders" (those traders could be HFT already active on Euronext and increasing their volumes

[17] Similar results were observed by analyzing transactions at both the exact microsecond precision and at the second level, revealing comparable patterns. This approach closely aligns with Eurex's public statistics on the proportion of volumes traded aggressively by proprietary clients against passive liquidity. However, differentiating between aggressive and passive fast trading remains unfeasible.

since competition and opportunities increased, or new HFT coming from Eurex, but also from off-book volumes from either platform).

4. Direct impact

The speed bump's implementation on all French equity options traded on Eurex initiates an initial assessment. This assessment aims to evaluate the mechanism's short-term effect by utilizing a difference-in-difference model on French equity options available for trading on both platforms.[18] The analysis compares liquidity variables for all French options traded on Eurex (treatment group) with those traded on Euronext (control group). However, this approach assumes that implementing the mechanism on Eurex did not impact options traded on Euronext (i.e., no spillover effect). To observe potential differentiated effects following the mechanism's introduction, we restrict the time window to 2 months before and after implementation.[19]

We evaluate short-term liquidity dynamics by analyzing the liquidity indicators: quoted spreads, effective spreads, realized spreads, price impact and market depth, alongside traded volumes and HFT presence. Quoted spread refers to the bid-ask spread displayed at the time of trade, calculated by dividing the difference between the ask and bid price by the mid-price. Effective spread, computed at each trade, approximates the actual liquidity cost paid by investors, calculated as the absolute value between the transaction price and the mid-price. In order to estimate the net revenue or loss of liquidity providers, realized spreads and price impact measures are calculated. [20] To gauge HFT

[18] This regression includes only options traded on both platforms, excluding those traded solely on Eurex or Euronext.

[19] Various time windows were considered, and did not yield different results.

[20] See Appendix B: Definition of the variables for the definitions of all the variables of the analysis.

presence, the HFT variable introduced in Section 3 generates a dummy, signalling for each contract if transactions with HFT occurred during the day. Alongside these, several control variables, including price level, maturity, and moneyness of the option contracts, are introduced into the regression. Following Etling and Miller (2000), moneyness is defined as the absolute value of the difference between the option's strike price and the price of its underlying equity at the end of the day. All these variables can influence the contract liquidity level: notably, options trading volumes have exhibited concentration in contracts with shorter maturities and strike prices closer to the underlying equity price (Cho and Engle (1999), Etling and Miller (2000)). Moreover, contracts with higher trading volumes often demonstrate narrower spreads (Mayhew (2002)). Consequently, all these factors are integrated as control variables in the regression analysis.[21]

The regression can be written as follows:

$$Y_{\mathrm{it}} = \alpha_t + \beta_1 Post_t + \beta_2 Speedbump_i + \beta_3 Post_t \times Speedbump_i + \gamma_1 X_{\mathrm{it}} + \varepsilon_{\mathrm{it}}$$

Where Y_{it} is the dependent variables for which the regression is estimated at date t for each contract i, i.e. quoted spread, effective spread, realized spread, price impact, market depth, volumes traded or HFT presence; POST_t is a dummy variable equal to 1 after the implementation of the speed bump; $\mathrm{Speedbump}_i$ a variable equal to 1 if the contract is affected by the speed bump (i.e. traded on Eurex); and X_{it} is the set of control variables, calculated on a daily basis for each contract. The coefficient of interest here is the difference-in-difference coefficient β_3, which captures the isolated effect of the mechanism.

[21] For detailed calculation methods of each indicator, refer to Appendix B: Definition of the variables.

Standard errors are clustered at the day level, and fixed effects are added per date and contract, which capture the systematic effects that affect all options on each day t. To avoid an overcontrolling bias, we also report the results without any controls in Table 7 in Appendix C: Other summary statistics, which shows that the overall results are not different.

Table 2: Direct impact analysis

	Indicators						
	Spread	Eff. spread	Real. spread	Price impact	Depth	Volumes	HFT
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Post*Speedbump	-0.0159^{***}	-0.0060^{***}	-0.0117^{**}	0.0053	-0.0109	-0.0500^{**}	-0.0387^{***}
	(0.0037)	(0.0021)	(0.0050)	(0.0043)	(0.0176)	(0.0236)	(0.0085)
Price level	-0.0005^{***}	-0.0003^{***}	-0.0004^{***}	-0.0002^{**}	0.0178^{***}	-0.0120^{***}	-0.0005^{***}
	(0.0002)	(0.0001)	(0.0001)	(0.0001)	(0.0013)	(0.0010)	(0.0002)
Maturity	-0.0003^{***}	-0.0002^{***}	-0.0001^{***}	-0.0001^{***}	0.0013^{***}	-0.0002^{***}	-0.0001^{***}
	(0.00001)	(0.00001)	(0.00001)	(0.00001)	(0.00004)	(0.0001)	(0.00001)
Moneyness	-0.3807^{***}	-0.2374^{***}	-0.1179^{***}	-0.0921^{***}	1.9976^{***}	0.1627^{***}	-0.0175
	(0.0152)	(0.0098)	(0.0110)	(0.0108)	(0.0600)	(0.0505)	(0.0120)
Fixed effects	Date	Date	Date	Date	Date	Date	Date
Observations	94,360	94,360	39,972	39,972	94,360	94,360	94,360
R^2	0.1711	0.0954	0.0111	0.0081	0.3682	0.0223	0.0232
Adjusted R^2	0.1703	0.0945	0.0089	0.0059	0.3676	0.0214	0.0223

Note: Standard errors are clustered at the date and contract level. Spread, effective, realized spread and price impact in percent, depth and volumes are transformed in log. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

The results in Table 2 show that the mechanism had an important impact on Eurex market liquidity, with a significant quoted spread decrease of -15.9%, a decrease in effective spread (-0.6%), and in realized spread (-1.2%). A significant decrease of HFT presence is observed in the short term, as well as a decrease in traded volumes (-5.1%). The mechanism did not yield a significant impact on price impact nor market depth. These results confirm the anticipated effects of the mechanism, demonstrating a short-term enhancement in Eurex's liquidity. This improvement could be attributed to the diminished adverse selection, with lower volumes and a decrease in HFT presence, that led to an improvement of the overall platform's liquidity. Moreover, these outcomes align with the expectations presented in Brolley and Cimon (2020), which anticipate enhanced liquidity and increased volumes on the delayed exchange, primarily driven by less latency-sensitive investors who might otherwise conduct off-book trading.

5. Spillover analysis

We now test whether the mechanism had an effect on the characteristics of trading on Eurex, but also on its competing platform, Euronext, expecting interconnectedness between the two competing platforms to have an effect on the observed impacts. The primary hypotheses we aim to test through this analysis are:

- Differentiated impact of the speed bump on liquidity: Following Baldauf and Mollner (2020), we hypothesize that the introduction of the speed bump mechanism on Eurex has an impact on the liquidity dynamics on both Eurex and Euronext. Specifically, we anticipate changes in bid-ask spreads, trading volumes, and depth following the implementation of the speed bump, with differential effects between options contracts traded exclusively on Eurex and those exclusively on Euronext. On Eurex, we expect an increase in the liquidity offered by market makers on the platform, resulting in a reduction in spreads and/or an increase in available depth and/or volumes traded. On Euronext, the expected effects and their effects on the liquidity offered and on volumes traded are difficult to predict. If the Eurex speed bump causes the increase in volumes traded and the competition for arbitrage opportunities can have beneficial effects, by decreasing spreads, and/or increasing market depth on Euronext. However, the externalities of latency arbitrage can also lead to adverse selection of marker-makers to protect themselves against the newcomers, which would then increase spreads and/or decrease the average depth available.

- HFT presence: We also expect the speed bump to influence the attractiveness of both Eurex and Euronext for HFT. This hypothesis is supported by the observed fluctuations in HFT presence on both platforms post-implementation. Since the volumes are equally distributed between the two venues and both a decrease on Eurex and an increase on Euronext of the presence of HFT is observed (figure 2b), indicating a potential shift in the preferences or strategies of HFT firms due to the new trading mechanism. This hypothesis follows Francioni and Gomber (2017) that gives insight into HFT strategies and their responsiveness to market structure modifications. The expected effect of implementing the speed bump is a reduction in the presence of HFT on Eurex, with and increase in HFT activity on the Paris-based platform.

- Spillover effect: Drawing from Biais et al. (2015), we anticipate that the coexistence of a “slow-friendly market” (i.e. which limits the ability of fast traders to adversely select slower ones) alongside a fast market will result in spillover effects. We hypothesize that changes in liquidity observed on one platform due to the speed bump might spill over and affect liquidity on the other platform, indicating interconnectedness in market behavior. Over time, market participants might shift between platforms due to this dynamic. Biais et al. (2015) suggests an equilibrium wherein slower institutions migrate to slower markets where there is no adverse selection. The choice of settling in either the slow or fast market depends on the threshold of speed technological costs for these institutions.

Thus, to measure the impact of Eurex implementation of the speed bump on changes in liquidity in the French equity options market, we add to the direct impact analysis a difference-in-difference analysis that can take into account possible carry-over or spillover effects. Thus, we propose an additional regression differentiating options contracts that are traded on both platforms or only on one. This differentiation enables us to observe the impact of the speed bump on the liquidity on Euronext as well. To make sure both samples of options traded on each platform are comparable, the database is matched with the most appropriate matching technique. In this set-up, to be able to assess potential spillover effects, the timeframe of observations is extended to 3 months before and after the introduction of the mechanism.[22] To confirms the observed results, a final robustness check compares the options that are only traded on Eurex with the options only traded on Euronext, meaning without possible spillover effect. As a final hypothesis, we expect the observed changes in liquidity metrics to hold, confirming the validity of the findings while isolating potential confounding factors or spillover effects.

5.1. Matching Process

If the initial sample is adequately balanced, employing sample matching serves as an extra layer of precaution. As anticipated, only a minimal number of options are removed through the matching procedure (980, from 6 smaller companies for which options are available on one platform only). Nevertheless, this process ensures that the final sample comprises options that possess comparable attributes across both markets.

In order to proceed with the matching process, we compare the nearest neighbor matching (NNM) and coarsened exact matching (CEM) techniques to evaluate their impact on the sample.[23] Evaluation of both methods includes

[22]This time window captures the change in HFT presence observed in figure 2a that begins a month after the speed bump, while not taking into account other effects observed in the data, such as the increase in depth observed on both venues beginning in November (4e) that is not linked to the mechanism. However, various time windows (2, 3, 4, 5 months) have been tested without changing the results.

[23]Nearest neighbor matching selects the closest counterpart for each observation, using a weighted function of defined covariates, while coarsened exact matching performs exact match-

(i) assessing result balance by comparing standardized mean differences and variance ratios before and after matching, and (ii) examining the effect of each method on the sample size. Figure 3 in Appendix C: Other summary statistics, illustrates the standardized mean differences between treatment and control groups after conducting CEM and NNM matching in the database. Following Iacus et al. (2012), CEM with replacement is chosen for the matching procedure.[24]

Table 3: Description of the sample after matching

Indicator	No. options	p5	mean	median	p95	sd
Spread	44,317	0.04	0.16	0.08	0.19	0.21
Effective spread	44,317	0.03	0.10	0.05	0.11	0.15
Realized spread	26,998	0.00	0.05	0.03	0.08	0.25
Price impact	26,998	−0.01	0.03	0.01	0.06	0.24
Depth	44,317	23.37	175.86	64.26	170.12	492.32
Volumes	44,317	8.33	50.09	21.00	52.44	113.04
HFT	44,317	0.0	0.11	0.0	0.13	0.21
Price level	44,317	0.48	3.27	1.18	2.88	8.16
Moneyness	44,317	−0.07	−0.03	−0.02	0.02	0.17
Underlying - volumes	44,317	643	2,433	1,408	3,801	2,484
Underlying - capitalisation	44,317	11,173	39,246	21,431	53,084	40,481
Underlying - volatility	44,317	0.35	1.02	0.70	1.30	1.10

Note: Spread, effective, realized spreads and price impact in percent, market depth in euros and volumes in quantity of contracts; HFT presence as the share of when the HFT proxy is observed; Underlying capitalisation in millions of euros; underlying volumes in thousands; p5/p95: percentiles; sd: standard deviation.

For the matching procedure, the variables considered rely on the option type (call or put) and the characteristics of the underlying option's share, since the level of liquidity of an option can be influenced by the characteristics of its underlying asset. We use here the market capitalisation, trading volumes and volatility of the underlying share as main characteristics. Thus, the chosen

ing on a set of covariates, grouping continuous covariates into strata and discarding strata lacking both treated and control observations. This method allows for multiple control observations to be matched to a single treated observation, with weights addressing any potential imbalance of observations.

[24] The CEM procedure involves grouping or coarsening options by variables, ensuring balance between options available for trading solely on one platform (Eurex or Euronext) and options available for trading on both platforms.

covariates are not affected by the treatment. Table 5 in Appendix C shows the differences between the treatment and control groups, before and after the matching, and Table 3 below presents a description of the main variables of the sample, after the matching process. After matching, all standardized mean differences for the covariates were below the standard of 0.1 and all standardized mean differences for squares and two-way interactions between covariates were below 0.15, indicating adequate balance (see figure 3 in Appendix C). In the end, the matching process takes out 1.5% of daily observations and the number of options in the matched database is 44,516 (45,597 before). Table 6 in Appendix C details the number of observations, options and underlying shares on both platforms before and after the matching process.

Finally, to fulfil the parallel trends assumption needed for the difference-in-difference analysis, figures 4 in Appendix C present the graphical evolution of the main dependent variables weighted averages for both places after the matching process, and show that the distance between the treatment and control group remains constant over time. The trend graphs confirm that both groups have comparable levels of spreads, effective spreads, market depth and volumes traded before the introduction of the speed bump on 3 June 2019.

5.2. Spillover analysis - Regression set-up

The difference-in-difference model can be written as follows, separating between the places where the options are traded:

$$Y_{\text{it}} = \alpha_t + \beta_1 Eurex\_only_i + \beta_2 Eurex\_both_i + \beta_3 Euronext\_only_i + \beta_4 Euronext\_both_i$$

$$+ c_1 Eurex\_only_i \times \text{Post}_t + c_2 Eurex\_both_i \times \text{Post}_t + c_4 Euronext\_both_i \times \text{Post}_t$$

$$+ \delta_1 \text{Post}_t + \gamma_1 X_{\text{it}} + \varepsilon_{\text{it}}$$

where Y_{it} is the dependent variables for which the regression is estimated at date t for each contract i, i.e. quoted spread, effective spread, market depth or volumes traded, Post_t is a dummy variable equal to 1 after implementation of the speed bump; and X_{it} is the set of control variables, calculated on a daily basis for each contract.

In order to separate contracts by their potential trading platform, $\_only_i$ and $\_both_i$ indicate those characteristics. In fact, if the contracts on Eurex and Euronext are not fungible, most of them have the same underlying shares (62 common underlying shares) and the same features. Thus, a market participant can choose to trade either on Eurex or on Euronext platform for those contracts. $Eurex\_both_i$ are the variables observed on Eurex for contracts traded on Eurex, but for contracts that can be traded on both platforms. Similarly $Euronext\_both_i$ concern the variables for contracts that can be traded on both platforms but are traded on Euronext. Only a limited number of contracts are proposed on one platform only: those options have underlyings of specific French stocks popular on the French market (options only proposed on Euronext) or on the German one (only on Eurex). An option is considered to be traded only on Euronext when no contract with the same underlying share is traded on Eurex (and vice versa for options available only on Eurex). $Eurex\_only_i$ and $Euronext\_only_i$ are the variables observed for contracts that are respectively only traded on Eurex, and similarly the variables observed on contracts traded only on Euronext. The options only traded on Euronext are used as a benchmark in the model since they are the only ones for which no spillover effect can be expected. Figures 5 and 6 clarify graphically the perimeter of each notation for all the options in the data set.

Similarly to the previous regression, standard errors are clustered at the date and contract level, and fixed effects are added per date, which captures the systematic effects that affect all matched options on each day t. To avoid an overcontrolling bias, we also report the results without any controls in Table 8 in Appendix C, which shows that the overall results are not different. Furthermore, other time windows have been tested (restricting the period pre- and post-speed bump), without an observable change in the results.

An improvement in the liquidity offered following the introduction of speed bump will be marked in our set-up by an increase in the quantity offered at the best limits with a constant (or decreasing) spread or by a decrease in the spread with a constant (or increasing) market depth. However, if the presence of adverse selection persists, market-makers may choose to offer less liquidity, increasing spreads, decreasing market depth or even withdrawing from the order book to avoid being "overtaken" in speculative runs. Similarly, if adverse selection persists, volumes should decrease significantly as a sign of less market participation.

As a robustness check of the set-up, a simple difference-in-difference regression model is comparing contracts that are only offered for trading on Euronext with those offered exclusively for trading on Eurex, meaning without possible spillover between the two platforms, with the same window and set-up.

5.3. Spillover analysis - Results

Table 4 presents the results, showing that implementing the speed bump had a bigger impact on options that were available for trading on the two competing platforms. The biggest impact of the mechanism appears for options traded on Eurex that can be traded on Euronext (*Eurex both*), with a significant decrease of their spreads (-8.5%), an increase of their market depth (+17.3%), and a decrease of the presence of HFT on their transactions, but no significant effect on effective spreads or its composites or on volumes. The effects for the options

Table 4: Speed bump impact on French equity options, spillover analysis

	Indicators						
	Spread	Eff. spread	Real. spread	Price impact	Depth	Volumes	HFT
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Eurex only*Post	−0.0445	0.0176	0.2568	−0.2517	0.0312	0.0282	−0.0643**
	(0.0349)	(0.0242)	(0.2353)	(0.2287)	(0.0594)	(0.0851)	(0.0263)
Eurex both*Post	−0.0845***	−0.0230	0.2499	−0.2629	0.1597***	0.0318	−0.0729***
	(0.0315)	(0.0193)	(0.2430)	(0.2397)	(0.0434)	(0.0696)	(0.0239)
Euronext both*Post	−0.0723**	−0.0171	0.2598	−0.2657	0.2139***	0.0753	−0.0135
	(0.0318)	(0.0194)	(0.2430)	(0.2401)	(0.0459)	(0.0675)	(0.0236)
Price level	−0.0009***	−0.0005***	−0.0004***	−0.0003***	0.0214***	−0.0171***	−0.0010***
	(0.0002)	(0.0001)	(0.0001)	(0.0001)	(0.0013)	(0.0011)	(0.0002)
Maturity	−0.0003***	−0.0001***	−0.0001***	−0.0001***	0.0013***	−0.0001**	−0.0001***
	(0.00001)	(0.00001)	(0.00001)	(0.00001)	(0.00003)	(0.00005)	(0.00001)
Moneyness	−0.3694***	−0.2286***	−0.1179***	−0.0824***	1.9692***	0.2265***	−0.0039
	(0.0130)	(0.0081)	(0.0085)	(0.0078)	(0.0463)	(0.0435)	(0.0100)
Fixed effects	Date	Date	Date	Date	Date	Date	Date
Observations	141,599	141,599	58,908	58,908	141,599	141,599	141,599
Adjusted R^2	0.1677	0.1006	0.0089	0.0083	0.3812	0.0256	0.0295

Note: Standard errors are clustered at the date and contract level. Spread, effective, realized spread and price impact in percent, depth and volumes are transformed in log. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

traded only on Eurex is smaller (*Eurex only*), observing a significant decrease of the presence of HFT but not other significant effect. On Eurex, this confirms that HFT presence decreased significantly after the introduction of the mechanism, for all options traded. The results confirm the expectations of Baldauf and Mollner (2020), with a reduction of the spread and an improvement of the liquidity of the instruments under an asymmetric speed bump.

The mechanism had also a positive and significant effect on market liquidity for the options that are traded on Euronext but can be traded on the competing platform, with both a decrease of their spreads (-7.2%) and an increase of their proposed depth (+23.7%). However, no effect was observed on effective spread, volumes or HFT presence. This improvement in market liquidity differs from the expected results of Brolley and Cimon (2020), where liquidity at the competing exchange deteriorates after the delay implementation. Due to the nonsignificant evolution of HFT presence on Euronext, improved liquidity can be attributed to increased trading competition, with new market participants either coming from Eurex liquidity investors and speculators, as well as from either platforms' off-book. It can also mean that the length of the delay is inadequate to fully change the segmentation of informed trading on each exchange (Brolley and Cimon (2020)).

Finally, the results of the difference-in-difference analysis comparing options traded on Eurex only (treatment) and those traded on Euronext only (control) from the matched sample allow for a check of the effect of the speed bump without possible spillover effects for the selected options.[25] However, the number of options is limited and the Table 9 in Annex C shows that for those instruments the mechanism had an impact only on Eurex HFT presence, which decreased significantly. These results confirm that even for the options that are traded

[25] In this set-up, only $Eurex\_only_i$ and $Euronext\_only_i$ options are kept in the regression.

only on Eurex, which are less liquid by nature, the mechanism reduced the presence of liquidity speculators, even though it did not impact their liquidity.

6. Conclusion

With more and more concern about the potential externalities resulting from HFT latency arbitrage, the introduction for the first time on a European market of an asymmetric delay provides evidence of a small market design change with a large impact. Although targeting aggressive orders with a minor delay of 3 millisecond, this asymmetric speed bump resulted in a significant improvement in liquidity, not only on the platform where it was implemented, but also on its competing platform.

Using a matched database of all French equity option transactions during the year of the speed bump introduction, the changes in liquidity are assessed through a difference-in-difference and an extended model allowing to analyse possible spillover effects between the two platforms trading those options. In the short-term, the direct impact of the mechanism shows an increase in liquidity and a decrease in volumes traded on Eurex. Taking into account the possible redistribution of market participants after the implementation of the mechanism, the results show that implementing a speed bump had a bigger impact on options that were available for trading on the two competing platforms. On Eurex, the mechanism improved liquidity and decreased HFT presence on the venue, confirming the expected reduction in adverse selection against the liquidity providers. But the mechanism also improved Euronext market quality, potentially coming from increased trading competition coming from new market participants, but also from an inadequate delay length, that is not optimally distributing informed trading on each exchange (Brolley and Cimon (2020)).

Those results provide for the first time an empirical assessment of the theoretical models of Baldauf and Mollner (2020) and Brolley and Cimon (2020) for a platform implementing an asymmetric speed bump: by providing an advantage to market-makers, the speed bump reduces adverse selection against the liquidity providers, which leads to an increased participation of those market participants and smaller spreads. The paper contributes to the continuing debate on market structure design, by showing that the modest change of the asymmetric speed bumps may be beneficial for market quality, while alleviating the costs of latency arbitrage. With the extension of the speed bump experiment to many other options traded on Eurex in 2022, further analysis of the evolution of market quality, attractivity and the behaviour of market participants might be warranted.

References

Aquilina, M., Budish, E.B., O'Neill, P., 2021. Quantifying the High-Frequency Trading. BIS Working Papers URL: https://www.bis.org/publ/work955.htm, doi:10.2139/ssrn.3636323.

Baldauf, M., Mollner, J., 2020. High-Frequency Trading and Market Performance. The Journal of Finance 75, 1495–1526. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.12882, doi:10.1111/jofi.12882.

Benos, E., Huang, W., Menkveld, A.J., Vasios, M., 2019. The Cost of Clearing Fragmentation. BIS Working Papers URL: https://www.bis.org/publ/work826.htm.

Biais, B., Foucault, T., Moinas, S., 2015. Equilibrium fast trading. Journal of Financial Economics 116, 292–313. URL: https://www.sciencedirect.

com/science/article/pii/S0304405X15000288, doi:10.1016/j.jfineco.2015.03.004.

Bongaerts, D., Kong, L., Van Achter, M., 2016. Trading Speed Competition: Can the Arms Race Go Too Far? Working Paper URL: https://papers.ssrn.com/abstract=2779904, doi:10.2139/ssrn.2779904.

Breckenfelder, J., 2019. Competition Among High-Frequency Traders, and Market Quality. ECB Working Paper series doi:10.2139/ssrn.3402867.

Brogaard, J., Hagströmer, B., Nordén, L., Riordan, R., 2015. Trading Fast and Slow: Colocation and Liquidity. The Review of Financial Studies 28, 3407–3443. URL: https://doi.org/10.1093/rfs/hhv045, doi:10.1093/rfs/hhv045.

Brolley, M., Cimon, D.A., 2020. Order-Flow Segmentation, Liquidity, and Price Discovery: The Role of Latency Delays. Journal of Financial and Quantitative Analysis 55, 2555–2587. URL: https://www.cambridge.org/core/journals/journal-of-financial-and-quantitative-analysis/article/abs/orderflow-segmentation-liquidity-and-price-discovery-the-role-of-latency-delays/524F0B00D8E9BAC8723956A6E2187E57#, doi:10.1017/S002210901900067X.

Budish, E., Cramton, P., Shim, J., 2015. The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. The Quarterly Journal of Economics 130, 1547–1621. URL: https://doi.org/10.1093/qje/qjv027, doi:10.1093/qje/qjv027.

Budish, E.B., Cramton, P., Kyle, A.S., Lee, J., 2022. Flow Trading. Working Paper URL: https://papers.ssrn.com/abstract=4145013, doi:10.2139/ssrn.4145013.

Chen, H., Foley, S., Goldstein, M.A., Ruf, T., 2017. The Value of a Millisecond: Harnessing Information in Fast, Fragmented Markets. Working Paper URL: https://papers.ssrn.com/abstract=2860359, doi:10.2139/ssrn.2860359.

Cho, Y.H., Engle, R.F., 1999. Modeling the Impacts of Market Activity on Bid-Ask Spreads in the Option Market. National Bureau Of Economic Research URL: https://papers.ssrn.com/abstract=196372.

Derchu, J., Guillot, P., Mastrolia, T., Rosenbaum, M., 2020. AHEAD: Ad Hoc Electronic Auction Design. Technical Report. Autorité des Marchés Financiers. URL: https://www.amf-france.org/en/news-publications/publications/reports-research-and-analysis/study-new-market-model-ahead, doi:10.2139/ssrn.3705514.

Etling, C., Miller, Jr., T.W., 2000. The relationship between index option moneyness and relative liquidity. Journal of Futures Markets 20, 971–987. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/1096-9934%28200011%2920%3A10%3C971%3A%3AAID-FUT5%3E3.0.CO%3B2-8, doi:10.1002/1096-9934(200011)20:10<971::AID-FUT5>3.0.CO;2-8.

Farmer J. Doyne, S.S., 2012. Review of the benefits of a continuous market vs. randomised stop auctions and of alternative Priority Rules. Technical Report. Great Britain, Government Office for Science. URL: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/289050/12-1072-eia11-continuous-market-vs-randomised-stop-auctions.pdf. impact Assessment EIA11, 2012.

Foucault, T., Kozhan, R., Tham, W.W., 2017. Toxic Arbitrage. Review of

Financial Studies 30, 1053–1094. URL: https://ideas.repec.org/a/oup/rfinst/v30y2017i4p1053-1094..html.

Francioni, R., Gomber, P., 2017. High Frequency Trading: Market Structure Matters, Springer International Publishing, Cham, pp. 363–390. URL: https://doi.org/10.1007/978-3-319-45848-9_13, doi:10.1007/978-3-319-45848-9_13.

Hagströmer, B., 2021. Bias in the effective bid-ask spread. Journal of Financial Economics 142, 314–337. URL: https://www.sciencedirect.com/science/article/pii/S0304405X21001458, doi:10.1016/j.jfineco.2021.04.018.

Hasbrouck, J., Saar, G., 2013. Low-latency trading. Journal of Financial Markets 16, 646–679. URL: https://www.sciencedirect.com/science/article/pii/S1386418113000165, doi:10.1016/j.finmar.2013.05.003.

Hendershott, T., Jones, C.M., Menkveld, A.J., 2011. Does Algorithmic Trading Improve Liquidity? The Journal of Finance 66, 1–33. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.2010.01624.x, doi:10.1111/j.1540-6261.2010.01624.x.

Hoffmann, P., 2014. A dynamic limit order market with fast and slow traders. Journal of Financial Economics 113, 156–169. URL: https://www.sciencedirect.com/science/article/pii/S0304405X14000610, doi:10.1016/j.jfineco.2014.04.002.

Hu, E., 2019. Intentional Access Delays, Market Quality, and Price Discovery: Evidence from IEX Becoming an Exchange. Technical Report. SEC Working Paper. URL: https://papers.ssrn.com/abstract=3195001, doi:10.2139/ssrn.3195001.

Iacus, S.M., King, G., Porro, G., 2012. Causal Inference Without Balance Checking: Coarsened Exact Matching. Political Analysis 20, 1–24. URL: https://www.cambridge.org/core/journals/political-analysis/article/abs/causal-inference-without-balance-checking-coarsened-exact-matching/5ABCF5B3FC3089A87FD59CECBB3465C0, doi:https://doi.org/10.1093/pan/mpr013.

Khapko, M., Zoican, M., 2021. Do speed bumps curb low-latency investment? Evidence from a laboratory market. Journal of Financial Markets 55, 100601. URL: https://www.sciencedirect.com/science/article/pii/S1386418120300707, doi:10.1016/j.finmar.2020.100601.

Kirilenko, A., Kyle, A.S., Samadi, M., Tuzun, T., 2017. The Flash Crash: High-Frequency Trading in an Electronic Market. The Journal of Finance 72, 967–998. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.12498, doi:10.1111/jofi.12498.

Mayhew, S., 2002. Competition, Market Structure, and Bid-Ask Spreads in Stock Option Markets. The Journal of Finance 57, 931–958. URL: https://www.jstor.org/stable/2697763.

Megarbane, N., Saliba, P., Lehalle, C.A., Rosenbaum, M., 2017. The Behavior of High-Frequency Traders Under Different Market Stress Scenarios. Market Microstructure and Liquidity 03, 1850005. URL: https://www.worldscientific.com/doi/abs/10.1142/S2382626618500053, doi:10.1142/S2382626618500053.

O'Hara, M., 2015. High frequency market microstructure. Journal of Financial Economics 116, 257–270. doi:10.1016/j.jfineco.2015.01.003.

Pagnotta, E.S., Philippon, T., 2018. Competing on Speed. Econometrica 86, 1067–1115. URL: https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA10762, doi:10.3982/ECTA10762.

7. Appendix A: Institutional set-up of the main speed bumps programmes

A speed bump is part of the strategy adopted by trading venues to attract market-makers, traders and investors, not only through the implementation of market-making programmes that encourage the provision of liquidity and commercial proposals (such as lower transaction fees or clearing costs)[26], but also, for the derivatives market in particular, by modulating the size of possible off-book transactions.

While some newer platforms use speed bumps as an argument to increase their uniqueness and market share, other platforms appear to test this mechanism on market segments with lower volumes, to be able to assess its efficiency before taking further actions. Lastly, calibrating the time delay appropriately seems important, as it is clearly defined by the time it takes to send and receive orders. The time delay must also be proportionate to the tick size in force on the market in question. They can take several forms:

- Symmetrical speed bumps, where the minimum delay between two orders is the same for all order types and market participants. For instance,this

[26]For instance, in the case of the venues observed in the article, both Eurex and Euronext offer an integrated model where clearing is done by their associated clearing houses (Eurex Clearing and LCH SA respectively). However, having multiple clearing houses that clear similar products end up increasing the overall cost of clearing, since traders bringing liquidity to multiple platforms cannot "net" their positions and must use multiple clearing houses (Benos et al. (2019)).

mechanism was implemented on the American platform IEX to slow down the speed of trading;

- Asymmetric speed bumps, which apply or not depending on the order type. This mechanism slows down only aggressive orders, thus explicitly targeting high-frequency arbitrage strategies;
- Random speed bumps, which are designed to prevent market participants from anticipating the effect of the slowdown. This particular mechanism can encourage duplicate orders but also reduce market participation.

Eurex – Passive Liquidity Protection (PLP)

Mechanism: Aggressive orders, i.e. orders that would trigger a transaction if they entered the order book, are set aside for 1 or 3 milliseconds before being integrated into the order book. Passive orders, i.e. orders that are not immediately executable, are not affected. They are entered directly into the order book and can be processed immediately. The time delay is 3 milliseconds for French equity options (= 0.003 of a second) and 1 millisecond for German equity options. Calibrating the time delay appropriately is important, as it is clearly defined by the time it takes to send and receive orders, as well as being proportionate to the tick size in force on the market.[27] The order response from the platform informs the participant whether their order is subject to the delay or not.

Scope: All French and German equity options (single underlying equity), including weekly options traded on the Eurex platform.[28]

[27]Eurex explains the reason for the time difference between the two as the result of examining the reaction time and the time required for market participants' transactions to pass through its matching engine. This is related to the greater geographical distance of the underlying equity market (infrastructure located near London for French equities traded on Euronext) compared with Eurex, causing higher latency in adjusting option prices based on French equity prices.

[28]For more technical details on PLP, see the dedicated pages on the Eurex website.

Start of experiment: 3 June 2019. The possibility of putting in place such mechanism was announced by Eurex in October 2018, and then discussed through an extensive participant consultation. The final set-up information were announced on 8 March 2019. However, no announcement effect is expected to impact our model, since the anticipations were built for more than 9 months - and the behavior of market participants is expected to switch once the program is activated, as is usually the case for latency arbitrageurs. [29]

Duration: Initially announced as an experiment lasting a minimum of six months, Eurex announced that it would continue the experiment, and then in May 2020 extended it to all equity options traded on the platform (French, German, but also Swiss, Dutch and Italian), i.e. more than 750 options from 10 countries. In addition, option trading on the DAX index, the main German stock market, were the first index option concerned by an asymmetric speed bump of 1.5 milliseconds beginning in August 2020.[30] At the end of May 2021, the options under a speed bump were extended to options on the FTSE100 index. Since February 2022, the mechanism covers all equity and index options traded on the platform.

Objective: According to Eurex, the mechanism's objective is not to slow down market trading but to increase fairness in the price discovery mechanism and improve market functioning. Eurex considers that liquidity providers are disadvantaged by the speed of order updates in the event of underlying movements, and therefore do not display their best prices in the order book. For the platform, this means wider spreads as well as clients who may prefer to trade off-book. The role of market-makers appears is crucial, especially for less liquid instruments such as options, where the probability of both buyers and

[29] Circular 027/19.

[30] DAX options accounted for 33.5 million trades during 2019, or 8% of index options trading on the platform.

sellers being present on the order book for all contracts on offer is low. With the PLP, Eurex intends to generate a virtuous circle that will encourage liquidity providers to be more active in the order book (rather than trading over-the-counter), thereby increasing the average depth of the order book and narrowing the average spreads, which is an attractive factor for clients. The aim is to increase the number of clients on the order book, with the initial objective of having more liquidity offered by the market-makers.

TSX Alpha Exchange

Mechanism: An asymmetric speed bump implemented in September 2015 imposing delay, on a randomized basis between 1 and 3 milliseconds, for orders likely to consume liquidity. However, it is possible for participants not to be affected by the mechanism by paying an additional fee, and being subject to a minimum size requirement (the minimum size requirement is provided daily and for each symbol by TSX).

Scope: TSX Alpha exchange is an exchange for equity listed on Toronto Stock Exchange or Venture Exchange. The mechanism was still in place in 2023 - with TSX Alpha exchange observed 6.4 trillion volumes in 2022 (or 11% of all volumes traded on TSX exchanges).[31]

Objective: The platform's stated goal is to attract "slow" (i.e. non-HFT) liquidity providers. An analysis published by TSX, in December 2019[32], shows that the implementation of this mechanism allows the exchange, compared to other platforms, to position itself in first place with the highest rate of presence and posted volumes at the national best bid and offer NBBO, the largest average trade size. The exchange is explicitly implementing the mechanism to "be an attractive destination for active retail and institutional orders".

[31] TSX statistics for 2022.
[32] TSX resources website.

IEX (Investors' Exchange LLC)

Mechanism: A symmetrical speed bump imposing a 350-microsecond (= 0.00035 seconds) delay between the request and the execution of all incoming orders on the platform, since the approval of the mechanism by the Securities and Exchange Commission (SEC), in June 2016.[33]

Scope: IEX is a national platform and therefore allows all US stocks and ETFs to be traded in the same way as the NYSE or Nasdaq. According to the platform, its market share of US equity trading volumes was around 2% in 2021.

London Metal Exchange (LME)

Mechanism: A symmetrical speed bump applying an 8-millisecond delay to all new incoming orders, but not applying to cancellation messages for existing orders.[34] The mechanism was introduced in March 2020 for a one-year trial period.

Scope: Gold and silver futures contracts (LME Precious Metals).

Objective: Considering itself a market where liquidity is still being built, LME wishes to attract new liquidity providers, encourage market participation, increase liquidity and improve the quality of spreads.

Programmes completed, not implemented or under discussion:

NYSE American (the New York Stock Exchange's segment for small and medium-sized enterprises): Had implemented a 350-microsecond symmetric speed bump for all orders on 1 July 2017. Designed to benefit passive orders,

[33] The SEC approval was permitted by a new interpretation of Rule 611 of the National Market System regulation, which governs best execution and ensures that there is a single best bid or offer at the national level, allowing the approval of intentional *de minimis* delays to orders. See SEC release No. 34-78102; File No. S7-03-16. However, the SEC confirms that speed bump programmes must be subject to regulatory assessment and "fairly applied", potentially closing the door to asymmetric speed bumps, as illustrated by the ban on the asymmetric speed bump proposed by the CBOE.

[34] See LME (2020), "Technical change to LME select FIX message processing for the LME precious market to introduce a fixed minimum delay", 20/052.

NYSE hoped it would encourage more trading from institutional investors at the mid-price. However, NYSE's assessments proved disappointing: market share decreased by 0.6% in its segment and average daily volumes decreased by 8%. More importantly, the speed bump did not have the desired effect: average spreads increased by 40%. The platform therefore removed it on 18 November 2019.[35]

Chicago Board Options Exchange (CBOE)

Mechanism: The CBOE proposed the implementation of a 4-millisecond (= 0.004 seconds, i.e. more than 10 times that of IEX) asymmetric speed bump for the first time in the United States on its EDGA Equities platform, which in April 2020 accounted for 1.6% of volumes traded in US equities. As with Eurex, the mechanism delays orders that would be executed immediately, allowing orders that cannot be executed immediately to be added to the book without a time delay.

Opposition: The public consultation that opened in 2019[36] provided an opportunity for many market participants and competitors to publicly oppose the mechanism, and in particular its asymmetrical nature, calling into question the fairness of market participants, the right to innovation and adding complexity to the market. The SEC rejected the proposed mechanism in February 2020,[37] finding that CBOE had not provided sufficient evidence that its speed bump "would not unfairly discriminate" against high-frequency traders' orders. It also found that the platform had not demonstrated that the proposal was sufficiently tailored to its stated purpose of improving displayed liquidity by reducing the

[35] See the NYSE analysis on its website.

[36] See the position documents available on the SEC website, Comments on CBOE EDGA Rulemaking.

[37] See SEC Order Disapproving Proposed Rule Change to Introduce a Liquidity Provider Protection Delay Mechanism on EDGA, Release No. 34-88261.

risk of adverse selection for liquidity providers.

ICE (Intercontinental Exchange Inc.)

Mechanism: An asymmetric speed bump imposing a 3-millisecond (= 0.003-second) delay on all orders that are not modifications or passive orders.

Scope: The measure concerns gold and silver futures and could be extended to other contracts in the future.

Objective: As ICE is not the preferred platform for gold and silver futures, the aim is to allow its investors to incorporate information when market movements, mainly from CME, affect prices.

Start of the mechanism: Accepted by the Commodity Futures Trading Commission (CFTC) in May 2019, the mechanism has not yet been implemented.[38]

8. Appendix B: Definition of the variables

- The **quoted spread** is the bid-ask spread displayed at the time of the trade. The quoted spread here is relative, as it is normalised by the mid-price to obtain a comparable view between instruments. It is calculated by dividing the difference between the ask and the bid price by the mid-price, and is in percent. This indicator is influenced by liquidity providers trying to be present in the order book at competitive prices.

- The **effective spread** at each trade is an approximation of the cost of liquidity actually paid by investors at the time of a trade, or a measure of the cost of immediate execution. It is usually defined as twice the difference between the transaction price and the fundamental value, using the mid price as an imperfect proxy (Hagströmer (2021)). The effective

[38] See the CFTC decision Submission No. 19-119.

spread at each trade should not deviate from the quoted spread most of the time. The effective spread is only larger than the quoted spread when the trade size exceeds the quantity available at the best limit. On the options market, larger volume trades, which would exceed the best limit offered, can be expected to take place outside the order book.

$$Effective\ spread_{\mathrm{ij}} = 2 * D_j * \left(Price_j - mid\ price_j\right) / mid\ price_j$$

Here, j is the transactions for contract i during day t, and D is a direction of trade indicator taking the value +1 for buyer-initiated trades, and -1 for seller-initiated trades. The effective spread is normalised by the mid-price to obtain a comparable view between instruments, multiplied by two to be comparable to the quoted spread. The daily effective spread is the average of effective spreads during day t for each contract.

- The effective spread is frequently decomposed into the price impact and the realized spread (Hendershott et al. (2011)) in order to estimate the net revenue or loss of liquidity providers. Here the **realized spread** is calculated when multiple transactions on a similar contract are observed during the day. It is calculated for each transaction as twice the difference between the price and the mid price of the following transaction, taking into account the direction of trading and normalised by the mid price. The daily realized spread is the average of realized spreads during day t for each contract.

$$Realised\ spread_{\mathrm{ij}} = 2 * D_j * \left(Price_j - mid\ price_{j+1}\right) / mid\ price_j$$

- Similarly, the **price impact** for each transaction is calculated as twice the difference between the mid price at the following transaction, and the mid price at the transaction (a proxy of the change in fundamental value following a trade), normalised by the mid price. The daily price impact is the average of price impacts during day t for each contract.

$$Price\ impact_{\mathrm{ij}} = 2 * D_j * \left(mid\ price_{j+1} - mid\ price_j\right) / mid\ price_j$$

- The **market depth** at each transaction is the quantity available at the best limit at the time of the trade. It is expressed in euros and is equal to the sum of the quantity available at the bid price multiplied by the best bid price and the quantity available at the ask price multiplied by the best ask price, divided by two. This indicator is a way of quantifying the effective passive presence of market participants in the order book.
- The **volumes traded** indicator reflects the amount traded during the transaction, an indicator for market activity. An option trade consists of the purchase or sale of one or more option lots: in most cases, one lot gives the right to buy or sell the equivalent of 100 shares. Thus, "volumes" are the number of lots traded and are not multiplied by 100. Both market depth and volumes are transformed in log in the regressions. In order to assess the representativeness of the resulting database of transactions, a comparison of monthly volumes for each underlying share was assessed, comparing the traded volumes published by each platform on their website (Euronext and Eurex) and the traded volumes available in the database. No particular underlying share appeared sub-represented in the database, and the monthly coverage of volumes remained mostly stable.
- **HFT presence** is approximated here by making use of the timestamp for each transaction, supposing that transactions reported at the nearest ten of microseconds are trades that include HFT (see Section 3). For

Table 5: Differences in the main variables used for the matching and the regressions, for the treatment and control groups, before and after the matching procedure

	Before matching				After matching			
	Control group		Treatment group		Control group		Treatment group	
	mean	std	mean	std	mean	std	mean	std
Underlying capitalisation	36,845	39,991	41,447	40,261	37,206	40,228	41,005	40,617
Underlying volumes	2,323	2,470	2,469	2,475	2,330	2,486	2,522	2,479
Underlying volatility	1.01	1.08	1.18	1.57	1.02	1.09	1.01	1.11
Spread	0.22	0.25	0.11	0.16	0.21	0.24	0.11	0.16
Effective spread	0.09	0.13	0.11	0.16	0.09	0.13	0.11	0.16
Realized spread	0.06	0.20	0.04	0.29	0.06	0.20	0.04	0.29
Price impact	0.03	0.19	0.04	0.29	0.02	0.19	0.04	0.29
Depth	215	659	145	278	218	665	139	259
Volumes	49	119	50	106	49	119	51	107

Note: Underlying capitalisation in millions of euros; underlying volumes in thousands. Spread, effective, realized spreads and price impact in percent, market depth in euros and volumes in quantity of contracts.

the regressions, the HFT variable is a dummy created for each contract, signalling if there was transactions that included HFT during the day.

- **Moneyness** is defined as the likelihood that an option will expire in the money - a proxy of the profitability of an option to the investors, if the option is exercised right away. Following Etling and Miller (2000) definition, it is calculated daily, for call options as the difference between the closing price of the underlying equity and the strike price of the option, divided by the strike price; and for put options as the difference between the strike price of the option and the closing price of the underlying equity, divided by the strike price. Thus, moneyness is positive for in-the-money call or put options, negative for out-of-the-money call or put options, and near zero for call or put options that are exactly at the money.

9. Appendix C: Other summary statistics

Table 6: Number of observations, options and underlyings before and after matching

	Before matching			After matching		
	No obs	No options	No underlying	No obs	No options	No underlying
Eurex only	2,385	1,259	17	2,349	1,229	15
Eurex both	147,380	23,131	63	144,000	22 555	62
Euronext both	128,812	19,592	62	128,812	19,592	62
Euronext only	2,865	1,313	25	2,005	942	22
Total	281,442	45,295	105	277,166	44,318	99

Table 7: Direct impact analysis, regression results without control variables

	Indicators						
	Spread	Eff. spread	Real. spread	Price impact	Depth	Volumes	HFT
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Post*PLP	-0.0157^{***}	-0.0061^{***}	-0.0117^{**}	0.0053	-0.0126	-0.0459^{*}	-0.0373^{***}
	(0.0040)	(0.0022)	(0.0050)	(0.0044)	(0.0194)	(0.0235)	(0.0085)
Fixed effects	Date	Date	Date	Date	Date	Date	Date
Observations	94,360	94,360	39,972	39,972	94,360	94,360	94,360
Adjusted R^2	0.0637	0.0026	0.0027	0.0016	0.0074	0.0096	0.0190

Note: Standard errors are clustered at the date and contract level. Spread, effective, realized spread and price impact in percent, depth and volumes are transformed in log. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Table 8: Spillover analysis, regression results without control variables

	Indicators						
	Spread	Eff. spread	Real. spread	Price impact	Depth	Volumes	HFT
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Eurex only*PLP	-0.0463	0.0163	0.2542	-0.2530	0.0351	0.0399	-0.0623^{**}
	(0.0361)	(0.0249)	(0.2356)	(0.2287)	(0.0700)	(0.0857)	(0.0263)
Eurex both*PLP	-0.0851^{**}	-0.0236	0.2487	-0.2631	0.1611^{***}	0.0414	-0.0701^{***}
	(0.0325)	(0.0199)	(0.2432)	(0.2397)	(0.0562)	(0.0699)	(0.0237)
Euronext both*PLP	-0.0724^{**}	-0.0173	0.2586	-0.2660	0.2134^{***}	0.0804	-0.0121
	(0.0329)	(0.0199)	(0.2432)	(0.2400)	(0.0578)	(0.0677)	(0.0235)
Fixed effects	Date	Date	Date	Date	Date	Date	Date
Observations	141,599	141,599	58,908	58,908	141,599	141,599	141,599
Adjusted R^2	0.0627	0.0136	0.0035	0.0045	0.0368	0.0113	0.0257

Note: Standard errors are clustered at the date and contract level. Spread, effective, realized spread and price impact in percent, depth and volumes are transformed in log. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.

Figure 1: Evolution of volumes on each venue

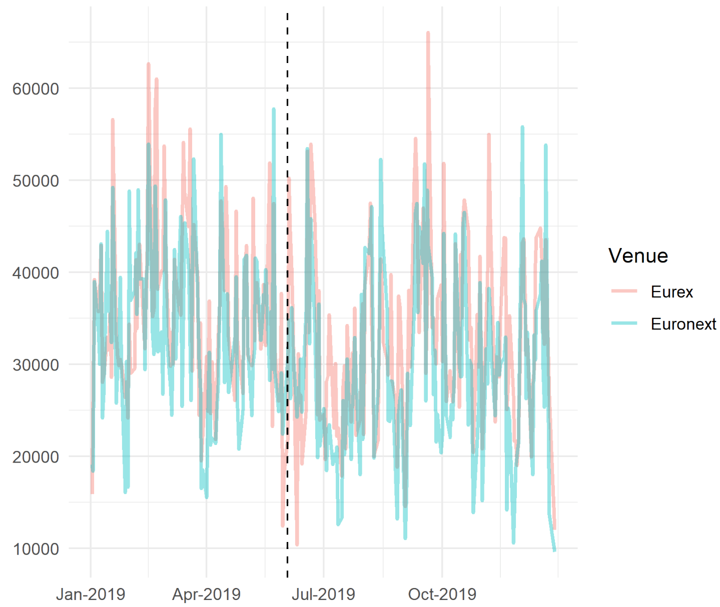

Figure 2: HFT presence

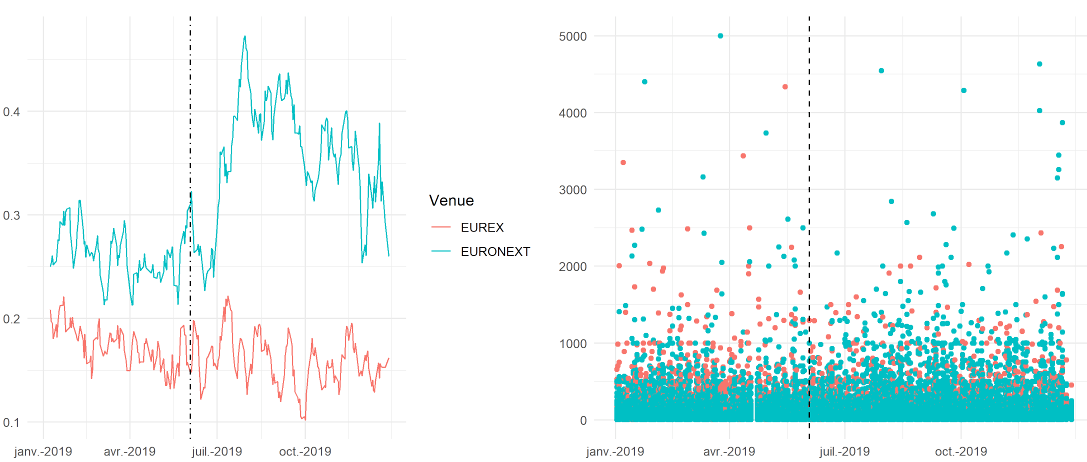

(a) Proxy of HFT volumes

(b) Distribution of HFT volumes

Figure 3: Standardized mean differences between the different matching methods

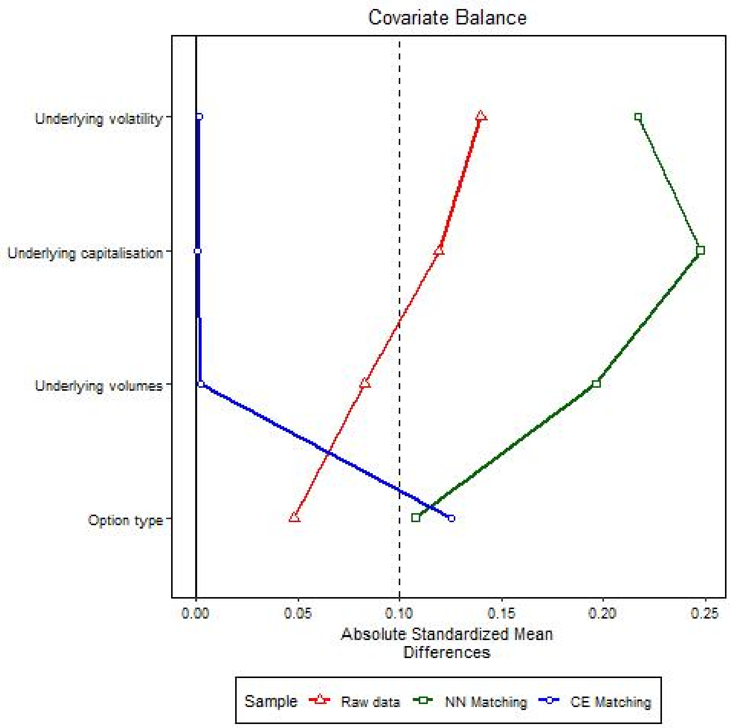

Table 9: Speed bump impact for the restricted sample

	Indicators						
	Spread	Eff. spread	Real. spread	Price impact	Depth	Volumes	HFT
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Post*PLP	−0.0454	0.0047	−0.0611	0.0780	0.0229	0.0332	-0.0451^{*}
	(0.0347)	(0.0213)	(0.1096)	(0.0996)	(0.0553)	(0.0746)	(0.0230)
Price level	-0.0086^{***}	-0.0045^{**}	-0.0156^{***}	0.0039	0.0385^{***}	-0.0139^{*}	−0.0015
	(0.0026)	(0.0019)	(0.0056)	(0.0063)	(0.0108)	(0.0075)	(0.0017)
Maturity	-0.0008^{***}	-0.0004^{***}	0.0001	−0.0003	0.0018^{***}	0.0007^{**}	0.00001
	(0.0001)	(0.0001)	(0.0002)	(0.0003)	(0.0002)	(0.0003)	(0.0001)
Moneyness	-0.4457^{***}	-0.3623^{***}	0.0683	−0.3579	1.3125^{***}	0.3271	0.0041
	(0.0982)	(0.0758)	(0.2079)	(0.2983)	(0.2321)	(0.2077)	(0.0467)
Fixed effects	Date	Date	Date	Date	Date	Date	Date
Observations	2,106	2,106	502	502	2,106	2,106	2,106
Adjusted R^2	0.2470	0.0962	0.1014	0.1224	0.3810	0.0374	0.0141

Note: Standard errors are clustered at the date and contract level. Spread, effective, realized spread and price impact in percent, depth and volumes are transformed in log. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$.

Figure 4: Parallel trends

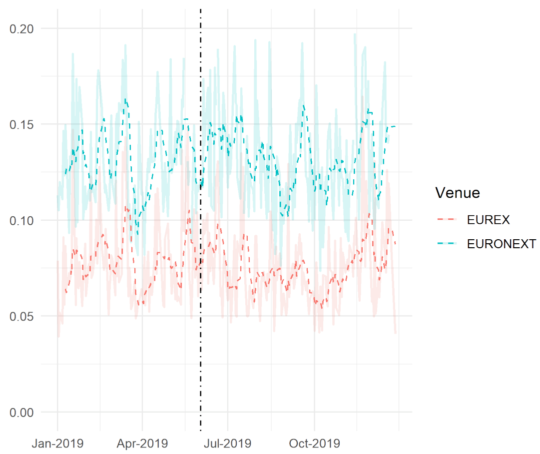

(a) Average spread by venue

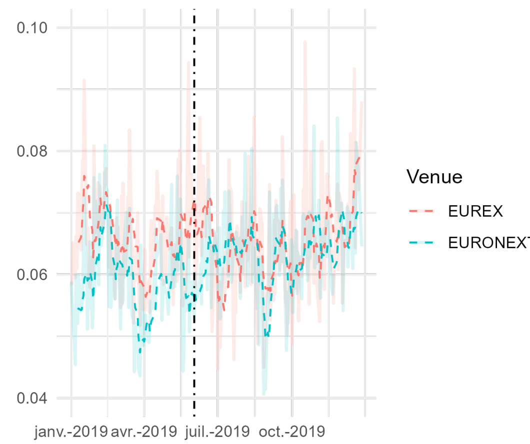

(b) Average effective spread by venue

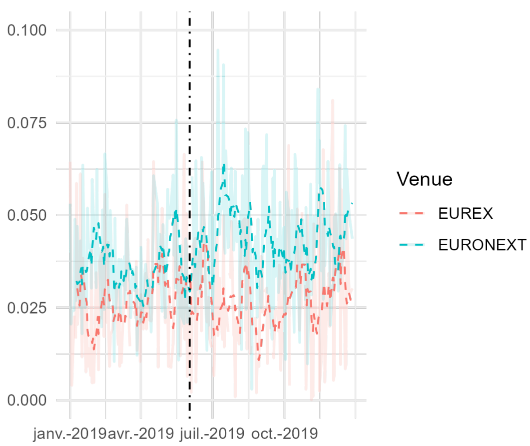

(c) Average realized spread by venue

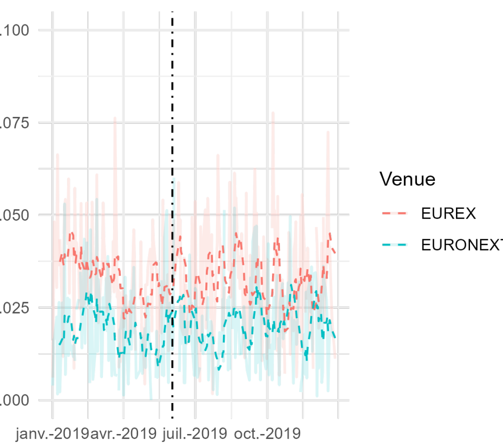

(d) Average price impact by venue

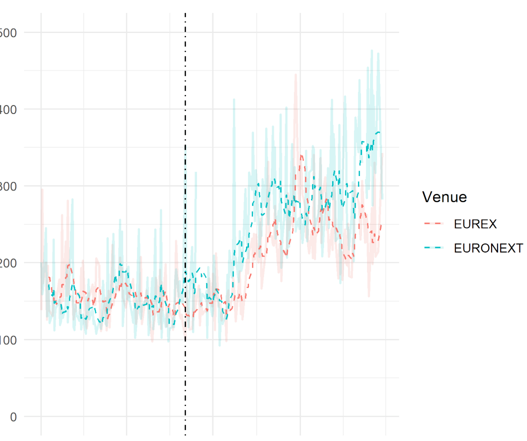

(e) Average depth by venue

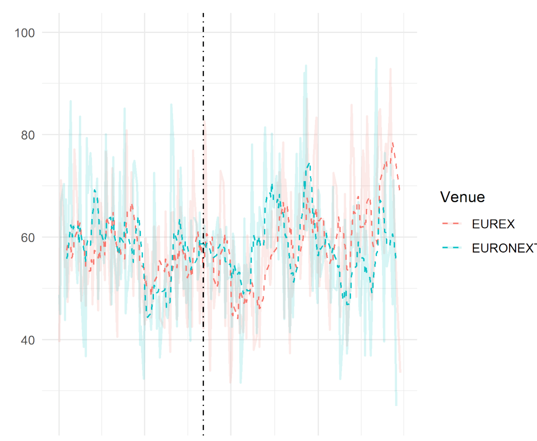

(f) Average volume by venue

Figure 5: Number of French equity options underlying shares, traded on Eurex and Euronext

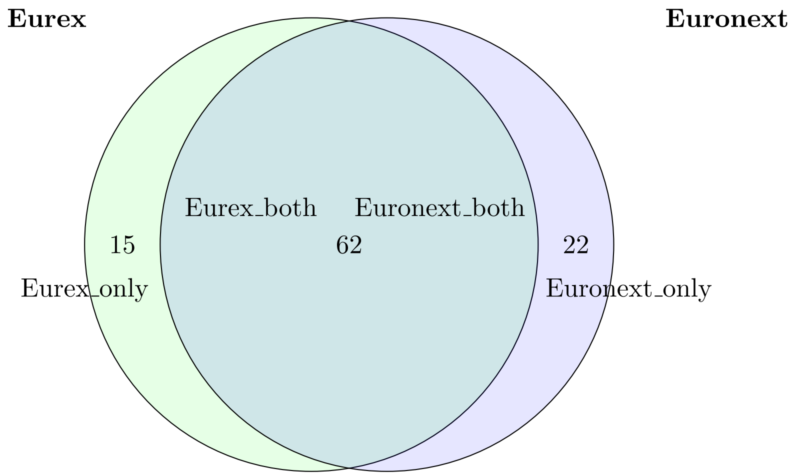

Note: *Eurex_both* (resp. *Euronext_both*) are contracts traded on Eurex (Euronext), but that can be traded on both platforms. *Eurex_only* (resp. *Euronext_only*) are contracts that are only traded on Eurex (Euronext).

Figure 6: Number of French equity options traded on Eurex and Euronext, matched sample

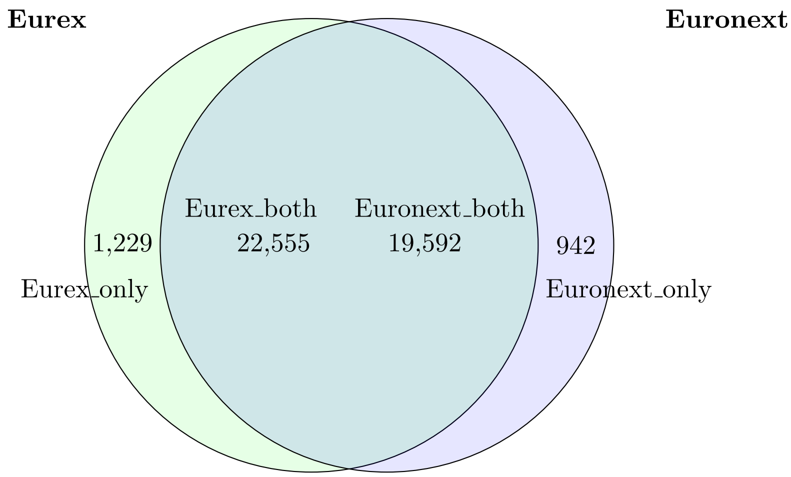

Note: *Eurex_both* (resp. *Euronext_both*) are contracts traded on Eurex (Euronext), but that can be traded on both platforms. *Eurex_only* (resp. *Euronext_only*) are contracts that are only traded on Eurex (Euronext).