# A BLUEPRINT FOR EUROPE?

# Investigating a Financial Transaction Tax on the French Market


*Authors*

Markus Lithell

Johan Ljungkvist

*Advisor*

Laurent Bach


*Date*

May, 2013


**ABSTRACT**

This paper investigates the market consequences of financial transaction tax (FTT) implementation. We use data on volume and volatility for securities subject to a recent French FTT, as well as comparable European securities that were not. We find that the FTT has a significant negative effect on both volume and volatility. Furthermore, we attempt to isolate the effects of the FTT's two components: a securities transaction tax (STT) and a high-frequency trading tax (HFTT). No volume or volatility effects are found when the STT is scrutinized on its own. However, we are unable to find any statistical significance in the difference between the diff-in-diff results from the two tests, thus preventing attribution of the negative volume and volatility effects to either HFTT or STT.



*Keywords: financial transaction tax, securities transaction tax, high-frequency trading tax*

## INTRODUCTION

Since the debate on financial transaction taxes (FTTs) was sparked in the early 1970's (with the introduction of the Tobin tax), the issue has been contested to no clear consensus. In the wake of the recent financial crisis, interest in FTTs has resurged; a multitude of governments have set to appraising the tax's potential, partly to pacify financial markets that are seen as turbulent and unpredictable, partly to nourish financially-troubled governments. At the moment, eleven European governments are in favor of implementing FTTs or have recently done so. The French government, one member of this group, recently instituted an FTT that went into effect on August 1st, 2012. This particular case forms the basis of our study.

Although several previous empirical studies have been conducted on the topic of FTTs, few have had access to a natural experiment as favorable as the recent French case. The French case provides the ideal springboard for FTT analysis, as it features a "true" FTT, where many previous studies have been forced to gauge the effects of an FTT using proxy events such as commission deregulations. Furthermore, the French FTT is holistically formulated, which makes it difficult for traders to circumvent. Some previous studies have investigated cases involving poorly designed FTTs; bizarre investor migration ensued. While such a study might be informative on the consequences of a lopsided FTT, it has little relevance for policymakers interested in the general application of an FTT. Finally, the French case features clear treatment-control configurations. The case has two distinct FTT targets: domestically listed French equities and French ADRs. The tax goes into effect on separate dates for these two groups. As such, the French case offers variety of temporal and unitary treatment-control configurations. This is a particularly favorable circumstance in comparison to earlier papers that have been forced to adopt second-best comparison groups. Altogether, these three reasons differentiate the French case from many of its predecessors.

Furthermore, there is no established consensus among empirical studies of FTTs to date; an essay presented by the *Österreichisches Institut für Wirtschaftsforschung* tabulates the results of 21 empirical studies on the subject from the past 30 years (Schulmeister *et al.*, 2008). Of these 21 studies, 10 found a positive connection between FTTs and price volatility, 5 found no such relationship, and 6 found this link to be negative. As such, we feel that a study on the French case is in an optimal position to contribute to the debate.

In this paper, we attempt to evaluate the market-relevant effects of an FTT using the French case. We conduct three experiments to this end. First, we compare the volume and volatility behaviors of tax-subject French firms with equivalent European equities. This stage is dubbed "FTT testing". We find that the FTT reduces both volume and volatility on aggregate. Second, in the "STT testing" experiment, we attempt to analyze the FTT in a bifurcated form by segregating the tax's two components: a securities transaction tax (STT) and a high-frequency trading tax (HFTT). To do this, we utilize the ADR equivalents of the equities from FTT testing, as ADRs are unaffected by the HFTT. In doing so, we find that the STT has noticeable effects on neither volume nor volatility. Third, we set out to determine whether the difference in results between FTT and STT testing is statistically significant. In other words, the "combined testing" stage checks whether the STT and FTT truly have different effects, thus implying that the HFTT accounts for a majority of the FTT's effect, or if these effects are, in fact, indistinguishable. The test yields results that indicate the latter to be the case; we cannot prove that the FTT and STT have significantly different effects.

The paper is laid out as follows. First, the theoretical literature on FTTs is summarized ("Previous Literature"). An overview of the French FTT follows ("Case Details"). Data and methodology are documented in the third and fourth sections ("Data" and "Method"). After that, the results of our study are recorded, the analysis of which follows ("Results" and "Analysis"). Finally, our concluding thoughts on the topic are presented ("Concluding Remarks"). To supplement the body of our paper, two appendices are provided, elaborating on our model and case ("Appendix A" and "Appendix B").

## PREVIOUS LITERATURE

Academics tend to view FTTs in one of two ways. One side purports that FTTs are undesirable, as they impede with the market's ability to operate efficiently. The other touts that markets are, in fact, too efficient in unregulated circumstances; thus, FTTs can improve market efficiency by curtailing noise trading. We will now analyze each of these standpoints.

### Arguments in Favor of an FTT

J.M Keynes argued in favor of FTTs in *The General Theory of Employment, Interest and Money* (1936). Keynes's argumentation is still relevant today, not least because he touches upon two points that stand at the very center of the current FTT debate. Firstly, Keynes distinguished between two different approaches to investing: speculation, which he defined as "the activity of forecasting the psychology of the market", and enterprise, defined as "the activity of forecasting the prospective yield of assets over their whole life". In short, Keynes argued that speculation could be detrimental to the functioning of financial markets if it became too prevalent. The following oft-cited passage illustrates this point:

> *"Speculators may do no harm as bubbles on a steady stream of enterprise. But the position is serious when enterprise becomes the bubble on a whirlpool of speculation. When the capital development of a country becomes a by-product of the activities of a casino, the job is likely to be ill-done."*

Keynes's preferred method of curtailing speculation (which he indeed saw as a whirlpool-in-development) was by way of FTT. This brings us to what we consider to be the second still-relevant part of Keynes's argument. While favoring FTTs, Keynes conceded the existence of a trade-off between transaction costs and market liquidity (in other words, an FTT would come at the cost of lower liquidity). As such, the investor confidence gained from stabilized prices may be partly counteracted by increased liquidity uncertainty.

In addressing the desirability of a securities transactions tax (STT), which is an FTT applied specifically to the securities market, Stiglitz (1989) and Summers & Summers (1989) agree with the premise that transactions taxes are a viable way of countering speculation. Referring to French and Roll (1986), they support this conjecture by providing empirical evidence indicating that market volatility was reduced as a result of the market closing for one day. Stiglitz argues that the closing of the market can be seen as an extreme form of a transaction tax (in practice equivalent to the introduction of a 100% transaction tax). He continues by saying that if the market consisted purely of what Keynes would call enterprising investors (i.e. if the market responded to fundamentals alone), then the closing of the market for one day should not affect the difference between the closing price on the day before and the opening price on the day after. If, on the other hand, the market to a large part consisted of speculators, then the closing of the market would reduce their ability to add noise to the market and thus reduce the deviation between the closing price on the day before and the opening price on the day after. According to Stiglitz, the reduced volatility observed by French and Roll thus supports his claim that speculators indeed are a source of market volatility.

At the heart of both papers, therefore, lies the argument that an FTT would decrease speculation and thereby also market volatility, while simultaneously leaving desirable fundamentals-based investing largely unaffected; over longer investment horizons, percentage returns are likely to be substantially larger than over the course of a few days or even a few hours, rendering such a (small) tax negligible. Furthermore, speculators, by their very nature, trade more frequently than long-term, "enterprising", investors. Consequently, they will be hit harder by such a tax.

In addition to the central argument outlined above, the two papers provide a number of other arguments in favor of an STT. For example, both argue that an STT could contribute toward lowering the profitability of the financial industry, incentivizing the most talented students to pursue careers in other industries, in which they can do society more good. Summers & Summers also argue that in curbing speculation, an STT would effectively lead to the emergence of a less myopic stock market, which in turn would allow firm managers to adopt a longer-term perspective in running their companies. Finally, an STT would, of course, also generate revenue to the government, which could also justify implementation.

**Arguments against an FTT**

To a large extent, the whole debate about the desirability of an FTT hinges on the validity of the efficient market hypothesis. Summers & Summers make this point when saying:

> *"The belief that facilitating trading improves the social functioning of financial markets is premised on the acceptance of the efficient market hypothesis. If prices in unfettered financial markets closely track fundamental values, then they will provide proper economic signals, guide investment appropriately, and facilitate the spreading of risks."*

In other words, this argument touts that if the market, on aggregate, correctly prices assets in spite of speculation (or maybe even because of it), then a tax to dissuade such behavior will be fruitless and likely even harmful. Summers & Summers, who argue in favor of an STT, consequently dedicate a part of their paper to making a case against the efficient market hypothesis.

Another common objection is to question the notion that FTTs push down stock market volatility. Although both sides agree that the tax will cause trading volume to decrease, they disagree when it comes to the resultant impact on market volatility. Whereas proponents argue that this lost volume largely consists of speculators and that volatility is reduced accordingly, opponents contend that the drop in volume will cause prices to be updated less frequently, giving rise to sharper swings in stock prices. An empirical paper by Farmer *et al.* (2004) gives support to this claim. After studying trading in AstraZeneca, they find that the stock's price jumped £374m in response to a small buy order of just £16,000 in a period during which the stock was unusually illiquid.

An outline of other potential problems associated with FTT implementation is provided by Schwert & Seguin (1993). The arguments emphasize that higher transaction costs give rise to higher required rates of return (i.e. higher costs of capital); that the burden of the tax, although designed to fall upon so-called speculators, hits a far wider group of investors; that the implementation costs for the government as well as the compliance and potential avoidance costs for investors are significant; and that capital structures are distorted (i.e. trading simply moves away from the securities being taxed to alternative investment vehicles not covered by the tax).

In summary, the key point of contention in the FTT debate is whether or not a transactions tax will lower market volatility. Virtually all academics agree that market volume will decrease as a result of a tax, but there is disagreement over how this drop in volume translates into volatility changes. The theoretical literature's failure to reach a consensus has spurred a variety of empirical studies. In the next section, we examine the backdrop of our own paper and its relationship to existing studies.

## CASE DETAILS

### Relationship to Existing Studies

Before we proceed with addressing the specifics of the French case, it is useful to establish how our case relates to the existing empirical literature on FTTs. As previously mentioned, our study differs from many of its predecessors. These differences are a consequence of circumstance; the 2012 French FTT has several features that make it ideal for empirical analysis.

First, the French FTT is holistic and thusly relatively devoid of loopholes (as described in the Case section). As this FTT is comparatively hard to circumvent, the case allows us to clearly isolate the theoretical taxation effect. This is important if conclusions are to have general applicability. The benefits of analyzing such an FTT become clearer when contrasted with a poorly formulated FTT. One notable example is the Swedish FTT that was applied (and intermittently modified) to Swedish stocks from 1984 to 1991, but only affected transactions conducted by domestic brokers. Such a faulty formulation allowed the tax to be easily avoided, making it a poor vehicle for FTT analysis (as it represented poor taxation methodology). Umlauf's (1993) empirical study is unfortunately hampered by such circumstances. Ironically, the Swedish case is probably the most-cited evidence against FTTs to this day.

Second, the French FTT case allows for the use of distinct and closely related treatment and control groups. The eligible candidates vary foremost in the unitary dimension (with listings in both Europe and the United States), but also in the temporal dimension (with several valuation points and implementation dates). This variety improves our chances for robust results. Several (otherwise successful) studies suffer from a lack of clear control groups. One example is Liu & Zhu's (2009) study on commission taxation on the Japanese stock market. In this case, taxation was applied to all domestic stocks. The pair uses Japanese ADRs as a control group. This has, as they admit, a problem of size; Japanese ADRs will only be representative of the largest firms affected by the tax. Since no better control options are available, they choose to apply the "'second best' control portfolios" (Liu & Zhu, 2009).

Third, the French case features a "true" FTT, in that the tax above all applies to financial transactions. This minimizes the extent to which our study requires the drafting of proxies. Accordingly, our scenario provides us with the opportunity to closely approximate the "true" effect of FTTs. Additionally, the case involves both an STT and an HFTT, which should be separable by varying time and control-treatment group configurations, as is argued in the Method section. One study which is troubled by the need to proxy for FTT effects is that of Jones and Seguin (1997). The study, which deals with commission deregulation on the NYSE, uses a similar proxy to the aforementioned Japanese study. In both cases, the researchers use the removal of minimum commission rates to approximate the removal of a true FTT. Although this is a viable approach, it does make results slightly more opaque.

For these three reasons, we think that our case is in an ideal position to contribute to the FTT debate. Since the empirical evidence has thus far been relatively mixed (as indicated by Schulmeister *et al.*, 2008), we feel that such a contribution is especially relevant.

**Background**

The implementation of an FTT in France was first proposed by President Nicholas Sarkozy in January 2012, an election year, in a climate where polls were showing overwhelming support across the political spectrum for such a tax. In this context, the tax was marketed as an instrument that would (a) help restore justice by forcing the financial sector to contribute towards paying for the very crisis that it had caused, (b) generate revenue to help balance the government's accounts, and (c) demonstrate the viability of an FTT to other European countries. During the presidential election campaign, would-be winner Francois Hollande declared that he too would carry through with the tax. The tax that came into force on August 1st, 2012, was thus largely a construct of the Sarkozy presidency, with the main exception that the tax rate was increased from 0.1% to 0.2% under Hollande.

The FTT was far less well-received in other quarters. Following Sarkozy's announcement, France's financial sector vehemently opposed the FTT, voicing concerns that the tax would impede growth, damage competitiveness, and vex the financing of the economy. Another repeated critic is the British Prime Minister David Cameron, who has called the tax "mad".

**The Tax**

Articles *235 TER ZD*, *235 TER ZD BIS*, and *235 TER ZD TER* of the French tax code constitute a bundle of components that together form the FTT. The third item is a tax on the purchase of credit default swaps (CDSs) on sovereign debt, and thus lies outside the scope of this study. The other two both impact equity trading however, meaning that they are relevant and deserve additional explanation.

*235 TER ZD:* This article deals with a 0.2% tax that applies to equity securities transactions (which is why we henceforth refer to it as an STT) in companies that (a) have their registered office in France and (b) had a market capitalization in excess of €1bn on

December 1st of the year preceding the year of taxation (with the exception of 2012, when this date is set to January 1st of the same year). Worthy of note is the fact that only securities acquisitions (thus, only the buyer is tax-liable) resulting in a daily net ownership change of one of these companies are subject to the tax.[1] The tax applies regardless of where the security is traded, regardless of the residence of the parties involved in the transaction, and regardless of where the contract is entered into. Apart from equity securities, the tax also applies to investment certificates, voting right certificates, and depositary receipts; the latter category encompasses American depositary receipts (ADRs), which are subject to the tax as of December 1st, 2012.

A number of exemptions are also contained in the article. Apart from the convertible bonds mentioned earlier, these exemptions cover the primary issuance of equity instruments, transactions by clearing houses in the course of their clearing activities, employee incentive plans, temporary transfers, and acquisitions by market makers (the last three exemptions have been simplified; see Appendix B for further details).

*235 TER ZD BIS:* This article outlines a 0.01% tax (hereafter referred to as an HFTT) on high frequency trading (HFT) that is levied on traders with a cancellation and/or modification ratio with respect to orders in excess of an 80% threshold (see Appendix B for specifics). The behavior targeted by this tax is known as "quote stuffing", and has to do with traders placing a huge number of orders, only to immediately thereafter cancel most of them, misleading competitors in the process.

The HFTT only applies to companies operating in France. In other words, firms engaging in HFT activities on the French market from abroad are not subject to the tax (this includes foreign branches of French firms). Additionally, market-making activities, as defined in Appendix B, are exempt. Importantly, all eligible parties have to pay the tax not only when trading in French stocks with a market capitalization greater than €1bn, but regardless of which security they trade in.

**Differentiating between FTT Components**

Since both the STT and the HFTT target the equity market and were introduced on the same day, trying to gauge the isolated effect of either one by studying the aggregate effect of the FTT is problematic. However, for purposes of general applicability, we nevertheless deem

[1] For example, if during the course of one day someone buys 2,000 shares and sells 500 shares in Sanofi, then the tax is calculated on the basis of a transaction of 1,500 shares (2,000-500) and not 2,500. In other words, only acquisitions causing a change in ownership, not the volume of trade per se, determine the tax. Derivative products resulting in the transfer of ownership are also affected. Convertible bonds are exempt, however.

this to be desirable (few, if any, other countries are likely to introduce a combined STT and HFTT; they are more likely to opt for only an STT instead). Fortunately, we would like to argue that there exist viable, alternative, procedures that manage to give a very good indication of the true effect of the STT on market volume and volatility. We will now describe two of these procedures, each of which depends on a different peculiarity of the French FTT to differentiate between the STT and the HFTT.

*Recategorization:* On January 1st, 2013, the list of stocks targeted by the STT was updated on the basis of their market capitalizations on December 1st, 2012 (for an overview of the relevant dates, refer to Figure 1). This meant that two stocks that had previously been targeted ceased to be so, and seven stocks that previously had not been targeted now were. Since the HFTT remained unchanged during this period, a study of these stocks as they passed from being non-targets to targets and vice versa would catch the STT effect.

Unfortunately this procedure suffers from an extremely small sample size. After data cleaning (see Data section for details on this procedure), only one stock is dropped and six stocks added on the recategorization date. Furthermore, momentum biases may cause volume and volatility to develop differently for stocks that are dropping in price and those that are climbing in price. Since this would lead to in uncertain results, we rely instead on the second test, which we will now describe.

*ADRs:* As of December 1st, 2012 the STT affects the trading of ADRs. The HFTT, on the other hand, will not impact ADRs. How can we draw this conclusion? As the HFTT only applies to high-frequency traders operating out of France, such actors would be at a competitive disadvantage when trading on the American market due to the large physical

**Figure 1: Relevant Dates**

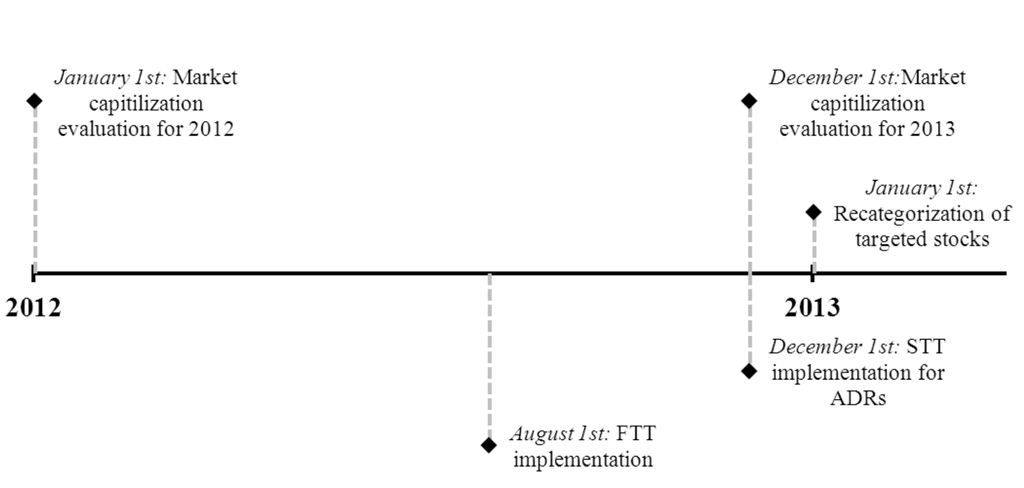

distance (which plays an integral role in high-frequency trading); if they wished to engage in such a practice, they would have to do so through a subsidiary in the US, which would be HFTT exempt. Hence, a study of the STT effect on ADRs would not be influenced by the HFTT.

We thusly regard a study on French ADRs as the best way to determine the effect of the STT on market volume and volatility. However, we have found anecdotal evidence indicating that some brokers in the US are not passing on the tax to their clients, but rather using commission fees to cover the cost. Such a practice will shift the tax burden away from the actual investors behind the trade to the brokers, eliminating the tax altogether from the investors' point of view. We were unable to glean exactly how prevalent this is, but it is obvious that if common enough, it would substantially reduce the effect (assuming that there is one) of the tax. On the whole, however, this practice seems to be relatively rare.

**Enforcement**

The institution legally obliged to both provide the declarations for an STT-liable trade and subsequently pay the tax (i.e. the accountable party) is either the investment service provider (broker) that executed the transaction on behalf of its client, or the securities account holder when a broker does not execute the transaction. The accountable party then sends the tax declaration and payment to Euroclear France, which has been given a mandate by the French government to collect the tax. Failure to declare and pay on time will trigger Euroclear France to report the matter to the tax authorities.

As regards the situation in the United States, the very legality of the STT on ADRs is being questioned owing to the existence of a US income tax treaty with France from 1994 that contains a clause forbidding extraterritorial STTs. It is believed that discussions are ongoing between French and US authorities on whether the current STT is covered by that tax treaty, but that the US Treasury is assisting in its enforcement in the interim. In any case, it is clear that US investors are having to pay the tax for now.

# DATA

## Data Overview

Data are collected using Thomson Reuters' DataStream. They cover equity stock markets in the European Union as well as their American Depositary Receipt (ADR) counterparts. We group stocks into categories depending on treatment and control group designation. For each security, several data types are collected: turnover volume, closing price, opening price, highest price, lowest price, and market capitalization. These data span from January 1st, 2012, to February 28th, 2013. The event date depends on a stock's classification. For EU equities, the event date is August 1st, 2012, while the corresponding date for ADRs is December 1st, 2012. For more details, refer back to the Case section.

For every security, a cleaning procedure is conducted to eliminate poor data, focused on closing price, opening price, highest daily price, and lowest daily price. Once data are sorted by stock and date, the number of consecutive 5-trading-day periods featuring static prices is tallied. If any stock features more than 20 such periods (including overlap), the stock is excluded from testing. For an in-depth discussion on the cleaning protocol's construction, rationale, and purpose, refer to Appendix A.

Table 1 features stock count as well as mean and median market capitalization values for reference and comparison purposes. We now proceed to explain each data group in turn, as designated by the "note" column of this table. Two distinct, but complementary, experiments are conducted using the "FTT testing" and "STT testing" datasets. Their usage is explained in the next section, Method.

*Note 1:* The treatment group for FTT testing consists of Euronext Paris-listed stocks with market caps of over €1bn on both of the evaluation dates: January 1st, 2012, and December 1st, 2012. Firms that are listed on Euronext Paris but have headquarters abroad are excluded (as they are not tax-subject). Equities that only exceed the market cap threshold on one of the two valuation dates are omitted; they are only FTT-subject for a limited portion the post-event window. As the number of such stocks only totals seven after cleaning, the sample is relatively unaffected by their exclusion.

*Note 2:* The control group for FTT testing consists of stocks with market caps of over €1bn on January 1st, 2012, from Belgium, Italy, the Netherlands, and Spain. These countries were chosen because they belonged to the largest Eurozone trading partners with France, and met a certain 10-year national index correlation threshold. Since these countries are all Eurozone members, we avoid any complications due to currency effects. See Appendix A for details.

*Note 3:* The treatment group for STT testing consists of the ADR counterparts of French FTT-subject firms, i.e. the equities from the group examined in Note 1. As a majority of French firms are unavailable as ADRs, only a limited portion (about one half) of the stock data is available for analysis. Consequently, one might worry that these ADRs might have some common characteristic (such as large size, as the market values in Table 1 would seem to indicate) that causes their sample to be relationally biased towards the EU equities. While a sample selection bias may occur, it will not affect testing; the control group, as described next, will be subject to the exact same biases. Thus, the treatment-control relationship will be intact, whether or not a bias exerts influence on the ADR data.

The same can be argued of the extensive data cleaning that occurs for ADR data. The ADRs undergo noticeably more trimming as a result of generally low turnover volumes. One might be concerned that the removal of such illiquid equities might induce a bias in our overall results. However, such concern is unwarranted; as both the treatment and control groups undergo the same extensive data cleaning, the treatment-control relationship, for all testing purposes, will still be intact. The situation does become slightly more complicated for combined testing, as will be discussed in the Results section.

*Note 4:* The control group for ADR testing represents firms from the same four countries as in Note 2, with a listed market cap of over €1b on January 1st, 2012. All ADRs fulfilling these criteria are included.

**Treatment-Control Quality Assessment**

Are the control groups, as outlined above, suitable for their role? To find out, we calculate the event window treatment-control correlation for equally weighted daily averages of each of our measures (as described in the following section). The results are presented in Table 2.

On the whole, the treatment and control groups appear to behave very similarly during the event window. This is especially true for volatility, which features solid treatment-control correlations for both FTT testing and STT testing. Volume is also fine for FTT testing, but is noticeably weaker for STT testing. Thus, we should keep in mind that the volume results for this set of regressions may be slightly unreliable. For the most part, however, we can be confident that our control groups are well-calibrated for their purpose.

TABLE 1: DATA OVERVIEW

*Stock count before and after data cleaning, market capitalization figures*

| Category | Pre-clean count | Post-clean count* | Mean MV** | Median MV** | Note |
|---|---|---|---|---|---|
| *FTT testing (EU equities)* | | | | | |
| Treatment | 107 | 106 | 9782 | 3963 | *1* |
| France | 107 | 106 | 9782 | 3963 | |
| Control | 139 | 134 | 8041 | 3145 | *2* |
| Belgium | 22 | 21 | 6924 | 3048 | |
| Italy | 43 | 42 | 6804 | 2674 | |
| Netherlands | 32 | 30 | 8273 | 4687 | |
| Spain | 42 | 41 | 9710 | 3453 | |
| *STT testing (ADRs)* | | | | | |
| Treatment | 60 | 39 | 26044 | 14039 | *3* |
| France | 60 | 39 | 26044 | 14039 | |
| Control | 94 | 45 | 24902 | 13427 | *4* |
| Belgium | 12 | 4 | 29884 | 8483 | |
| Italy | 29 | 11 | 20510 | 13426 | |
| Netherlands | 26 | 21 | 21080 | 11793 | |
| Spain | 27 | 9 | 36976 | 38066 | |

Above, some characteristics of our data sample are outlined. The data are presented in two categories: FTT testing and STT testing. The former consists of equities traded on European markets, while the latter is made up of ADR counterparts. Further, each category is divided into a treatment and control group, the compositions of which are determined by nationality (and are explained more fully in Appendix A). Note that treatment and control groups tend to have very similar market capitalization characteristics within each testing group. Furthermore, note that the STT testing data belong to a higher market cap echelon than the FTT testing data; thus, it would seem that there is some positive correlation between size and ADR listing. The implications of this are subject to further discussion in Note 3. Finally, one should be aware of the fact that the market cap values for EU equities are presented in EUR, while those of ADRs are shown in USD.

**Uses 5-20 data cleaning procedure*

***Post-cleaning figures, valuation on Jan 1, 2012*

TABLE 2: TREATMENT-CONTROL CORRELATIONS

*Event window correlation using daily averages of volume and volatility measures*

| Measure | FTT testing | STT testing |
|---|---|---|
| Volume (non-logarithmic)<br>*Event window standard* | 0.833 | 0.327 |
| Volume (logarithmic)<br>*Event window standard* | 0.923 | 0.676 |
| Volatility (daily basis)<br>*Jones-Seguin* | 0.887 | 0.883 |
| Volatility (weekly basis)<br>*Yang-Zhang* | 0.900 | 0.960 |
| Volatility (monthly basis)<br>*Yang-Zhang* | 0.900 | 0.978 |
| Volatility (weekly basis)<br>*"Common"* | 0.932 | 0.953 |
| Volatility (monthly basis)<br>*"Common"* | 0.888 | 0.941 |

In order to assess the suitability of our control groups, the event window correlation between treatment and control group data is calculated. The process employs daily averages for each of our volume and volatility measurements, the usage of which will receive further explanation in Method. With the exception of STT testing volume, the correlations indicate sound control-treatment relationships.

## METHOD

### Experimental Setup

In the main body of our essay, three experiments are conducted. The first experiment, dubbed "FTT testing", uses a treatment group that consists of tax-subject French equities and a control group of European equities (as described in the Data section). The experiment compares the volume and volatility measurements for each group before and after the implementation of the FTT on August 1st, 2012. The regressions conducted for this test do not distinguish between the HFTT and STT; both are relevant for this sample at this time.

The second set of testing, labeled "STT testing", uses the ADR counterparts of the treatment and control groups from primary testing (whenever these are available) under a similar set of testing parameters. This time, the event day is December 1st, 2012 – the date on which the ADR tax takes effect. These regressions will only analyze the effects of the STT, as we do not expect French high-frequency traders to conduct any transactions with ADRs (as explained in the Case section).

For these two experiments, we employ identical experimental setups. In both cases we were faced with the combined facts that (a) we were dealing with a natural experiment and that (b) we could construct a comprehensive panel data set with relative ease. Given these circumstances, we regarded the use of a difference-in-differences (diff-in-diff) estimator as the most auspicious approach with regard to achieving respective experiments' goals. In general, our regressions thus take the following form:

$$E(y_{it}|i,t) = \alpha_0 + \alpha_1 dT_i + \alpha_2 d2_t + \beta_1 dT_i \cdot d2_t + a_i + u_{it} \qquad \textit{Eq. 1}$$

where $\beta_1$ is the diff-in-diff estimator. As regards the independent variables, $d2_t$ is a dummy representing the transactions tax introduction and $dT_i$ is a dummy used to differentiate between the control and treatment group. The last two terms constitute the model's error term, $a_i$ being the time-invariant component (i.e. a stock fixed effect) and $u_{it}$ varying over time. Finally, $y_{it}$ is the dependent variable, in our case either volume or volatility.

The third experiment makes use of all the data underlying experiments one and two in order to statistically test the difference between the FTT and the STT. Since this setup consists of two treatment and control groups as well as two different event dates, the model represented by Equation 1 is too limited. We therefore construct a different model specifically for this experiment:

$$E(y_{it}|i,t) = \delta_0 + \underbrace{\delta_1 dF_i + \delta_2 dADR_i + \delta_3 dF_i \cdot dADR_i}_{Level\,\text{Effects}} + \underbrace{\delta_4 dT_{1,t} + \delta_5 dT_{2,t}}_{Time\,Effects} + ...$$

$$\underbrace{\delta_6 dF_i \cdot dT_{1,t} + \delta_7 dF_i \cdot dT_{2,t}}_{French\,Time\,\text{Effects}} + \underbrace{\delta_8 dADR_i \cdot dT_{1,t} + \delta_9 dADR_i \cdot dT_{2,t}}_{ADR\,\text{Time Effects}} + ...$$

$$\underbrace{\varepsilon_1 dF_i \cdot (1 - dADR_i) \cdot dT_{1,t}}_{FTT\,d\text{iff-in-diff}} + \underbrace{\varepsilon_2 dF_i \cdot dADR_i \cdot dT_{2,t}}_{STT\,\text{diff-in-diff}} + a_i + u_{it} \qquad \textit{Eq. 2}$$

where $\varepsilon_1$ and $\varepsilon_2$ are the diff-in-diff estimators. Regarding the independent variables, $dF_i$ is a dummy representing whether or not a stock is French (French ADRs are encompassed by this definition), $dADR_i$ is a dummy representing ADRs, $dT_{1,t}$ is a dummy equal to one after August 1st, 2012 (the FTT event date), and $dT_{2,t}$ is a dummy equal to one after December 1st, 2012 (the STT event date). Lastly, the error terms, $a_i$ and $u_{it}$, as well as the dependent variable $y_{it}$, all have the same interpretation as in the previous model.

Finally, armed with all the regression coefficients from the regression just described, we use an F-test to test the null hypothesis that $\varepsilon_1$ and $\varepsilon_2$ are equal to one another (put differently, we test whether we can reject that the FTT and the STT have the same effect on market volume and volatility).

**Specification Tests**

As we are dealing with panel data, a natural question to ask is whether to use fixed or random effects regressions. To answer this question we make use of the Hausman test, which indicates that fixed effects regressions are the appropriate choice. Accordingly, the stock fixed effect and time-invariant dummies are removed prior to estimation. We also run a modified Wald test, as preinstalled in Stata, to test for panel-level heteroskedasticity. We find heteroskedasticity to be an issue, which is why we elect to use clustered robust standard errors, a procedure that also corrects for serial correlation.

In summary, we opt to use fixed effects transformation and clustered robust standard errors in our regressions.

**Volume Standardization**

As a consequence of our samples containing stocks with market caps ranging from €1bn to roughly €100bn, we were faced with the problem of aggregating vastly different volumes for meaningful comparisons over time. If we were to simply sum all the volumes by day, for example, the largest stocks would be heavily over-weighted in our results, owing to their larger trading volumes. In effect, our results would not reflect the effect of the FTT on the purported group (i.e. the treatment and control group) as a whole, but would rather, to a large extent, mirror the effect on the largest stocks in that group. To get around this problem we choose to standardize trading volumes on a stock-by-stock basis, expressing volumes in terms of each stock's mean volume during the event window. We then aggregated these new figures by group designation, giving us equally-weighted volume figures for both the treatment and control groups.

**Robustness**

Returning to our choice of model (Equations 1 and 2), we argue that there are two major threats to the robustness of our results. First, a poor choice of control group could render our diff-in-diff estimator useless. Luckily, such deficiencies are absent from our groupings, as was illustrated in Table 2 in the Data Section. Second, since no convention exists regarding the measurement of daily stock price volatility, the choice of volatility measure could noticeably impact results. For this reason, we choose to re-run our regressions using five different measures of volatility, again in the name of robustness. All five measures are log-normally distributed; accordingly, we always use the logarithmic form of each measure when conducting regressions. We now explain these volatility measures one at a time.

*Jones-Seguin volatility (daily basis)*: Following Jones and Seguin (1997), this is an interday measure of volatility based on returns calculated using daily closing prices. It is an unbiased estimator of standard deviation if returns are normally distributed, and is defined as:

$$\hat{\sigma} = (\pi/2)^{1/2}|r_t| \qquad \textit{Eq. 3}$$

where $|r_t|$ is the daily stock return expressed in absolute terms. We opt to include this measure for two reasons. First, it has been used in a number of previous FTT studies (e.g. Jones & Seguin (1997); Liu & Zhu (2009)) and thus facilitates comparisons. Secondly, it relies on daily data, which effectively utilizes more unique inputs than weekly or monthly measures in any given timeframe.

*Yang-Zhang volatility (weekly[2] basis):* Proposed by Yang & Zhang (2000), this is a weekly measure of volatility based on intraday stock data – specifically daily high, daily low, daily opening, and daily closing prices. The calculation of the estimator requires data inputs

[2] Since the FTT was introduced on August 1st, 2012, a Wednesday, we use different event dummies whenever looking at the weekly volatility measures. These dummies ( $dT$ in Equation 1 and $dT_1$ in Equation 2) are equal to zero the week before the event and one the week after, but designed so as to exclude the observation during the event week itself. December 1st, 2012, is luckily a Saturday (so no adjustment is necessary for ADRs).

from multiple time periods, which is why we calculate it on a weekly basis, i.e. using five days' worth of data (the reported standard deviations are converted into daily values, but each day of the same week then features the same standard deviation). The measure is the minimum-variance unbiased variance estimator, independent of both drift and interday jumps with respect to the stock price. It is defined as:

$$V = V_O + kV_C + (1-k)V_{RS} \qquad \text{Eq. 4}$$

where V is the variance estimator (see Appendix A for component definitions). Note that we report standard deviations and not variances, meaning that we take the square root of V. We choose to use this measure because it uses a different set of inputs than our other measures, but also because it is one of the most efficient estimators available.

*Yang-Zhang volatility (monthly basis)*: This is the same measure as the weekly basis Yang-Zhang measure, except calculated using monthly segments rather than weekly segments. The standard deviations are again converted into daily values, implying that each day of the same month has the same standard deviation.

*"Common" volatility (weekly basis)*: This is simply the "common" measure of standard deviation calculated using stock returns on a week-by-week basis, but converted into daily values and reported as such. It is similar to the weekly-input Yang-Zhang measure, in the sense that each day of the same week features the same standard deviation. For completeness, it is defined as:

$$\hat{\sigma} = \sqrt{\sum_{i=1}^{n} \frac{(r_i - \bar{r})^2}{n-1}} \qquad \text{Eq. 5}$$

where $r_i$ is the stock return on day *i*. We include this measure because it is the simplest measure available and also the most widely used and understood.

*"Common" volatility (monthly basis)*: This is the same measure as the weekly basis "common" measure, except instead of being calculated on a weekly basis it is calculated on a monthly basis. The standard deviations are converted into daily values, meaning there is only one unique standard deviation per month.

## RESULTS

In order to holistically appraise the effects of the FTT, we investigate both volume and volatility developments. As previously laid out, our initial tests review the FTT's effects on aggregate, after which we attempt to isolate the effect of the STT. Finally, a comparative test is conducted to determine if the difference in FTT and STT effects is statistically significant.

### FTT Testing

The FTT's aggregate effect is gauged by contrasting French and European equities. Regression results are presented in Table 3. The $\beta_1$ coefficient estimates indicate that FTT-subject stocks experience economically and statistically significant reductions in volume and volatility. While the logarithmic form for volume (which allows for normally distributed hypothesis testing) is difficult to interpret, the non-logarithmic form has a rather simple interpretation: on average, the FTT caused daily turnover volume to be 14.9% lower during the post-event period than during the event window. Even though the statistical significance of this figure is dubious (because of non-normal distribution), the logarithmic volume measurement attests to a similarly negative effect. Additionally, four out of five volatility measures indicate concomitant diff-in-diff decreases with statistical significance. The estimated effect varies between 4.1% and 10.2%, and is thusly also economically significant. The Jones-Seguin measure yields statistically insignificant results, but has a fairly strong t-value of -1.34.

**Figure 2: FTT Testing**

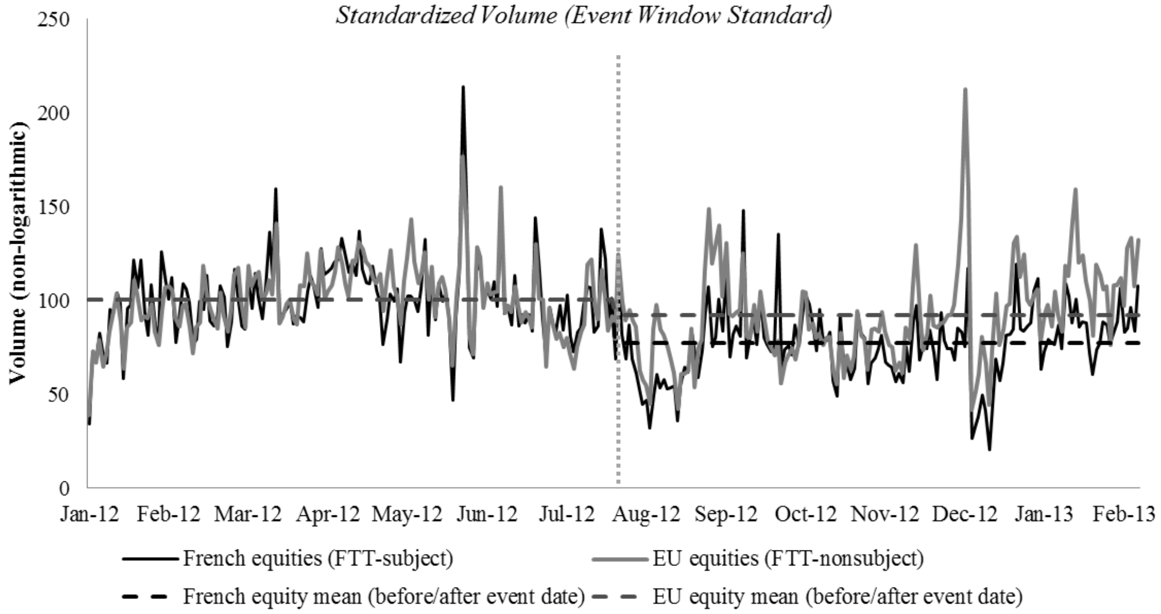

Figure 2 plots standardized non-logarithmic volume. The time series plot illustrates several points. Firstly, French "large" stocks (which account for the lion's share of the French market) experience a noticeable volume downturn during the month of August; average levels are 39% lower than during the event window. This development was widely portrayed by mainstream media sources as a consequence of the FTT's implementation. However, the one-month dip does not appear to be an exclusively French phenomenon; the European "large" stocks display similar decreases of approximately 24%. Furthermore, volume levels for both groups bounce back by mid-September. Thus, the volume reduction is unlikely to be an FTT-exclusive phenomenon. Instead, and secondly, the French treatment group appears to very closely track its selected Eurozone counterparts, both before and after implementation. The actual volume diff-in-diff stems from the fact that the French volume seems consistently lower during the post-event window than that of the control group. Thus, an FTT-spurred volume reduction, as the regressions indicate, seems more likely to manifest itself as an omnipresent and constant, albeit relatively small, effect, rather than a dramatic market shock.

**STT Testing**

Next, we attempt to isolate the STT effect from its HFTT counterpart, as explained in the Method section. We do this by employing data on ADRs, which should be unaffected by French high-frequency traders. We employ the same four countries as before to serve the function of control group. Regression results are displayed in Table 4.

**Figure 3: STT Testing**

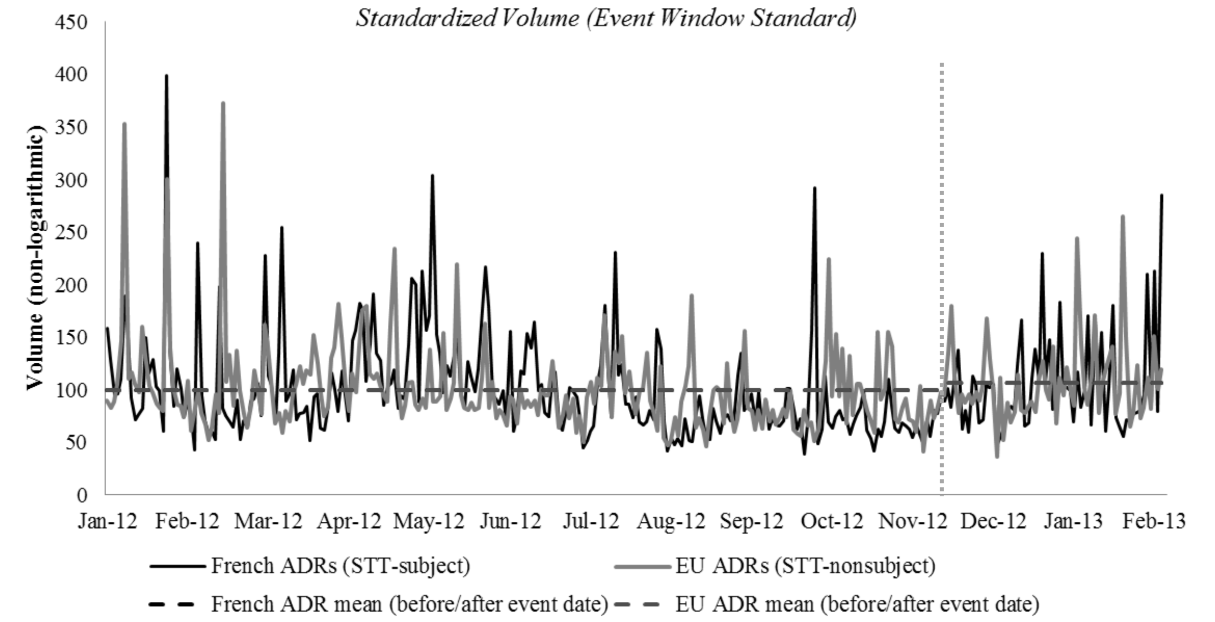

Unlike the Eurozone equities, the STT-subject ADRs show no sign of FTT influence. Volume changes are highly insignificant for both measures, rendering us unable to reject the null hypothesis of no diff-in-diff effect. The same applies to all five volatility measures, none of which produce statistically significant results. While some of the measures (such as volume) are too insignificant to gauge any sort of effect, others give some indication of a negative diff-in-diff effect. For example, the monthly "common" volatility has an economically significant effect and a t-value of -1.41. While STT testing faintly hints at a negative volatility diff-in-diff, the indicators are mostly indistinguishable from zero.

We now turn to Figure 3, which features the non-logarithmic standardized volume development over time. In accordance with the regression results, the plot shows no significant diff-in-diff effect. The graph also illustrates that the control group for this particular measure is of slightly lesser quality than its FTT testing counterpart. On the other hand, Table 2 in Data shows us that the treatment-control event window correlations for all other measures are very high; thus, the control group should still be considered reliable for volatility measures.

At this point, FTT testing has provided strong economically and statistically significant diff-in-diff effects. STT testing, on the other hand, has failed to produce any significant indicators of a diff-in-diff effect. We proceed with combined testing in order to determine whether this difference is statistically significant or not.

**Combined Testing**

The combined testing appends the data groups used in FTT testing and STT testing, in order to identify whether a statistically significant difference in the diff-in-diff effects of each sample exists. In this way, we can assess whether the effects of the FTT and STT are systematically different. However, the approach is not without its problems, as will be discussed later. Until then, results are presented in Table 5.

The two coefficients of interest are $\varepsilon_1$ and $\varepsilon_2$. These capture the effects of FTT testing and STT testing, respectively. In other words, $\varepsilon_1$ approximates the marginal effect experienced by French domestically-listed equities after August 1$^{st}$, 2012. The coefficient $\varepsilon_2$ estimates the marginal effect experienced by French ADRs after December 1$^{st}$, 2012. By presenting these estimates concomitantly, we can use an F-test to gauge the extent to which they differ. The results of such F-testing are available in Table 6.

Before we address the F-test results, let us discuss the regression results presented in Table 5. It should be immediately apparent that $\varepsilon_1$, which approximates the same effect as $\beta_1$

did during FTT testing (see Table 3), differs from this measure. While $\beta_1$ featured statistically significant results across the board, the $\varepsilon_1$ estimates are a far cry from statistical significance. Does this mean that the results from FTT testing are wrong? Probably not; the most likely explanation is that the combined testing features a diluted control group that yields less reliable results. More precisely, the French stocks that are FTT-subject are now compared to EU stocks, EU ADRs, and French ADRs. Furthermore, unpredictable effects that were previously ruled out, such as currency fluctuations and a potential size bias (as discussed in Data), will now have reentered the picture. The control group is now less suitable for its purpose, since its characteristics have drifted from those of the treatment group. Accordingly, statistically significant results are less likely to be obtained.

With this in mind, we proceed to the results of the F-test. As the p-values indicate, the test universally fails to reject the null hypothesis that $\varepsilon_1$ and $\varepsilon_2$ are equivalent. This implies that we cannot identify any systematic difference between the effects of the FTT and STT. While control group dilution casts some doubts on this conclusion, the result still stands: the FTT and STT diff-in-diffs are not statistically distinguishable.

TABLE 3: FTT TESTING RESULTS

*Fixed effects regression with clustered robust standard errors, two-tailed hypothesis tests*

| Measure | $\alpha_0$ | $\alpha_2$ | $\beta_1$ | Observations | Stock count | $R^2$ |
|---|---|---|---|---|---|---|
| Volume (non-logarithmic) | 100.3*** | -8.202 | -14.86** | 71,317 | 240 | 0.006 |
| *Event window standard* | (1.638) | (5.631) | (6.038) | | | |
| Volume (logarithmic) | 4.476*** | -0.228*** | -0.102** | 70,791 | 240 | 0.057 |
| *Event window standard* | (0.0106) | (0.0324) | (0.0408) | | | |
| Volatility (daily basis) | -4.376*** | -0.245*** | -0.0429 | 69,999 | 240 | 0.014 |
| *Jones-Seguin* | (0.00789) | (0.0212) | (0.0319) | | | |
| Volatility (weekly basis) | -3.994*** | -0.219*** | -0.0412* | 71,754 | 240 | 0.105 |
| *Yang-Zhang* | (0.00623) | (0.0177) | (0.0248) | | | |
| Volatility (monthly basis) | -3.938*** | -0.178*** | -0.0538** | 72,960 | 240 | 0.135 |
| *Yang-Zhang* | (0.00607) | (0.0173) | (0.0240) | | | |
| Volatility (weekly basis) | -4.087*** | -0.292*** | -0.0635** | 71,747 | 240 | 0.083 |
| *"Common"* | (0.00773) | (0.0209) | (0.0312) | | | |
| Volatility (monthly basis) | -3.953*** | -0.211*** | -0.0875*** | 72,960 | 240 | 0.135 |
| *"Common"* | (0.00757) | (0.0207) | (0.0303) | | | |

Each dependent variable of interest is regressed on independent variables as designated by Equation 1. The interaction dummy coefficient, which indicates the FTT diff-in-diff, is labeled $\beta_1$. Note that all measures except the non-logarithmic volume are normally distributed; the statistical significance of this particular indicator should be treated with some scepticism. Also note that $\alpha_1$ is omitted from the table above, as the French dummy variable is time-invariant.

**Indicates p-value of less than 10%*

***Indicates p-value of less than 5%*

****Indicates p-value of less than 1%*

TABLE 4: STT TESTING RESULTS

*Fixed effects regression with clustered robust standard errors, two-tailed hypothesis tests*

| Measure | $\alpha_0$ | $\alpha_2$ | $\beta_1$ | Observations | Stock count | $R^2$ |
|---|---|---|---|---|---|---|
| Volume (non-logarithmic) | 100.0*** | 6.880 | 8.657 | 24,050 | 84 | 0.000 |
| *Event window standard* | (2.100) | (8.772) | (21.21) | | | |
| Volume (logarithmic) | 4.050*** | 0.0384 | -0.0153 | 24,050 | 84 | 0.000 |
| *Event window standard* | (0.0120) | (0.0683) | (0.119) | | | |
| Volatility (daily basis) | -4.228*** | -0.175*** | -0.0405 | 23,464 | 84 | 0.006 |
| *Jones-Seguin* | (0.00464) | (0.0326) | (0.0444) | | | |
| Volatility (weekly basis) | -3.745*** | -0.219*** | -0.00551 | 25,393 | 84 | 0.049 |
| *Yang-Zhang* | (0.00415) | (0.0311) | (0.0384) | | | |
| Volatility (monthly basis) | -3.640*** | -0.209*** | -0.0446 | 25,452 | 84 | 0.132 |
| *Yang-Zhang* | (0.00491) | (0.0381) | (0.0451) | | | |
| Volatility (weekly basis) | -4.017*** | -0.192*** | -0.0312 | 25,449 | 84 | 0.026 |
| *"Common"* | (0.00411) | (0.0288) | (0.0385) | | | |
| Volatility (monthly basis) | -3.856*** | -0.195*** | -0.0603 | 25,452 | 84 | 0.099 |
| *"Common"* | (0.00464) | (0.0347) | (0.0429) | | | |

The regressions above are performed in the same manner as those of the FTT testing. The difference in results stems from the use of a wholly different treatment-control combination.

**Indicates p-value of less than 10%*

***Indicates p-value of less than 5%*

****Indicates p-value of less than 1%*

TABLE 5: COMBINED TESTING RESULTS

*Fixed effects regression with clustered robust standard errors, two-tailed hypothesis tests*

| Measure | $\delta_0$ | $\delta_4$ | $\delta_5$ | $\delta_6$ | $\delta_7$ | $\delta_8$ | $\delta_9$ | $\varepsilon_1$ | $\varepsilon_2$ | Observations | Stock count | $R^2$ |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Volume (non-logarithmic) | 100.3*** | -17.01*** | 21.60*** | -12.02 | -16.05*** | 5.407 | -5.003 | 3.691 | 45.78 | 95,371 | 324 | 0.003 |
| *Event window standard* | (1.718) | (5.624) | (4.547) | (12.29) | (5.574) | (8.969) | (9.091) | (13.71) | (31.85) | | | |
| Volume (logarithmic) | 4.371*** | -0.308*** | 0.197*** | -0.0288 | -0.163*** | 0.0739 | -0.00932 | -0.00643 | 0.165 | 94,845 | 324 | 0.041 |
| *Event window standard* | (0.0102) | (0.0338) | (0.0282) | (0.0897) | (0.0438) | (0.0645) | (0.0645) | (0.0992) | (0.103) | | | |
| Volatility (daily basis) | -4.318*** | -0.204*** | -0.101*** | -0.0460 | 0.00359 | 0.0183 | 0.0406 | 0.00192 | -0.0159 | 93,545 | 324 | 0.015 |
| *Jones-Seguin* | (0.00634) | (0.0212) | (0.0217) | (0.0381) | (0.0354) | (0.0366) | (0.0396) | (0.0509) | (0.0577) | | | |
| Volatility (weekly basis) | -3.910*** | -0.183*** | -0.0825*** | -0.00993 | 0.0309 | -0.0313 | 0.00277 | -0.0446 | -0.0290 | 96,811 | 324 | 0.107 |
| *Yang-Zhang* | (0.00521) | (0.0173) | (0.0207) | (0.0382) | (0.0278) | (0.0330) | (0.0351) | (0.0466) | (0.0463) | | | |
| Volatility (monthly basis) | -3.845*** | -0.149*** | -0.0698*** | -0.0221 | 0.0140 | -0.00507 | -0.0437 | -0.0377 | -0.0445 | 98,496 | 324 | 0.157 |
| *Yang-Zhang* | (0.00499) | (0.0169) | (0.0220) | (0.0299) | (0.0298) | (0.0278) | (0.0423) | (0.0391) | (0.0528) | | | |
| Volatility (weekly basis) | -4.043*** | -0.243*** | -0.114*** | -0.0294 | 0.0394 | -0.0208 | 0.0924*** | -0.0510 | -0.0552 | 96,860 | 324 | 0.086 |
| *"Common"* | (0.00615) | (0.0209) | (0.0237) | (0.0371) | (0.0358) | (0.0327) | (0.0353) | (0.0499) | (0.0535) | | | |
| Volatility (monthly basis) | -3.910*** | -0.154*** | -0.136*** | -0.0287 | 0.00632 | -0.00674 | 0.0396 | -0.0615 | -0.0486 | 98,496 | 324 | 0.156 |
| *"Common"* | (0.00594) | (0.0198) | (0.0235) | (0.0293) | (0.0353) | (0.0280) | (0.0418) | (0.0433) | (0.0567) | | | |

Each of our seven regressands is regressed on the set of regressors (including dummy variables and interaction variables) outlined in Equation 2. Our two coefficients of interest are $\varepsilon_1$ and $\varepsilon_2$. The former encapsulates the diff-in-diff effect first inspected during FTT testing, i.e. the marginal effect for measures pertaining to French non-ADRs after August 1st, 2012. The latter captures the diff-in-diff effect from STT testing, i.e. the marginal measurement effect for French ADRs after December 1st, 2012. Three coefficients ($\delta_1$, $\delta_2$, and $\delta_3$) are omitted, as they correspond to time-invariant independent variables.

**Indicates p-value of less than 10%*

***Indicates p-value of less than 5%*

****Indicates p-value of less than 1%*

TABLE 6: F-TEST RESULTS

*F-test of the coefficients $\varepsilon_1$ and $\varepsilon_2$*

| Measure | F(2, 323) | Prob > F |
|---|---|---|
| Volume (non-logarithmic)<br>*Event window standard* | 1.22 | 0.2953 |
| Volume (logarithmic)<br>*Event window standard* | 1.29 | 0.2756 |
| Volatility (daily basis)<br>*Jones-Seguin* | 0.05 | 0.9541 |
| Volatility (weekly basis)<br>*Yang-Zhang* | 0.51 | 0.5982 |
| Volatility (monthly basis)<br>*Yang-Zhang* | 0.67 | 0.5108 |
| Volatility (weekly basis)<br>*"Common"* | 0.77 | 0.4640 |
| Volatility (monthly basis)<br>*"Common"* | 1.11 | 0.3323 |

For each of the regression in Table 5, an F-test is conducted on the coefficient estimates of $\varepsilon_1$ and $\varepsilon_2$. The significance level required to reject the null hypothesis ($H_0$: $\varepsilon_1 = \varepsilon_2$) is provided in the right-hand column.

**ANALYSIS**

Before we proceed with our analysis, let us first summarize our results. Our first set of regressions, which compared tax-subject French firms with international firms of equivalent size, seems to indicate that the FTT, when taken as a whole, has a contractionary effect on both volume *and* volatility. In order to separate out STT and HFTT effects, we enlisted the help of ADRs. As French high-frequency traders were unlikely to have any sort of active relationship with ADRs, we reasoned that any volume or volatility effects would stem from the STT. From these regressions, we found signs of STT effects for neither volume *nor* volatility. However, when a combined test was performed to determine if the difference in results was statistically significant, we found that no significant difference could be proved.

How do we interpret these results? We draw a couple of conclusions. Firstly, it seems that an FTT, engineered in the French style, results in volume decreases and volatility decreases among tax-subject stocks. Secondly, as we are unable to disprove the presence of STT effects, we deduce that one or both of the STT and HFTT contribute to the overall volume and volatility effects.

The first conclusion supports the arguments of Stiglitz and Summers & Summers, as presented earlier in this paper. It confirms the hypothesis that a tax on transactions reduces the volume of trade, as well as reducing the volatility of stock.

The second conclusion can be interpreted in several ways, depending on how one interprets the role of each sub-tax. Since we know that volatility decreases from FTT implementation, at least one of the sub-taxes *has* to be exerting downwards pressure. We would argue that the effect is most likely attributable to the HFTT. This builds on the assumption that quote stuffing, the high-frequency trading behavior targeted by the HFTT, qualifies as a sort of "speculation". While some may balk at such a notion, is it really so unreasonable to infer that a technique *deliberately* used to misinform the market might contribute to greater market uncertainty, and thus larger price fluctuations? We do not think so.

Of course, it could be argued that the HFTT increases volatility. But that would mean that the STT decreases volatility, which is a much more contentious claim. In fact, the role of the STT is much more uncertain, and connects back to the heart of the debate. If the STT were shown to mollify the HFTT's negative influence, FTT opponents would be right; informative trading activity would have been suppressed. If, on the other hand, the STT were shown to also exert downwards pressure, FTT proponents would be right; some of the traders affected by the STT would qualify as "speculators". Unfortunately, we have been unable to provide concrete insight into the effects of this STT. The only possible contribution on STTs that might be perceived from this paper is that the STT effect, at low tax percentages, is unlikely to be dramatic (as our STT testing indicates). However, this is hardly an original claim; the general inconclusiveness of the empirical evidence on the issue attests to the difficulty in confidently identifying STT effects. Furthermore, the STT analyzed in this case is of a very particular design. Thus, it would hardly be fitting to see these indicators as befitting all STTs. These peculiarities are also manifested in a number of exemptions, the presence of which could explain the overall absence of any significant effects for this particular STT.

These inferences lead us to two concluding suggestions. First, additional research will be necessary to confidently identify the effect of FTTs, to the extent that they can be generalized. Second, care has to be taken to not regard all FTTs as identical; the formulation of each tax and sub-tax has an important impact on the overall effect. This, of course, amplifies the need for more empirical research. These notions form the crux of our concluding thoughts, which follow.

## CONCLUDING REMARKS

Our study set out to diagnose the market consequences of the implementation of a financial transaction tax (FTT). To accomplish this, we used data on volume and volatility for securities subject to a recent French FTT and comparable European securities that were not. Furthermore, we attempted to isolate the effects of the FTT's two components: a securities transaction tax (STT) and a high-frequency trading tax (HFTT). This attempt used ADR-equivalents to the subjects of our initial tests. Overall, we found that the FTT as a whole had a significant negative effect on both volume and volatility. No such effect was found when the HFTT was absent in practice. However, a third set of tests indicated that the difference in results was not statistically significant; thus, we were unable to conclude whether the negative relationship between FTT implementation and market volume/volatility was attributable to the HFTT, STT, or both.

Two implications of this study are generally relevant. First, it is difficult, but essential, to dedicate particular attention to the specificities of each FTT. An FTT is no more than the sum of its parts, and these can have markedly different effects depending on their targets (such as high-frequency trading, securities transactions, or derivatives transactions). Moreover, even though any two sub-FTTs may address the same target, their design may differ significantly. Such divergences have a noticeable influence on the effect of the FTT as a whole. The proverbial devil is in the details.

As specificity is so crucial, the need for a breadth of empirical evidence becomes all the more important. This motivates our second implication: more empirical studies are required to catalogue FTTs and sub-FTTs. Luckily (for researchers, at least), the immediate future promises a variety of FTT-related natural experiments. Most recent is the STT that was implemented on Italian markets on March 1st, 2013, soon to be bolstered by an HFTT and a derivatives tax on July 1st, 2013. Further down the road, several EU governments are considering an FTT on equities, bonds, and derivatives, as proposed by the European Commission. The tax is intended to take effect in January 2014, and covers eleven member states (including Belgium, France, Germany, Italy, Portugal, and Spain) that have agreed to go ahead with the tax using the so-called "enhanced cooperation" procedure. Consequently, the tax will be cooperatively implementable in each country without the approval of other EU members. Together, these nations account for two thirds of the EU's GDP.

The results of our study should not be seen as a reliable augur of what the wider European FTT experience would be like; the differences in design are vast. However, this should not be seen as a bad thing. Such differences are, in fact, precisely the prerequisites for furthering our understanding of FTT design. For example, the EU FTT lacks an HFT component. Even if we limit ourselves to the STT component, the French tax stands out as being very cautious. The EU tax, for example, has fewer exemptions and, perhaps most importantly, is levied on all intraday trade orders in conjunction with their execution (recall that the French STT only affects *acquisitions* leading to a *net change in ownership*, evaluated over the course of an entire day). Furthermore, the tax will require adopters to abandon all pre-existing FTTs. Thus, researchers will be able to compare the new FTT's influence to market environments with and without old (and different) FTTs. For such reasons, the new European tax should provide an excellent opportunity for a comprehensive FTT study.

On the whole, then, more empirical studies are necessary to map FTTs. The need for such taxonomy is a testament in and of itself to the intricacies involved in FTT design. Hopefully, the reemergence of the FTT into public policy debate will galvanize interest in such research. It is our hope that this paper can provide some meager contribution to FTT taxonomy. In this way, our own ambitions mirror those of the French policymakers looking to contribute to the blueprint for European FTTs. With a sufficient number of such alignments, researchers and policymakers can work towards refining this blueprint.

## REFERENCES


Farmer J.D., Gillemot, L., Lillo, F., Szabolcs, M., and Sen, A. (2004). "What Really Causes Large Price Changes?" *Quantitative Finance*, 4(4), pp. 383-397.

French, K. and Roll, R. (1986). "Stock Return Variances: The Arrival of Information and the Reaction of Traders." *Journal of Financial Economics*, 17(1), pp. 5-26.

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

Keynes, J.M. (1936). The General Theory of Employment, Interest and Money. London: Macmillan.

Liu, S. and Zhu, Z. (2009). "Transaction Costs and Price Volatility: New Evidence from the Tokyo Stock Exchange." *Journal of Financial Services Research*, 36(1), pp. 65-83.

Schulmeister, S., Schratzenstaller, M., and Picek, O. (2008). "A General Financial Transaction Tax: Motives, Revenues, Feasibility and Effects." Österreichisches Institut für Wirtschaftsforschung (WIFO) working paper.

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

Stiglitz, J.E. (1989). "Using Tax Policy to Curb Speculative Short-Term Trading." *Journal of Financial Services Research*, 3(2-3), pp. 101-115.

Summers, L.H. and Summers, V.P. (1989). "When Financial Markets Work Too Well: A Cautious Case for a Securities Transactions Tax." *Journal of Financial Services Research*, 3, pp. 261-286.

Umlauf, S.R. (1993). "The Transaction Taxes and the Behaviour of the Swedish Stock Market." *Journal of Financial Economics*, 33(2), pp. 227-240.

Yang, D. and Zhang, Q. (2000). "Drift Independent Volatility Estimation Based on High, Low, Open, and Close Prices." *The Journal of Business*, 73(3), pp. 477-492.

## APPENDIX A: MODELLING CLARIFICATIONS AND JUSTIFICATIONS

### Calculating the Yang-Zhang Variance Estimator

The variance estimator is defined as:

$$V = V_O + kV_C + (1-k)V_{RS} \qquad \textit{Eq. 6}$$

where

$$V_O = \frac{1}{n-1}\sum_{i=1}^{n}(o_i - \bar{o})^2 \qquad \textit{Eq. 7}$$

$$V_C = \frac{1}{n-1}\sum_{i=1}^{n}(c_i - \bar{c})^2 \qquad \textit{Eq. 8}$$

$$V_{RS} = \frac{1}{n}\sum_{i=1}^{n}\left[u_i\left(u_i - c_i\right) + d_i\left(d_i - c_i\right)\right] \qquad \textit{Eq. 9}$$

$$k = \frac{0.34}{1.34 + \frac{n+1}{n-1}} \qquad \textit{Eq. 10}$$

$$o_i = \ln O_i - \ln C_{i-1} \qquad \textit{Eq. 11}$$

$$u_i = \ln H_i - \ln O_i \qquad \textit{Eq. 12}$$

$$d_i = \ln L_i - \ln O_i \qquad \textit{Eq. 13}$$

$$c_i = \ln C_i - \ln O_i \qquad \textit{Eq. 14}$$

and O is the opening price, C the closing price, H the daily high price, and L the daily low price of the stock on a given day. Finally, n is the number of periods used in the estimation process. We choose to define a period as a day, so for our weekly measure n is generally 5. For monthly measures, n ranges from 28 to 31.

### International Control Group Selection

To form our international control group, we used a country-based correlation threshold. The 10-year correlation between the French CAC 40 index and the respective Eurozone national index was gauged against a 0.85 threshold. The time span from Dec 31st, 2001, to Dec 30th, 2011, was chosen for modernity, and does not feature any overlap with our event window. Major Eurozone trading partners were the primary candidates, and correlation results are featured in Table 7.

Note that the correlation with Germany, despite being a significant trading partner, is much lower than with the other countries. Since more countries would decrease the relevance of the control group (for economic significance), but fewer would yield less data for analysis (for statistical significance), we found the 0.85 mark to be a reasonable, though somewhat arbitrary, cut-off point. Since these countries are all Eurozone countries, we also avoid any complications due to currency effects.

**Cleaning Procedure Design**

The intent of the cleaning protocol is to filter out "bad" data, i.e. data that is infrequently updated. Accordingly, the number of 5-day periods during which no change occurs in closing, opening, highest, or lowest prices is tallied, including overlap (thus, four different counts are made for each date). If any count exceeds 20, the stock is dropped. For this reason, the process is referred to as a "5-20" cleaning protocol. Altering these two criteria has a negligible effect on our samples, as Table 8 shows. This would seem to indicate that the stocks that are flagged really are "bad", as even the most draconian of cleaning processes barely increases the number of stocks dropped (i.e. "good" stocks almost never get dropped).

It is impossible for us to know with certainty whether or not the data that are filtered out are inaccurately reported or accurately reported, but illiquid. However, the use of cleaning measures is appropriate in either case, as can be proved by ratiocination. If the data are inaccurate, then they should not be included in testing, as they will distort results. If the data are accurately reported, but feature price-stasis due to illiquidity, they should also be removed; the control and treatment groups should have as similar characteristics as possible (except the event effect, of course), including liquidity-based characteristics. As such, filtering out illiquid stocks will only make our testing groups more alike, and accordingly allow testing to produce better results. Furthermore, the effects of a tax on illiquid stocks are likely to be erratic, which would make accurate assessment more difficult.

TABLE 7: CONTROL GROUP SELECTION

*Import/export partnership with France, "main" index correlation*

| Characteristic | Belgium | Germany | Italy | Netherlands | Spain |
|---|---|---|---|---|---|
| Main index | BEL 20 | DAX | FTSE MIB | AEX Index | IBEX 35 |
| French export allocation | 7.4% | 16.7% | 8.3% | 4.3% | 7.4% |
| French import allocation | 11.3% | 19.1% | 7.7% | 7.5% | 6.6% |
| Index correlation | 0.980 | 0.709 | 0.901 | 0.937 | 0.865 |

This table features some of the candidates for the international control group. In order to qualify, a country had to qualify as a major French trading partner and Eurozone member. Once the initial candidates were identified, the correlation between each country's "main" stock index and the French CAC 40 was calculated based on the 10-year period preceeding the event window. The correlation threshold was set a 0.850. Thus, the control group ended up consisting of Belgium, Italy, the Netherlands, and Spain.

TABLE 8: CLEANING PROCEDURE VARIATION

*Stock count with varied data cleaning criteria*

| Category | Pre-clean count | "3-10" | "4-15" | "5-20" | "6-25" | "7-30" |
|---|---|---|---|---|---|---|
| *FTT testing (EU equities)* | | | | | | |
| Treatment | 107 | 100 | 105 | 106 | 106 | 106 |
| France | 107 | 100 | 105 | 106 | 106 | 106 |
| Control | 139 | 134 | 134 | 134 | 135 | 135 |
| Belgium | 22 | 21 | 21 | 21 | 21 | 21 |
| Italy | 43 | 42 | 42 | 42 | 42 | 42 |
| Netherlands | 32 | 30 | 30 | 30 | 30 | 30 |
| Spain | 42 | 40 | 41 | 41 | 42 | 42 |
| *STT testing (ADRs)* | | | | | | |
| Treatment | 60 | 37 | 39 | 39 | 39 | 39 |
| France | 60 | 37 | 39 | 39 | 39 | 39 |
| Control | 94 | 42 | 43 | 45 | 45 | 46 |
| Belgium | 12 | 4 | 4 | 4 | 4 | 4 |
| Italy | 29 | 10 | 10 | 11 | 11 | 11 |
| Netherlands | 26 | 19 | 20 | 21 | 21 | 21 |
| Spain | 27 | 9 | 9 | 9 | 9 | 10 |

This table features the results from a variety of data cleaning procedure variations. The first number designates the number of consecutive trading days with static prices used for each "bad" stock tally. The second number features the tally required for the stock to be dropped. Thus, the "3-10" procedure is the most stringent and the "7-30" procedure the most lenient. The "5-20" criteria are used by default throughout this paper.

## APPENDIX B: CASE CLARIFICATIONS

### Article 235 TER ZD

***Exemptions***

Some of the exemptions are fairly straightforward and require no clarification. Others are more nuanced and thus deserve some elaboration:

*Employee incentive plans:* Covered herein are not merely acquisitions of equity instruments made by employees through their employee savings plans, but also acquisitions by employee funds. To the extent that the shares acquired through a share buy-back are sold to the issuer's own employees through their employee savings plan, these too are exempt from the tax.

*Temporary transfers:* This exemption covers the lending of securities, sale and repurchase agreements (repos), as well as buy/resell and sell/buy transactions.

*Market-making activities:* Market-making activities are defined as either (a) the quoting of bid and ask prices and the provision of liquidity to the market on a regular and continual basis, (b) the execution of orders issued by clients, or (c) the hedging of positions with regard to (a) or (b). Hence market makers not only enjoy the common restriction of only having to pay taxes on the acquisitions leading to a change in ownership; they avoid the tax whenever their actions fall within the legal bounds of what is considered market-making activities.

### Article 235 TER ZD BIS

***HFT Definition***

For the purposes of the HFTT, high-frequency trading is defined as "the habitual addressing of orders for own account using an automated mechanism." To qualify, orders must be spaced at intervals smaller than 0.5 seconds.

***Tax Calculation***

The tax rate of 0.01% applies to the amount of orders cancelled or modified in excess of the 80% threshold. In other words, the number of securities that were either cancelled or modified in excess of the threshold constitutes the basis for the tax. The following example provided in the French tax code helps to illustrate how the tax works:

A trading desk identified as having engaged in high-frequency trading places initial orders to buy or sell 40,000 securities in a stock which had an average value of €45 during that trading day. Subsequently, 200 of these orders are modified and 35,000 are cancelled.

The modification/cancellation ratio thus becomes: ((200+35,000)/(200+40,000))*100 = 87.56%. Since this rate exceeds the 80% threshold, the tax is activated. The basis of assessment is calculated as: 35,200 [modified/cancelled orders] – (40,200 [initial and modified orders]*80%) = 35,200 – 32,160= 3,040*€45 [daily average stock price] = €136,800. The amount of tax due thus becomes €136,800 * 0.01% = €13.68.