Financial Transaction Taxes, Market Composition, and Liquidity

Jean-Edouard Colliard* and Peter Hoffmann†

First version: December 2012

This version: December 28, 2015

New version. Preliminary.

Abstract

We use the introduction of a financial transactions tax (FTT) in France in 2012 to test competing theories on the impact of FTTs. We find no support for the idea that an FTT improves market quality by affecting the composition of trading volume. Instead, our results are in line with the idea that a lower trading volume reduces liquidity, and thereby market quality. Moreover, short-term institutional investors sell 8% of their holdings of French stocks, which are bought by more long-term investors. This effect on shareholder composition suggests that the debate on FTTs should include their indirect effects on corporate governance and not only on market quality.

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

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

*HEC Paris. E-mail: colliard@hec.fr .

†Contact author. European Central Bank, Financial Research Division, Sonnemannstrasse 20, D-60314 Frankfurt am Main, Germany. E-mail: Peter.Hoffmann@ecb.int .

We are grateful to EUROFIDAI for access to the BEDOFIH database. We would like to thank Bruno Biais, Fany Declerck, Hans Degryse, Patrice Fontaine, Thierry Foucault, Luc Goupil, Laurent Grillet-Aubert, Philippe Guillot, Philipp Hartmann, Frank de Jong, Julien Leprun, Laurence Lescourret, Simone Manganelli, Sophie Moinas, Alper Odabasioglu, Richard Payne, Rémy Praz, Paolo Sodini, Elvira Sojli, David Thesmar, as well as seminar and conference participants at the French Treasury, Autorité des Marchés Financiers (AMF), ECB, VU University Amsterdam, University of Vienna, Toulouse School of Economics, the Banque de France Workshop on Algorithmic and High Frequency Trading, the 6th Erasmus Liquidity Conference, the Arne Ryde Workshop, a workshop of the FBF-IDEI chair "Investment Banking and Financial Markets Value Chain", the 2014 Conference on "The Future of Securities Markets Regulation" at the University of Geneva, the 2014 Frontiers of Finance Conference, the 2014 CREDIT Conference, the 10th Annual Central Bank Workshop on Market Microstructure, the 2014 Paris December International Finance Meeting, the 2014 conference "Market Microstructure: Confronting Many Viewpoints", and the 2015 Bocconi-CONSOB conference "The Development of Securities Markets Trends, Risks and Policies" for useful 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.

1 Introduction

The global financial crisis has renewed interest in financial transaction taxes (FTT), a development that has been fueled by the combination of strapped public finances and public discontent with the financial industry.[1] Since suggested by Keynes (1936), the idea of taxing trading activity has been at the center of a wider debate on whether there is too much trading in financial markets.

Indeed, proponents of FTTs like Stiglitz (1989) argue that markets are populated by too many noise traders whose actions are not based on information and thus generate "excess volatility". According to this view, an FTT improves market quality by reducing the proportion of such "futile" trading. Opponents to FTTs question the relevance of this *composition effect*. Moreover, they argue that it is dominated by a *liquidity effect*: Reduced participation, even by uninformed agents, has a negative effect on liquidity. This prevents the correction of mispricings by arbitrageurs, thereby increasing volatility and reducing price efficiency.

The theoretical literature, surveyed in Section 2, has rationalized these different arguments. In line with this tension, empirical work has delivered mixed conclusions, albeit with a slight overweight towards papers documenting negative consequences for market quality.[2] However, existing studies focus on the aggregate impact of FTTs, mainly due to a lack of granular data. Accordingly, they fail to offer a more detailed investigation into the different economic mechanisms highlighted in the theoretical debate. In particular, there is no evidence on composition and liquidity effects.

This paper aims to bridge this gap between the existing theoretical and empirical work by examining the introduction of a 20 bps tax on the purchase of French equities on August 1st, 2012. We assess the FTT's causal impact using a difference-in-differences framework. Importantly, we estimate treatment effects for different types of market participants, which allows us to shed light on the role of composition and liquidity effects. To this end, we use

[1] In October 2012, 11 EU countries committed to the introduction of a harmonized tax on financial transactions, currently planned to be launched in 2016. Transaction taxes and administrative charges on trading activity (e.g. the SEC's Section 31 fee) are internationally widespread. See Matheson (2011) for an overview.

[2] There are numerous empirical studies estimating the impact of FTTs and transaction costs in general. A non-exhaustive list includes Roll (1989), Umlauf (1993), Jones and Seguin (1997), Baltagi, Li, and Li (2006), Hau (2006), Liu and Zhu (2009), or Pomeranets and Weaver (2012). For a more complete overview, see Matheson (2011).

two different datasets which i) distinguish among investors with different trading technologies (HFT, Mixed, non-HFT), and ii) provide information on a number of investor characteristics (portfolio turnover, size, investment style, etc.) for a large set of institutional portfolios.

We clearly reject the hypothesis that an FTT improves market quality through a composition effect. Instead, we find evidence in favor of sizeable liquidity effects. In particular, we show that high-frequency traders significantly reduce their trading activity although they are effectively exempt from the tax. We provide evidence suggesting that this indirect effect is a response to both increased trading costs as well as a lower demand for liquidity. Finally, we show that investors use portfolio turnover and portfolio holdings as two alternative margins of adjustment to the increase in transaction costs. In particular, we document evidence for high-turnover investors selling some of the affected securities to low-turnover investors. The underlying changes in the ownership structure of affected companies can have important consequences for corporate governance.

More in detail, our analysis starts out with an overview of the French FTT's aggregate impact. We document that trading volume declines on average by around 10%, accompanied by a moderate decline in market quality. However, the overall impact masks some significant heterogeneity across stocks. We find that stocks that are part of Euronext's "Supplemental Liquidity Provider" (SLP) programme, a rebate scheme aimed at incentivizing high-frequency traders to provide liquidity in the most active stocks, are largely unaffected by the FTT. In contrast, the remaining stocks display a strong reduction in trading volume (-20%) and a significant decline in market quality (higher bid-ask spreads, larger price impacts, and increased volatility). This discrepancy suggests that the liquidity effects of a tax may depend on the affected stock's level of liquidity.

We proceed to analyzing the FTT's impact on different trader types using a dataset that groups market participants into three different categories.[3] We find that, despite being effectively tax-exempt, HFTs constitute the most affected trader type with a decline in trading activity of 35%. This indirect impact constitutes strong evidence in favor of a significant liquidity effect. We further provide evidence suggesting that HFTs are affected due to i) an increase in bid-ask spreads hurting the profitability of trading opportunities, as well as ii) an

[3] Due to data availability, this analysis is restricted to a subset of the affected securities.

overall decline in trading activity decreasing opportunities for arbitrage or intermediation.

Finally, we use theories of portfolio choice under transaction costs in order to explore the two different mechanisms through which FTTs can affect trading volume. First, different types of investors can have different elasticities to transaction costs, and adjust their trading volume accordingly (intensive margin). Second, the tax can encourage some investors to sell their stocks to other market participants with potentially different trading patterns (extensive margin). In order to separate these two effects, we examine the trading volume and holdings of institutional investors and find support for both channels. Consistent with the model of Amihud and Mendelson (1986), investors with a high portfolio turnover reduce their holdings of French stocks, whereas investors with a low turnover increase them. Moreover, in line with Constantinides (1986) and Vayanos (1998), investors also significantly decrease their trading volume. This effect is stronger for investors with a high portfolio turnover, and robust to controlling for the mechanical decrease due to reduced holdings.

Overall, our evidence shows that the impact of an FTT, and transaction costs in general, is more complex than a decrease in trading volume by the "marginal investor". Its effects, both direct and indirect ones, are heterogeneous across investors. Importantly, they affect both trading activity as well as portfolio choice.

While our evidence is important for understanding the impact of transaction costs in securities markets in general, it also has rich implications for the debate on transactions taxes. FTTs are frequently motivated by a mix of fiscal (raising revenue) and Pigovian (correcting externalities) motives. We find no evidence for the Pigovian rationale, but we neither observe a significant decrease in market quality for the most liquid stocks. Moreover, the FTT has an important effect on the clientele of French stocks that is not apparent when looking only at aggregate measures of market quality. It reallocates the ownership of French stocks towards more long-term institutional investors, which may actually improve the corporate governance of the affected companies.[4] In contrast, the tax seems to have a more negative impact for less liquid stocks.

This paper coincides with a number of other empirical studies of the French FTT which,

[4]See Stein (1989) for a theoretical argument, and Bushee (2001) and Derrien, Kecskes, and Thesmar (2013) for recent evidence.

despite using a variety of different control groups, reach remarkably similar conclusions concerning the FTT's aggregate impact.[5] In addition, Meyer, Wagener, and Weinhardt (2014) show that the FTT's impact is roughly comparable across different trading venues. Haferkorn and Zimmermann (2013) focus on price discovery and document a decrease in market integration following the policy experiment. Coelho (2014) provides additional evidence from the introduction of the Italian FTT in 2013. However, our paper has an entirely different focus, as the granularity of our data allows us to examine the impact on different types of market participants, and thus to shed light on competing theories of FTTs.

The remainder of this paper is organized as follows. Section 2 details the policy experiment, our main testable hypotheses, and our identification strategy. Section 3 presents our results on aggregate market quality. Section 4 details the impact of the tax on different trader types depending on their speed, while Section 5 considers the impact on the portfolios and trading volumes of different institutional investors. Section 6 concludes.

2 Hypothesis development and methodology

2.1 The policy experiment

On August 1st, 2012, France introduced an FTT of 20 bps on stock purchases.[6] 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,[7] to trades on any trading platform as well as in the over-the-counter market, and to all investors, irrespective of their country of residence.[8]

Importantly, the tax is payable on daily net position changes (i.e. ownership transfers),

[5]See Meyer, Wagener, and Weinhardt (2014), Becchetti, Ferrari, and Trenta (2014), Capelle-Blancard and Havrylchyk (2013), Coelho (2014), and Haferkorn and Zimmermann (2013).

[6]The exact description of the experiment can be found in articles 235 ter ZD and 235 ter ZD bis of the French tax law code.

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

[8]The French FTT relies on the so-called "issuance principle", under which taxation is based on a security's country of registration and not on the residence of the counterparties involved in the transaction. American Depository Receipts (ADRs) were not subject to the tax at the time of the study. Throughout our sample period, there are six French stocks whose ADRs are traded on US exchanges. In the Online Appendix, we present some evidence that refutes the idea that these instruments were used actively in an effort to evade the French FTT.

which implies that pure intraday trading is *de facto* exempted (the same restriction applies to most real-world FTTs, e.g., the British stamp duty). In addition, there are a number of important exemptions: The tax does not apply to newly emitted shares, to transactions by clearing houses, to employee stock ownership plans and, most importantly, to market makers.[9]

Simultaneously, the government introduced a tax of 1 bp on the notional amount of modified or cancelled messages by HFTs exceeding an order-to-trade ratio of 5:1.[10] Unlike the FTT, this tax applies to trading in all French stocks. However, it is only levered on HFTs residing in France, which excludes all major HFT firms. Moreover, message traffic due to market-making is exempt. Accordingly, the scope of this policy is extremely limited, and the French securities markets regulator itself describes its impact as "minimal".[11] This is also corroborated by the fact that this tax generated zero revenues.[12] In Section B.3 of the Online Appendix, we provide additional evidence that confirms this view. Thus, we henceforth consider the policy experiment to only consist of the FTT.

2.2 Hypothesis development

Many arguments in favor and against FTTs have been made since the idea was floated by Keynes (1936), either through the lens of theoretical models or more informally in a number of essays. Arguments in favor of such a tax broadly rely on the idea that it will correct existing market inefficiencies by changing the composition of the trader population. However, opponents of FTTs contend the economic significance as well as the general desirability of this "composition effect". Instead, they warn of a negative "liquidity effect" associated with a widespread decrease in trading activity due to an increase in transaction costs. Most of

[9] Market-making is defined as either quoting competitive bid and ask prices and/or providing liquidity on a regular and continuous basis, or executing orders on the behalf of clients, or hedging positions due to these activities.

[10] In this context, HFT is defined as the regular submission of orders with a resting time of less than 0.5 seconds.

[11] See "Risk and Trend Mapping No14 - Macrofinancial risks for 2014: Market experts' analyses", downloadable at `http://www.amf-france.org/en_US/Publications/Lettres-et-cahiers/Risques-et-tendances/Archives.html?docId=workspace%3A%2F%2FSpacesStore%2F42715c14-a052-43ff-9580-1ea0f52e6c38`.

[12] See the report of the French parliament's finance committee, available at: http://www.assemblee-nationale.fr/14/rap-info/i1328.asp

the debate revolves around the relative importance of these two countervailing mechanisms and their combined impact. In the following, we briefly contrast both perspectives and the competing hypotheses they imply concerning the effects of FTTs on market quality.

Composition vs. liquidity effects. Loosely speaking, proponents of a tax argue that the trading activity of some market participants constitutes a negative externality. At the same time, these agents are assumed to be particularly sensitive to transaction costs. Accordingly, an FTT will affect them disproportionately and thus help to reduce the existing "pollution" in the market. For example, Stiglitz (1989) postulates that an FTT has a stronger effect on noise traders than on agents with accurate fundamental information. The resulting change in the trader population is expected to decrease (non-fundamental) volatility and improve price efficiency.[13] Similar in spirit, Summers and Summers (1989) argue that frequent trading is the essence of "positive feedback traders" (e.g., trend followers) who tend to amplify price swings.[14]. Hence, an FTT will decrease their activity to a larger extent than that of traders who base their decisions on fundamentals.

The opposing view relies on the existence of strong participation externalities in financial markets (see, e.g., Pagano (1989)), resulting in indirect effects of FTTs. For example, Schwert and Seguin (1993) suggest that such a policy will lead to wider bid-ask spreads due to decreased turnover, increased hedging costs, and a possible increase in adverse selection. Likewise, Ross (1989) argues that the absence of short-term traders will hurt liquidity by increasing the role of costly inventory positions for the provision of liquidity. As a consequence, arbitrageurs will find it more costly to correct mispricings, whose magnitude and persistence are therefore expected to increase. This implies that prices become more volatile and less efficient.

Song and Zhang (2005) present an equilibrium model in which both effects are present,

[13] Foucault, Sraer, and Thesmar (2011) show that an increase in transaction costs for retail investors only indeed causes a decrease in volatility, but this is a different experiment from increasing transaction costs for all market participants simultaneously. Indeed, experimental evidence shown in Bloomfield, O'Hara, and Saar (2009) casts doubt on the idea that noise traders are more affected by a general tax than other agents. Deng, Liu, and Wei (2014) suggest that a tax on all transactions can have the desired composition effect when the proportion of noise traders on the market is large enough, but they cannot test the hypothesis that these traders are disproportionately affected by the tax.

[14] See De Long, Shleifer, Summers, and Waldmann (1990) for a model in which the existence of positive feedback traders leads to "excess" volatility.

and the net effect of an FTT depends on their relative strength.[15] If, for example, the "liquidity effect" is sufficiently strong, informed traders reduce their trading activity in a response to lower liquidity. In this case, volatility increases and price efficiency declines.

Based on these two polar views, we arrive at the following competing hypotheses concerning FTT's impact on aggregate market quality.

Hypothesis 1A (Composition effect dominates)**.** *The French FTT leads to a decrease in trading volume and volatility, and an increase in price efficiency.*

Hypothesis 1B (Liquidity effect dominates)**.** *The French FTT leads to a decrease in trading volume, liquidity, and price efficiency, and an increase in volatility.*

Estimating an FTT's impact on aggregate variables provides useful information on the relative strength of these two effects, but it fails to gauge the magnitudes of the individual effects. However, the institutional design of the French FTT allows us to directly test for the presence of liquidity effects. As mentioned in the previous section, the French FTT only applies to ownership transfers and thus implicitly exempts HFTs. Accordingly, observing an impact of the FTT on their trading activity can only be due to a liquidity effect.

Hypothesis 2 (Liquidity effect)**.** *The French FTT leads to a decrease in tax-exempt HFT trading volume.*

Clientele effects. While both Hypotheses 1A and 1B predict that an FTT reduces trading volume, it is important to note that this can occur through two different mechanisms. First, investors can simply scale back their trading activity and thus reduce their portfolio turnover ("turnover adjustment"). This mechanism is at the center of equilibrium models with transaction costs such as Constantinides (1986) and Vayanos (1998). Second, investors with a high portfolio turnover can reduce some of their holdings of affected securities by selling them to market participants with a lower portfolio turnover ("holdings adjustment"). This "clientele effect" features prominently in the literature on FTTs. For example, Keynes

[15] Other equilibrium models also conclude that the effect of an FTT is ambiguous. In Dow and Rahi (2000), the impact of an FTT on price informativeness is positive if and only if informed traders are more risk averse than uninformed traders. Similarly, Subrahmanyam (1998) presents a setting in which an FTT decreases liquidity if and only if the number of informed traders is sufficiently large.

(1936) suggests that the introduction of a tax on trading will affect corporate decisions by giving less weight to shareholders engaging in short-term trading ("speculation") relative to those making a long-term investment ("enterprise"). These considerations are also echoed by Stiglitz (1989) and Summers and Summers (1989). More generally, Amihud and Mendelson (1986) develop a model of portfolio choice under transaction costs and show that, in equilibrium, assets with higher bid-ask spreads are held by investors with shorter average holding periods. We summarize this discussion in the following empirical hypotheses, which will be tested using data on institutional portfolio snapshots.

Hypothesis 3A (Holdings adjustment)**.** *Investors react to the French FTT primarily by adjusting their portfolio holdings in affected securities. In particular, market participants with a long investment horizon increase their holdings in affected securities relative to short-term investors.*

Hypothesis 3B (Turnover adjustment)**.** *Investors react to the French FTT primarily by reducing their portfolio turnover in affected securities.*

Importantly, Hypotheses 3A and 3B are not to be taken as competing hypotheses, but rather as complementary. In practice, investors are likely to use both margins of adjustment, and their relative importance is ultimately an empirical question that we aim to shed light on. Notably, both mechanisms differ in their policy implications. For example, a strong holdings adjustment implies that an FTT has a direct effect on corporate ownership, and thus potentially also on managerial decisions.

2.3 Identification Strategy and Data

We adopt a simple difference-in-differences approach for identifying the FTT's causal impact on market quality and the trading volume of different market participants. To this end, we compare treated French stocks to a group of non-treated control stocks that are otherwise as similar as possible.

As much of our work relies on high-frequency data, it is important to ensure that the data for both groups stem from the same microstructural environment, including the trading protocol, the tick size regime, and the fee structure. The last point is of particular importance because part of our analysis makes use of a rebate scheme for limit orders offered by Euronext, the primary market for French stocks. Fortunately, Euronext also constitutes the main trading venue for Belgian, Dutch and Portuguese stocks. Moreover, the Luxembourg Stock Exchange also uses Euronext's Universal Trading Platform (UTP) as part of a cross-membership cooperation. Accordingly, non-French stocks listed on Euronext (as of the cutoff date January 1st 2012) form a natural pool of control stocks for our diff-in-diff analysis.

We define our final sample of stocks as follows. We start by collecting all the constituents of the Euronext 100 and Euronext Next 150 Indices, which represent the 250 most liquid stocks listed on Euronext. Because Belgium increased its pre-existing FTT on August 1st 2012 and Portugal was heavily affected by the sovereign debt crisis, we restrict the control group to stocks registered in Luxembourg and the Netherlands. Further, we require stocks to trade at least 20 times a day in order to ensure a minimum level of liquidity. Finally, we drop two stocks due to takeover activity, so that our final sample comprises of 172 stocks.[16]

Our first specification consists in comparing all French stocks with a market capitalization above 1 bln EUR and thus affected by the FTT (91 stocks) to the non-French stocks above the same threshold (32 stocks). In a second step, we separately estimate the FTT's impact on stocks that are part of Euronext's SLP programme and those that are not. We have 53 French stocks that are subject to the FTT and at the same time part of the SLP programme, for which the 27 non-French SLP stocks form a natural control group. For the 38 French non-SLP stocks affected by the tax, we form a control group composed of 5 non-French non-SLP stocks above the 1 bln EUR threshold and 47 non-SLP stocks below the 1 bln EUR (30 French and 17 non-French). Table 1 summarizes the different groups of stocks we are using.

[Insert Table 1 here.]

Importantly, practitioners, government officials and regulators advised us in private conversations that the trading activity in August is unlikely to correctly reflect the impact of the policy change because of i) temporary (legal) uncertainty among investors on whether they are subject to the tax or not and ii) a seasonal decline in trading activity for French stocks

[16] All stocks are listed in the Online Appendix B.1.

due to country-wide summer holidays. In order to take such a possibility into account, we use a 5-month sample period from June to October 2012 (109 trading days) and opt for a flexible framework that allows the treatment effect in the first month after the policy change (i.e., August) to be potentially different from the impact in September and October.[17] Formally, the assumption underlying our approach is that for each stock i and date t the variable of interest, $y_{i,t}$, satisfies the following equation:

$$\begin{aligned} \mathbb{E}(y_{i,t} \mid i,t) &= \alpha + \gamma Treated_i + \delta^{Aug} Post_t^{Aug} + \delta^{Sep/Oct} Post_t^{Sep/Oct} \\ &+ \beta^{Aug} Treated_i * Post_t^{Aug} + \beta^{Sep/Oct} Treated_i * Post_t^{Sep/Oct}, \end{aligned}$$

where $Treated_i$ is a dummy variable equal to one for treated stocks and zero otherwise, and $Post_t^{Aug}$ and $Post_t^{Sep/Oct}$ similarly mark days in August and September/October, respectively. The coefficient $\beta^{Sep/Oct}$ captures the long-term impact of the treatment.

This specification relies on the standard common trends assumption that the variables of interest for both groups of stocks should co-move closely absent any treatment. A common issue concerning the difference-in-difference methodology is that this assumption cannot be tested formally. However, the Online Appendix B.5 provides some "placebo" tests that have become customary in the literature on policy evaluation. Together with visual inspection of the data series shown in A.2 and the high correlations between the treated and the control group in the pre-event window (reported in the Online Appendix B.2), these tests confirm the validity of our control group.

Data sources. Our empirical tests are based on data from several sources. We obtain millisecond-stamped intraday data for the market activity (trades and quotes) on Euronext from Thomson Reuters Tick History (TRTH), which we use to compute a wide range of microstructure variables at the stock-day level (see Section 3). Further, we were granted access to the BEDOFIH database, which covers stocks for which Euronext Paris constitutes the primary market center. This database assigns each side (i.e. limit and market order) of

[17] In the Online Appendix B.4, 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 are free from seasonal influences.

every transaction to one out of three trader categories: HFT, non-HFT, and Mixed Traders (MTs). This data is described in more detail in Section 4. Finally, we obtain data on institutional investment portfolios from Factset. Given that this last dataset is only available at the quarterly frequency, we use the four quarterly snapshots of the calendar year 2012 in our analysis of institutional trading.

3 The effect on market quality

In this section, we examine the FTT's impact on aggregate market quality on the background of the competing Hypotheses 1A and 1B. To this end, we focus on the measures emphasized by the theoretical literature: trading activity, volatility, price efficiency, and liquidity.

3.1 Measures of market quality

First, we briefly detail the different variables that we construct in order to test the predictions from the literature. All measures are computed at the stock-day level using intraday data from Euronext's limit order book.[18] We discard trades that are executed off-book, during call auctions and the "trading-at-last" period.

Trading activity. We measure trading activity by the natural logarithm of the EUR value traded for stock i on day t, denoted $logvolume_{i,t}$.

Volatility. We use two different variables in order to assess intraday price volatility. Realized volatility, $RV_{i,t}$, is defined as the sum of squared returns based on the final mid-quote of 5-min intervals. This measure is annualized through multiplication with a factor of $\sqrt{262}$. Alternatively, we compute $range_{i,t}$ as the intraday price range (maximum minus minimum price) across all trades, normalized to the average traded price. Both measures are expressed in percentage points.

Price efficiency. We measure price inefficiency as the absolute value of first-order return

[18] While most stocks are also traded on a number of competing trading platforms, Euronext clearly dominates trading in French and Dutch stocks. Moreover, public data available on the webpage of BATS Chi-X Europe shows that its market share was basically unaffected by the FTT: 68.6% during June-July 2012, 69.2% in August 2012, and 68.4% during September-October 2012. It is therefore reasonable to assume that the observed changes in market quality on Euronext are representative for the changes occurring on other trading venues. This is also consistent with the results reported by Meyer, Wagener, and Weinhardt (2014).

autocorrelations, based on the final mid-quote of 5-min intervals. Intuitively, efficient prices should be unpredictable, and both positive and negative autocorrelations indicate deviations from a random walk process.

Liquidity. We use a number of variables in order to capture all three dimensions of liquidity as defined by Kyle (1985): Tightness, depth, and resiliency. A standard measure of tightness is the quoted relative half-spread, which we compute as a weighted average across all time intervals on day t and denote by $QS_{i,t}$. The relative effective spread, $ES_{i,t}$, is a more robust measure because it measures the spread at the time of a transaction and thus reflects the actual trading cost incurred. We compute this measure as an equal-weighted average across all trades in a given stock across the entire trading day. The effective spread for the τ-th trade in stock j is defined as

$$ES_{\tau} = q_{\tau}\frac{p_{\tau} - mid_{\tau}}{mid_{\tau}},$$

where p_{τ} denote the transaction price, mid_{τ} refers to the contemporaneously prevailing mid-quote, and q_{τ} is a buy-sell indicator taking the value of 1 (-1) for buys (sells).[19] It is useful to decompose the effective spread into two separate components, the price impact, PI, and the realized spread, RS. These are defined as:

$$PI_{\tau} = q_{\tau}\frac{mid_{\tau+5min} - mid_{\tau}}{mid_{\tau}}$$
$$RS_{\tau} = q_{\tau}\frac{p_{\tau} - mid_{\tau+5min}}{mid_{\tau}},$$

where $mid_{\tau+5min}$ denotes the mid-quote 5 minutes after the transaction. The price impact is often interpreted as a measure of adverse selection, while the realized spread is often taken as a proxy for the revenues of liquidity providers (assuming they are closing their position after 5 minutes).

While spreads are relevant to estimate the trading costs of small transactions, a larger

[19] Trades are signed using the Lee and Ready (1991) method, and we aggregate individual orders that are executed simultaneously into one single transaction.

trader will also take into account the quoted depth, denoted $depth_{i,t}$. This measure is computed as the sum of the available liquidity at the inside spread (both bid and ask side), weighted by time and measured in thousands of EUR.

Finally, resiliency measures the speed at which depth is replenished over time following a shock to liquidity. This dimension of liquidity is particularly relevant to larger investors, who try to minimize the impact of their orders by splitting them into numerous smaller transactions. For measuring resiliency, we follow Kempf, Mayston, and Yadav (2009) and estimate the following dynamic model for each stock-day,

$$\Delta depth_\tau = \alpha - \kappa \cdot depth_\tau + \sum_{k=1}^{K} \gamma_k \Delta depth_{\tau+K} + \epsilon_\tau,$$

where $depth_\tau$ denotes quoted depth at the end of the τ-th time interval, and Δ is the first-difference operator. Our estimation is based on 1-minute intraday intervals, and we set $K = 5$. The resulting estimate of κ for stock i on day t, denoted $res_{i,t}$, is our measure of resiliency, as a larger coefficient estimate indicates faster mean-reversion. Economically, $(\ln 2)/\kappa$ measures the half-life of a shock to market depth.[20]

Table 2 presents the mean and standard deviation of each variable for French and non-French stocks during the pre-event period. Overall, the differences in terms of the overall level of market quality are rather small across both groups of stocks. We also include statistics for the log market capitalization and the inverse of the stock price, two variables that we subsequently use as potential control variables. Likewise, there is no notable difference for these variables across both groups.

[Insert Table 2 here.]

3.2 Aggregate impact

Column (1) of Table 3 contains the estimated treatment effects from our baseline diff-in-diff analysis based on stocks above the 1 bln EUR threshold. The t-statistics are based on

[20]Formally, this is true under the assumption that depth follows an arithmetic Ornstein-Uhlenbeck process, see Kempf, Mayston, and Yadav (2009).

standard errors clustered by stock and time.[21] In addition, we graphically illustrate these estimates in Figure 1 by plotting the cross-sectional averages of the variables for both groups of stocks minus their respective pre-event average over time. The dashed lines indicate the sub-period averages for June/July, August and September/October. The difference between the two dashed lines in September/October is equal to the diff-in-diff estimate of the causal impact of the tax.

[Insert Table 3 and Fig. 1 here.]

In line with Hypotheses 1A and 1B, the results show that the French FTT leads to a significant reduction in trading activity. Compared to the control group, trading volume in French stocks declines by approximately 10% in September and October 2012. This underlines that the FTT was an important policy event with considerable impact.

We next turn to the results on volatility and price efficiency, for which Hypotheses 1A and 1B have opposite predictions. First, we observe a decline in the informational efficiency of prices following the introduction of the FTT. Relative to the control group, the absolute value of 5-minute midquote return autocorrelations increase by 0.007 for French stocks (significant at the 5% level). While the economic magnitude of this effect is relatively modest, it still is noteworthy because one would assume that the exemption of intraday trading allows short-term arbitrageurs to continue eliminating price inefficiencies quickly. This result supports Hypothesis 1B. Second, we find that the FTT has no impact on intraday volatility. This result obtains for both measures of price variability, $RV_{i,t}$ and $range_{i,t}$. This is at odds with both Hypotheses 1A and 1B, and is indicative of an overall moderate impact of the tax.

Finally, we turn to the FTT's impact on liquidity. Our estimates show that both quoted and effective spreads were essentially unaffected by the policy experiment. Similarly, we also observe no meaningful variation in price impacts and realized spreads. In contrast, we find a sizeable and strongly significant decrease in quoted depth of approximately 10,700 EUR. This corresponds to a decline of roughly 20%, compared to the pre-event average. Moreover, we also document a small and statistically significant reduction in market resiliency. Given

[21] To alleviate any concern of the results being driven by autocorrelation in the time-series, we alternatively follow Bertrand, Duflo, and Mullainathan (2004) and estimate the model after collapsing all the data in only one "pre" and one "post" periods. These results are reported in the Online Appendix B.6.

the pre-event average speed of mean reversion $\kappa = 0.50$ (based on a 1-minute frequency), the estimated treatment effect of -0.017 implies an increase in the expected half-life of shocks to market depth of approximately 3 seconds (from 83s to 86s).

In sum, our estimates suggest that the French FTT had an overall negative impact on market quality. This result is consistent with the views from the theoretical literature emphasizing the importance of liquidity effects (Hypothesis 1B). In contrast, we reject the predictions associated with a beneficial effect of such a policy through composition effects (Hypothesis 1A). However, the aggregate impact is economically rather small, which suggests that some of the FTT's features such as the market-making exemption and the restriction to ownership transfers were at least partially successful in alleviating its adverse side effects. Yet, the broad picture may eventually mask more significant effects for specific groups of stocks. We next turn to investigating this issue in more detail.

3.3 Heterogeneity across stocks

When evaluating the benefits and concerns associated with FTTs, one important question is whether its effects are heterogeneous across different types of stocks. This is a particularly relevant issue in the presence of liquidity effects, because a decrease in market participation can be expected to have a more pronounced effect on stocks with a relatively low level of liquidity.

In order to investigate this issue, we exploit Euronext's "Supplemental Liquidity Provider" (SLP) programme. This incentive scheme grants rebates on limit orders to a set of market participants in exchange for a commitment to providing additional liquidity.[22] The programme aims at enhancing liquidity in an effort to mitigate competition from alternative trading venues, which is fiercest in the most active stocks.[23] Its structure is particularly geared towards high-frequency market-makers, who have been shown to rely heavily on liq-

[22]More specifically, liquidity providers are required to post two-sided quotes with a minimum size of 5,000 EUR during 95% of the trading day. In addition, they are required to commit to a fixed percentage of time during which they are present at the inside quote, which is determined through competitive bidding during the application process. See https://www.euronext.com/sites/www.euronext.com/files/launch_of_a_supplemental_liquidity_provider_programme_on_european_blue_chips.pdf and https://www.euronext.com/sites/www.euronext.com/files/ifca120326.pdf for further details.

[23]See "Euronext launches DMM-style programme in Europe", Financial Times (Online Edition), April 17, 2011.

uidity rebates (e.g., Brogaard, Hendershott, and Riordan (2014)).

Given that the SLP programme effectively enhances liquidity in the stocks with the highest ex-ante liquidity, it gives rise to a natural partitioning of our sample. Thus, in order to examine whether the FTT's impact varies with the level of liquidity, we estimate separate treatment effects for SLP and non-SLP stocks.

For the SLP stocks, our difference-in-difference procedure is based, as before, on comparing treated French stocks to non-French control stocks. For non-SLP stocks, however, we need to modify our identification strategy because there are only 5 control stocks with a market capitalization of more than 1 billion EUR. Accordingly, we additionally use the 30 French and 17 non-French non-SLP stocks below the 1 bln EUR threshold to form a control group of sufficient size (52 stocks) for the 38 treated French non-SLP stocks. Columns (1) and (2) of Table 4 contain the diff-in-diff estimates.

[Insert Table 4 here.]

Our estimates show that the impact of the FTT varies substantially with SLP membership. While trading volume is broadly unchanged for SLP stocks, non-SLP stocks display a decrease in market activity of about 20% due to the FTT.[24] Moreover, non-SLP stocks exhibit significant increases in intraday volatility, quoted and effective spreads, price impacts, as well as a decrease in resiliency. All these negative effects on market quality are absent in SLP stocks, which only display a reduction in market depth and a slight decrease in price efficiency.

Our results for non-SLP stocks are thus all in line with Hypothesis 1B, and clearly reject Hypothesis 1A. The decrease in volume associated with the FTT hurts liquidity and crowds out "useful" trades, so that volatility increases and price efficiency decreases. Moreover, the relatively modest aggregate impact is actually hiding a large impact on non-SLP stocks and a minor impact on SLP stocks. It thus seems that the strength of the liquidity effect depends on the actual level of liquidity. The most liquid stocks are largely insulated against the adverse effects (except depth), while the rest of the market experiences a more pronounced decline

[24]The interpretation of a coefficient β in a semi-log specification as a percentage change is only valid if its magnitude is sufficiently small. In the text we always report the correct percentage change, given by $\exp(\beta) - 1$ (up to a Jensen error), see Halvorsen and Palmquist (1980).

in the overall level of liquidity, as well as a more significant decrease in trading activity. In Section 4, we turn to the potential role of liquidity effects in explaining these differences by examining the impact on different groups of market participants.

The role of tick size constraints. One potential concern associated with the observation that non-SLP stocks are significantly more affected by the FTT than non-SLP stocks pertains to the role of the minimum tick size. Given their higher liquidity, it is possible that the observed non-result for SLP stocks in terms of spreads is a mechanical result of the minimum tick size representing a binding constraint for these stocks. In order to investigate this alternative explanation for the FTT's differential impact between SLP and non-SLP stocks, we estimate a new specification where we allow the treatment effect to vary with the extent to which the pricing grid represents a binding constraint. The associated results, which are provided in the Online Appendix, show that the difference in the FTT's impact across SLP and non-SLP stocks is not explained by the tick size. This finding is also robust to testing for non-linear threshold effects.

4 Liquidity effects

Like basically all previously implemented FTTs (e.g., the UK stamp duty), the French tax effectively exempts intraday trading activity due to a restriction to ownership transfers. As indicated in Section 2.2, we can exploit this structure in order to shed light on the liquidity effect. In particular, we use data that allows for the identification of different trader types in order to estimate group-specific treatment effects. Given that high-frequency traders exclusively engage in tax-exempt intraday trading, they can only be affected indirectly by the FTT through the liquidity effect.

To examine this issue, we rely on the BEDOFIH database, which assigns both sides of each transaction (limit and market order) to one of three distinct categories: high-frequency traders (HFTs), mixed traders (MTs), and non-high-frequency traders (non-HFTs). This classification was conducted by AMF, the French securities markets regulator, and is based on the median order lifetimes of individual exchange members as well as additional expert knowledge concerning their trading strategies and business model. The first category covers firms that can unambiguously be identified as pure-play HFT outlets trading on their own account. The MT category is composed of exchange members whose order flow is a blend of HFT and non-HFT. According to the providers of the database, this group mainly comprises of large banks and brokers that either have some proprietary HFT activities or offer direct market access to HFT firms.[25] Finally, the remaining category (non-HFTs) includes smaller banks and retail brokerage firms.

Table 5 details the market shares for each of the three different trader categories across the set of French securities subject to the FTT. We report separate statistics for SLP and non-SLP stocks given that we expect these two groups to differ significantly with respect to the importance of HFT activity. A glance at the numbers confirms this, as we find HFTs to account for 27.9% of the trading volume in SLP stocks, but only around 16.9% in non-SLP stocks. We further report each trader type's market share for market and limit orders in order to see whether some market participants are more likely to trade with a particular order type. This decomposition reveals a striking difference between both groups of stocks in terms of the type of HFT activity. While HFTs roughly split their trades equally among both order types in SLP stocks, they almost exclusively trade via market orders in non-SLP stocks. Instead, we observe that both MTs and non-HFTs display a higher share of trading via limit orders in the latter group of stocks. Liquidity provision is thus structured very differently in both groups: while in SLP stocks HFTs provide liquidity in 27.8% of the trades, in non-SLP stocks liquidity is almost exclusively provided by MTs (65.2%) and non-HFTs (31.0%).

Table 6 provides trader-type level information on price impacts and realized spreads. For robustness, we report the results for different frequencies. These statistics allow us to gauge differences in informed trading, as well as agents' ability to avoid being picked off when submitting limit orders. Consistent with previous studies (Brogaard, Hendershott, and Riordan (2014), Carrion (2013)), we find that HFTs' market orders have the largest permanent price impact. Moreover, HFTs also earn the largest realized spreads, in line with a superior management of outstanding limit orders. In contrast, non-HFTs suffer the largest adverse selection costs, being reflected in the smallest realized spreads.

[25] In the Nasdaq dataset used, e.g., by Brogaard, Hendershott, and Riordan (2014), this group of traders is included in the non-HFT category.

[Insert Tables 5 and 6 here.]

Importantly, the BEDOFIH database only covers stocks for which AMF is the competent authority, that is, French securities. We thus cannot compare HFT volume in French stocks to HFT volume in non-French stocks. Accordingly, we restrict our analysis in this section to non-SLP stocks, for which we can use other French stocks below the 1 bln EUR threshold as a control group. Table 7 contains the diff-in-diff estimates for trading volume, as well as the impact on the share of market and limit orders attributable to each of the three different trader groups.[26]

[Insert Table 7 here.]

We find that, in line with Hypothesis 2, HFTs are strongly impacted by the FTT despite the effective exemption of intraday trading. Their activity decreases by 35% (significant at the 1% level), which is in fact the largest decline across all three trader groups. While the trading volume of MTs declines by around 22%, non-HFTs are only marginally affected, and the estimated treatment effect of -3% is statistically insignificant.

Interestingly, we observe that the FTT has not only affected the composition of the overall trading volume, but also triggered changes in the relative use of market and limit orders across trader types. For example, we observe that MTs strongly reduce their use of limit orders in favour of market orders. In contrast, the relative increase in non-HFT activity is particularly driven by an increased use of limit orders, indicating that this trader type has become more important for liquidity provision due to the FTT.

The fact that the FTT led to a large decrease in HFT activity despite the exemption of intraday activity indicates the presence of a strong liquidity effect. Our data allows us to further investigate the mechanisms behind this effect. In particular, two possible explanations come to mind. First, HFTs act as intermediaries or arbitrageurs, such that their trading activity varies with that of end-investors such as, e.g., asset managers. Accordingly, a decrease in the overall trading volume should also imply a decrease in HFT activity. Second, as shown

[26] In order to avoid estimation errors due to a small number of missing values, we follow Bertrand, Duflo, and Mullainathan (2004) and collapse all the pre-treatment data into a single "pre period" and the post-treatment data into a single "post period".

previously, HFTs mainly trade via market orders in non-SLP stocks. This implies that their trading activity is sensitive to an increase in transaction costs such as the effective spread. Accordingly, the reduced level of market liquidity can force them to scale back their trading because some previously profitable trading strategies are now unprofitable. As we will show, both effects contribute to the decrease in HFT volume.

In order to explore the first mechanism, we re-estimate the treatment effect for HFTs when simultaneously controlling for the FTT-induced reduction in the trading activity of other market participants (MTs and non-HFTs). If the entire decrease in HFT volume is due to the overall decline in market activity, the resulting estimate for the FTT's causal impact should be equal to zero. However, Table 7 reveals that we continue to obtain a significant and negative treatment effect of around -24% after controlling for the overall decline in trading volume. This corresponds to roughly two thirds of our initial estimate. Our conclusions are qualitatively unchanged when additionally adding a squared volume term, or controlling for intraday volatility.

[Insert Table 8 here.]

We next turn to examining the potential effects of increased effective spreads on the profitability of HFTs' market orders. We use the negative of the realized spread as a proxy for the profits earned on market orders, given that this measure is widely used to gauge the revenues of liquidity providers. The first column of Table 9 contains the estimated profits on HFT market orders in treated non-SLP stocks for June and July, averaged across all stock-days. For robustness, we present the figures for different time horizons. In line with their large price impact, we generally find these trades to be profitable, with the largest profits accruing for longer holding periods. In order to gauge the effects of a decline in liquidity on HFT profits, we then apply a mechanical increase in effective spreads of 1.331 bps to these figures, in line with the estimated treatment effect (see Table 4). We observe that this increase in trading costs effectively renders HFT in non-SLP stocks a money-losing business, irrespectively of the holding period. To quantify the possible effects on HFTs' trading behaviour, we compute the fraction of stock-days that would turn from profitable into unprofitable, based on our measure of HFT profits. The third column reveals that the resulting effect is quantitatively very large. The increase in effective spreads turns between 17% and 44% of all stock-days from profit-making enterprises into loss-generating ones. This is consistent with the increase in trading costs through reduced market liquidity being able to generate large indirect effects on tax-exempt HFTs.

[Insert Table 9 here.]

This evidence suggests the following explanation for the impact of the FTT on non-SLP stocks: the FTT has a direct effect on MTs, who reduce their trading volume. As a result, a higher proportion of limit orders are submitted by non-HFTs. As pointed out in Table 6, the limit orders of non-HFTs are more adversely selected than the MTs'. Non-HFTs thus ask for a higher compensation when submitting limit orders, and spreads are on average higher. This indirect effect hurts the HFTs, who are mostly using market orders, and they also reduce their trading volume. Both effects interact in equilibrium. Theoretically, the fact that less markets orders are submitted by the very informed HFTs should limit the increase in spreads, but not cancel it. This is exactly what we observe.

The important driver of this liquidity effect is the fact that the tax has a larger direct impact on liquidity providers than on liquidity consumers. This mechanism, which has been neglected so far by the theoretical literature, seems to be in line with the idea of exempting liquidity provision. However, this is difficult in smaller stocks, because liquidity is to a large extent provided by end investors themselves and not by specialized market-makers.

5 Two margins of adjustment

We now turn to testing Hypotheses 3A and 3B, that is, we examine the relative importance of portfolio holdings and portfolio turnover for investors' reaction to the policy experiment. Beyond shedding light on existing theories on FTTs, and transaction costs in general, this exercise is also important from a policy perspective. Indeed, a reshuffling of stocks from high-turnover to low-turnover investors can have important long-term effects on ownership structure and, ultimately, managerial decision-making (as suggested, e.g., by Stiglitz (1989)).

We base our analysis on portfolio snapshots of a rich set of institutional investors, obtained from Factset. This database also contains useful information on fund characteristics, in particular their portfolio turnover. An additional advantage of this data is that it pertains to the "buy-side", that is investors with a relatively long investment horizon that usually do not engage in intraday trading. Accordingly, all institutions in this sample can be expected to be fully exposed to the tax.

We start out by screening the database for investment funds holding any of our sample securities throughout the calendar year 2012.[27] Most funds report at the monthly or quarterly frequency, but not always at the quarter end. In order to bring all data to the same frequency, we only consider the last report in a given calendar quarter and assume it is filed at the quarter end.[28] We then restrict our sample to funds reporting at least once per quarter and with non-zero holdings of at least one French and one control stock throughout the entire period 2012:Q1-2012:Q4. We also limit the sample to Closed-end Funds, Hedge Funds, Non-Public Funds, Open-end Funds, Pension Funds, and Offshore Funds. This leaves us with 3,241 funds.[29]

Given our interest in the cross-sectional variation of the FTT's impact, we need to measure the treatment effect at the fund level. Based on the hypotheses developed in Section 2.2, we focus on the FTT's impact on portfolio holdings and trading volume. In line with the extensive literature on institutional investment, we compute fund-specific treatment effects based on changes in portfolio holdings of treated and control stocks.[30] Here, the set of treated and control stocks corresponds to that of Section 3.2, that is we only consider stocks above the 1 bln EUR threshold.

Fund-specific treatment effects. We use the following two measures of fund-specific

[27]Factset uses a variety of sources to collect information on holdings in non-U.S. securities, see Ferreira and Matos (2008). The information extracted from 13-F filings (referred to as "Institutions" in the database) only pertains to securities listed on U.S. exchanges, such that this part of the database is not relevant for our analysis.

[28]We discard all reports filed in July 2012 in order to ensure that all reports allocated to the third quarter are filed after the launch of the FTT. This reduces the number of funds in our sample by around 3%, and does not affect our results.

[29]We verify that our sample of funds represents a significant part of the overall tax base by applying a 20 bps surcharge to all purchases of French securities taking place in Q4:2012 across all funds. This yields a tax revenue estimate of 22.65 million EUR, which can be linearly extrapolated to 37.75 million EUR for the 5 months of 2012 under the FTT regime. This corresponds to roughly 19% of the total 2012 revenue of 198 million EUR.

[30]It is common practice to measure the trading activity of institutional investors by changes in their portfolio holdings. See, e.g., Grinblatt, Titman, and Wermers (1995) for an analysis of herding in the fund industry, and Manconi, Massa, and Yasuda (2012) for a more recent application in the context of contagion.

treatment effects. First, did_f^H, quantifies the FTT's impact on security holdings. It is computed as the log change in fund f's total holdings of treated stocks between the end of the second quarter (Q2) and the end of the third quarter (Q3) of 2012, minus the contemporaneous log change in holdings of control stocks. In order to avoid picking up effects related to price changes, we value all holdings using the stock prices prevailing at the end of Q2.

The second variable, did_f^V, measures how funds' trading activity changes in response to the implementation of the FTT. We define fund-specific trading volume for quarter t as the absolute value of the change in holdings between quarter $t-1$ and quarter t, again evaluated at the closing prices at $t-1$. Given that the computation of trading volume uses data from two adjacent quarters, we compute did_f^V as the log change of trading volume in treated stocks between Q4 and Q2, minus the contemporaneous log change for control stocks. Notice that did_f^V is only defined for the 2,436 funds that trade both in treated and control stocks during Q2 and Q4, and we henceforth restrict our analysis to these funds.[31]

Explanatory variables. As already mentioned, Factset complements the data on portfolio holdings with a small set of fund characteristics. Based on this information, we construct the following explanatory variables, which we will subsequently use to explain cross-sectional variation in the FTT's impact. $Logsize_f$ is defined as the natural logarithm of a fund's total assets under management (in USD). $Price\_to\_book_f$ is a proxy for the investment style of fund f. A high (low) value implies that a fund predominantly invests into growth (value) stocks. The measure is computed by Factset as the average price-to-book ratio across all portfolio constituents. $Turnover_f$ is a discrete variable that measures a fund's portfolio turnover, i.e. its trading intensity. It takes values ranging from -2 ("Very Low") to +2 ("Very High"), and is based on a classification provided by Factset.[32] Finally, $Index_f$ is a dummy variable which takes a value of one for index funds, and zero otherwise.

Table 10 contains some summary statistics for these explanatory variables. The average fund has a price-to-book ratio of slightly above 3 and approximately 1.2 billion USD assets under management. Around 13% of the funds are index funds. The average fund is cat-

[31] All our results for did_f^H are qualitatively unchanged when including funds that do not trade in either Q4 and Q2 in at least one group of stocks, i.e. funds for which did_f^V is not defined.

[32] This classification contains the following values: "Very Low", "Low", "Medium", "High", and "Very High".

egorized to have a low portfolio turnover (-0.89) but there is considerable cross-sectional variation across funds with a standard deviation of 1.14.

[Insert Tables 10 and 11 here.]

Results. Panel A of Table 11 reports the coefficients from cross-sectional regressions of the fund-level treatment effects on our explanatory variables, separately for changes in portfolio holdings (did_f^H) and changes in trading volume (did_f^V). In all regressions, t-statistics are based on White standard errors robust to heteroskedasticty.

We first discuss the impact on portfolio holdings. Recall that, based on Hypothesis 3A in Section 2.2, we particularly expect the treatment effect to be negatively related to portfolio turnover. Before examining the cross-sectional variation and its determinants, we start out with a simple regression on only a constant in order to gauge the average treatment effect across funds. The resulting coefficient estimate in column (1) is very small and statistically insignificant. This indicates that the FTT does not induce the average institutional investor in our sample to increase or decrease her holdings of French stocks relative to the control group.

We proceed to examining how the cross-sectional variation in funds' reaction to the FTT can be explained by differences in fund characteristics. The resulting estimates are tabulated in column (2). In order to facilitate the interpretation of the intercept, we demean both $Logsize_f$ and $Price\_to\_book_f$ before the estimation. Accordingly, the constant represents the treatment effect for a non-index fund of average size, average investment style, and with a medium turnover.

In line with Hypothesis 3A, the coefficient on $Turnover_f$ is negative and strongly statistically significant, indicating that investors with shorter (longer) average holding periods have reduced (increased) their holdings in French stocks subject to the FTT, relative to control stocks. Notably, the economic magnitude of this effect is large. The coefficient estimate of -0.038 implies that investors with a very high turnover are predicted to sell 7.9% of their holdings of French stocks, while investors with a very low turnover increase them by 7.3%. This provides strong support for the "clientele effect" in Amihud and Mendelson (1986) and the associated prediction that an exogenous increase in transaction costs will induce agents with a short investment horizon to sell some of their holdings to long-term investors.[33]

We now turn our attention to the results on funds' trading volume, that is, Hypothesis 3B. We proceed as before and start by a simple regression on a constant, which yields a highly significant coefficient estimate of -0.180, see column (3). This implies that the introduction of the FTT induced the average fund to reduce its trading activity in French stocks by approximately 16%, relative to the control group. Hence, there is evidence in favor of a significant role for both margins of adjustment, holdings and turnover. Moreover, the magnitude of the treatment effect for turnover is more than one and a half times the impact on the aggregate trading volume and thus consistent with the view that buy-side investors are disproportionately affected by the tax.

A number of interesting findings emerge when turning to the role of individual fund characteristics, as shown in column (4). In particular, the coefficient on portfolio turnover is negative and strongly significant, which suggests that the FTT leads to a stronger (weaker) reduction in trading volume for the most (least) active investors. While intuitive, we are not aware of any theoretical model that predicts this relationship, as standard theories of transaction costs usually assume either homogeneity across investors (see Constantinides (1986) and Vayanos (1998)) or exogenous investment horizons (as in Amihud and Mendelson (1986)). Notice that the effect is economically large, as the coefficient estimate of -0.148 suggests that the funds with the highest portfolio turnover reduce their trading in French stocks by around 42%, while those with the least reshuffling activity marginally increase their trading.[34]

One potential issue when measuring the FTT's impact on trading volume is the fact that a reduction in the holdings of French stocks will lead to a mechanical reduction in trading volume even if agents keep their portfolio turnover constant. Intuitively, a fund that reduces its holdings by 10% and keeps the same turnover will mechanically trade 10% less, everything

[33] Further, we observe that funds with a high price-to-book ratio as well as index funds responded to the FTT by reducing their holdings in French stocks relative to the control group. The latter result may seem surprising, as index funds should not have a lot of flexibility in their portfolio choice. We explain this result in the Online Appendix B.9: The impact on index funds is almost entirely driven by investors who rely on the use of synthetic replication strategies via total return swaps. In contrast, the FTT's impact on index funds using physical replication strategies is close to the average treatment effect.

[34] Given an intercept of -0.251, the treatment effects for investors with very high (very low) turnover are $-0.251 \pm 2 \times (-0.148)$. We then apply the transformation explained in footnote 24.

else equal. We therefore include did_f^H as an additional control variable in our cross-sectional regression in order to disentangle a mechanical reduction in trading activity from a change in actual behaviour. The results are tabulated in column (5). As expected, the coefficient on did_f^H is positive and significant at the 1% level, confirming the mechanical relationship just discussed. However, the coefficient estimate is considerably smaller than 1, which confirms that funds use both margins of adjustment. Accordingly, none of our previous results changes materially once controlling for the FTT-induced change in portfolio holdings.

We proceed by repeating the above analysis separately for SLP and non-SLP stocks. We follow our approach from Section 3.3 and use both French and non-French stocks below the 1 billion EUR threshold as control group for French non-SLP stocks subject to the tax. For each subsample, we only include funds that have positive holdings of at least one treated and one control stocks throughout the entire year 2012. Moreover, we discard funds for which did_f^V is not defined due to zero trading volume. The results are depicted in Table 11, where Panel B refers to SLP stocks and Panel C to non-SLP stocks.

Interestingly, we find that the holdings adjustment is limited to SLP stocks. For non-SLP stocks, the coefficient on $Turnover_f$ is statistically insignificant. In turn, we find a more pronounced decrease in trading volume, which declines by around 19.5% for SLP stocks and 26.5% for non-SLP stocks. This observation is consistent with our results concerning the differences in the FTT's impact on market quality across both groups of stocks (Section 3.3) being driven by liquidity effects. For the less liquid stocks, investors exclusively adjust via the intensive margin, thereby hurting market liquidity disproportionately. This finding is also in line with the more prominent role of MTs and non-HFTs for liquidity provision in non-SLP stocks documented in Section 4.

In sum, our results provide evidence consistent with investors using both portfolio holdings as well as turnover as complementary margins of adjustment to an exogenous increase in transaction costs through the adoption of the French FTT. In particular, we provide evidence supporting the idea that even a moderate FTT can significantly alter the shareholder composition of affected stocks (Hypothesis 3A), consistent with arguments frequently made by proponents of such a policy. Additionally, short-term investors adjust their trading activity to a larger extent than long-term investors, which implies that the net effect on portfolio holdings is weaker than a model with constant turnover would suggest. Moreover, reduced participation by institutional investors can have negative side-effects on liquidity provision, in particular absent specialized market makers (in line with our results in Section 4).[35]

6 Conclusion

This paper uses the French FTT launched in 2012 in order to shed light on the main economic mechanisms mentioned in the debate on FTTs. We find no evidence for the composition effect through which an FTT is supposed to improve market quality. Instead, we find evidence for a liquidity effect, through which the tax worsens market quality and indirectly affects even exempted traders. The impact is strongest in less liquid stocks, in line with the idea that participation externalities are less important at the margin for the most actively traded stocks. More generally, our results show that a small increase in transaction costs does not only affect "futile" marginal investors, as claimed by some proponents of FTTs. Instead, we observe that the tax induces many different types of market participants to scale back their trades.

The idea to impose a Pigovian tax on trading volume, seen as a negative externality, thus does not seem to apply well to modern markets. The policy debate on FTTs, which has been revived recently with the project of a pan-European FTT, should therefore rather focus on two alternative motivations, for which we provide more mixed evidence.

First, as for any tax, the resulting revenues need to be compared to the associated economic distortions. With its exemptions and safeguards for liquidity provision, the French design attempts to minimize the negative side effects.[36] Nevertheless, our results suggest that the French FTT is to a large extent a disguised tax on savers, whose costs and benefits

[35] Keim and Madhavan (1997) provide evidence consistent with institutional investors acting as important liquidity providers. This is especially true in today's markets where they have access to sophisticated execution algorithms that allow them to trade with both limit and market orders.

[36] In particular, the French design compares rather well to the European project. Indeed, we estimate that the French design would yield approximately 2.7 billion EUR if implemented at the EU level, based on an extrapolation of the actual 2012 revenues of 198 Million for August-December 2012, see `http://www.assemblee-nationale.fr/14/rap-info/i1328.asp`. This compares to 4.8-6.5 billion EUR that the European Commission expects to raise with a tax on *all* equity transactions (including intraday activity), and without any safeguards for liquidity provision. Details on the computations can be found in the Online Appendix.

require a general equilibrium analysis.[37]

Second, we find evidence that the tax has caused an important change in the shareholder composition of the affected companies. Proponents of FTTs would expect the associated increase in shareholders' investment horizon to foster a reduction of managerial myopia or short-termism, as argued by Stein (1989). However, we find that agents react to the tax by adjusting not only their holdings but also their turnover, which dampens this effect significantly. Indeed, an FTT is a rather indirect tool to allocate more influence to long-term shareholders.

This second point illustrates a more general phenomenon. The theoretical literature suggests that market participants adjust their behavior so as to minimize the impact of the tax (e.g. Constantinides (1986)). The impact of a tax on aggregate market outcomes should thus be second order compared to changes in the affected investors' portfolios and trading strategies. To our knowledge, the French FTT is the first experiment of this kind for which the necessary disaggregated data is available, opening new avenues for research on FTTs and transaction costs in general.

[37]See Diamond and Saez (2011) for an up-to-date discussion of this literature.

A Appendix

A.1 Tables

Table 1: Treated and control groups.

Treated group	Control group
FTT for all stocks	
91 French stocks above 1 bln EUR	32 non-French stocks above 1 bln EUR
FTT for SLP stocks	
53 French SLP stocks above 1 bln EUR	27 non-French SLP stocks above 1 bln EUR
FTT for non-SLP stocks	
38 French non-SLP stocks above 1 bln EUR	5 non-French non-SLP stocks above 1 bln EUR 47 non-SLP stocks below 1 bln EUR

Table 2: Summary statistics for treated and control stocks.

This table contains the empirical averages and standard deviations of market quality variables for French and non-French stocks, both below and above the 1 bln EUR threshold, over the period June-July 2012. All figures are computed at the stock-day level.

	French	Non-French
logvolume	16.23	16.71
	(1.49)	(1.14)
RV	28.12	24.83
	(11.49)	(10.43)
range	2.83	2.49
	(1.52)	(1.40)
QS	5.71	4.19
	(4.15)	(2.14)
ES	4.38	3.32
	(3.02)	(1.56)
RS	0.38	0.15
	(2.04)	(1.03)
PI	4.00	3.18
	(2.47)	(1.71)
depth	58.66	80.61
	(50.66)	(47.17)
res	0.50	0.53
	(0.15)	(0.13)
$\|AR\|$	0.11	0.11
	(0.09)	(0.09)
invprice	0.06	0.07
	(0.11)	(0.06)
logmktcap	22.43	22.50
	(1.09)	(1.09)
# Stocks	91	32
# Obs.	3,913	1,376

Table 3: Causal impact of the FTT on all stocks.

This table contains the estimates for the coefficient $\beta^{Sep/Oct}$ from the regression equation

$$\begin{aligned} y_{i,t} &= \alpha + \gamma Treated_i + \delta^{Aug} Post_t^{Aug} + \delta^{Sep/Oct} Post_t^{Sep/Oct} \\ &+ \beta^{Aug} Treated_i * Post_t^{Aug} + \beta^{Sep/Oct} Treated_i * Post_t^{Sep/Oct}, \end{aligned}$$

where $y_{i,t}$ denotes one of the ten market quality variables defined in Section 3. We consider the sample of stocks with a market capitalization of more than 1 billion EUR and $\beta^{Sep/Oct}$ identifies the average impact of the FTT. T-statistics based on standard errors clustered at the stock and day level are given in parentheses. $***$, $**$, and $*$ denote statistical significance at the 1%, 5%, and 10% level, respectively.

Variable	$\beta^{Sep/Oct}$
logvolume	-0.094**
	(-2.17)
RV	0.200
	(0.22)
range	-0.070
	(-0.66)
QS	-0.031
	(-0.13)
ES	0.023
	(0.13)
PI	0.187
	(1.16)
RS	-0.164
	(-1.25)
depth	-10.718***
	(-2.82)
res	-0.017*
	(-1.87)
$\|AR\|$	0.007*
	(1.91)
# Treated	91
# Control	32
# Obs.	13,407

Table 4: Causal impact of the FTT on SLP and non-SLP stocks.

This table contains the estimates for the coefficient $\beta^{Sep/Oct}$ from the regression equation

$$\begin{aligned} y_{i,t} &= \alpha + \gamma Treated_i + \delta^{Aug} Post_t^{Aug} + \delta^{Sep/Oct} Post_t^{Sep/Oct} \\ &+ \beta^{Aug} Treated_i * Post_t^{Aug} + \beta^{Sep/Oct} Treated_i * Post_t^{Sep/Oct}, \end{aligned}$$

where $y_{i,t}$ denotes one of the ten market quality variables defined in Section 3. Column (1) refers to the stocks pertaining to Euronext's SLP programme and uses non-French SLP stocks as control group. Column (2) refers to the remaining stocks, and uses both non-French non-SLP stocks and French non-SLP stocks below the 1 billion EUR threshold as control group. T-statistics based on standard errors clustered at the stock and day level are given in parentheses. $***$, $**$, and $*$ denote statistical significance at the 1%, 5%, and 10% level, respectively.

Variable	(1) SLP stocks	(2) Non-SLP stocks
$logvolume$	-0.015	-0.224***
	(-0.32)	(-3.29)
RV	0.207	2.468**
	(0.21)	(2.03)
$range$	-0.101	0.256*
	(-0.88)	(1.87)
QS	0.148	1.254*
	(1.21)	(1.85)
ES	0.160	1.088**
	(1.44)	(2.08)
PI	0.161	1.835***
	(1.15)	(5.34)
RS	-0.001	-0.746**
	(-0.01)	(-2.14)
$depth$	-12.646***	-2.112
	(-2.69)	(-1.13)
res	-0.008	-0.029***
	(-0.82)	(-3.24)
$\|AR\|$	0.009*	0.007
	(1.93)	(1.20)
# Treated	53	38
# Control	27	52
# Obs.	8,720	9,810

Table 5: Breakdown of trading activity by trader type for SLP and non-SLP stocks.

This table contains the cross-sectional averages for the proportion of trading volume (Share Volume) as well as the proportion of limit orders (Share Limit) and market orders (Share Market) attributable to each of the three trade type categories in the BEDOFIH database (HFT, MT, non-HFT). The estimates are tabulated separately for SLP and non-SLP stocks above the 1 billion EUR threshold. Standard errors computed across stock-days are given in parentheses. All figures are based on the months of June and July only.

Variable/Group	SLP			Non-SLP > 1 bln		
	Share Volume	Share Limit	Share Market	Share Volume	Share Limit	Share Market
HFT	27.9	27.8	27.9	16.9	3.7	30.1
	(0.2)	(0.2)	(0.2)	(0.2)	(0.1)	(0.3)
Mixed	55.9	54.7	57.1	55.7	65.2	46.3
	(0.2)	(0.2)	(0.2)	(0.3)	(0.3)	(0.3)
Non HFT	16.2	17.5	15.0	27.3	31.0	23.7
	(0.1)	(0.1)	(0.1)	(0.3)	(0.3)	(0.3)

Table 6: Price impacts and realized spreads, by trader type, for SLP and non-SLP stocks.

This table contains the cross-sectional averages for price impacts and realized spreads at different time horizons. Both measures are tabulated separately for SLP and non-SLP stocks above the 1 billion EUR threshold, and each trader type available in the BEDOFIH database (HFT, MT, non-HFT). In the computation, stock-day observations with a missing value for at least one trader category were discarded. Standard errors computed across stock-days are given in parentheses. All figures are based on the months of June and July only.

Variable/Group		SLP			Non-SLP > 1 bln		
	Horizon	HFT	MT	Non-HFT	HFT	MT	Non-HFT
Price impact	10s	3.31	2.48	1.37	5.46	3.62	3.14
		(0.02)	(0.02)	(0.02)	(0.07)	(0.05)	(0.06)
	5min	3.32	3.05	1.72	6.12	5.16	4.26
		(0.03)	(0.04)	(0.05)	(0.09)	(0.08)	(0.10)
	30min	3.35	3.00	1.56	6.53	5.95	4.59
		(0.05)	(0.08)	(0.11)	(0.14)	(0.17)	(0.22)
Realized spread	10s	0.35	-0.07	-0.16	4.25	2.15	2.06
		(0.01)	(0.01)	(0.02)	(0.17)	(0.06)	(0.08)
	5min	0.01	-0.36	-0.46	3.19	1.17	0.64
		(0.03)	(0.02)	(0.05)	(0.19)	(0.07)	(0.10)
	30min	-0.07	-0.18	-0.38	2.51	0.85	-0.08
		(0.05)	(0.05)	(0.13)	(0.30)	(0.11)	(0.21)

Table 7: Causal impact of the FTT on trading volume and order flow composition for different trader types. Non-SLP stocks.

This table contains the estimates for the coefficient $\beta^{Sep/Oct}$ from the regression equation

$$\Delta y_i = \alpha + \beta^{Sep/Oct} \times Treated_i,$$

where $\Delta y_i = y_i^{Sep/Oct} - y_i^{Jun/Jul}$, and $y_i^{Jun/Jul}$ and $y_i^{Sep/Oct}$ denote, for a specific trader type (HFT, MT, non-HFT) either the log of the average daily trading volume, the average share of limit order volume, or the average share of market order volume, across all trading days in June-July and September-October, respectively. This procedure corresponds to the estimation of treatment effects based on a cross-sectional regression with time-series collapsed information suggested in Bertrand, Duflo, and Mullainathan (2004). T-statistics based on White standard errors robust to heteroskedasticity are given in parentheses. $***$, $**$, and $*$ denote statistical significance at the 1%, 5%, and 10% level, respectively.

Group/Variable	Log volume	Share Limit (%)	Share Market (%)
HFT	-0.434***	-0.654**	-7.109***
	(-3.05)	(-2.06)	(-4.76)
MT	-0.253***	-4.933***	3.776***
	(-3.00)	(-3.66)	(2.69)
Non-HFT	-0.029	5.586***	3.332**
	(-0.32)	(3.96)	(2.43)
# Treated	38	38	38
# Control	29	29	29

Table 8: Test for a mechanical decrease in HFT activity. Non-SLP stocks.

This table contains the estimates for the coefficient $\beta^{Sep/Oct}$ from the regression equation

$$\Delta logvolume_i^{HFT} \quad = \quad \alpha + \beta^{Sep/Oct} \times Treated_i + \gamma(\Delta X_i),$$

where $\Delta logvalue_i^{HFT} = logvolume_i^{HFT,Sep/Oct} - logvolume_i^{HFT,Jun/Jul}$, and $logvolume_i^{HFT,Jun/Jul}$ and $logvolume_i^{HFT,Sep/Oct}$ denote the log of average daily HFT trading volume for stock i, and $(\Delta X_i = X_i^{Sep/Oct} - X_i^{Jun/Jul})$ is a vector of control variables based on the difference of sub-period averages. These variables are i) the log of the daily average non-HFT and MT volume, its square, and the average daily realized volatility. This procedure corresponds to the estimation of treatment effects based on a cross-sectional regression with time-series collapsed information suggested in Bertrand, Duflo, and Mullainathan (2004). T-statistics based on White standard errors robust to heteroskedasticity are given in parentheses. $***$, $**$, and $*$ denote statistical significance at the 1%, 5%, and 10% level, respectively.

Variable	Coefficient		
$Treated_i$	-0.272**	-0.266**	-0.346**
	(-2.25)	(-2.38)	(-2.46)
$\Delta logvalue^{Other}$	0.860***	0.198	0.666***
	(4.71)	(0.07)	(2.89)
$\Delta logvalue^{Other\ 2}$		0.024	
		(0.23)	
$\Delta vola\_5$			0.029*
			(1.82)
# Treated	38		
# Control	29		

Table 9: HFT gains on market orders, before and after the increase in spreads.

This table contains the cross-sectional averages at different time horizons of the following variables: (1) the negative of the realized spread (in bps) on HFTs' market orders, representing HFTs' trading profits; (2) the same figure minus 1.331 bps, the estiamted treatment effect for effective spreads in non-SLP stocks, thus representing the HFTs' counterfactual trading gains; (3) the proportion of stock-days that turn from profitable into unprofitable after the 1.331 increase in effective spreads. All figures are based on June-July only.

Horizon	HFT profits (before)	HFT profits (after)	% stock-days becoming unprofitable
10s	0.18	-1.15	44
5min	0.83	-0.50	31
30min	1.25	-0.08	17

Table 10: Summary statistics of fund characteristics.

This table contains summary statistics on fund characteristics for the 2,435 investment funds used in our analysis. *Size* denotes total assets under management in billion USD. *Turnover* is a discrete variable ranging from −2 ("Very Low" portfolio turnover) to +2 ("Very High" portfolio turnover). *Price_to_book* denotes the fund's average price-to-book ratio based on its portfolio holdings, and *Index* is a binary variable equal to one for Index Funds, and zero otherwise. Standard deviations are reported in parentheses. All variables are provided by Factset.

Variable	Mean
size	1.20
	(6.27)
Turnover	-0.89
	(1.14)
Price_to_book	3.20
	(1.17)
Index	0.13
	(0.34)

Table 11: Causal impact of the FTT on investment funds' portfolio holdings and trading volume, for all stocks, SLP and non-SLP stocks.

This table contains coefficient estimates from a linear regression of fund-specific treatment effects in terms of portfolio holdings (columns (1) and (2)) and trading volume (columns (3) - (5)) on investment fund characteristics. Panel A refers to all stocks, and the control group consists of all non-French stocks. Panel B refers to SLP stocks, and the control group consists of all non-French SLP stocks. Panel C refers to non-SLP stocks, and the control group consists of non-French non-SLP stocks as well French non-SLP stocks below the 1 bln EUR threshold. The dependent variables are defined as

$$\begin{aligned} did_f^H &= [log(p^T_{Q2}.x^T_{f,Q3}) - log(p^T_{Q2}.x^T_{f,Q2})] - [log(p^C_{Q2}.x^C_{f,Q3}) - log(p^C_{Q2}.x^C_{f,Q2})] \\ \text{and } did_f^V &= [log(p^T_{Q3}.|\Delta x^T_{f,Q4}|) - log(p^T_{Q1}.|\Delta x^T_{f,Q2}|)] - [log(p^C_{Q3}.|\Delta x^C_{f,Q4}|) - log(p^C_{Q1}.|\Delta x^C_{f,Q2}|)], \end{aligned}$$

where $x^T_{f,t}$ and $x^C_{f,t}$ denote the (column) vectors of holdings in treated and control stocks by fund f at time $t \in \{Q2, Q3\}$ and p^T_t and p^C_t denote the associated (row) price vectors. $\Delta x^T_{f,t}$ denotes $x^T_{f,t} - x^T_{f,t-1}$, with $\Delta x^C_{f,t}$ defined accordingly, and the notation $|.|$ is used for the element-wise absolute value of a vector. *size* denotes total assets under management in million USD. *Turnover* is a discrete variable ranging from -2 ("Very Low" portfolio turnover) to $+2$ ("Very High" portfolio turnover). *Price_to_book* denotes the fund's average price-to-book ratio based on its portfolio holdings, and *Index* is a binary variable equal to one for Index Funds, and zero otherwise. All variables are provided by Factset. 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.

Panel A: all stocks.

	Holdings		Trading volume		
Expl. variable	(1)	(2)	(3)	(4)	(5)
Cons.	0.001	-0.009	-0.180***	-0.251***	-0.247***
	(0.12)	(-0.55)	(-4.87)	(-4.77)	(-4.75)
Logsize		-0.002		-0.050**	-0.049**
		(-0.41)		(-2.48)	(-2.45)
Turnover		-0.038***		-0.148***	-0.131***
		(-3.40)		(-4.20)	(-3.75)
Price_to_book		-0.025**		-0.096***	-0.085**
		(-2.53)		(-2.68)	(-2.35)
Index		-0.183***		-0.464***	-0.383***
		(-5.18)		(-5.14)	(-4.21)
did_f^H					0.442***
					(2.78)
R^2	0.000	0.023	0.000	0.020	0.035
# Obs.	2,436	2,436	2,436	2,436	2,436

Panel B: SLP stocks.

	Holdings		Trading volume		
Expl. variable	(1)	(2)	(3)	(4)	(5)
Cons.	-0.001	-0.013	-0.195***	-0.290***	-0.286***
	(-0.12)	(-0.83)	(-4.77)	(-4.98)	(-4.93)
Logsize		-0.012**		-0.067***	-0.063***
		(-2.41)		(-3.07)	(-2.89)
Turnover		-0.034***		-0.172***	-0.160***
		(-3.47)		(-4.50)	(-4.21)
Price_to_book		0.011		-0.077*	-0.081*
		(1.08)		(-1.85)	(-1.95)
Index		-0.142***		-0.438***	-0.388***
		(-4.28)		(-4.51)	(-4.05)
did_f^H					0.347**
					(2.22)
R^2	0.000	0.016	0.000	0.021	0.029
# Obs.	2,189	2,189	2,189	2,189	2,189

Panel C: Non-SLP stocks.

	Holdings		Trading volume		
Expl. variable	(1)	(2)	(3)	(4)	(5)
Cons.	-0.007	-0.011	-0.265***	-0.150*	-0.144
	(-0.41)	(-0.39)	(-4.36)	(-1.65)	(-1.62)
Logsize		-0.006		-0.108***	-0.105***
		(-0.68)		(-3.08)	(-3.03)
Turnover		0.003		-0.010	-0.012
		(0.13)		(-0.16)	(-0.20)
Price_to_book		-0.037***		-0.036	-0.013
		(-2.68)		(-0.63)	(-0.23)
Index		0.032		-0.604***	-0.623***
		(0.51)		(-4.16)	(-4.46)
did_f^H					0.607***
					(2.84)
R^2	0.000	0.008	0.000	0.037	0.067
# Obs.	818	818	818	818	818

A.2 Figures

Figure 1: Graphical illustration of the causal impact of the FTT.

This figure illustrates our difference-in-difference estimates for the causal impact of the FTT on the market quality variables defined in 3. For each variable, we plot the cross-sectional average for treated (in black) and control (light red) stocks with a market capitalization of more than 1 billion EUR, minus the respective pre-event average. For improved readability, we use 3-day moving averages. The dashed lines indicate the sub-period averages for June/July, August, and September/October. The difference between the two dashed lines in September/October is equal to the diff-in-diff estimate of the causal impact of the tax.

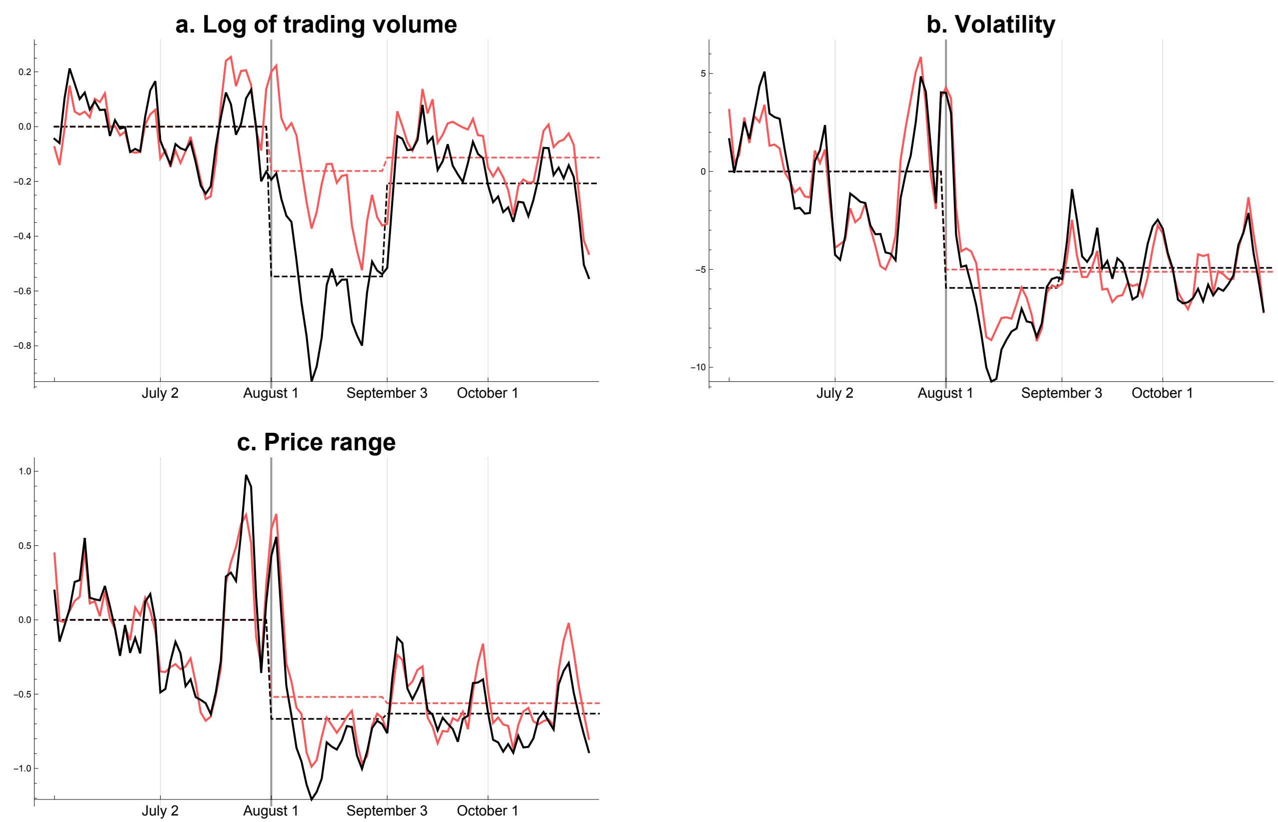

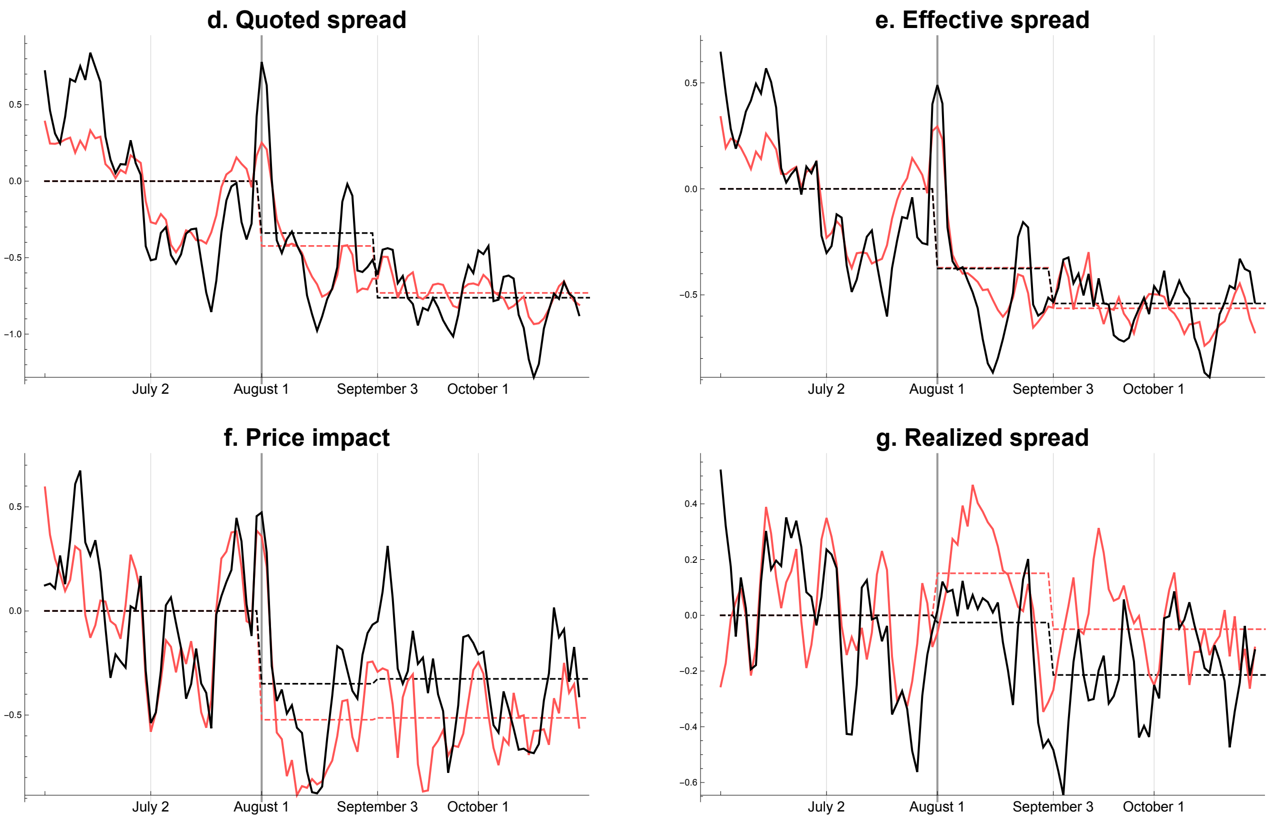

d. Quoted spread
0.5
0.0
−0.5
−1.0
July 2
August 1
September 3
October 1
e. Effective spread
0.5
0.0
−0.5
July 2
August 1
September 3
October 1
f. Price impact
0.5
0.0
−0.5
July 2
August 1
September 3
October 1
g. Realized spread
0.4
0.2
0.0
−0.2
−0.4
−0.6
July 2
August 1
September 3
October 1

h. Depth

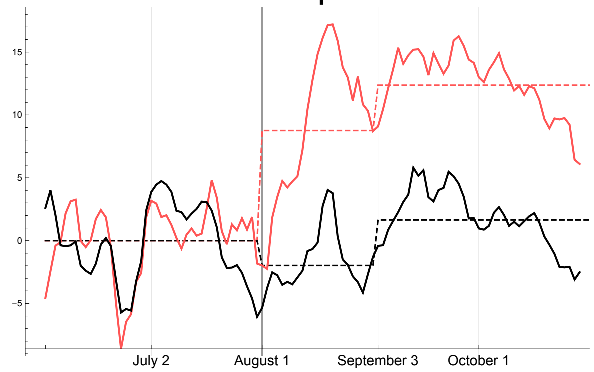

i. Resiliency

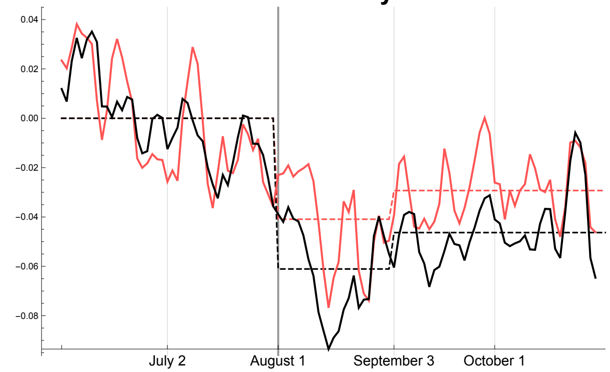

j. Autocorrelation

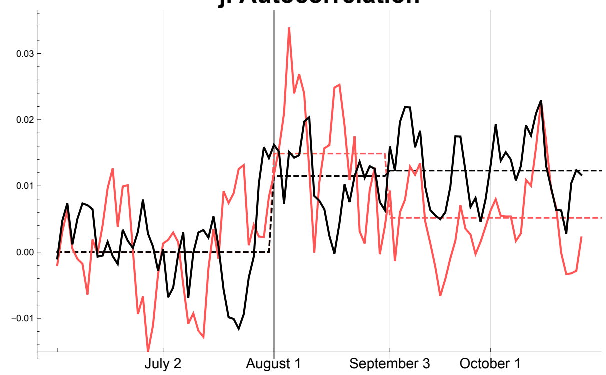

References

Amihud, Y., and H. Mendelson (1986): "Asset pricing and the bid-ask spread," *Journal of Financial Economics*, 17(2), 223 – 249. 3, 8, 24, 25

Baltagi, B., D. Li, and Q. Li (2006): "Transaction tax and stock market behavior: evidence from an emerging market," *Empirical Economics*, 31, 393–408. 1

Becchetti, L., M. Ferrari, and U. Trenta (2014): "The impact of the French Tobin tax," *Journal of Financial Stability*, 15(0), 127 – 148. 4

Bertrand, M., E. Duflo, and S. Mullainathan (2004): "How Much Should We Trust Differences-In-Differences Estimates?," *The Quarterly Journal of Economics*, 119(1), 249–275. 14, 19, 34, 35

Bloomfield, R., M. O'Hara, and G. Saar (2009): "How Noise Trading Affects Markets: An Experimental Analysis," *Review of Financial Studies*, 22(6), 2275–2302. 6

Brogaard, J., T. Hendershott, and R. Riordan (2014): "High-Frequency Trading and Price Discovery," *Review of Financial Studies*, 27(8), 2267–2306. 16, 18

Bushee, B. J. (2001): "Do Institutional Investors Prefer Near-Term Earnings over Long-Run Value?," *Contemporary Accounting Research*, 18(2), 207–246. 3

Capelle-Blancard, G., and O. Havrylchyk (2013): "The Impact of the French Securities Transaction Tax on Market Liquidity and Volatility," Working paper. 4

Carrion, A. (2013): "Very fast money: High-frequency trading on the NASDAQ," *Journal of Financial Markets*, 16(4), 680 – 711, High-Frequency Trading. 18

Coelho, M. (2014): "Dodging Robin Hood: Responses to France and Italy's Financial Transaction Taxes," Discussion paper. 4

Constantinides, G. M. (1986): "Capital Market Equilibrium with Transaction Costs," *Journal of Political Economy*, 94(4), 842–62. 3, 7, 25, 28

De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann (1990): "Positive Feedback Investment Strategies and Destabilizing Rational Speculation," *The Journal of Finance*, 45(2), pp. 379–395. 6

Deng, Y., X. Liu, and S.-J. Wei (2014): "One Fundamental and Two Taxes: When Does a Tobin Tax Reduce Financial Price Volatility?," Working Paper 19974, National Bureau of Economic Research. 6

Derrien, F., A. Kecskes, and D. Thesmar (2013): "Investor Horizons and Corporate Policies," *Journal of Financial and Quantitative Analysis*, 48, 1755–1780. 3

Diamond, P., and E. Saez (2011): "The Case for a Progressive Tax: From Basic Research to Policy Recommendations," *Journal of Economic Perspectives*, 25(4), 165–90. 28

Dow, J., and R. Rahi (2000): "Should Speculators Be Taxed?," *The Journal of Business*, 73(1), 89–107. 7

Ferreira, M. A., and P. Matos (2008): "The colors of investors' money: The role of institutional investors around the world," *Journal of Financial Economics*, 88(3), 499 – 533, Darden - {JFE} Conference Volume: Capital Raising in Emerging Economies. 22

Foucault, T., D. Sraer, and D. Thesmar (2011): "Individual Investors and Volatility," *The Journal of Finance*, 66(4), 1369–1406. 6

Grinblatt, M., S. Titman, and R. Wermers (1995): "Momentum Investment Strategies, Portfolio Performance, and Herding: A Study of Mutual Fund Behavior," *American Economic Review*, 85(5), 1088–1105. 22

Haferkorn, M., and K. Zimmermann (2013): "Securities Transaction Tax and Market Quality - The Case of France," Working paper. 4

Halvorsen, R., and R. Palmquist (1980): "The Interpretation of Dummy Variables in Semilogarithmic Equations," *American Economic Review*, 70(3), 474–75. 16

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. 1

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. 1

Keim, D. B., and A. Madhavan (1997): "Transactions costs and investment style: an inter-exchange analysis of institutional equity trades," *Journal of Financial Economics*, 46(3), 265–292. 27

Kempf, A., D. Mayston, and P. Yadav (2009): "Resiliency in limit order markets: a dynamic view of liquidity," Working paper. 13

Keynes, J. M. (1936): *The General Theory of Employment Interest and Money*. Palgrave MacMillan. 1, 5, 7

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. 12

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. 1

Manconi, A., M. Massa, and A. Yasuda (2012): "The role of institutional investors in propagating the crisis of 2007-2008," *Journal of Financial Economics*, 104(3), 491 – 518, Market Institutions, Financial Market Risks and Financial Crisis. 22

Matheson, T. (2011): "Taxing Financial Transactions; Issues and Evidence," IMF Working Papers 11/54, International Monetary Fund. 1

Meyer, S., M. Wagener, and C. Weinhardt (2014): "Politically Motivated Taxes in Financial Markets: The Case of the French Financial Transaction Tax," *Journal of Financial Services Research*. 4, 11

Pagano, M. (1989): "Trading Volume and Asset Liquidity," *The Quarterly Journal of Economics*, 104(2), 255–274. 6

Pomeranets, A., and D. Weaver (2012): "Security transaction taxes and market quality," Working paper 2011-26, Bank of Canada. 1

Roll, R. (1989): "Price volatility, international market links, and their implications for regulatory policies," *Journal of Financial Services Research*, 3, 211–246. 1

Ross, S. A. (1989): "Commentary: Using tax policy to curb speculative short-term trading," *Journal of Financial Services Research*, 3. 6

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. 6

Song, F. M., and J. Zhang (2005): "Securities Transaction Tax and Market Volatility," *The Economic Journal*, 115(506), 1103–1120. 6

Stein, J. C. (1989): "Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior," *The Quarterly Journal of Economics*, 104(4), 655–69. 3, 28

Stiglitz, J. (1989): "Using tax policy to curb speculative short-term trading," *Journal of Financial Services Research*, 3, 101–115. 1, 6, 8, 21

Subrahmanyam, A. (1998): "Transaction Taxes and Financial Market Equilibrium," *The Journal of Business*, 71(1), pp. 81–118. 7

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. 6, 8

Umlauf, S. R. (1993): "Transaction taxes and the behavior of the Swedish stock market," *Journal of Financial Economics*, 33(2), 227–240. 1

Vayanos, D. (1998): "Transaction costs and asset prices: a dynamic equilibrium model," *Review of Financial Studies*, 11(1), 1–58. 3, 7, 25