Sand in the Chips? Evidence on Taxing Transactions in Modern Markets*

Jean-Edouard Colliard and Peter Hoffmann†

First version: December 2012

This version: July 31, 2013

Abstract

We present evidence on the causal impact of financial transaction taxes on market quality in a modern market structure by exploiting the introduction of such a levy in France on August 1st, 2012. Our evidence suggests that the substantial changes in market structure over the past decades play an important role in reassessing the long-standing idea of the FTT. While we document a surprisingly mild impact on exchange-based trading due to exemptions for liquidity provision, off-exchange trading declined by 40%, and the largest OTC trades virtually disappeared. This suggests that market segmentation poses a considerable challenge to current policy proposals.

Journal of Economic Literature Classification Number: G10, G14, G18, H32.

Keywords: Financial transaction tax, OTC markets, liquidity, high-frequency trading.

---

*We would like to thank Bruno Biais, Fany Declerck, Hans Degryse, Laurent Grillet-Aubert, Philipp Hartmann, Frank de Jong, Simone Manganelli, Elvira Sojli as well as seminar and conference participants at the French Treasury, the Autorité des Marchés Financiers (AMF), the ECB, and the Arne Ryde Workshop for comments and suggestions. The views expressed in this paper are the authors' and do not necessarily reflect those of the European Central Bank or the Eurosystem.

†European Central Bank, Financial Research Division. E-mail: `Jean-Edouard.Colliard@ecb.int` and `Peter.Hoffmann@ecb.int`. Contact author: Peter Hoffmann, ECB, Kaiserstrasse 29, D-60311 Frankfurt am Main, Germany.

1 Introduction

The long-debated idea of a financial transaction tax (FTT) has regained the favour of many governments since the beginning of the financial crisis. Many countries, particularly in Europe, are under pressure to increase tax revenues in order to revamp public finances, and the political attractiveness of imposing levies on financial activities has increased significantly in light of the widespread public discontent with the financial sector. Consequently, 11 European countries[1] formally committed in October 2012 to introducing an FTT under enhanced cooperation, and the details are currently being debated among European leaders. Yet, even prior to that, France had unilaterally introduced a levy of 0.2% (20 bps) on the purchase of shares of French companies with a market capitalization of more than 1 billion EUR on August 1st, 2012. Importantly, the tax only applies to transfers of ownership and thus implicitly exempts intraday trading activity. Simultaneously, a tax on high-frequency trading (HFT) by domestic entities was introduced.

The French experiment is unique due to the radical changes which have transformed the traditional trading floor into a purely electronic limit order book that competes with other trading venues and mechanisms such as multi-lateral trading facilities, dark pools, and the over-the-counter market. Thus, to our knowledge, it is the first time that such a levy is introduced in a modern market structure and therefore a good opportunity to assess how such a tax affects the marketplace of the 21st century. As the FTT has reappeared on the political agenda, it would be questionable to base policy decisions on empirical evidence that dates back to the 1990s or relates to emerging markets. Throwing "sand in the chips" of modern markets may imply both new benefits and new risks compared to previous experiences.

In order to identify the causal impact of this policy change, we exploit the fact that Euronext as a pan-European stock exchange is the primary listing venue not only for French stocks but also for those of other European countries. Hence we are able to implement a difference-in-differences approach based on a control group of non-French equities traded in the same market under identical rules.

Overall, trading volume in French stocks on Euronext dropped by 10% as a consequence

[1] See "Financial transaction tax gains approval", Financial Times Online Edition, October 9th, 2012.

of the policy change, and this development was accompanied by a rather muted decrease in liquidity. While bid-ask spreads and intraday volatility were not affected, we document significant decreases in market depth (driven by the most liquid stocks), resiliency, and price efficiency. We further document that the policy change has led to a significant decrease in measures of low-latency activity.

Although the overall impact of the FTT is certainly negative, our evidence suggests that market quality on Euronext was quite resilient. It is particularly remarkable that spreads did not increase significantly despite a measurable impact on HFT activity. This is most likely due to a cautious implementation that restricts taxation to ownership transfers and additionally protects market-making activities.

Given that trading on Euronext only represents a fraction (albeit an important one) of trading in the fragmented marketplace for European equities, we additionally examine how the policy change has affected the activity outside the primary market. While other "lit" markets are affected to a similar degree, we find that trading volume reported for off-exchange trading mechanisms such as dark pools and over-the-counter decreases by a staggering 40%. We show that this effect is largely due to a reduction in trade sizes, with the market for large block trades disappearing almost completely. When summing across all market segments, the French FTT reduced trading volume by 30%.

In sum, our results highlight the importance of accounting for recent evolutions in market structure when debating the long-standing idea of an FTT. In particular, we show that failing to consider the complete spectrum of trading mechanisms in an increasingly segmented marketplace leads to incomplete conclusions. Focusing exclusively on exchange-based trading suggests that today's electronic limit-order markets are surprisingly resilient to such a tax. As these markets mainly rely on liquidity provision by high-frequency traders, the restriction to taxing only ownership transfers together with an explicit protection of market-making was successful at safeguarding market quality. Due to its focus on domestic market participants, the additional HFT tax did not offset these mitigating effects. In contrast, the adopted safeguards have failed in protecting off-exchange trading, as they were evidently tailored to the specifics of lit markets. Clearly, it may be difficult to design measures that are effective in preserving market functioning in less transparent and more informally organized market segments.

Notably, the current draft for a pan-European FTT[2] by the European Commission does not include any liquidity-protecting exemptions and furthermore applies homogeneously to a wide array of different markets, which raises some concerns about its potential impact. Moreover, we estimate the revenues that could be raised by extending the French implementation of the FTT to the entire European Union at around 3.4 billion EUR. In comparison, the European Commission expects that between 4.8 and 6.5 billion EUR could be raised by taxing *all* European equities, without any exemptions. This rather narrow gap suggests that liquidity protecting measures need not be overly costly in terms of foregone revenues.

The remainder of this paper is organized as follows. The next section relates our paper to the existing literature. Section 3 details the legal background and provides an outline of our identification strategy. Section 4 presents our results on lit market quality, while Section 5 looks at the impact on other trading venues, in particular the OTC market. Section 6 details the impact of the tax in the cross-section of stocks, followed by the Conclusion.

2 Related literature

The idea of using an FTT to curb short-term speculation and noise trading with the aim of reducing excessive volatility dates back to Keynes (1936). Many years later, Tobin (1978) suggested to "throw sand in the wheels of our excessively efficient international money markets". Stiglitz (1989) and Summers and Summers (1989) argue that additional benefits include the reduction of managerial myopia and excessive investment in research activities. From their point of view, the gains of taxing transactions outweigh its potentially detrimental impact on market liquidity.

The general equilibrium framework of Kupiec (1996) challenges this view. He shows that an FTT may actually increase return volatility because the drop in asset prices offsets the volatility-dampening decrease in noise trading. Schwert and Seguin (1993) provide an overview of the trade-off between benefits and costs in terms of liquidity and efficiency, and conclude that the latter are likely to be large. Song and Zhang (2005) model this trade-off

[2]See http://ec.europa.eu/taxation_customs/resources/documents/taxation/other_taxes/financial_sector/com(2011)594_en.pdf.

by assuming that an FTT discourages noise trading, which can lead either to an increase or a decrease in volatility, depending on whether fundamental uncertainty or uncertainty on supply is more important.

Several papers have empirically examined the impact of various FTTs. Roll (1989) studies a cross-section of 23 countries and finds a negative, but insignificant relationship between stock market volatility and transaction taxes. Umlauf (1993) studies the Swedish experiment during the 1980s, when an FTT led to a large part of trading migrating to London, while the stocks whose activity remained in Stockholm displayed a statistically insignificant increase in volatility. Baltagi, Li, and Li (2006) examine the impact of a tax increase for trading Chinese shares in 1997 and document a decrease in volume but an increase in volatility. Pomeranets and Weaver (2012) study a multitude of changes in New York State Security Transaction Taxes between 1932 and 1981 and conclude that the tax decreased trading, led to wider bid-ask spreads, and higher price impacts.

More evidence can be collected by examining changes in other fixed costs of trading that are similar in effect to taxes. Jones and Seguin (1997) study the introduction of negotiated, lower commissions on US national exchanges in 1975 and conclude that this change led to a reduction in volatility. In contrast, Liu and Zhu (2009) document an increase in volatility on the Tokyo stock exchange following a similar event in 1999. Hau (2006) examines migrations between different tick-size regimes on the Paris Bourse and concludes that an increase in fixed transaction costs leads to an increase in volatility. The only recent evidence stemming from a modern market structure is provided by Malinova and Park (2011), who consider the impact of different fee changes on the Toronto Stock Exchange in 2005, but their focus in on the breakdown between make and take fees.

It is important to notice that financial markets have undergone a tremendous change over the past decades. Chordia, Roll, and Subrahmanyam (2011) document that trading volume in US equities has increased roughly fivefold between 1993 and 2008. This development has been accompanied by the rise of HFT in the mid 2000's, which now accounts for the majority of overall trading in US and European equities. Simultaneously, the number of available trading venues, both lit and dark, has multiplied. The trade-offs associated with an FTT may be entirely different in today's markets and our paper contributes to this literature by updating the available evidence with data from a fragmented and computerized market.

The controversies over the impact of HFT nicely illustrate the new potential benefits and costs of FTTs in modern markets. HFTs can play a useful role when they behave as competitive market-makers (Jovanovic and Menkveld (2012)), but may also use their speed to take advantage of other participants (Biais, Foucault, and Moinas (2013), Hoffmann (2013), Cartea and Penalva (2012)). Even when they trade based on their quick processing of information, their contribution to price discovery is not necessarily positive (compare Martinez and Rosu (2013) and Foucault, Hombert, and Rosu (2012)). Empirical studies also offer a mixed picture, suggesting that algorithmic trading improves liquidity (e.g. Hendershott, Jones, and Menkveld (2011)), but at the same time also increases volatility (Boehmer, Fong, and Wu (2012)). Brogaard, Hendershott, and Riordan (2012) find that HFT contributes to price discovery, and Malinova, Park, and Riordan (2012) show that HFT activity seems to reduce trading costs for other market participants, including retail investors. In line with an ambiguous impact of HFT, we find that the policy change led to a decrease in low-latency activity without an impact of similar magnitude on market quality.

To our knowledge, there is no recent empirical evidence on OTC markets for large stocks that are simultaneously listed on lit trading venues. Our results suggest that these markets play an important role, in particular for large block trades. In this sense, our paper relates to the literature on "upstairs markets", as studied for instance by Bessembinder and Venkataraman (2004) at the time of the Paris Bourse. A widespread view of the role of these markets is that they allow large uninformed trades to be executed without having a large price impact or being front-run (see Smith, Turnbull, and White (2001), and Madhavan (2000) for a review), a concern that may have re-emerged with certain HFT strategies (see Hirschey (2013)).

A number of recent empirical papers such as e.g. Degryse, de Jong, and van Kervel (2013), Weaver (2011) or Comerton-Forde and Putnins (2012) point at the importance of externalities between competing trading mechanisms. In fact, the evidence collected here suggests that potential asymmetries in the effects of liquidity preserving measures across market segments could have triggered some substitution from dark to lit trading venues.

3 Methodology and data

3.1 The policy experiment

The different measures introduced on August 1st are the following:

- Article 235 ter ZD of the *Code Général des Impôts* (the French tax law code) introduces a "stamp duty" of 0.2% of the transaction price, payable by the buyer on daily net position changes (i.e. ownership transfers). This tax applies to shares of all listed companies incorporated in France with a market capitalization above one billion euros on December 1st of the previous year[3], and to trades on any platform or OTC. American Depository Receipts and Global Depository Receipts were not subject to the tax at the time of the study.

There are a number of important exemptions from the tax: It does not apply to newly emitted shares, to transactions by clearing houses, to employee stock ownership plans and, most importantly, to transactions due to market-making. Market-making is legally defined as either quoting competitive bid and ask prices and providing liquidity on a regular and continuous basis, or executing orders on the behalf of clients, or hedging positions due to these two activities. Notice that a market-maker ending each day with a balanced inventory is already exempt from the tax since only the net position change is taxed. This additional exemption means that in the case where a market-maker has a positive overnight position[4] he can be exempted if he shows that this position was accumulated as the result of market-making activities. Due to all these exemptions, the main agents directly affected by the tax are buy-side investors. Notice in particular that non-French investors are taxed as well as French investors.

- Article 235 ter ZD bis of the *Code Général des Impôts* introduces a tax on HFT. It applies only to HFTs based in France, transacting on their own account. HFT is defined as the regular submission of orders with a resting time of less than 0.5 seconds. Market-making, smart-order routing and automated execution of large orders are exempted. When the order-to-trade ratio exceeds 5 : 1, a tax of 0.01% (1 bp) has to be paid on the notional of each additional order that is modified or cancelled. Finally, and importantly, the tax on high-

[3]With the exception of the first year of implementation, for which the relevant date was January 1st 2012.

[4]Market-makers may for example decide to hedge positive inventories in the derivatives markets instead of closing the position on the cash market.

frequency trading applies to all stocks, even those not affected by the tax on transactions.

- Article 235 ter ZD ter of the *Code Général des Impôts* taxes transactions on CDS on European sovereign bonds. These products are outside the scope of this paper.

3.2 Data and Identification Strategy

The goal of our analysis is to identify the causal impact of a policy experiment consisting in the simultaneous introduction of an FTT and a tax on domestic HFT activity. It is important to stress that the actual implementation of these levies makes it impossible to disentangle the effects of the individual measures: Smaller stocks (i.e. below 1 billion EUR market capitalization) are not subject to the FTT, but at the same time they are less likely to be affected by the HFT tax, because those types of traders are known to play a more prominent role in the most active securities (see e.g. Brogaard, Hendershott, and Riordan (2012)). Therefore we just aim at identifying the joint effect of both taxes (henceforth referred to as the "policy change").

We adopt a difference-in-differences (DiD) approach in order to measure the causal impact of the policy change on trading activity and a wide array of measures of market quality. As usual, identification requires the availability of a suitable control group, i.e. a set of stocks that are not subject to the FTT but otherwise as similar to the treated stocks as possible. Because we are looking at high-frequency data, it is important to ensure that the data for both treated and non-treated stocks stems from the same microstructural environment, which includes the trading protocol, the fee structure, the tick size regime, etc. Fortunately, the primary market for French stocks, Euronext, also constitutes the main trading venue for Belgian, Dutch and Portuguese stocks which therefore form a natural control group. Moreover, the Luxembourg Stock Exchange also uses Euronext's Universal Trading Platform (UTP) as part of a cross-membership cooperation. Because Belgium increased its existing FTT on August 1st 2012 (only levied on Belgian residents) and Portugal was heavily affected by the sovereign debt crisis we restrict the control group to stocks registered in Luxembourg and the Netherlands.

We define an initial pool of securities by collecting all constituents of the Euronext 100 and Euronext Next 150 Indices with a market capitalization of more than 1 billion EUR (as of the cutoff date January 1st 2012). A small number of French stocks are not included in these indices but are still subject to the FTT. These firms do not fulfill Euronext's minimum requirements in terms of free float and/or liquidity so that we decide not to include them in our sample. Additionally, we discard all stocks registered in Portugal and Belgium, as well as all bank stocks in order to avoid picking up large idiosyncratic effects due to the sovereign debt crisis. Moreover, we drop shares of CFAO which were subject to a public tender offer in August 2012. Finally, we require a stock to trade on each day of our sample period, which spans the period June 1st - October 31st (109 trading days), so that our sample consists of 117 stocks (86 treated and 31 controls, see Table 1 in the Appendix). We obtain millisecond-stamped intraday data for the market activity (trades, quotes, and the state of the limit order book at each point in time) on Euronext from Thomson Reuters Tick History (TRTH), from which we compute a wide range of microstructure variables at the stock-day level (see Section 4). Additionally, we obtain data on trades executed on MTFs, dark pools and over-the-counter using the TRTH Cross-Market Data Service.

The successful identification of the causal impact of the policy change with a DiD approach requires that the variable of interest, $y_{i,t}$, satisfies the so-called common trends assumption, which is formally given by

$$E(y_{i,t} \mid i,t) = \alpha_i + \gamma_t + \beta D_{i,t} \tag{1}$$

where the α_i and γ_t are stock and time fixed effects and D is a dummy variable equal to one for treated stocks on days *after* the policy change and zero otherwise. Intuitively, this assumption requires that the variables of interest for both groups of stocks comove closely absent any treatment, and the coefficient β captures the treatment effect.

While many applications of the difference-in-differences estimator rely on just two observations per individual (one before and one after the treatment), the abundance of financial data allows for repeated observations in both the pre- and post-event periods. This can be particularly useful if one wants to allow the treatment effect to vary across sub-periods, for example in order to disentangle temporary effects from permanent ones. In fact, we were advised in private conversations with practitioners, government officials and securities markets regulators that the trading activity in August is unlikely to correctly reflect the impact of the policy change because of i) (legal) uncertainty among investors whether they are subject to the tax or not and ii) a seasonal decline in trading activity for French stocks due to country-wide summer holidays. In order to take such a possibility into account, we opt for a more flexible framework and allow the treatment effect in the first month after the policy change (i.e. August) to be potentially different than the impact in September and October. Hence our specification reads as

$$E(y_{i,t} \mid i, t) = \alpha_i + \gamma_t + \beta^{Aug} D_{i,t}^{Aug} + \beta^{Sep/Oct} D_{i,t}^{Sep/Oct} \tag{2}$$

where $D_{i,t}^{Aug}$ and $D_{i,t}^{Sep/Oct}$ are dummy variables that take a value of 1 for treated stocks on trading days in August and September/October, respectively, and zero otherwise. In Section 4.3 we provide empirical support for our choice of specification by confirming the suspicion that trading activity in French stocks is generally subject to a slowdown in August, while both September and October generally appear to be free from seasonal influences. While the common trends assumption cannot be tested formally, we use this "placebo" DiD (together with visual inspection) to confirm the validity of our control group, as is customary in the literature on policy evaluation. Throughout the remainder of the paper, we will refer to $\beta^{Sep/Oct}$ as the permanent impact of the policy change.

4 The impact on the primary market

4.1 Variables

We begin our analysis by examining the causal impact of the policy change on several traditional measures of market quality such as trading volume, intraday volatility, price efficiency, and liquidity. All variables are computed for each stock-day in our sample (12,753 observations) using trades and quotes from the continuous trading session on Euronext, i.e. we discard trades that are executed off-book, during call auctions and the "trading-at-last" period. Trades are signed using the Lee and Ready (1991) method and we aggregate individual orders that are executed simultaneously into one single transaction. The variable definitions are standard and as follows.

$logvolume_{i,t}$: Natural logarithm of order-book traded value (in EUR, continuous session).

$RV_{i,t}$: Realized volatility based on the final midquote of 5-min intervals (in %, annualized).

$range_{i,t}$: Intra-day price range normalized by median trade price (in %).

$QS_{i,t}$: Time-weighted relative quoted half-spread (in bps).

$ES_{i,t}$: Transaction-weighted relative effective half-spread (in bps).

$depth_{i,t}$: Time-weighted depth at the inside quotes (in 1,000 EUR).

$res_{i,t}$: Speed of mean reversion for market depth, based on 1-min intervals with 5 lags (see Kempf, Mayston, and Yadav (2009)).

$|AR|_{i,t}$: Absolute value of first-order return autocorrelations, based on the final midquote of 5-min intervals.

Table 2 in the Appendix gives some sample statistics for each of these variables for both treated and control stocks.

4.2 Results

In order to estimate the causal impact of the policy change, we operationalize equation (2) by estimating the following regression:

$$\begin{aligned} y_{i,t} &= \alpha + \gamma \mathbf{1}(i \in \Theta) + \delta^{Aug}\mathbf{1}(t \in Aug) + \delta^{Sep/Oct}\mathbf{1}(t \in Sep/Oct) \\ &+ \beta^{Aug}\mathbf{1}(i \in \Theta) \times \mathbf{1}(t \in Aug) + \beta^{Sep/Oct}\mathbf{1}(i \in \Theta) \times \mathbf{1}(t \in Sep/Oct) + \varepsilon_{i,t} \end{aligned} \tag{3}$$

where $\mathbf{1}(x)$ is an indicator variable that takes the value of 1 if the statement x is true and zero otherwise, and Θ denotes the set of treated stocks. Table 3 in the Appendix reports the estimates for the coefficients β^{Aug} and $\beta^{Sep/Oct}$ with t-statistics based on standard errors clustered at the stock level. In addition, Figure 1 graphically illustrates the diff-in-diff estimates for each variable by plotting the cross-sectional averages for both groups of stocks minus their respective pre-event averages over time[5]. In addition, the dashed lines indicate

[5] We use 3-day moving averages in all figures for better readability.

the subperiod averages for June/July, August and September/October. The high degree of comovement between both series in all cases suggests that the common trends assumption is indeed satisfied.

We begin our discussion by looking at the impact of the policy change on trading volume, which is a natural starting point because it is not only relevant from the perspective of market quality but also is of paramount interest to the tax authorities. Panel a. of Figure 1 plots the evolution of the log of trading volume (from the continuous trading session) for both groups of stocks. We observe that August displays a considerable decrease in trading activity across both markets compared to June/July, a fact that is most likely due to fading market tensions in connection with a public statement by the ECB's President to do "whatever it takes" to preserve the Euro. Subsequently, we observe a rebound in trading activity. Turning to the differences across groups, we see that stocks subject to the tax display a decrease in trading volume of around 32% in August relative to the control sample, while this difference shrinks considerably towards the end of August and then remains remarkably stable at roughly 10%. The associated regression coefficients β^{Aug} and $\beta^{Sep/Oct}$ (reported in Table 3) are statistically significant at the 1% and 5% level, respectively.[6] While the results suggest that the policy change has led to a reduction in trading volume, a large part of the effect appears to have been transitory and the permanent impact is approximately 10%. Section 4.3 provides evidence that a large proportion of the transitory impact is due to seasonality in trading activity.

We next turn our attention to volatility, which we measure by realized volatility at the 5-min frequency or alternatively by the intraday price range. As mentioned in the introduction, the main economic rationale behind an FTT (beyond raising revenue) is the reduction of asset price volatility through the curbing of "excessive" trading activity. Opponents argue that such a levy will in fact lead to the opposite result because it becomes more costly to trade against mispricings. Panels b. and c. of Figure 1 illustrate the DiD estimates and show that the implementation of the FTT did not have a sizeable impact on market volatility. Looking at the regression coefficients in Table 3, we find that the permanent treatment effect $\beta^{Sep/Oct}$ is statistically insignificant for both variables. For August, we find a slight decrease

[6] The interpretation of a coefficient in a semi-log specification as a percentage change is only valid if its magnitude is sufficiently small. The correct percentage change is given by $\exp(\beta) - 1$ (up to a Jensen error).

in the intraday price range (β^{Aug} is significant at the 10% level), which one may interpret as weak support for the arguments of the FTT's proponents. Nevertheless, this effect is rather small compared to the large temporary reduction in trading activity, suggesting a rather unfavorable volume-volatility trade-off at best. We will further discuss this issue in Section 4.3.

We proceed by examining the impact of the policy change on all three dimensions of market liquidity as defined by Kyle (1985): Tightness (spreads), depth, and resiliency. While the French authorities explicitly exempted market-making activities from the FTT in an effort to preserve market quality, liquidity providers may still be affected indirectly by the decreased market activity that e.g. increases the expected turnaround time for inventory positions. Panels d. to g. of Figure 1 depict the impact on quoted and effective half-spreads, market depth and resiliency. Interestingly, we find that quoted half-spreads were not affected by the policy change, and this conclusion does not change if we alternatively consider the effective half-spread. This finding is consistent with the market-making exemption being successful at protecting liquidity providers. On the other hand, we find that both market depth and resiliency decrease significantly. On average, quoted depth at the inside quote is reduced by about 11,000 EUR (corresponding to slightly less than 20% of the pre-event average of 58,000 EUR), and its resiliency (the speed of mean reversion after shocks) is lowered by 0.017 (compared to a pre-event average of 0.49) . The coefficient estimates are significant at the 1% and 10% level, respectively, and interestingly there is no difference between the temporary effect β^{Aug} and the permanent impact $\beta^{Sep/Oct}$.

Finally, we turn to the informational efficiency of prices as measured by deviations from a pure random walk. The results (Panel h. of Figure 1 and Table 3) indicate that the adoption of the FTT has made prices less efficient as we observe a permanent increase in the absolute value of return autocorrelations of about 0.007 (significant at the 5% level). Compared to a pre-event average of 0.11 for treated stocks, this represents an increase of around 7%. This effect was absent in August, as β^{Aug} is statistically not different from zero.

While we have allowed for a differential impact of the policy change in August, our specification assumes that the impact is constant across the months of September and October. A look at Figure 1 reveals that this assumption appears to be valid for all variables. Additionally, we estimate a regression that allows for a different treatment effect in each month and report the difference between September and October in the last column of Table 3. As can be seen readily, none of the differences is statistically significant, and the magnitudes are very small across the board.

Overall, we find that the FTT has permanently reduced trading volume, market depth and resiliency as well as price efficiency. On the other hand, intraday volatility and quoted (effective) half-spreads are unaffected by the policy change. While the overall effect is certainly negative, its magnitude definitely falls short of a complete erosion of market quality. In fact, with the exception of market depth, the negative effects on liquidity are rather muted in terms of economic significance. In Section 6 we provide evidence that the result on market depth is driven by the most liquid stocks only. We also observe some temporary effects in August, an issue that we turn to examine next.

4.3 Accounting for seasonality

As mentioned in Section 3.2, practitioners suggested to us that the trading activity in August may not properly reflect the permanent impact of the policy change due to short-run (legal) uncertainty and seasonality in trading activity. The results from the previous subsection somewhat confirm these suspicions as we find some temporary effects in August for trading volume and intraday volatility. While it is close to impossible to measure the extent of uncertainty among investors, it is relatively straightforward to verify whether certain variables are subject to seasonal factors based on past data. To this end, we collect data for the months June - October for the three years prior to our sample period (2009 - 2011) and estimate a placebo-DiD, where we allow a different treatment effect for each calendar month. For this analysis only, we discard three stocks from our initial sample because of incomplete data for this period. Moreover, we drop October 23rd 2009 due to some missing observations. The resulting estimates (see Table 4 in the Appendix) strongly confirm the existence of seasonal factors in trading activity and volatility. In line with Hong and Yu (2009), we hypothesize that this effect is due to different vacation patterns across France and the Netherlands.[7]

[7] While it is common in France to take off most or even the entire month of August for summer holidays, this pattern is less prevalent in the Netherlands.

Compared to the control group, French stocks generally display a drop in traded volume of roughly 15% during the month of August, accompanied by a slight (and statistically significant) decrease in intraday volatility. None of the remaining variables appears to be subject to a seasonal influence during August. The coefficients for September and October are statistically insignificant with the exception of quoted half-spreads, which are somewhat lower in France during these months. Nevertheless, this result does not appear to be robust as it is not present for effective half-spreads. Overall, this placebo-DiD additionally supports the validity of our control group (see e.g. Autor (2003)).

Now we may de-seasonalize the treatment effect estimates for trading volume and volatility during August 2012 via a diff-in-diff-in-diff procedure. Given that the remaining variables are not subject to seasonality, applying this procedure to them would only (incorrectly) decrease the precision of the estimates. Hence we estimate the following equation for the months June - August and years 2009 - 2012 exclusively for trading volume and our two measures of volatility:

$$\begin{aligned} y_{i,t} &= \alpha^{09-11} + \alpha^{12}\mathbf{1}(t \in 2012) + \gamma^{09-11}\mathbf{1}(i \in \Theta) + \gamma^{12}\mathbf{1}(t \in 2012) \times \mathbf{1}(i \in \Theta) \\ &+ \delta^{09-11}\mathbf{1}(t \in Aug) + \delta^{12}\mathbf{1}(t \in 2012) \times \mathbf{1}(t \in Aug) \\ &+ \beta^{09-11}\mathbf{1}(t \in Aug) \times \mathbf{1}(i \in \Theta) + \beta^{12}\mathbf{1}(t \in 2012) \times \mathbf{1}(t \in Aug) \times \mathbf{1}(i \in \Theta) + \varepsilon_{i,t} \end{aligned} \tag{4}$$

where the seasonally adjusted treatment effect for August is given by β^{12}. The results in Table 5 show that the de-seasonalized impact on trading volume is now roughly 19%, which is considerably closer to the permanent impact of around 10%. The remaining discrepancy may be due to the mentioned (legal) uncertainty or other short-run factors. In terms of volatility, we conclude that the FTT did not have a statistically significant impact even in the short run.

4.4 High-frequency activity

A key difference between the market environment considered here and previous studies on securities transactions taxes is the presence of high-frequency traders. This type of market participants may have been affected by the policy change for three reasons: (i) in their role as "middlemen" they can be indirectly affected by a reduction in volume due to the stamp duty, (ii) French HFTs were directly impacted by a specific levy on order cancelations, (iii) other HFTs may indirectly be affected via changes in the competitive environment.

Unfortunately, our data does not allow for the direct identification of high-frequency traders, so that we have to resort to indirect measurement. As the HFT surcharge was targeted at fast order cancelations, we try to examine whether the policy change has led to systematic changes in the distribution of the order resting times, something that could plausibly be attributed to changes in HFT activity (see e.g. Hasbrouck and Saar (2009)).

Our data contains the state of the first 10 levels of the limit order book at each point in time. It is important to notice that we only observe a truncated version of the limit order book and not the entire message traffic, which does not allow us to track all orders until cancelation/execution. For example, one does not observe cancelations or modifications of orders that move beyond the 10th level as a result of changes in the mid-quote. This additionally complicates the tracking of time priority across orders.

In order to circumvent this issue, we therefore focus our analysis on orders where we are able to observe both the addition and the cancelation messages with a sufficiently high degree of certainty. To reduce the computational burden, we restrict our analysis to the bid side. We begin by collecting all messages that add liquidity, and then search for messages with a higher timestamp on the same day that led to a removal of the same amount of liquidity (depth) at the same limit price. If the removal was a trade (which we can identify from the trades file), we drop the order, and in case there are multiple liquidity-removing messages of the same type, we take the one that occurs first. Using this procedure, we are able to retrieve the lifetime of more than 80% of all buy limit orders, and based on this data we compute the following measures for each stock-day:

$median.lapse_{i,t}$: Median cancelation time of orders that were at the inside quote at the time of their submission, in seconds.

$pct.inside_{i,t}$: Proportion of total messages that either add or remove liquidity to/from the inside quote.

$cancel.0 - 500_{i,t}$: Proportion of limit orders submitted at the inside quote that are cancelled within 500 ms.

As reported in Table 6, treated and untreated stocks have similar averages for all three variables before the policy change. For illustration, panel a. of Figure 2 plots the difference in the empirical frequencies of cancelation times (up to 100ms) between September/October and June/July for both treated and untreated stocks.

This graph nicely illustrates the high quality of our control group, as the spikes in the distributions due to periodic cancelations perfectly coincide, which is due to the fact that the orders for both groups of stocks are submitted to the same matching engine. We see an increase in cancelations at ultra-high frequencies for Dutch stocks and at the same time a decrease for French stocks, implying that the policy change increased the time orders for French stocks spent in the limit order book. While the effect is more pronounced for extremely short cancelation times, it in fact persists until 400ms. Not reported on this graph, cancelations below 1ms represented 8.43% of the total for treated stocks before the treatment, compared to 7.13% after, while the proportion for control stocks slightly increased. This is direct evidence that the policy change has significantly affected low-latency activity.

While panel a. of Figure 2 provides evidence that our control group is even able to identify changes in a low-latency environment, the distributions may in fact be dominated by more active days and stocks. In order to avoid potential biases, panel c. of Figure 2 plots the evolution of the average median cancelation times[8] for each group and confirms the previously drawn conclusions. As a result of the policy change, the average median order-resting time for French stocks increased by 3.8 seconds in August and 2.4 seconds in September and October (both effects are significant at the 1% level, see Table 6), which is consistent with a significant reduction in low-latency activity. Compared to a pre-event average of 3.6 seconds for treated stocks, this means that the median cancelation time more than doubled in August, and increased by 67% in September.

We proceed by looking specifically at the proportion of orders that are canceled in less than half a second, because this is the minimum time span between two messages that French law defines as HFT activity. Panel b. reveals that, on average, the probability of observing cancelations within 500 ms declined by 6.6% and 5.4% in August and September/October,

[8] The mean cancelation time is contaminated by few orders resting extremely long in the book.

respectively. Again, both effects are statistically significant at the 1% level, as indicated in Table 6.

Finally, order aggressiveness also declined significantly, as the proportion of orders submitted at the inside quote decreased substantially (−5.2% in August and −7.3% in September/October), which is in line with reduced competition at the inside quote. Although market-makers are exempt from both the FTT and the HFT tax, other French HFTs are likely to be heavily impacted by the surcharge on cancelations.

Overall, these results strongly confirm that the policy change had a significant and permanent impact on high-frequency activity. The evidence suggests that two things may have happened (simultaneously or not): (i) affected French high-frequency traders may have reacted to the surcharge by canceling their orders less frequently, or alternatively by increasing their usage of market orders (as those always decrease the order-to-trade ratio); (ii) French high-frequency traders may have reduced their trading activity, or even dropped out of the market, thus leading to a "slowdown" unless their activity was replaced by HFTs located outside France.

The large impact of the policy change on measures of high-frequency trading is remarkable when compared to the rather muted impact on liquidity: Median cancelation times have increased by 67% without any significant impact on volatility or spreads for instance. While this experiment does not allow us to propose a causal interpretation, the evidence suggests that electronic markets can withstand a significant amount of "sand in the chips" of their participants, at least when some safeguards for liquidity provision are in place. The next section will further show that volume was much less resilient on OTC and dark trading venues.

5 Other markets

5.1 Trading volume

So far, we have restricted our analysis to the market activity generated during the continuous trading session on Euronext. While this is likely to be the center of price formation, an investor who wants to trade any of these stocks can choose among several alternative venues and mechanisms. First, a non-negligible proportion of trading volume is generated in the call auctions that are used to determine the opening and closing prices. Moreover, all stocks in our sample are also traded on other European exchanges as well as on a number of multilateral trading facilities (MTFs) such as Chi-X, BATS and Turquoise. While these new trading venues have managed to acquire a sizeable market share since the adoption of MiFID in 2007, the actual trading mechanism employed mirrors that of Euronext (open limit order book). Besides these "lit" markets, investors have the option to trade off-exchange, either in a dark pool or alternatively via private negotiation in the over-the-counter (OTC) market.

Given that a significant proportion of the on-exchange activity stems from high-frequency traders, the observed volume provides a rather noisy picture of "end user" activity. In our view, off-exchange trading activity as well as Euronext's call auctions are likely to be more representative of the activity of investors who are subject to the FTT due to the accumulation of net positions. For example, passive index funds may prefer to trade in the closing auction in order to minimize the tracking error of their portfolios. Similarly, large block trades usually take place in the OTC market.

In order to deepen our understanding of the impact of the policy change on trading activity we obtain additional data from TRTH Cross-market Data Service which compiles trades reported on any European venue. We delete trades that execute more than 20% away from the current VWAP as data errors and collapse simultaneously executed orders into single trades as before. We additionally omit trades that are not in EUR. The data on Euronext's call auctions (including the "trading-at-last" period) is taken from the trade and quote file used to construct the variables in Section 4.

In order to minimize the number of stock-day observations with zero volume due to infrequent trading on particular venues, we pool the resulting data into the following five categories: Euronext LOB Volume, Euronext Auction Volume, Other LOB Volume (order-book volume on other exchanges and MTFs), Dark Pool Volume, and OTC Volume. Table 7 in the Appendix provides summary statistics for both groups of stocks. Across all stock-days, trading on Euronext's limit order book averages about 1/3 of the total reported trading volume across all trading venues. Another 30-40% is attributed to the OTC market, while the remaining share is split between the other categories, with the majority pertaining to other lit venues (20-23%) and Euronext's call auctions (around 7%). Overall, this distribution is quite similar across both groups and roughly representative of European equity markets. As expected, the average trade size in the OTC market is considerably higher than on Euronext, which in turn exceeds the size of transactions executed elsewhere.

In order to examine the causal impact of the policy change on the different types of trading volumes, we estimate equation (3) with the log of each of our five volume categories as dependent variable.[9] We additionally estimate the impact for the sum of trading volume across all categories. The results are given in Table 8 and clearly demonstrate that it is highly important not to miss off-exchange trading. Restricting our attention to the causal impact for September/October, we find that the implementation of the FTT has caused OTC trading to decrease by around 42% relative to the group of control stocks. Dark trading has been similarly affected with a decrease of around 38%, while the reduction in activity on Euronext's call auctions as well as on other lit markets is only approximately 15% and thus relatively similar to the 10% decrease documented previously for the continuous trading session on the primary market. When summing across all five categories, we conclude that overall reported trading volume has decreased by almost one third.

In principle, a decrease in trading volume can occur either via a lower average trade size or a reduction in the number of trades (or both). We thus decompose the log of trading volume into the sum of the log of the average trade size and the log of the number of trades. The last two columns of Table 8 report the permanent treatment effect for each of those two components[10]. Interestingly, the reduction in on-exchange volume (either on Euronext or on other lit markets) is due to small decreases of roughly equal magnitudes (about 5%) in the number of trades and the average trade size. This picture changes considerably once we turn our attention to the OTC market. The coefficients on the average trade size and the number of trades are equal to -0.39 and -0.15, respectively, both significant at the 1% level and summing up to our -0.54 estimate for the logarithm of OTC volume. In other words, the 42% decline in OTC volume on the French OTC market is the combination of a

[9] Because there are a small number of stock-days with zero trading volume for some categories we use $y_{i,t} = \log(1 + volume_{i,t})$.

[10] This analysis is not relevant for the volume in the Euronext auctions, so that we do not decompose the impact on total volume either.

32% decline in average trade size and a 14% decline in the number of trades. For dark pool volume, the decomposition reveals a very similar picture. This observation points towards a disproportionately large impact of the policy change on large off-exchange transactions.

This result also suggests that trading venues with less (or without) HFT activity were impacted to a much larger extent by the policy change. It is thus likely that the contribution of HFT to the total decrease in trading volume on Euronext is small, and that the main impact of the policy change on trading volume was due to the FTT itself, and not to the HFT surcharge.

Notice that so far we have simply adopted the point of view of the tax authorities, for which total volume is a proxy for the tax base. Given that there is a variety of different types of transactions in the OTC market, the reported volume is potentially difficult to compare with that negotiated in lit markets. Therefore, OTC volume will be studied in greater detail in the next subsection.

5.2 A closer look at the OTC market

Trade Sizes: To advance our understanding of the FTT's impact on OTC activity, we first classify trades into different buckets according to size. Small (large) trades have a traded value of less (more) than $25,000$ ($250,000$) EUR, and trades in between these two thresholds are attributed to the medium category[11]. We then compute the trading volume in each category for each stock-day.

The upper panel of Table 9 contains the permanent treatment effects for the trading volume due to small, medium, and large trades. The overall picture is consistent with a sharp decline in large trades for French stocks, as the volume due to this trade category has declined by 60%. While small trades have been affected only slightly, the turnover due to medium size trades appears to have actually *increased* by 15%.

We furthermore define the indicator variable $\mathbf{1}(MaxSize_{i,t} \geq X)$ to take a value of 1 for stock i on day t if there was at least one trade whose value exceeded X EUR, and zero

[11] In order to ensure that our results are robust to differences in trade sizes across stocks, we alternatively classify trades using the "large-in-scale" (LIS) definition set forth in the MiFID regulatory framework, see `http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2006:241:0001:0025:EN:PDF`. The corresponding results are virtually identical.

otherwise. We then use this variable to estimate a linear probability model based on equation (3) for different values of X beyond the threshold for large trades. Because this regression is equivalent to estimating averages of $\mathbf{1}(MaxSize_{i,t} \geq X)$ for sub-periods and/or subsets of stocks, the right-hand side is always guaranteed to be inside the unit interval such that OLS estimates are unbiased and consistent, as discussed in Horrace and Oaxaca (2006). Table 9 contains the estimates which show that the market for extremely large transactions effectively disappeared. For example, the probability of observing a trade above 10,000,000 EUR on a given stock-day decreases by 25 percentage points, which represents more than two thirds of the pre-event probability of 31%.

Execution costs for small trades: In the absence of data on quotes for the OTC market, we can estimate execution costs only based on actual transactions. One issue that arises in this context is that only small transactions are observed with sufficient regularity (there are only four stocks with missing observations for small trades). Given that MiFID allows larger trades to be reported with a considerable delay (in some cases up to several days), the determination of the associated costs appears rather difficult in any case. We therefore compute the average effective half-spread (defined as the relative absolute difference between the transaction price and the contemporaneous[12] midquote on Euronext) for each stock-day based only on transactions classified as small and use it as dependent variable in our DiD framework. For robustness, we also report results for very small trades that do not exceed $10,000$ EUR.

The results[13] are also reported in Table 9. The estimated permanent treatment effect is positive and statistically significant at the 1% level, thus indicating that the introduction of the FTT has significantly increased execution costs for small trades in the OTC market. For example, for transactions with a value smaller than $25,000$ EUR, the introduction of the FTT has increased effective half-spreads by 1.07 bps, which represents 14% of the pre-event average of 7.48 bps. The economic significance is somewhat smaller for transactions below $10,000$ EUR, but still remains considerable with a treatment effect of 0.74 bps compared to a

[12]We omit trades that are reported before the market opening, which constitute a negligible amount of volume. Trades reported after the market close are matched to the last midquote of continuous trading.

[13]In each of the estimations, a small number of stocks have to be discarded due to some days with missing observations.

pre-event average of 6.97 bps. Interestingly, we observe that the impact observed for August is generally smaller, although it is not statistically different from the permanent effect.

In order to check the robustness of these results, we repeat the above calculations after excluding trades that are reported after the market close (5:35 p.m.). The resulting estimates in Table 9 show that the increase in execution costs for stocks subject to the FTT becomes considerably smaller and statistically insignificant once we discard after-hours trading. After the close, liquidity providers do no longer have the option of rebalancing their inventory through the exchange, which therefore may increase the risk of accumulating taxable overnight positions unless they enjoy a market-making exemption. In addition, the scope for higher (effective) trading costs for small trades during the day is limited as investors may always turn to the lit market as an alternative (this is not the case for large trades).

Discussion: Overall, the evidence points towards the policy change having had a more pronounced impact on dark trading compared to lit market segments, so that focusing exclusively on the latter would lead to a significant underestimation of the impact of the tax. Before we discuss some potential explanations for this observation, we briefly comment on some issues that may arise when comparing volume reported on- and off-exchange.

A recent study[14] by the Association for Financial Markets in Europe (AFME) claims that certain types of OTC transactions such as e.g. broker give-ups do not actually constitute "true" liquidity but are rather to be seen as "technical trades" that artificially inflate the reported trading activity. Moreover, some agency trades may eventually lead to double-counting, e.g. in cases where a broker purchases a block on the lit market and the subsequent transfer to the client shows up as an OTC trade. Based on this view, one could be concerned that part of the observed decrease in OTC activity is due to the disappearance of such trades as agents adjust their reporting behaviour to avoid being taxed multiple times for a single transaction. Nevertheless, market participants confirmed to us that the tax authorities have ensured that transactions are only taxed once so that the observed changes in volume are unlikely to be due to strategic considerations. In line with this argument, we find that the impact on dark pool trading is very similar to the treatment effect found for the OTC market,

[14] AFME, "The Nature and Scale of OTC Trading in Europe", 2011.

which suggests that we are successful at capturing changes in off-exchange activity.

While the heterogeneous nature of OTC trading makes it difficult to compare to lit trading volume, it is important to stress that the rise of HFT over the past decade has also led to a transformation of on-exchange liquidity. In particular, recent research advances the idea that trading volume has become an increasingly poor indicator of market liquidity due to increased intermediation activity. Menkveld and Yueshen (2013) provide an interesting argument in this direction by suggesting that high trading volume may signal that middlemen fail to find end-investors and thus are forced to pass their inventory on to other HFTs, akin to the hot-potato effect identified in Kirilenko *et al.* (2011). While we acknowledge that our estimates for the impact on overall trading volume may be somewhat influenced by differences in the nature of (reported) volumes across trading mechanisms, this heterogeneity should not affect the estimated effects for the individual marketplaces.

One potential explanation for a larger impact on OTC trading is the possibility that volume reported off-exchange is more representative for investors that are subject to the FTT. Given that HFT activity accounts for roughly 40% of the total trading volume in European equities according to recent estimates[15], it seems plausible that on-exchange volume is a rather noisy proxy for the activity of investors that engage in the accumulation of net positions. While we have shown in Section 4.4 that low-latency activity was affected significantly, a large chunk of the HFT community is based in the United Kingdom and thus not directly affected by the policy change.

Although the FTT also exempts market-making activities that take place away from regulated markets such as exchanges and MTFs, it is important to point out that market participants are required to clearly separate market-making from other trading desks in order to enjoy this privilege. Clearly, this may impair the provision of liquidity by agents that do so on an irregular, informal, or opportunistic basis (e.g. hedge funds), in particular for large block trades that are too big to be unwound in a single trading day (see e.g. Duffie (2012)) and thus lead to the accumulation of taxable overnight positions. The virtual disappearance of very large OTC transactions is consistent with this intuition, as they are likely to involve considerable inventory positions. Our results on trade sizes suggest that due

[15] "Understanding high-frequency trading", World Federation of Exchanges

to lacking liquidity for large blocks, liquidity consumers may prefer to split up their orders, leading to more medium-sized trades. Similarly, the observation on elevated trading costs when including transactions occurring after the official market close points into the direction of supply effects, because in these instances liquidity providers are not able to use the lit market to close out net positions before the end of the day and thus may face an increased risk of ending up with a taxable overnight position.

One additional explanation of the differential impact between on and off-exchange trading suggested to us by practitioners is the role played by mutual funds eligible for French "PEA" accounts. These popular investment vehicles offer a tax advantage to retail investors if they invest more than 75% of their portfolio in European equities. In the past few years, an increasing number of funds have been using total return swaps to offer exposure to other assets while still meeting this requirement[16]. As a consequence of the FTT, these funds may have substituted domestic stocks with other European stocks, thus contributing to a drop in OTC transactions for French equities.

Overall, our results suggest that the OTC market was significantly more impacted than on-exchange trading, and at least some of this discrepancy appears to be rooted in liquidity safeguards having a heterogeneous effect across different market segments. In addition to this, this differential impact may have been amplified via substitution effects between lit and dark trading. Consistent with this interpretation, Degryse, de Jong, and van Kervel (2013) provide evidence that a decrease in OTC activity is associated with an increase in liquidity in electronic limit order markets.

5.3 Tax revenues and trade-offs

Equipped with a more complete picture of the FTT's impact on trading volume, we can now turn to contrasting the costs and revenues from the implementation of this policy with some simplified calculations. While no official figures have been released yet, an unofficial revenue of 250 million EUR for the period August-December 2012 has circulated in the French

[16]See e.g. http://www.challenges.fr/patrimoine/20120906.CHA0432/comment-internationaliser-son-pea.html.

press[17]. Extrapolating this figure to a full year yields an annual revenue of 600 million EUR. Importantly, this figure falls considerably short of the French authorities' estimate of 1.6 billion EUR per year, which points at an underestimation of the impact on revenue-generating market activity.

While the observed decrease in trading volume is rather large, a rigorous evaluation of the associated welfare losses is an extremely complex task. We can at least roughly estimate losses to market participants under some assumptions. A transaction T taking place in market segment P prior to the imposition of the tax implied gains from trade that were higher than the effective half-spread $ES_{T,P}$ times the volume of the transaction $volume_{T,P}$. Because it is rather difficult to obtain reliable cost estimates for large OTC transactions (e.g. due to delayed reporting for large trades) and those in other lit markets are likely to be of a similar magnitude, we assume that the value-weighted average half-spread on Euronext (5.12 bps for the pre-event period) is a reasonable lower bound on the average transaction cost. Given annual trading volume for mechanism P[18], the associated deadweight loss is then given by $volume_P \times 5.12bps \times (1 - exp(-\beta_P^{Sep/Oct}))$. Summing up over our 5 possible trading mechanisms, our estimated lower bound for deadweight losses to market participants is equal to 473 million EUR per year.

Together with the 600 million EUR paid in taxes, these unrealized gains from trade yield a total cost to market participants of around 1.1 billion EUR per year, or 1.79 times the tax revenue. Even though many transactions may involve traders simply taking opposite bets, a zero-sum situation with no trading gain, our simple computation probably underestimates trading gains, especially on the OTC market where even the spreads on small transactions are already above those assumed here. This estimate is to be interpreted rather as measuring market participants' resistance to such a tax, and not necessarily as a welfare loss. The most significant costs in terms of welfare would additionally come from spillovers to the real sector, something that we cannot take into account here.

Given the current debate around the foreseen implementation of a pan-European FTT,

[17] See e.g. http://www.lefigaro.fr/conjoncture/2013/03/03/20002-20130303ARTFIG00128-la-recette-decevante-de-la-taxe-tobin.php.

[18] This can be easily computed by multiplying the stock-day average trading volumes in Table 7 with the number of stock-days in the pre-event period June-July (3,698) and annualizing with a factor of 6.

a relevant question to ask is how much one would raise by applying the French model to the remaining countries of the EU, or at least to the 11 states that have committed to implementing an FTT. To obtain such an estimate, we begin by assuming that a country's tax revenue would be proportional to its share in the total turnover of European stocks with a market capitalization above 1 billion EUR (the tax base). If we call T_i this total turnover in country i and $i = F$ for France, we thus assume that the yearly revenue of the tax is equal to $600(T_i/T_F)$ million EUR. Table 10 presents the results for several European countries, where the stock turnover for country i is based on stocks contained in the Stoxx 600 index[19]. An interesting cross-check consists in comparing our estimate for the UK, 724 million EUR, with the actual revenue obtained from the UK stamp duty. Over the period August-December 2012[20], HM Treasury received 977 million GBP, which corresponds to 2.9 billion EUR for a full year based on an average exchange rate of 1.2475 EUR/GBP. This is about 4 times our estimate, but given that the stamp duty (50 bps) is 2.5 times higher than the French FTT and is also levied on smaller stocks, our estimates appear rather reasonable. For all EU-countries, we arrive at an estimated revenue of 3.4 billion EUR, and restricting the set to the group of 11 that has committed to the FTT yields 2.1 billion EUR.

In its impact assessment of the European FTT, the European Commission uses a simple formula with various parameterizations to estimate potential revenues of a 10 bps tax on all equity transactions in the EU-27 (payable by both sellers and buyers, no exemption from market-making, no limitation to the largest companies). The estimates lie between 4.8 and 6.5 billion EUR[21] and are probably quite optimistic. The comparison suggests that the potentially lost revenue through an exemption of market-making and a restriction to stocks of large companies may be small compared to the potential inefficiencies arising from a tougher implementation.

[19] We compute country shares based on the turnover in the 6 months prior to the FTT's adoption. Note that several smaller EU countries are not represented in the index, but this is due to their low stock market capitalization. Given the 1 billion EUR threshold, this is unlikely to affect our results (the smallest ES600 component has a market capitalization below this threshold).

[20] See http://www.hmrc.gov.uk/statistics/receipts/info-analysis.pdf.

[21] The documentation is available on the Commission's website: http://ec.europa.eu/taxation_customs/taxation/other_taxes/financial_sector/.

6 Heterogeneity in the cross-section

We further investigate the impact of the policy change by examining potential heterogeneity in the treatment effect across stocks. To this end we sort stocks into terciles according to market capitalization (as of May 31st 2012), liquidity (proxied by the inverse of the Amihud (2002) illiquidity measure), and stock price volatility (defined as the standard deviation of daily returns). The latter two measures are computed using daily data from Datastream for the 6 months preceding our sample period. We then estimate the following type of regression for each characteristic, where G_k denotes the set of stocks in the kth tercile:

$$\begin{aligned} y_{i,t} &= \sum_{k=1}^{3} \alpha_k \mathbf{1}(i \in G_k) + \sum_{k=1}^{3} \gamma_k \mathbf{1}(i \in G_k \cap \Theta) + \sum_{k=1}^{3} \delta_k^{Aug} \mathbf{1}(i \in G_k, t \in \text{Aug}) + \sum_{k=1}^{3} \delta_k^{Sep/Oct} \mathbf{1}(i \in G_k, t \in \text{Sep/Oct}) \\ &+ \sum_{k=1}^{3} \beta_k^{Aug} \mathbf{1}(i \in G_k \cap \Theta, t \in \text{Aug}) + \sum_{k=1}^{3} \beta_k^{Sep/Oct} \mathbf{1}(i \in G_k \cap \Theta, t \in \text{Sep/Oct}) \end{aligned} \tag{5}$$

We additionally compare the treatment effect for the constituents of the EuroStoxx50 index (hereafter ES50) and the remaining stocks.

Impact on the lit market: Table 11 in the Appendix contains the coefficient estimates for the FTT's impact on lit market quality across the different terciles for each characteristic. Overall, there is relatively little heterogeneity in the cross-section. One notable exception is trading volume, which decreases most for relatively small and illiquid stocks as well as for stocks with either high or low volatility. More importantly, there are also some differences with respect to the impact on market depth. In particular, it is only the largest and most liquid stocks that experience a significant decline in the available liquidity at the inside quotes, while the remaining stocks are at most marginally affected. Clearly, the overall negative impact on this dimension of market quality is driven by this group of stocks.

Interestingly, we find the most pronounced differences across stocks when sorting according to ES50 membership. Index members experience no significant decrease in trading activity, but at the same time display an increase in quoted and effective spreads as well as a decrease in resiliency. Realized volatility also increases slightly, although the coefficient is statistically insignificant.

Turning to our measures of low-latency activity, we find that the previously documented decrease is mostly concentrated in the smallest and most illiquid stocks, in particular if we restrict our attention to the median order lifetime. Given that the HFT tax is only levied on French traders, this effect is most likely due to local HFT players being relatively more important in smaller and less liquid stocks (notably the effect is again strongest when sorting on ES50 membership) and the policy change inducing them to either submitting more market orders or alternatively dropping out of the market. For the largest stocks, HFT activity seems to have been only marginally affected.

Impact on other trading venues: Turning to the results on trading volume across different venues / market segments (Table 12), we find that the largest and most liquid stocks display the strongest decrease in OTC trading, which is in stark contrast to the findings for on-exchange trading volume (Euronext LOB and Other LOB), where they constitute the least affected group. This is consistent with a possible substitution effect between on and off-exchange trading already mentioned earlier. In fact it seems plausible that the magnitude of this effect depends on lit-market liquidity, i.e. on-exchange trading is a relatively more attractive alternative for liquid stocks. Interestingly, dark pool activity and call auction trading were most strongly affected for illiquid and small stocks, with decreases of up to 72% and 33%, respectively. In terms of execution quality on the OTC market, we observe that the largest and most liquid stocks are the least affected, which is in stark contrast to them displaying the largest decreases in OTC volume. A glance at Table 13 reveals that very large/liquid and small/illiquid stocks display the most pronounced decrease in trading volume for very large trades, which therefore suggests that the segment of large trades was relatively more important for the former group of stocks. Notice that this is not inconsistent with the large differences across the two groups for the estimated impact on the probability to observe large trades, as the pre-event probabilities are very different.

7 Conclusion

This paper has examined the causal impact of the French FTT on market quality. While the effect on exchange-based trading on Euronext is relatively modest with a decrease of 10% in volume, a slight worsening of market quality and a reduction in low-latency activity, we show that the OTC market has been affected to a much larger degree. Off-exchange trading has declined by more than 40%, with large block trades disappearing almost completely.

Our evidence suggests that modern electronic markets may be much more resilient to the introduction of an FTT compared to what is suggested by previous episodes. In particular, the exemptions of intraday trading and market-making activities appear to have been successful in avoiding a larger drop in market liquidity for exchange-based trading. In contrast, dark trading has proven much more sensitive to the FTT, which is likely due to a more informal liquidity provision and a greater risk of accumulating overnight positions when taking on large block trades. This differential impact was apparently not expected by the French authorities, as suggested by a considerable shortfall in preliminary revenue figures for the period August - December 2012.

The mentioned discrepancy across trading mechanisms underlines the importance of protecting market functioning by adopting (possibly market-specific) safeguards to liquidity provision in order to mitigate the adverse impact of an FTT. Notably, the current draft of the pan-European FTT as put forth by the European Commission does not foresee any exemptions for market-making or other related activities and applies homogenously to different market structures, which suggests that one may expect a much more severe impact.

A Appendix

A.1 Tables

Table 1: List of stocks in the treated and the control groups.

Treated stocks				Control stocks	
Accor	Faurecia	Lafarge	Safran	Aalberts Industries	SBM Offshore
Aeroports de Paris	Eramet	Lagardere	Sanofi	Aegon	SES
Air France-KLM	Essilor International	Legrand	Schneider Electric	Ahold	Eurocommercial Properties
Air Liquide	Eutelsat Communications	Klepierre	SCOR	Akzo Nobel	STMicroelectronics
Arkema	Eurazeo	LVMH Moet Hennessy Louis Vuitton	SEB	ASM International	TNT Express
Alstom	Havas	Etablissements Maurel et Prom	Suez Environnement Co	ASML Holding	Unilever
Alcatel-Lucent	Sodexo	Mercialys	Vinci	Boskalis Westminster	Vopak
AXA	Eiffage	Cie Generale des Etablissements Michelin	Cie de St-Gobain	Corio	Wereldhave
Societe BIC	France Telecom	Metropole Television	Silic	Delta Lloyd	Wolters Kluwer
BioMerieux	Rallye	Wendel	Thales	DSM	
Bouygues	Cie Generale de Geophysique - Veritas	Nexans	Technip	EADS	
Bureau Veritas	Groupe Eurotunnel	Neopost	Societe Television Francaise 1	Reed Elsevier	
Cap Gemini	Gecina	L'Oreal	Total	Fugro	
Carrefour	Bourbon	Orpea	Unibail-Rodamco	Heineken	
Casino Guichard Perrachon	GDF Suez	Pernod-Ricard	Veolia Environnement	Royal Imtech	
CNP Assurances	Gemalto	Peugeot	Virbac	ING Groep	
Danone	ICADE	PPR	Vivendi	ArcelorMittal	
Dassault Systemes	Iliad	Publicis Groupe	Vallourec	KPN	
Christian Dior	Imerys	Remy Cointreau	Valeo	Nutreco	
Edenred	Ingenico	Renault	Zodiac Aerospace	Philips Electronics	
Electricite de France	Ipsen	Rubis		Randstad Holding	
Euler Hermes	JCDecaux	Rexel		Royal Dutch Shell	

Table 2: Summary statistics for treated and control stocks.

This tables contains sample statistics across all stock-days (12,753) observations for each of the market quality measures described in Section 4.1.

	Treated Group			Control Group		
	Mean	Median	Std. dev.	Mean	Median	Std. dev.
$volume_{i,t}$ (million EUR)	26	11	36	34	17	43
$RV_{i,t}$ (%)	27	26	11	25	23	11
$range_{i,t}$ (%)	2.8	2.4	1.5	2.5	2.2	1.4
$QS_{i,t}$ (bps)	5.7	4.3	4.2	4.2	3.8	2.2
$ES_{i,t}$ (bps)	4.4	3.4	3.1	3.3	3.0	1.6
$depth_{i,t}$ (1,000 EUR)	58	50	51	81	73	48
$res_{i,t}$	0.49	0.50	0.15	0.53	0.53	0.13
$\|AR\|_{i,t}$	0.11	0.090	0.086	0.11	0.088	0.091
Number of stocks		86			31	

Table 3: Impact of the policy change on market quality.

This table contains the regression coefficient from estimating equation (3) via OLS for all market quality measures described in Section 4.1. The last column contains the difference between the treatment effects for September and October from an augmented model with month-specific coefficients. T-statistics based on standard errors clustered at the stock level are given in parentheses. $***$, $**$, and $*$ denote statistical significance at the 1%, 5%, and 10% level, respectively.

Variable/Coefficient	β^{Aug}	$\beta^{Sep/Oct}$	$\beta^{Oct} - \beta^{Sep}$
$logvolume_{i,t}$	-0.38^{***}	-0.10^{**}	0.022
	(−7.91)	(−2.31)	(0.46)
$RV_{i,t}$	−0.88	0.52	−1.01
	(−1.36)	(0.59)	(−1.22)
$range_{i,t}$	-0.14^{*}	−0.040	−0.086
	(−1.74)	(−0.37)	(−1.05)
$QS_{i,t}$	0.069	−0.019	0.039
	(0.37)	(−0.08)	(0.27)
$ES_{i,t}$	−0.020	0.029	0.056
	(−0.13)	(0.16)	(0.56)
$depth_{i,t}$	-11.4^{***}	-11.0^{***}	0.058
	(−2.73)	(−2.84)	(0.03)
$res_{i,t}$	-0.020^{*}	-0.017^{*}	0.0082
	(−1.74)	(−1.74)	(0.83)
$\|AR\|_{i,t}$	−0.0039	0.0073^{**}	−0.00095
	(−0.77)	(2.08)	(−0.20)

Table 4: Test for seasonal effects 2009-2011.

This table contains the regression coefficients from estimating a placebo diff-in diff based on equation (3) for all market quality measures described in Section 4.1 with month-specific treatment effects using data from 2009-2011. 3 Stocks are dropped from the initial sample in Table 1 due to incomplete data. T-statistics based on standard errors clustered at the stock level are given in parentheses. $***$, $**$, and $*$ denote statistical significance at the 1%, 5%, and 10% level, respectively.

Variable/Coefficient	β^{Aug}	β^{Sep}	β^{Oct}
$logvolume_{i,t}$	-0.17^{***}	-0.0054	-0.036
	(-3.84)	(-0.13)	(-0.85)
$RV_{i,t}$	-1.49^{***}	-0.86	-0.51
	(-2.81)	(-1.23)	(-1.06)
$range_{i,t}$	-0.15^{**}	0.053	-0.041
	(-2.36)	(1.13)	(-0.80)
$QS_{i,t}$	-0.096	-0.19^{*}	-0.22^{**}
	(-0.77)	(-1.74)	(-2.14)
$ES_{i,t}$	-0.056	-0.00	-0.023
	(-0.58)	(-0.00)	(-0.28)
$depth_{i,t}$	0.61	-1.56	0.27
	(0.44)	(-1.02)	(0.18)
$res_{i,t}$	-0.0085	0.0069	-0.0025
	(-1.62)	(1.39)	(-0.50)
$\|AR\|_{i,t}$	-0.00027	-0.0019	0.0036
	(-0.09)	(-0.67)	(1.27)

Table 5: Diff-in-diff-in-diff estimates for August.

This table contains the regression coefficients from estimating the diff-in-diff-in-diff model in equation (4). 3 Stocks are dropped from the initial sample due to incomplete data. T-statistics based on standard errors clustered at the stock level are given in parentheses. $***$, $**$, and $*$ denote statistical significance at the 1%, 5%, and 10% level, respectively.

Variable/Coefficient	β^{12}
$logvolume_{i,t}$	-0.21^{***}
	(-4.26)
$RV_{i,t}$	0.59
	(0.56)
$range_{i,t}$	-0.0030
	(-0.02)

Table 6: Policy impact on message traffic.

The first two columns contain the regression coefficients from estimating equation (3) via OLS for the three measure of HFT activity described in Section 4.4. T-statistics based on standard errors clustered at the stock level are given in parentheses. $***$, $**$, and $*$ denote statistical significance at the 1%, 5%, and 10% level, respectively. The last two columns report the averages (across stock-days) for each variable and group of stocks.

Variable/Coefficient	β^{Aug}	$\beta^{Sep/Oct}$	Pre-tax average	
			Treated	Untreated
$median.lapse_{i,t}$	3.8^{***}	2.4^{***}	3.6	3.3
	(3.98)	(3.83)		
$cancel.0-500_{i,t}$	-0.066^{***}	-0.054^{***}	0.34	0.32
	(−5.34)	(−4.13)		
$pct.inside_{i,t}$	-0.052^{***}	-0.073^{***}	0.37	0.38
	(−4.44)	(−4.98)		

Table 7: Summary statistics on trading volume.

This table contains sample statistics for trading volume, the number of trades and trade sizes across both groups of stocks and for each individual category of trading mechanisms. All figures are based on stock-days.

	Treated stocks											
	Value (in million EUR)			No. of trades (in 1,000)			Size (in 1,000 EUR)			Share (%)		
	Mean	Median	StDev	Mean	Median	StDev	Mean	Median	StDev	Mean	Median	StDev
Euronext LOB	25.8	11.5	35.7	1.6	1.2	1.4	11.6	9.8	7.8	36.4	14.3	31.4
Euronext Auction	6.9	2.3	14.6	-	-	-	-	-	-	6.8	4.6	6.4
Other LOB	16.5	7.2	23.6	3.8	2.5	3.8	3.3	3.0	1.7	23.7	8.6	20.0
OTC	84.4	13.5	323.9	0.6	0.4	0.7	109.2	37.5	320.9	32.0	22.5	36.2
Dark	1.2	0.4	2.2	0.1	0.1	0.2	7.1	4.7	17.3	1.2	2.3	0.6
All	134.8	41.7	364.8	6.2	4.2	5.9	14.3	8.5	20.3	-	-	-

	Control stocks											
	Value (in million EUR)			No. of trades (in 1,000)			Size (in 1,000 EUR)			Share (%)		
	Mean	Median	StDev	Mean	Median	StDev	Mean	Median	StDev	Mean	Median	StDev
Euronext LOB	34.1	17.4	43.0	1.8	1.5	1.3	15.1	12.6	8.2	31.2	11.4	36.4
Euronext Auction	6.7	3.0	16.3	-	-	-	-	-	-	7.3	3.5	6.2
Other LOB	22.8	10.7	28.0	4.7	3.2	4.2	3.9	3.7	1.6	19.5	8.0	23.2
OTC	47.0	13.8	100.0	0.7	0.5	0.7	53.7	27.6	87.6	40.9	16.0	29.3
Dark	1.3	0.5	2.1	0.2	0.1	0.2	7.2	4.7	18.8	1.2	1.4	0.8
All	112.1	51.5	160.8	7.4	5.5	6.1	11.2	8.6	8.5	-	-	-

Table 8: Impact on different categories of trading volume.

This table contains the regression coefficients from estimating equation (3) via OLS for the log of trading volume for different trading mechanisms. The last two columns decompose the impact on the log of trading volume into the impact on the log of average trade size and the log of the number of trades. T-statistics based on standard errors clustered at the stock level are given in parentheses. $***$, $**$, and $*$ denote statistical significance at the 1%, 5%, and 10% level, respectively.

	Log Trading volume		Log Avg. trade size	Log No. of trades
Trading mechanism/Coefficient	β^{Aug}	$\beta^{Sep/Oct}$	$\beta^{Sep/Oct}$	$\beta^{Sep/Oct}$
Euronext LOB	-0.38^{***}	-0.10^{**}	-0.043	-0.056^{*}
	(-7.91)	(-2.31)	(-1.51)	(-1.67)
Euronext Auction	-0.30^{***}	-0.17^{***}	–	–
	(-7.04)	(-3.84)	–	–
Other LOB	-0.34^{***}	-0.16^{***}	-0.061^{**}	-0.095^{**}
	(-7.67)	(-3.41)	(-2.12)	(-2.11)
OTC	-0.91^{***}	-0.54^{***}	-0.39^{***}	-0.15^{***}
	(-7.54)	(-6.18)	(-4.61)	(-3.34)
Dark	-1.04^{***}	-0.48^{***}	-0.34^{***}	-0.14
	(-5.76)	(-2.72)	(-3.65)	(-1.52)
All	-0.68^{***}	-0.36^{***}	–	–
	(-8.32)	(-6.21)	–	–

Table 9: Impact on the OTC market for different trade sizes.

The top panel of this table contains the regression coefficients from estimating equation (3) via OLS for the log of trading volume based on different trade sizes. The middle panel contains the regression coefficients from a linear probability model based on equation (3), where the dependent variable takes a value of 1 for a given stock-day if at least one trade above a certain threshold was observed, and zero otherwise. The bottom panel coefficients come from estimating equation (3) via OLS for OTC execution costs based on trades with a size below a given threshold. T-statistics based on standard errors clustered at the stock level are given in parentheses. $***$, $**$, and $*$ denote statistical significance at the 1%, 5%, and 10% level, respectively.

Threshold	β^{Aug}	$\beta^{Sep/Oct}$	Pre-tax average	
			Treated	Untreated
Impact on Log volume for trades of size X EUR:				
$X \leq 25,000$ EUR	-0.25^{***}	-0.092^{*}	13.7	14.3
	(-3.45)	(-1.73)		
$25,000 \leq X \leq 250,000$EUR	-0.11	0.16^{**}	14.0	14.4
	(-1.45)	(2.21)		
$X \geq 250,000$ EUR	-1.91^{***}	-0.93^{***}	14.9	15.5
	(-7.51)	(-3.67)		
Impact on probability to observe a trade of size at least X EUR:				
$250,000$ EUR	-0.073^{***}	-0.024	0.90	0.95
	(-4.63)	(-1.58)		
$1,000,000$ EUR	-0.099^{***}	-0.073^{***}	0.70	0.74
	(-4.15)	(-2.87)		
$2,500,000$ EUR	-0.18^{***}	-0.18^{***}	0.53	0.50
	(-7.63)	(-7.72)		
$10,000,000$ EUR	-0.22^{***}	-0.25^{***}	0.31	0.24
	(-5.79)	(-8.23)		
Impact on execution costs, trades below X EUR:				
$10,000$ EUR	0.50^{**}	0.74^{**}	6.97	5.93
	(2.06)	(2.52)		
$25,000$ EUR	0.65^{**}	1.07^{***}	7.48	6.23
	(2.60)	(3.43)		
Impact on execution costs, trades during market hours, below X EUR:				
$10,000$ EUR	0.32	0.30	6.24	5.45
	(1.45)	(1.13)		
$25,000$ EUR	0.36	0.34	6.60	5.62
	(1.50)	(1.21)		

Table 10: Country share of EuroStoxx600 turnover for companies with market capitalization above 1 billion EUR, and imputed potential tax revenues.

This table contains the country-specific shares in Euro Stoxx 600 turnover (in percentage points) as well as the resulting revenue estimates for an implementation of the French FTT. Details can be found in Section 5.3. The underlined countries are among the 11 that have committed to the implementation of an FTT.

Country	Turnover (share)	Estimated revenues (Million EUR)
Austria	0.3	11.7
Belgium	1.3	46.1
Denmark	1.4	49.8
Finland	2.0	68.2
France	17.5	600.0
Germany	19.0	651.7
Greece	0.1	3.9
Ireland	0.7	23.4
Italy	9.6	328.6
Luxembourg	0.9	31.2
Netherlands	6.0	205.1
Portugal	0.3	10.6
Spain	13.8	475.1
Sweden	6.0	207.8
United Kingdom	21.1	724.3
Total - EU	100.0	3437.6
Total - 11	61.9	2127.7

Table 11: Impact of the policy change on the Euronext limit order book, in the cross-section.

This table contains the regression coefficients from estimating equation (5) via OLS for the variables considered in Tables 3 and 6, allowing for a heterogeneous impact for different terciles of market capitalization, liquidity and volatility, and for stocks part of the EuroStoxx50 index vs. others. $\beta_1^{Sep/Oct}$ represents the impact for the largest, most liquid, or most volatile stocks, respectively. T-statistics based on standard errors clustered at the stock level are given in parentheses. $***$, $**$, and $*$ denote statistical significance at the 1%, 5%, and 10% level, respectively.

Variable/Sort	Market capitalization			Liquidity			Volatility			EuroStoxx50	
	$\beta_1^{Sep/Oct}$	$\beta_2^{Sep/Oct}$	$\beta_3^{Sep/Oct}$	$\beta_1^{Sep/Oct}$	$\beta_2^{Sep/Oct}$	$\beta_3^{Sep/Oct}$	$\beta_1^{Sep/Oct}$	$\beta_2^{Sep/Oct}$	$\beta_3^{Sep/Oct}$	$\beta_{Non-ES50}^{Sep/Oct}$	$\beta_{ES50}^{Sep/Oct}$
$logvolume_{i,t}$	−0.078	0.013	−0.17*	−0.11	0.028	−0.18*	−0.16***	−0.037	−0.11**	−0.11**	−0.083
	(−1.28)	(0.28)	(−1.73)	(−1.62)	(0.54)	(−1.83)	(−2.91)	(−0.38)	(−2.30)	(−2.18)	(−1.51)
$RV_{i,t}$	0.083	1.07	2.15	−0.66	2.85**	0.52	−0.35	1.14	1.07	0.12	2.77
	(0.08)	(0.59)	(1.38)	(−0.64)	(2.50)	(0.17)	(−0.17)	(0.82)	(1.40)	(0.12)	(1.59)
$range_{i,t}$	−0.094	0.088	0.13	−0.074	0.25**	−0.064	−0.25	0.058	0.10	−0.089	0.20
	(−0.75)	(0.40)	(0.67)	(−0.60)	(1.99)	(−0.18)	(−0.99)	(0.33)	(1.17)	(−0.74)	(0.95)
$QS_{i,t}$	0.099	−0.069	0.61	−0.015	0.29*	0.45	−0.20	0.44	−0.25	−0.11	0.38**
	(0.61)	(−0.33)	(1.03)	(−0.09)	(1.94)	(0.61)	(−0.49)	(0.90)	(−0.77)	(−0.40)	(2.50)
$ES_{i,t}$	0.050	0.079	0.50	−0.034	0.37***	0.32	−0.11	0.090	−0.008	−0.038	0.31**
	(0.37)	(0.42)	(1.30)	(−0.26)	(2.95)	(0.60)	(−0.37)	(0.38)	(−0.03)	(−0.18)	(2.53)
$depth_{i,t}$	−21***	−3.3	0.30	−23***	−4.9*	2.2	−8.8*	0.53	−25***	−8.6**	−26***
	(−2.75)	(−1.05)	(0.14)	(−2.75)	(−1.90)	(0.93)	(−1.72)	(0.18)	(−2.98)	(−2.09)	(−3.27)
$RES_{i,t}$	−0.013	−0.0068	−0.0076	−0.015	−0.0004	−0.013	−0.038**	−0.009	−0.006	−0.014	−0.047***
	(−0.85)	(−0.48)	(−0.53)	(−0.98)	(−0.03)	(−0.74)	(−2.31)	(−0.55)	(−0.37)	(−1.47)	(−3.01)
$\|AR\|_{i,t}$	0.019***	−0.007	0.012	0.013**	−0.002	0.016*	0.010*	0.009	0.002	0.0072*	0.0052
	(3.82)	(−1.43)	(1.49)	(2.38)	(−0.41)	(1.90)	(1.94)	(1.28)	(0.37)	(1.86)	(0.87)
$median.lapse_{i,t}$	0.27	2.0***	2.7*	0.32*	0.85	3.49**	2.27*	3.21***	4.35*	2.95***	0.19
	(1.53)	(4.23)	(1.94)	(1.66)	(1.30)	(2.18)	(1.85)	(2.88)	(1.77)	(4.09)	(1.46)
$cancel.0-500_{i,t}$	−0.028*	−0.034**	−0.058**	−0.033**	−0.021	−0.064**	−0.049**	−0.056***	−0.055**	−0.062***	−0.019*
	(−1.88)	(−2.54)	(−2.56)	(−2.10)	(−1.54)	(−2.25)	(−2.02)	(−3.11)	(−2.28)	(−4.12)	(−1.91)
$pct.inside_{i,t}$	−0.035**	−0.093***	−0.045*	−0.042***	−0.062***	−0.057**	−0.057*	−0.084***	−0.076***	−0.083***	−0.026
	(−2.33)	(−4.88)	(−1.83)	(−2.66)	(−2.98)	(−2.07)	(−1.88)	(−3.92)	(−3.28)	(−5.03)	(−1.27)

Table 12: Impact of the policy change on volume on other trading venues, in the cross-section.

This table contains the regression coefficients from estimating equation (5) via OLS for the variables considered in Table 8, allowing for a heterogeneous impact for different terciles of market capitalization, liquidity and volatility, and for stocks part of the EuroStoxx50 index vs. others. $\beta_1^{Sep/Oct}$ represents the impact for the largest, most liquid, or most volatile stocks, respectively. T-statistics based on standard errors clustered at the stock level are given in parentheses. $***$, $**$, and $*$ denote statistical significance at the 1%, 5%, and 10% level, respectively.

Variable/Sort	Market capitalization			Liquidity			Volatility			EuroStoxx50	
	$\beta_1^{Sep/Oct}$	$\beta_2^{Sep/Oct}$	$\beta_3^{Sep/Oct}$	$\beta_1^{Sep/Oct}$	$\beta_2^{Sep/Oct}$	$\beta_3^{Sep/Oct}$	$\beta_1^{Sep/Oct}$	$\beta_2^{Sep/Oct}$	$\beta_3^{Sep/Oct}$	$\beta_{Non-ES50}^{Sep/Oct}$	$\beta_{ES50}^{Sep/Oct}$
Log of trading volume											
Euronext LOB	−0.078	0.013	−0.17*	−0.11	0.028	−0.18*	−0.16***	−0.037	−0.11**	−0.11**	−0.083
	(−1.28)	(0.28)	(−1.73)	(−1.62)	(0.54)	(−1.83)	(−2.91)	(−0.38)	(−2.30)	(−2.18)	(−1.51)
Euronext Auction	−0.079	−0.007	−0.40***	−0.062	−0.099*	−0.33**	−0.24***	−0.14*	−0.16***	−0.18***	−0.15***
	(−1.48)	(−0.11)	(−4.09)	(−1.03)	(−1.72)	(−2.52)	(−2.76)	(−1.67)	(−3.42)	(−3.61)	(−2.96)
Other LOB	−0.10*	−0.024	−0.30***	−0.12*	−0.11	−0.17	−0.19**	−0.11	−0.16***	−0.17***	−0.11
	(−1.67)	(−0.32)	(−3.38)	(−1.73)	(−1.57)	(−1.66)	(−2.40)	(−1.29)	(−2.71)	(−3.32)	(−1.48)
OTC	−0.75***	−0.41***	−0.52***	−0.76***	−0.46***	−0.49***	−0.44***	−0.61***	−0.60***	−0.48***	−0.91***
	(−5.33)	(−3.70)	(−3.34)	(−4.93)	(−4.11)	(−2.92)	(−3.85)	(−3.02)	(−5.42)	(−4.80)	(−8.93)
Dark	−0.42***	0.18	−1.28**	−0.39***	−0.10	−1.19*	−0.69**	−0.39	−0.39*	−0.51**	−0.35***
	(−3.13)	(1.10)	(−2.32)	(−2.74)	(−0.59)	(−1.79)	(−2.10)	(−1.07)	(−1.95)	(−2.50)	(−2.61)
All	−0.43***	−0.24***	−0.37***	−0.45***	−0.27***	−0.34***	−0.34***	−0.31**	−0.43***	−0.33***	−0.52***
	(−4.51)	(−3.67)	(−3.67)	(−4.25)	(−4.04)	(−3.15)	(−5.09)	(−2.31)	(−6.48)	(−5.12)	(−7.73)

Table 13: Impact on the OTC market for different trade sizes, in the cross-section.

This table contains the regression coefficients from estimating equation (5) via OLS for the variables considered in Table 9, allowing for a heterogeneous impact for different terciles of market capitalization, liquidity and volatility, and for stocks part of the EuroStoxx50 index vs. others. $\beta_1^{Sep/Oct}$ represents the impact for the largest, most liquid, or most volatile stocks, respectively. T-statistics based on standard errors clustered at the stock level are given in parentheses. $***$, $**$, and $*$ denote statistical significance at the 1%, 5%, and 10% level, respectively.

Threshold/Sort	Market capitalization			Liquidity			Volatility			EuroStoxx50	
	$\beta_1^{Sep/Oct}$	$\beta_2^{Sep/Oct}$	$\beta_3^{Sep/Oct}$	$\beta_1^{Sep/Oct}$	$\beta_2^{Sep/Oct}$	$\beta_3^{Sep/Oct}$	$\beta_1^{Sep/Oct}$	$\beta_2^{Sep/Oct}$	$\beta_3^{Sep/Oct}$	$\beta_{Non-ES50}^{Sep/Oct}$	$\beta_{ES50}^{Sep/Oct}$
Impact on Log volume for trades of size X EUR:											
$X \leq 25,000$	−0.15**	0.050	−0.15	−0.17**	−0.10	0.075	−0.10	−0.070	−0.10	−0.078	−0.18**
	(−2.25)	(0.47)	(−1.45)	(−2.27)	(−1.54)	(0.47)	(−0.90)	(−0.76)	(−1.33)	(−1.29)	(−2.40)
$25,000 \leq X \leq 250,000$	0.039	0.29**	0.22	0.038	0.24***	0.27	0.13	0.095	0.23*	0.18**	0.033
	(0.57)	(4.41)	(1.07)	(0.52)	(3.74)	(1.25)	(1.56)	(0.65)	(1.90)	(2.17)	(0.38)
$X \geq 250,000$	−0.88***	−0.31	−1.67**	−0.85***	−0.22	−2.1**	−0.66***	−1.50**	−0.64**	−0.92***	−1.0***
	(−5.47)	(−1.19)	(−1.98)	(−5.06)	(−0.84)	(−2.30)	(−2.60)	(−2.37)	(−2.53)	(−3.13)	(−9.51)
Impact on probability to observe a trade of size at least X EUR:											
250,000	−0.003	0.015	−0.086	0.00	0.025**	−0.12**	−0.012	−0.067*	0.006	−0.029	0.00
	(−1.04)	(1.27)	(−1.53)	(*n.a.*)	(2.04)	(−1.96)	(−0.92)	(−1.68)	(0.40)	(−1.62)	(*n.a.*)
1,000,000	−0.006	−0.063	−0.16***	−0.008	−0.065	−0.16***	−0.037	−0.12***	−0.060**	−0.087*	0.00
	(−1.49)	(−1.00)	(−3.06)	(−1.57)	(−1.18)	(−3.98)	(−0.66)	(−2.60)	(−2.34)	(−1.95)	(*n.a.*)
2,500,000	−0.080***	−0.26***	−0.20***	−0.11***	−0.26***	−0.17***	−0.18***	−0.19***	−0.17***	−0.20***	−0.053***
	(−3.23)	(−5.43)	(−5.36)	(−3.76)	(−5.51)	(−4.02)	(−3.94)	(−5.00)	(−4.55)	(−7.69)	(−4.14)
10,000,000	−0.43***	−0.25***	−0.058***	−0.45***	−0.21***	−0.064***	−0.22***	−0.21***	−0.30***	−0.21***	−0.43***
	(−7.85)	(−6.16)	(−3.41)	(−7.96)	(−6.28)	(−3.04)	(−4.91)	(−4.04)	(−5.80)	(−6.54)	(−7.98)
Impact on execution costs, trades below X EUR:											
10,000	0.20	0.41	2.19***	0.06	1.17***	1.27	1.01	0.77	0.52***	0.72**	0.96**
	(0.75)	(0.74)	(3.32)	(0.26)	(3.72)	(1.25)	(1.56)	(1.32)	(3.41)	(2.10)	(2.39)
25,000	0.16	0.69	3.00***	0.04	1.42***	2.13**	1.42**	1.01*	0.84***	1.15***	0.87**
	(0.08)	(1.19)	(4.41)	(0.16)	(4.37)	(2.00)	(2.05)	(1.69)	(3.70)	(3.14)	(2.00)

A.2 Figures

The following figures illustrate our diff-in-diff estimates by plotting the cross-sectional averages for the treated group (in red) and the control group (in blue) minus their respective pre-event averages over time. We use 3-day moving averages in all figures for better readability. The dashed lines indicate the subperiod averages for June/July, August and September/October. We additionally report the coefficient estimates at the bottom of each graph.

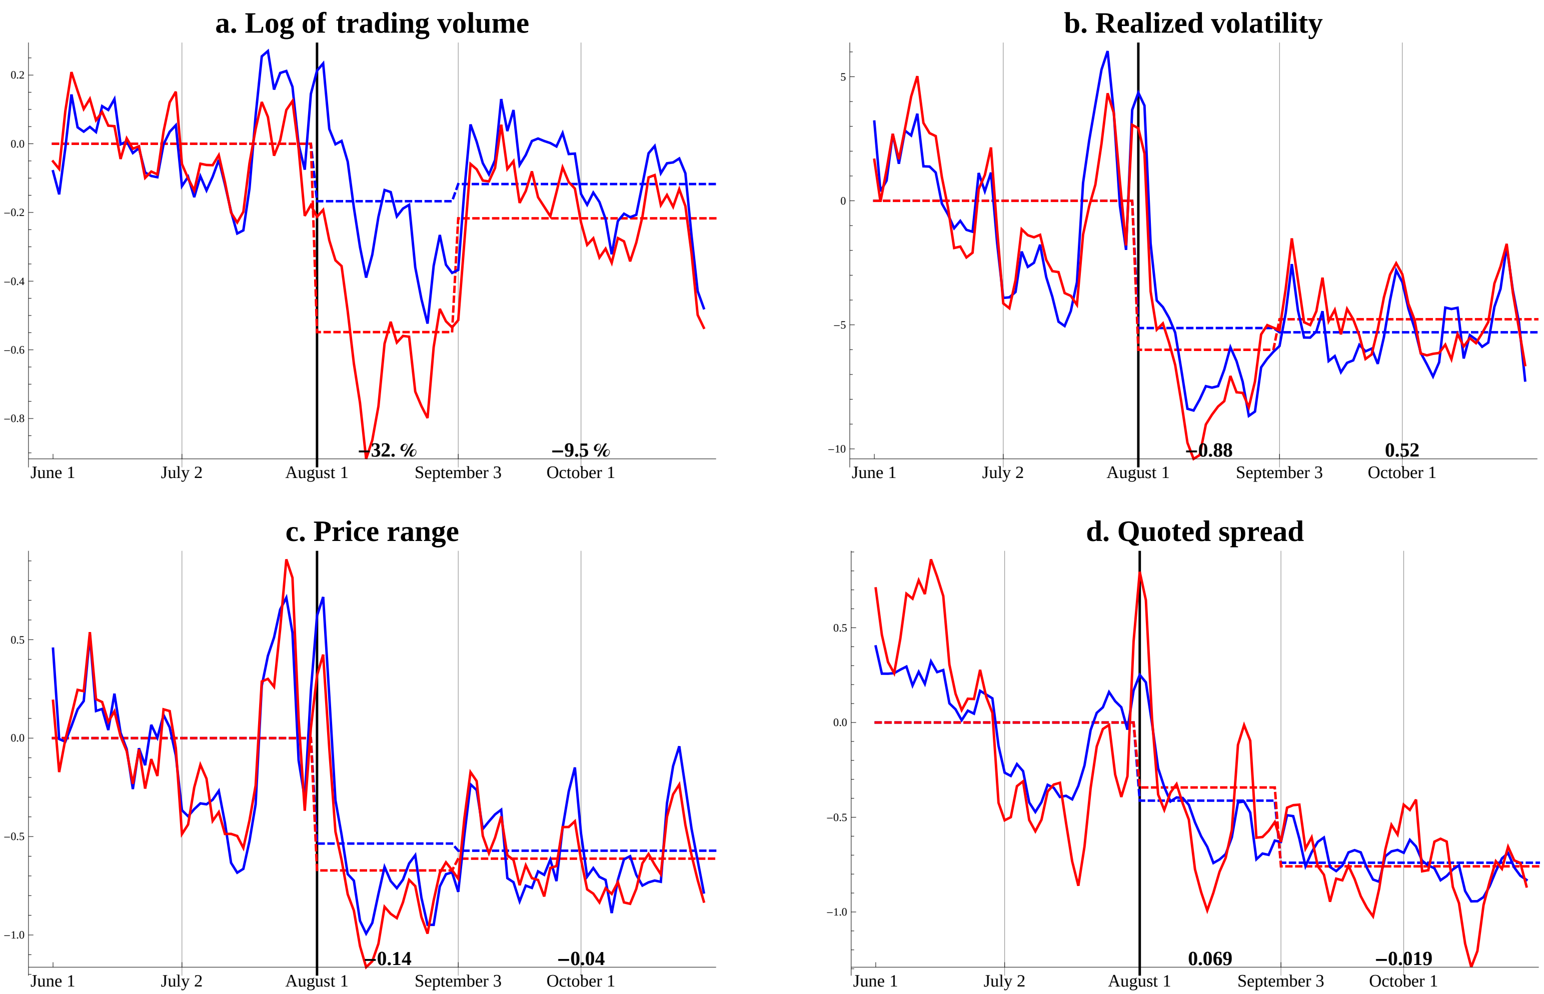

Figure 1: Impact of the tax on market quality variables.

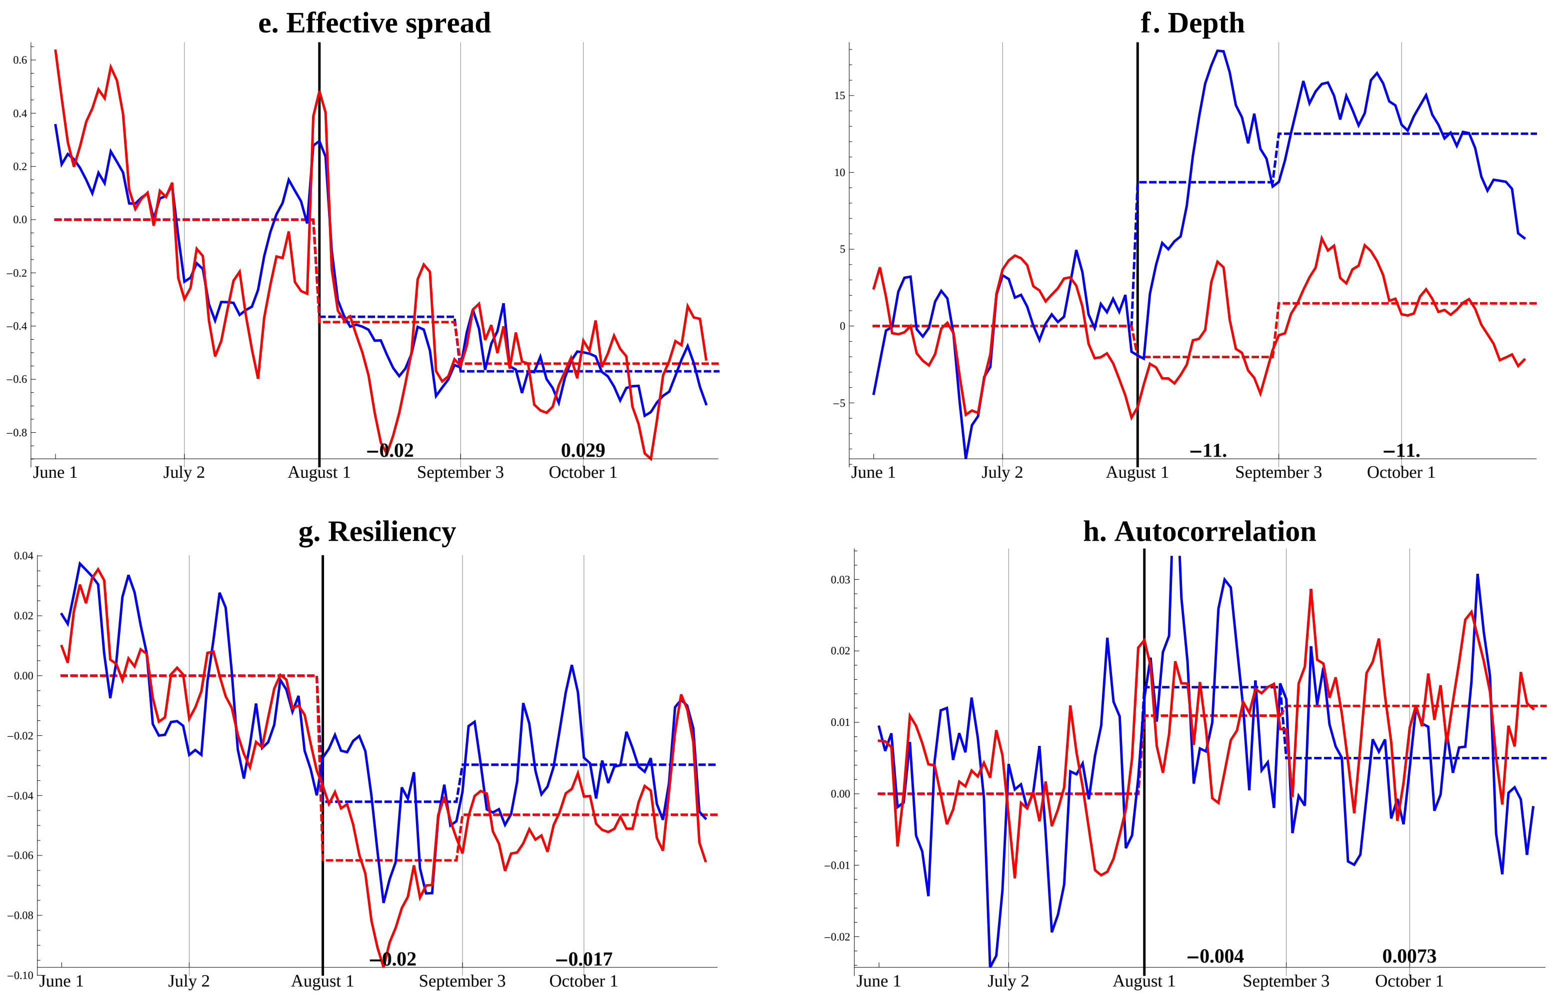
e. Effective spread
0.6
0.4
0.2
0.0
−0.2
−0.4
−0.6
−0.8
−0.02
0.029
June 1
July 2
August 1
September 3
October 1
f. Depth
15
10
5
0
−5
−11.
−11.
June 1
July 2
August 1
September 3
October 1
g. Resiliency
0.04
0.02
0.00
−0.02
−0.04
−0.06
−0.08
−0.10
−0.02
−0.017
June 1
July 2
August 1
September 3
October 1
h. Autocorrelation
0.03
0.02
0.01
0.00
−0.01
−0.02
−0.004
0.0073
June 1
July 2
August 1
September 3
October 1

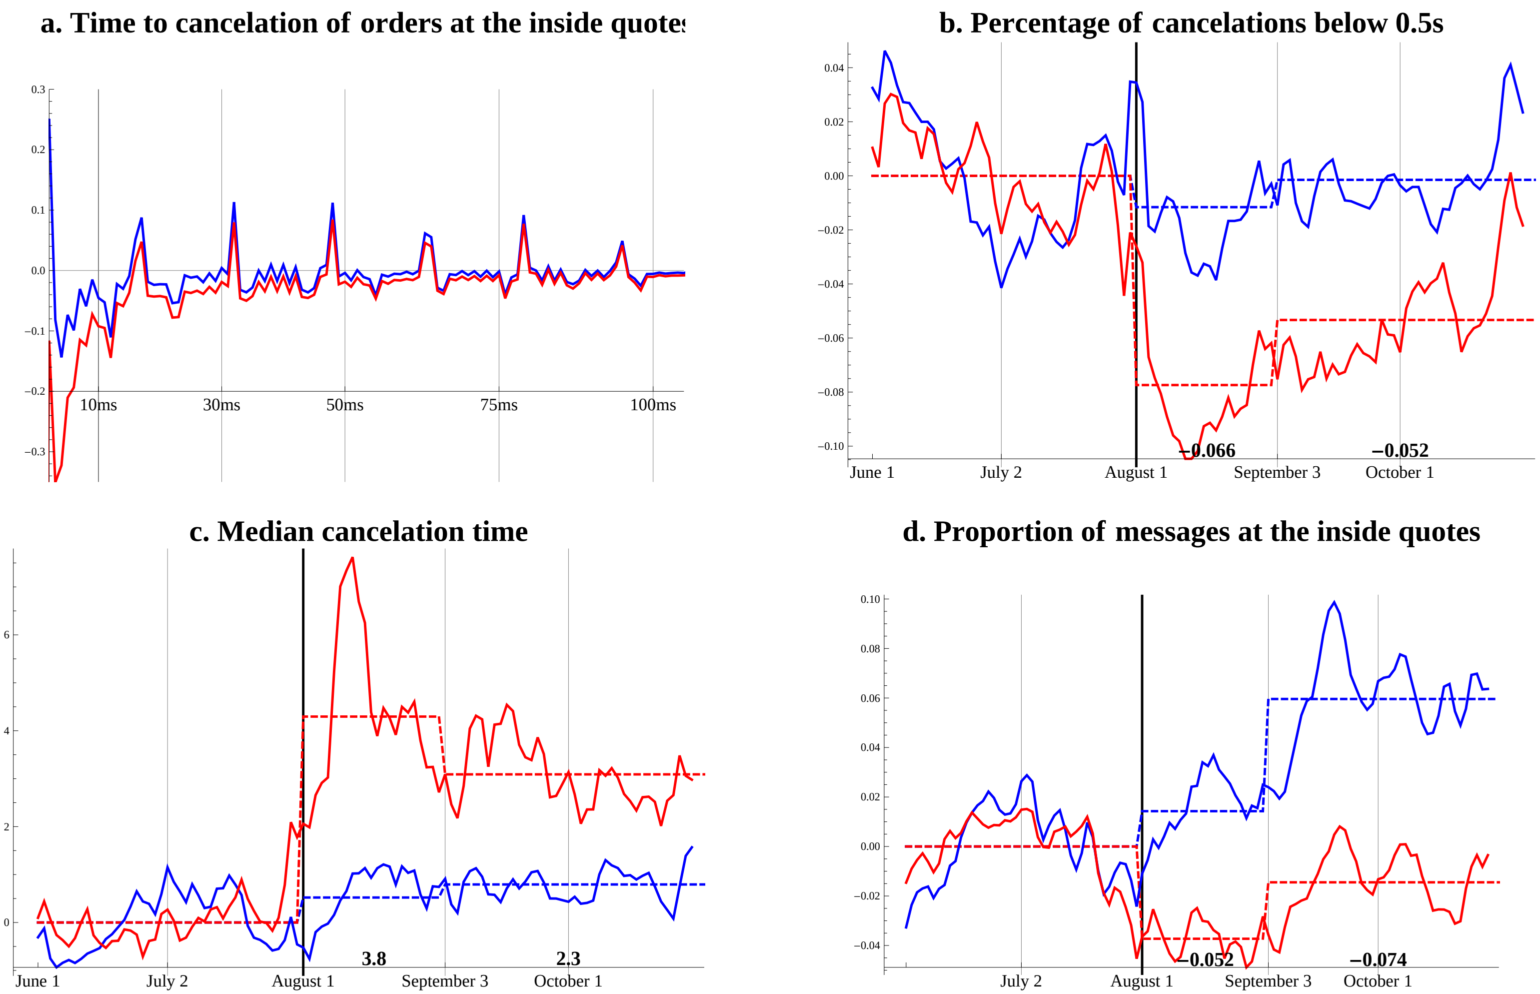

Figure 2: Impact of the tax on message traffic variables.

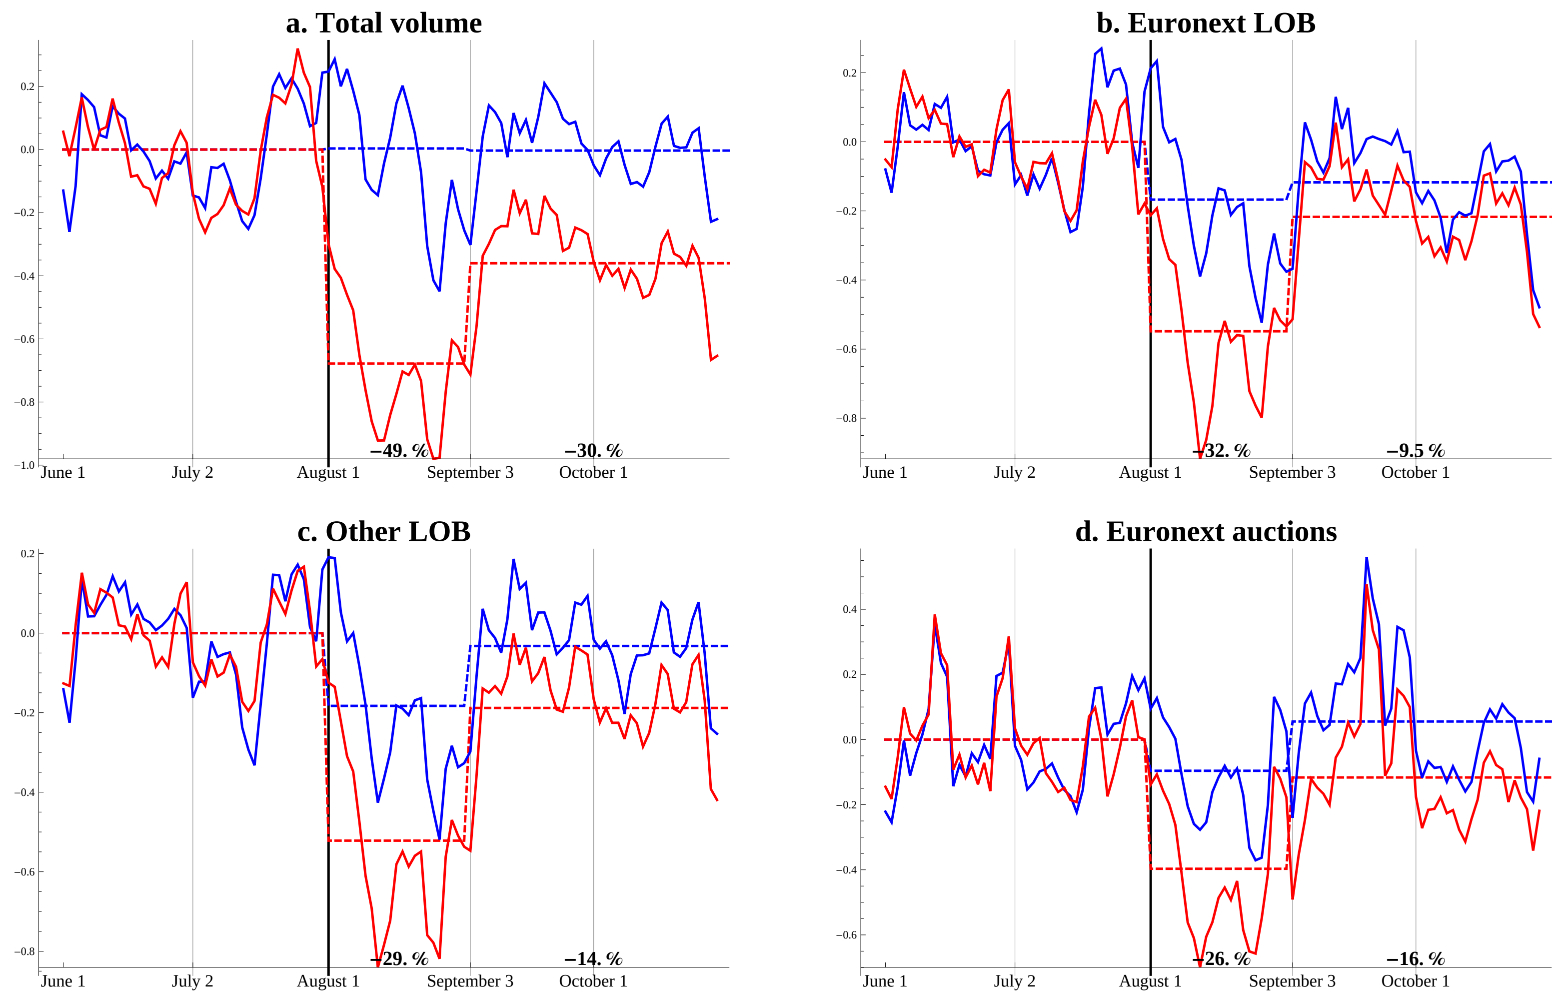

Figure 3: Impact of the tax on trading volume, for all trading venues.

e. Dark pools

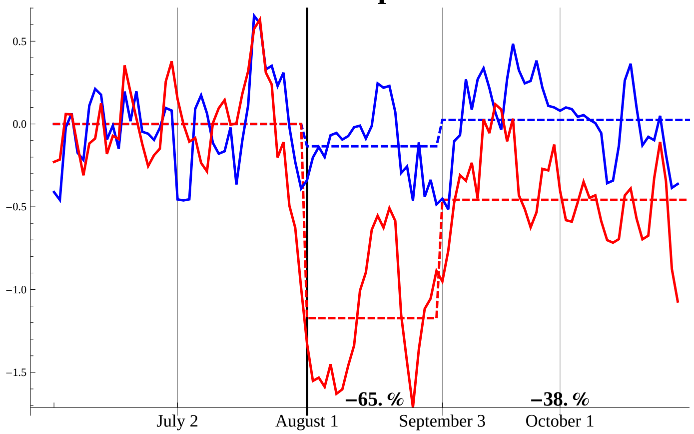

f. OTC

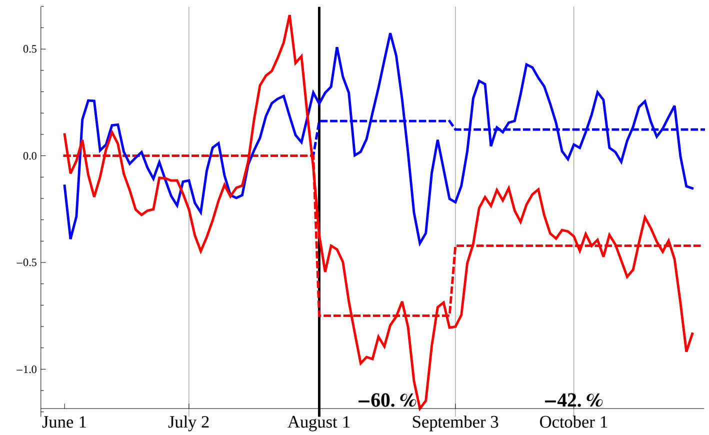

References

Amihud, Y. (2002): “Illiquidity and stock returns: cross-section and time-series effects,” *Journal of Financial Markets*, 5(1), 31–56. 27

Autor, D. H. (2003): “Outsourcing at Will: The Contribution of Unjust Dismissal Doctrine to the Growth of Employment Outsourcing,” *Journal of Labor Economics*, 21(1), 1–42. 14

Baltagi, B., D. Li, and Q. Li (2006): “Transaction tax and stock market behavior: evidence from an emerging market,” *Empirical Economics*, 31, 393–408. 4

Bessembinder, H., and K. Venkataraman (2004): “Does an electronic stock exchange need an upstairs market?,” *Journal of Financial Economics*, 73(1), 3–36. 5

Biais, B., T. Foucault, and S. Moinas (2013): “Equilibrium Fast Trading,” Working paper. 5

Boehmer, E., K. Fong, and J. Wu (2012): “International Evidence on Algorithmic Trading,” Working paper. 5

Brogaard, J., T. Hendershott, and R. Riordan (2012): “High-frequency trading and price discovery,” Working paper. 5, 7

Cartea, l., and J. Penalva (2012): “Where is the Value in High Frequency Trading?,” *Quarterly Journal of Finance*, 02(03). 5

Chordia, T., R. Roll, and A. Subrahmanyam (2011): “Recent trends in trading activity and market quality,” *Journal of Financial Economics*, 101(2), 243 – 263. 4

Comerton-Forde, C., and T. J. Putnins (2012): “Dark Trading and Price Discovery,” Working paper. 5

Degryse, H., F. de Jong, and V. van Kervel (2013): “The impact of dark trading and visible fragmentation on market quality,” Working paper. 5, 24

Duffie, D. (2012): “Market Making Under the Proposed Volcker Rule,” Working paper. 23

Foucault, T., J. Hombert, and I. Rosu (2012): “News Trading and Speed,” Working paper. 5

Hasbrouck, J., and G. Saar (2009): “Technology and liquidity provision: The blurring of traditional definitions,” *Journal of Financial Markets*, 12(2), 143–172. 15

Hau, H. (2006): “The Role of Transaction Costs for Financial Volatility: Evidence from the Paris Bourse,” *Journal of the European Economic Association*, 4(4), 862–890. 4

Hendershott, T., C. Jones, and A. Menkveld (2011): “Does algorithmic trading improve liquidity?,” *Journal of Finance*, 66(1), 1–33. 5

Hirschey, N. (2013): “Do High-Frequency Traders Anticipate Buying and Selling Pressure?,” Working paper. 5

Hoffmann, P. (2013): “A Dynamic Limit Order Market with Fast and Slow Traders,” Working paper 1526, European Central Bank. 5

Hong, H., and J. Yu (2009): “Gone fishin’: Seasonality in trading activity and asset prices,” *Journal of Financial Markets*, 12(4), 672–702. 13

Horrace, W. C., and R. L. Oaxaca (2006): “Results on the bias and inconsistency of ordinary least squares for the linear probability model,” *Economics Letters*, 90(3), 321–327. 21

Jones, C. M., and P. J. Seguin (1997): “Transaction Costs and Price Volatility: Evidence from Commission Deregulation,” *The American Economic Review*, 87(4), pp. 728–737. 4

Jovanovic, B., and A. J. Menkveld (2012): “Middlemen in Limit-Order Markets,” Working paper. 5

Kempf, A., D. Mayston, and P. Yadav (2009): “Resiliency in limit order markets: a dynamic view of liquidity,” Working paper. 10

Keynes, J. M. (1936): *The General Theory of Employment Interest and Money*. Palgrave MacMillan. 3

Kirilenko, A., A. Kyle, M. Samadi, and T. Tuzun (2011): “The Flash Crash: The Impact of High Frequency Trading on an Electronic Market,” Working paper. 23

Kupiec, P. (1996): “Noise traders, excess volatility, and a securities transactions tax,” *Journal of Financial Services Research*, 10, 115–129. 3

Kyle, A. S. (1985): “Continuous Auctions and Insider Trading,” *Econometrica*, 53(6), 1315–35. 12

Lee, C. M. C., and M. J. Ready (1991): “Inferring Trade Direction from Intraday Data,” *Journal of Finance*, 46(2), 733–46. 9

Liu, S., and Z. Zhu (2009): “Transaction Costs and Price Volatility: New Evidence from the Tokyo Stock Exchange,” *Journal of Financial Services Research*, 36, 65–83. 4

Madhavan, A. (2000): “Market microstructure: A survey,” *Journal of Financial Markets*, 3(3), 205–258. 5

Malinova, K., and A. Park (2011): “Subsidizing Liquidity: The Impact of Make/Take Fees on Market Quality,” Working paper. 4

Malinova, K., A. Park, and R. Riordan (2012): “Do Retail Traders Suffer from High Frequency Traders?,” Working paper. 5

Martinez, V. H., and I. Rosu (2013): “High Frequency Traders, News and Volatility,” Working paper. 5

Menkveld, A., and B. Yueshen (2013): “Middlemen Interaction and Its Effect on Market Quality,” Working paper. 23

Pomeranets, A., and D. Weaver (2012): “Security transaction taxes and market quality,” Working paper 2011-26, Bank of Canada. 4

Roll, R. (1989): “Price volatility, international market links, and their implications for regulatory policies,” *Journal of Financial Services Research*, 3, 211–246. 4

Schwert, G. W., and P. J. Seguin (1993): "Securities Transaction Taxes: An Overview of Costs, Benefits and Unresolved Questions," *Financial Analysts Journal*, 49(5), pp. 27–35. 3

Smith, B. F., D. A. S. Turnbull, and R. W. White (2001): "Upstairs Market for Principal and Agency Trades: Analysis of Adverse Information and Price Effects," *The Journal of Finance*, 56(5), 1723–1746. 5

Song, F. M., and J. Zhang (2005): "Securities Transaction Tax and Market Volatility," *The Economic Journal*, 115(506), 1103–1120. 3

Stiglitz, J. (1989): "Using tax policy to curb speculative short-term trading," *Journal of Financial Services Research*, 3, 101–115. 3

Summers, L., and V. Summers (1989): "When financial markets work too well: A cautious case for a securities transactions tax," *Journal of Financial Services Research*, 3, 261–286. 3

Tobin, J. (1978): "A Proposal for International Monetary Reform," *Eastern Economic Journal*, 4(3-4), 153–159. 3

Umlauf, S. R. (1993): "Transaction taxes and the behavior of the Swedish stock market," *Journal of Financial Economics*, 33(2), 227–240. 4

Weaver, D. G. (2011): "Internalization and Market Quality in a Fragmented Market Structure," Working paper. 5