The consequences of a Tobin Tax – An experimental analysis

Michael Hanke[†,*], Jürgen Huber[#,†], Michael Kirchler[†] and Matthias Sutter[‡]

[†] University of Innsbruck

[#] Yale University

[‡] University of Cologne

This version: 26th April 2006

Abstract

The idea of introducing a transactions tax on foreign exchange markets has floated around since the early 1970s when James Tobin first proposed it in order to cut back speculation and to reduce volatility. Recently, this so-called Tobin tax has become popular – mainly among politicians – as a possible source for tax revenues. The economic consequences of introducing a Tobin Tax are, however, completely unknown, as such a tax has not been introduced on any real foreign exchange market so far. We report an experimental study where a tax is introduced on a subset or all foreign exchange markets. The consequences are very clear: If introduced unilaterally, the Tobin tax causes a dramatic shift of trading volume to the untaxed market (thereby yielding almost no tax revenues), together with an increase in volatility and a decrease in efficiency in the taxed market. If the tax is introduced simultaneously on all markets, volatility is hardly affected, but trading volume and efficiency are still reduced substantially.

JEL-classification: E62; H2

Keywords: Tobin tax, Experiment, Volatility, Foreign exchange

[*] Corresponding author. E-mail: Michael.Hanke@uibk.ac.at

1 Introduction

In this paper we present an experimental test of the economic consequences of a Tobin Tax. Such a transactions tax on foreign exchange or stock markets has been first advocated by James Tobin in the early 1970ies and it has been controversial among economists and politicians since then (for a collection of articles on various aspects of this topic, see Haq et al., 1996). The actual implementation of a Tobin Tax on real-world markets would, of course, ultimately resolve the controversies over the alleged consequences of a Tobin Tax on market volume, price volatility, market efficiency, and tax revenues, to name but a few of the disputed issues. Yet, given that a Tobin Tax has not been implemented on any real foreign exchange market so far, we resort to examining its consequences in an economic experiment. In our experimental treatments we let subjects trade currencies on two distinct markets. Initially, there is no tax on either of these markets, but then a small transactions tax is either introduced in one of the two markets or in both. Finally, the tax is removed in order to study whether some of the effects of the tax persist even after its removal.

The experimental results provide unambiguous evidence of the consequences of a Tobin Tax. If introduced unilaterally, the tax causes a dramatic shift of trading volume to the untaxed market. Hence, the tax revenues are almost negligible. Volatility on the taxed market increases substantially, while market efficiency decreases markedly. The removal of a unilateral tax does not restore pre-tax levels, indicating that the effects of a Tobin Tax cannot be undone (completely) by lifting it after it has been introduced once. The simultaneous introduction of the tax on both markets reduces market volume and efficiency, but has hardly any effect on market volatility. Besides these aggregate effects we find through an analysis of individual trading patterns that a Tobin Tax reduces in particular the trading activity of high-frequency traders (who might be called the speculators on our experimental markets).

Although this was presumably one of the motivations for James Tobin's proposal (in order to benefit those who need to trade on foreign exchange markets for "bona fide" commercial reasons; Tobin, 1978; Eichengreen et al., 1995), the latter effect comes at the cost of increasing volatility if the tax does not cover all markets.

Our results should prove useful both for politicians and for economists. In recent years, the Tobin Tax has become popular among politicians as an instrument to fight speculation and stabilize foreign exchange markets. Although the tax revenues are often downplayed as "side-effects", it seems obvious that the fiscal benefits of the tax play an important role for its popularity. This has been explicitly acknowledged, for instance, by Wolfgang Schüssel, Austrian Chancellor and EU-president in the first half year of 2006. When taking over presidency in January 2006, he proposed the introduction of a Tobin Tax to provide a stable revenue basis for the EU budget. A Tobin Tax would have the advantage of neither relying on contributions from EU-member states' governments nor putting a burden on EU citizens directly. Our results suggest that such a tax cannot be considered a stable basis for tax revenues, especially when it is not introduced worldwide.

The academic debate about the consequences of a Tobin Tax has, somewhat surprisingly, only gained momentum from the 1990s onwards. So far, there is a series of theoretical and some empirical contributions that have not yet reached a consensus on the most likely effects of a Tobin Tax. Parts of the controversy are probably due to different approaches concerning the coverage of the tax, either uniformly across all markets or applying only to a subset of markets. The only experimental study that we are aware of (O'Hara et al. 2006) will be discussed in some detail below.

Assuming full coverage of the tax, Kupiec (1995) models the Tobin Tax as a special case of a Keynesian transactions tax. Relying partly on the empirical evidence concerning a transaction tax on stocks in Sweden (Umlauf, 1993), Kupiec concludes that a Tobin Tax would decrease informational efficiency and lead to lower liquidity. As far as volatility is

concerned, he does not arrive at a clear-cut conclusion, because a possible reduction in volatility might be wiped out by an increase in liquidity premia. An indirect examination of the possible consequences of a Tobin Tax is provided by Aliber et al. (2003). They consider the Tobin Tax as a particular type of transactions costs on currency markets. Therefore, they investigate in their empirical study the impact of the size of transactions costs on trading volume and volatility. Using an innovative approach to back out transactions costs implied by futures prices, they show that higher transactions costs are associated with higher volatility and lower trading volume on foreign exchange markets. Though in line with the main thrust of our results, Aliber et al. (2003) cannot provide evidence on the differential effects of introducing, respectively removing, a Tobin Tax on a subset of markets or all markets.

Palley (1999) presents a microeconomic model with two groups of risk-neutral traders (fundamentalists and noise traders). He shows that noise traders (speculators) lead to inefficiencies and higher costs for fundamentalists. Therefore, anything that reduces the volume of noise trading without harming fundamentalists would be considered positive. Palley (1999) then argues that although a Tobin Tax would hit fundamentalists and noise traders alike with respect to a single transaction, noise traders would be affected more heavily due to their higher trading frequency. As a consequence, noise trading would be reduced.

Contrary to Palley (1999) where the Tobin Tax is assumed to encompass all markets, there are also some papers which consider the existence of (Tobin-)tax havens. Though based on slightly different models, Mannaro et al. (2005) and Westerhoff and Dieci (2006) consider two markets where traders can choose on which market to trade and where a Tobin Tax is either implemented in both markets or in just one of them. Both the simulation approach of Mannaro et al. (2005) as well as the analysis in Westerhoff and Dieci (2006) show that introducing the tax on only one market leads to a strong decrease in trading volume and higher volatility in the taxed market. Westerhoff and Dieci (2006) also stress that the interplay between liquidity and volatility (via the price impact of orders) is difficult to assess in

practice. Our experiment supports the general results of Mannaro et al. (2005) and Westerhoff and Dieci (2006), as we will show that a one-sided introduction of a Tobin Tax leads, indeed, to a (partly dramatic) decline in trading volume and to an increase in the volatility of the taxed market. As such, our behavioural data cast doubt on Tobin's (1996) claim that the existence of tax havens would not lead to strong distortions between taxed and untaxed markets as long as the G7 and the big financial centers implemented the Tobin Tax synchronously and without exceptions.[1] Other valuable sources on the Tobin Tax include Amihud/Mendelson (1992), Dow and Rahi (2000), Habermeir/Kirilenko (2003), Stiglitz (1989), and Summers/Summers (1989).

In the only experimental study so far touching this topic, Bloomfield et al. (2006) explore the effects of the introduction of a Tobin Tax on different trader groups. They focus on noise traders and find that their market share is reduced by the introduction of the tax. However, they only have one market, so issues like tax havens and avoidance of the tax cannot be addressed.

The rest of our paper is organized as follows: In Section 2, we present our market model and the experimental design. Section 3 reports the experimental results. Section 4 concludes the paper by relating our findings to the previous literature and by discussing the practical implications of our results.

[1] In order to tackle the possible problems associated with the existence of tax havens, Kenen (1996) proposed two strategies to avoid a massive shift of transaction volume to the tax havens. One is to tax transfers of funds to or from such locations at penalty rates high enough to deter market participants to relocate their transactions, the other is to levy the tax in the country the deal is closed rather where the transaction occurs. Both strategies seem very difficult to implement in practice.

2 Market model and experimental design

2.1 Model description

There are two markets (called LEFT and RIGHT) on which two virtual currencies (Gulden and Taler) can be traded. Both markets are implemented as continuous double-auction markets with open order books. Traders can be active on both markets simultaneously. Apart from prohibiting short sales, there are no limitations to trading, meaning that traders are allowed to freely place limit and market orders. Buying a currency on one market and selling it on the other is possible, as is buying on both markets or selling on both markets. Orders have to include the number of Taler a participant wants to trade and the amount of Gulden (the "home currency") offered or asked for each Taler.

For the sake of simplicity we introduce a symmetric information structure and assume that traders know the exact future (fundamental) Gulden-value of the Taler. That means that we implement the fundamental value of the Taler (in Gulden) as a random walk without drift: $V_k = V_{k-1} + \varepsilon_k$, where V_k denotes the fundamental value in period *k*, and ε_k is a standard normal random variable. V_0 is set to 40.

2.2 Basic experimental design

We let groups of 20 subjects trade currencies on both markets, LEFT and RIGHT. Each trader is initially endowed with 200 Taler and 8,000 Gulden. That means that the total supply of Taler (in both markets combined) is fixed (20x200=4,000), while the amount of Gulden in the markets may rise and fall depending on the development of the price of the Taler. The experiment is divided into trading periods of 100 seconds each. In total there are 18 trading

periods[2], preceded by 5 trial periods in order to accustom subjects with the trading environment (the experimental instructions in the Appendix include several trading screens).

At the start of each trading period subjects are informed about the fundamental value of the Taler and the order books are empty. During a trading period subjects are continuously informed about open orders, their own holdings of both currencies and about their wealth. The latter is calculated as the sum of the Gulden holdings and the Taler holdings (number of Taler held multiplied by the current Taler price in Gulden).[3] Even if not involved, subjects also see the prices of all transactions in the current period.

When trading stops (after 100 seconds), subjects receive a summary of the trading activities of all previous periods in a "history screen". It contains, for each market, the closing price, total trading volume, (if applicable) the amount of taxes paid, and the trading volume of the subject on this market. Additionally, the current holdings of Taler and Gulden are displayed, as well as the subject's wealth.

2.3 Experimental treatments

Table 1 summarizes our experimental treatments. They differ with respect to when and on which market a (two-way) Tobin tax of 0.5% of the trading price is levied. A dash ("-") indicates that there is no tax. While traders are only informed that there will be 15 to 25 periods of trading, there are essentially three parts of the experiment (periods 1-6, 7-12, and 13-18). In the first part all treatments are identical. That means that all treatments start with the very same instructions and without any tax. We consider the absence of a tax as the most

[2] We told participants that each experiment would be randomly terminated between period 15 and 25. This was done to avoid strategic behaviour towards the end of the experiment.

[3] If the current Taler prices in the two markets deviate, the price with the higher volume is used to value the Taler.

realistic starting condition. Only after period 6 there is a change announced, however, it is not revealed that there will be another change after period 12.

The treatment abbreviations in Table 1 indicate the number of markets with a tax in each of the three parts. The numbers 0, 1, and 2 specify the number of taxed markets in the three parts of the experiment. 000 is our control treatment, which is used to correct for systematic changes in some of our measures used (details on this are provided in Section 3). The final letter (L or R) in the treatment notation indicates on which market (LEFT or RIGHT) the tax is levied if it only applies to one market, but not to both. For instance, in treatment 021R the tax is introduced on both markets from period 7 to period 12, but from period 13 on it is only sustained on the RIGHT market (and, hence, removed on the LEFT market).

From Table 1 it becomes clear that in the second part of the experiment the tax is either introduced on one market (010L, 010R, 012L, 012R) or on both markets (021L, 021R). In the third part the tax is either removed from one previously taxed market (010L, 010R, 021L, 021R) or it is introduced on one hitherto untaxed market (012L, 012R).

Table 1. Experimental treatments

	Periods 1-6		Periods 7-12		Periods 13-18	
Treatment	LEFT	RIGHT	LEFT	RIGHT	LEFT	RIGHT
000	-	-	-	-	-	-
010L	-	-	0.5%	-	-	-
010R	-	-	-	0.5%	-	-
012L	-	-	0.5%	-	0.5%	0.5%
012R	-	-	-	0.5%	0.5%	0.5%
021L	-	-	0.5%	0.5%	0.5%	-
021R	-	-	0.5%	0.5%	-	0.5%

2.4 Experimental implementation

For each treatment we ran two sessions á 20 subjects. The 240 participants were business students at the University of Innsbruck. Sessions were computerized (using zTree by Fischbacher, 1999) and lasted on average 75 minutes. Traders were paid according to their wealth, benchmarked by the average wealth of all traders, and earned on average 17 €.

3 Experimental results

3.1 Trading volume

We start our analysis by looking at the trading volume. In Figure 1 we present the average volume of Taler traded per period. We pool the L- and R-treatments in order not to overload the figure. Hence, "010" in Figure 1 captures the two sessions from 010L and the two sessions from 010R.

Figure 1. Trading volume in Taler

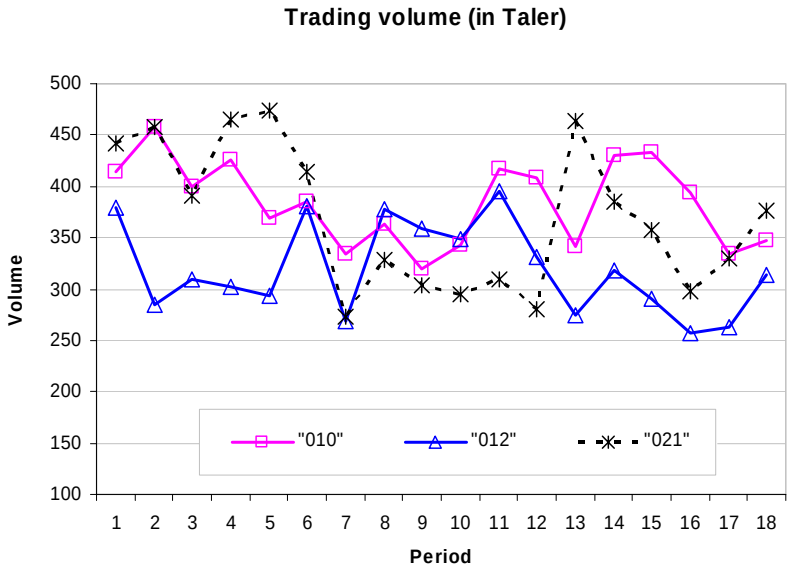

The overall trading volume in "010" does not seem to be much influenced by the introduction of a tax in periods 7-12. However, the detailed analysis of market shares and trading volumes on taxed, respectively untaxed, markets will reveal that the introduction of the tax has substantial – and immediate – consequences which are hardly detectable when looking at the aggregate trading volume.

When the tax is levied on both markets (see periods 7-12 in "021" and periods 13-18 in "012") the overall trading volume is clearly reduced in comparison to the levels prevailing before introducing the tax on both markets. Trading volume falls by 32% from periods 1-6 to periods 7-12 in "021" and by 17% from periods 7-12 to periods 13-18 in "012". This indicates that the trading volume is clearly negatively affected if the tax is encompassing, i.e. when there are no tax havens.

Table 2: Trading volume in Taler

The first value in each cell represents the average trading volume per period. Below we calculate the changes in volume per period to the prior part of the experiment in percent. The value in parentheses shows the *p*-value (one-sided) of a non-parametric Wilcoxon-Test of trading volume compared to the trading volume of the previous part of the experiment. In the last column we calculate the number of observations.

Treatment	1-6	7-12	13-18
	817.1	728.3	760.7
010		-10.9%	+4.4%
		(0.060)	(0.314)
	650.2	693.4	572.5
012		+6.6%	-17.4%
		(0.133)	** (0.005)
	880.9	596.8	737
021		-32.3%	+23.5%
		** (0.000)	* (0.007)
N	24	24	24

In Figures 2 to 4 we take a closer look at the effects of a Tobin Tax on trading volume in the taxed market. In particular, these figures show the market share of the LEFT market, i.e. the relation of trading volume on the LEFT market to the trading volume of both markets

together.[4] The first thing to notice from Figures 2 to 4 is the fact that in periods 1-6 the LEFT market has an average market share of about 68%, which implies that the trading volume on LEFT is more than double the one on RIGHT. Hence, trading has a strong bias in favour of the LEFT market (in fact the market which is on the left-hand-side of subjects' trading screens). This asymmetry in trading volumes allows us to consider the consequences of a Tobin Tax, depending upon whether it is introduced in a large market (LEFT) or in a small market (RIGHT).

Figure 2. Trading volume in 010L and 010R

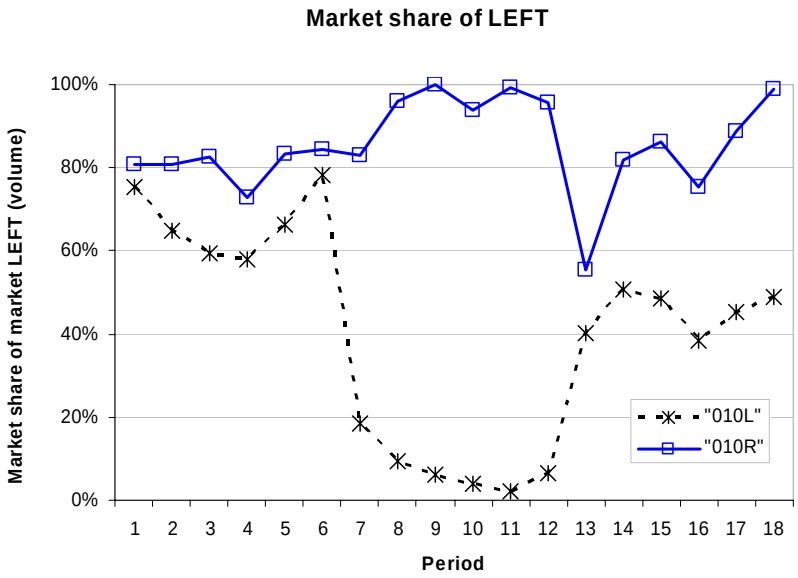

The introduction of the tax in 010L causes a drop in the market share of the LEFT market from almost 80% in period 6 to less than 20% in period 7. Hence, the shifts in trading volume as a consequence of taxation are very rapid and very strong. Whereas the average market share of LEFT in periods 1-6 was 68%, it drops to 8% in periods 7-12. Removing the tax after period 12 leads to an increase in market share (to about 47% in periods 13-18), but the pre-tax

[4] The market share for RIGHT is of course 1-(market share LEFT).

levels are not reached any more. Hence, a removal of the tax cannot fully undo the negative effects that have been caused when it is introduced in a large (LEFT) market.

Table 3: Market share of market LEFT

The first value in each cell represents the average market share of market LEFT per period. Below we calculate the changes in market share per period to the prior part of the experiment in percent. The value in parentheses shows the *p*-value (one-sided) of a non-parametric Wilcoxon-Test of market share compared to the market share of the previous part of the experiment. In the last column we calculate the number of observations.

Treatment	1-6	7-12	13-18
	67.8	8.2	47.3
010L		-87.9%	+476.8%
		** (0.001)	** (0.001)
	78.7	92.7	77.4
010R		+17.8%	-16.5%
		** (0.002)	* (0.007)
N	12	12	12

Considering treatment 010R we note that LEFT gains market share when the tax is levied on the RIGHT market from periods 7-12. The RIGHT market's share drops from 21.3% in periods 1-6 to 7.3% in periods 7-12. Note that in relative terms the loss in market share for RIGHT (loss of 66%) is less dramatic than when the tax is levied on the LEFT market (loss of 88%). Interestingly, removing the tax from the small RIGHT market brings back the market share almost exactly (22.6%) to pre-tax levels, which means that the harm on trading volume can be undone almost completely when a Tobin Tax is introduced and later abolished on a small market.

Figure 3. Trading volume in 012L and 012R

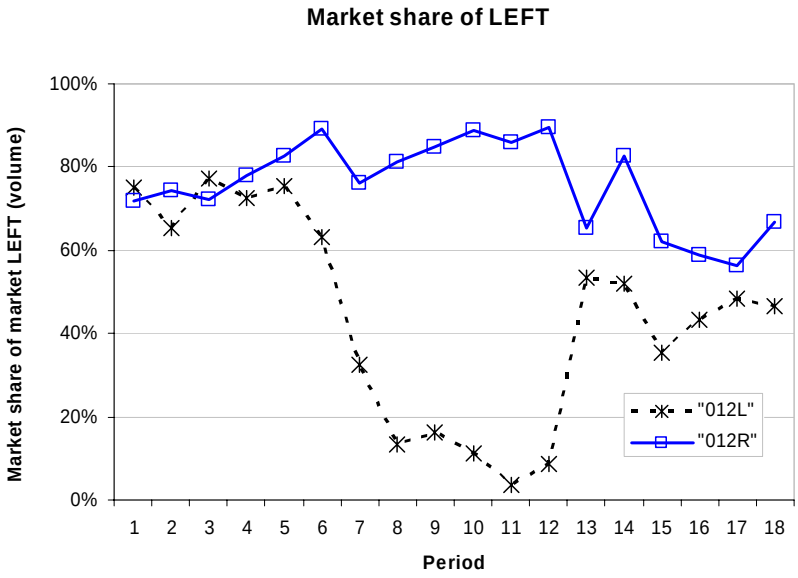

Figure 3 carries basically the same message as Figure 2. When the tax is introduced on the large LEFT market it leads to a very strong loss of market shares (of 82% when comparing periods 7-12 with periods 1-6). However, these losses can not be fully regained when the RIGHT market is also taxed from period 13 on.

Table 3: Market share of market LEFT

The first value in each cell represents the average market share of market LEFT per period. Below we calculate the changes in market share per period to the prior part of the experiment in percent. The value in parentheses shows the *p*-value (one-sided) of a non-parametric Wilcoxon-Test of market share compared to the market share of the previous part of the experiment. In the last column we calculate the number of observations.

Treatment	1-6	7-12	13-18
	72.9	13.1	44.6
012L		-82.0%	+240.5%
		** (0.001)	** (0.001)
	78.1	84.2	65.0
012R		+7.8%	-22.8%
		* (0.025)	** (0.003)
N	12	12	12

If the small RIGHT market is taxed first (in 012R) there are again only rather small effects of the tax. The RIGHT market loses only about 28% of its pre-tax market share when

comparing periods 7-12 to periods 1-6. These losses are more than regained when the tax is also introduced in the large LEFT market from period 13 on. Whereas RIGHT has a market share of 22% in periods 1-6, the share is 35% in periods 13-18. Even though this increase might seem surprising at first sight, it can be explained by the persistent observation that the large market is much stronger affected by the introduction of the Tobin Tax than the small one. Hence, a lot of trading activity swaps from LEFT to RIGHT in period 13.

Figure 4. Trading volume in 021L and 021R

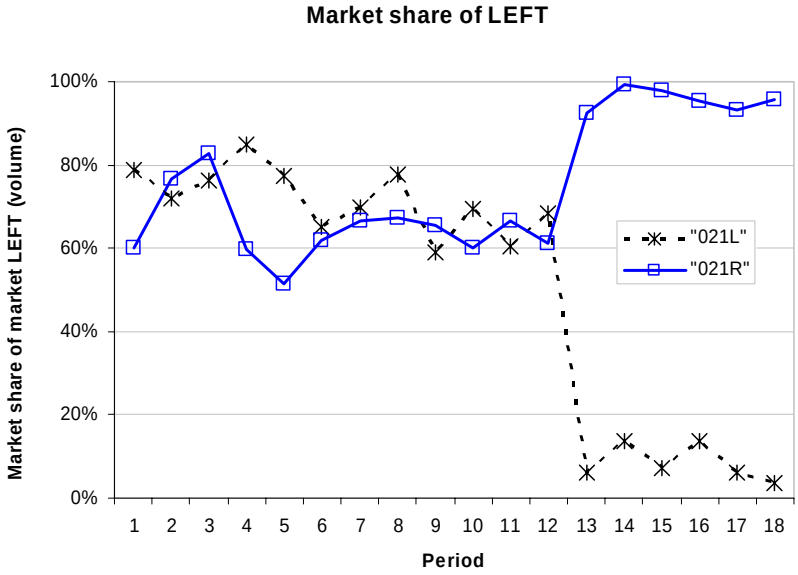

Finally, Figure 4 presents the two treatments where the Tobin Tax is levied, first, simultaneously on both markets, and only afterwards abolished on one of them. The introduction of the tax in period 7 shows no marked effect on the market share of both markets. But if the tax is abolished in one of the markets after period 12, there is a huge jump in market shares in favour of the untaxed market. For instance, when the tax remains valid in LEFT, but is abolished in RIGHT (021L), the market share of LEFT drops from 65% in period 12 to 6% in period 13.

Table 4: Market share of market LEFT

The first value in each cell represents the average market share of market LEFT per period. Below we calculate the changes in market share per period to the prior part of the experiment in percent. The value in parentheses shows the *p*-value (one-sided) of a non-parametric Wilcoxon-Test of market share compared to the market share of the previous part of the experiment. In the last column we calculate the number of observations.

Treatment	1-6	7-12	13-18
	75.7	67.6	8.6
021L		-10.7%	-87.3%
		* (0.042)	** (0.001)
	66.8	60.6	95.1
021R		-9.3%	+56.9%
		(0.120)	** (0.001)
N	12	12	12

3.2 Tax Revenues

The likely fiscal effects of a Tobin Tax have been discussed very controversially. Looking at the tax revenues in the different scenarios the evidence from our experimental markets implies that there is massive tax evasion if the tax is levied on one market only. Naïve estimates of the revenues would multiply total turnover by the tax rate (in our case, 1%, since both buyer and seller pay a rate of 0.5%). This ignores, however, the reduction in trading volume induced by the tax. This reduction is most pronounced when the tax is introduced unilaterally. The slanted figures in Table 5 show "hypothetical" tax returns calculated in this manner for all treatment/period combinations where no market is taxed. For combinations where both markets are taxed, upright figures show actual tax revenues paid by market participants in our experiments. Below those numbers, in parentheses, we show the ratio of tax revenues over total market turnover. This is always equal to one when both markets are taxed, since there is no possibility for tax evasion. When comparing actual tax revenues for these markets to hypothetical levels in previous periods (when no market was taxed), however, we see a marked decrease due to the effect of tax on turnover. For treatment/period combinations where only one market is taxed, the picture is very clear. First,

actual tax revenues are only a small fraction of hypothetical revenues. Second, due to tax evasion via trading in the untaxed market, the ratios of tax revenues to total market turnover are around 0.1%, only about one tenth of the value when both markets are taxed.

Table 5: Tax revenues in Taler

Treatment	Period 1-6	Period 7-12	Period 13-18
010L	*2845.7*	212.3	*3523.2*
		(0.07%)	
010R	*4820.8*	196.7	*3840.5*
		(0.05%)	
012L	*2553.2*	352.4	2439.7
		(0.13%)	(1.00%)
012R	*3774.6*	637.5	3250.1
		(0.16%)	(1.00%)
021L	*4077.3*	2518.5	266.9
		(1.00%)	(0.08%)
021R	*4362.3*	3253.3	167.2
		(1.00%)	(0.04%)

3.3 Frequency of transactions

Figures 5 through 7 show how the frequency of transactions is influenced by the introduction, respectively removal, of a Tobin Tax. When the tax is levied on the large LEFT market only in treatments 010L and 012L (see Figure 5), there is a very strong shift of transactions from the LEFT market to the RIGHT market from period 7 on. When the tax is either abolished (in 010L) or the RIGHT market is also taxed (in 012L) from period 13 on much of the trading volume activity floats back to the LEFT market, without reaching pre-tax levels, though.

Figure 5. Transactions per period when LEFT is taxed only

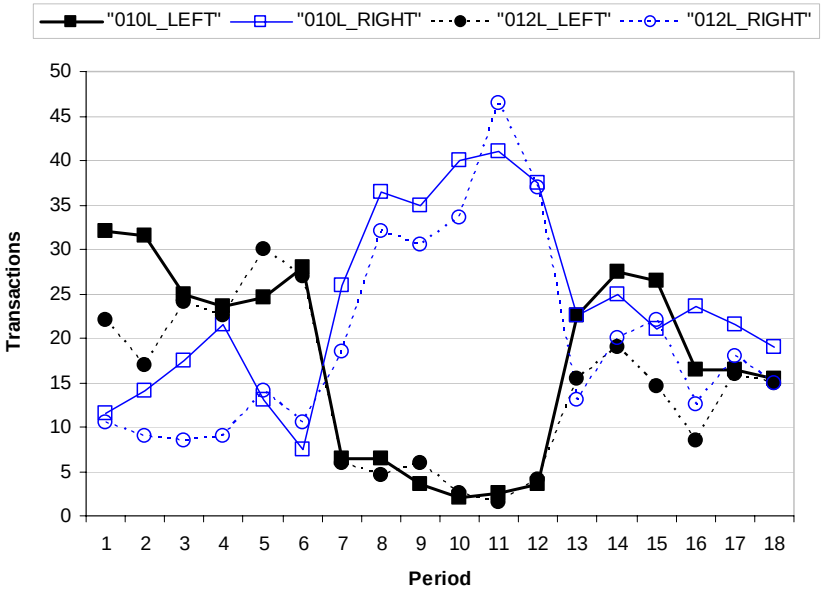

Figure 6 shows what happens when the RIGHT market is taxed solely. Like with the total trading volume, the number of transactions on RIGHT is less affected by the tax (from period 7-12) than if the tax is levied on LEFT (see the sharp kinks in Figure 5 which are hardly detectable in Figure 6). This is another confirmation that the market size (small or large) interacts with the introduction of a Tobin Tax.

Figure 6. Transactions per period when RIGHT is taxed only

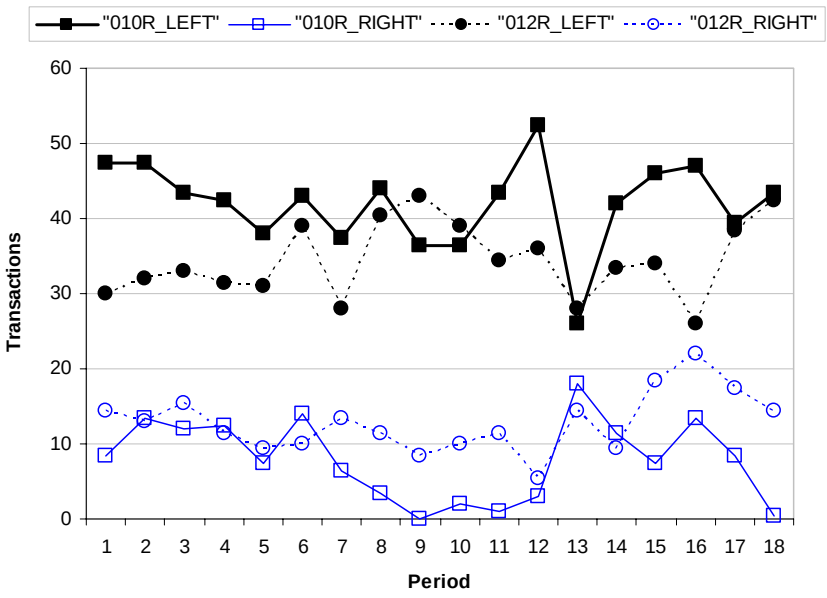

Figure 7 presents the number of transactions in treatments 021L and 021R. Abolishing the tax on one of the markets after period 12 leads to a strong increase in transactions in the untaxed market.

Figure 7. Transactions per period when tax is introduced in both markets at first

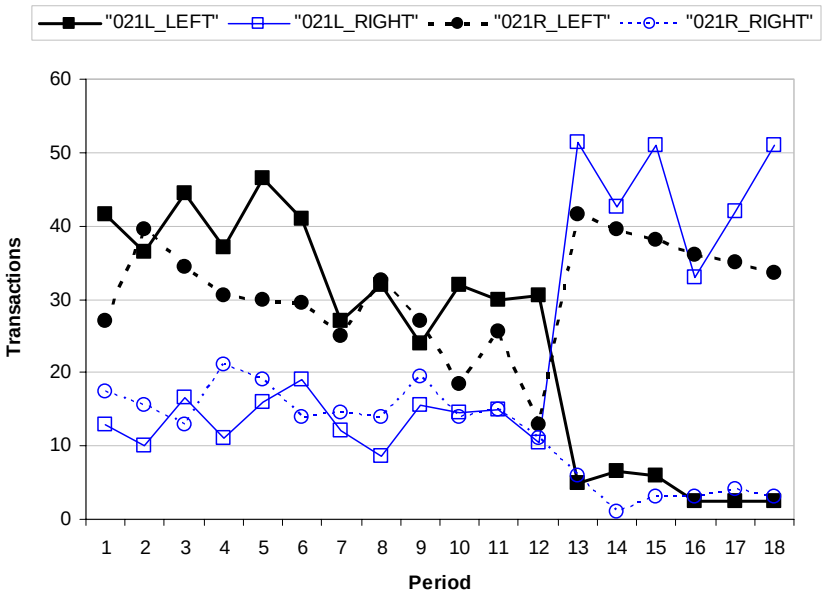

From Figures 5 through 7 we can calculate the number of transactions for the different cases. In untaxed markets, the average number of transactions per period (of 100 seconds) is 27.3. If both markets are taxed simultaneously, the number of transactions drops to 20.0, but if the tax is levied on only one market, there are only 4.7 transactions per period on this market.

The tax has also a similarly negative effect on traders' willingness to accept orders and, thus, make a transaction. 82% of orders lead to a transaction on untaxed markets. If both markets are taxed, it is only 65% of orders, and even less orders, 44%, turn into transactions on taxed markets when the other market is not taxed.

Table 6: Transactions/Limit Orders

In the first line we calculate the percentage of limit orders that yield a transaction. Below we conduct Mann-Whitney U-Tests between the different regimes (*p*-values are given in parentheses). As data points we compute the ratio (Transactions/Limit Orders) for each part of each experimental session. We can see that with exception of the comparison between 'double no tax regimes' and 'no tax regimes, when the other market is taxed' all ratios are significantly different from each other on a 1% level.

	Double no tax, [0T]	No tax, other market taxed [0T']	Single tax [1T]	Double tax [2T]
Mean	81.6	89.3	43.5	65.3
No tax, other market taxed [0T']	(0.368)			
Single tax [1T]	** (0.000)	** (0.000)		
Double tax [2T]	** (0.004)	** (0.000)	** (0.003)	

The volume of a transaction also differs with respect to the (non-)existence of a Tobin Tax. There are 17.1 Talers traded per transaction on untaxed markets. When both markets are taxed, only 14.7 Talers are traded per transaction. And in case one market is taxed, but not the other, only 13.1 Talers are traded per transaction on the taxed market.

Table 7: Volume per transaction

In the first line we calculate the average volume per transaction in the corresponding tax regimes. Below we conduct Mann-Whitney U-Tests between the different regimes (*p*-values are given in parentheses). As data points we compute the average volume per transaction for each part of each experimental session. We can see that in single taxed markets the average trading volume is significantly lower than in 'double no taxed markets' and in 'no taxed markets, while the other market is taxed'.

	Double no tax, [0T]	No tax, other market taxed [0T']	Single tax [1T]	Double tax [2T]
Mean	16.3	17.0	13.1	14.7
No tax, other market taxed [0T']	(0.306)			
Single tax [1T]	** (0.010)	* (0.012)		
Double tax [2T]	(0.060)	(0.080)	(0.151)	

3.4 Market volatility

As has been observed previously with continuous double auction markets (Sunder, 1995) the volatility of transaction prices, typically measured by the standard deviation of prices, is generally decreasing in time (as measured by the number of trading periods). Since we have introduced a Tobin Tax only in periods 7 or later, it is inadequate to compare the standard deviation of prices before and after introducing the tax. In order to determine the effects of the tax on standard deviations we compare the taxed market with the untaxed one. More precisely, we calculate the standard deviation of prices within a given part of the experiment (e.g. in periods 7-12) on the taxed market (e.g. the LEFT market in 010L) and compare it with the standard deviation of prices on the untaxed market (e.g. on the RIGHT market in 010L) during the same periods. There are 12 such matched observations (6 treatments á 2 sessions). Averaging across all observations, the standard deviation of prices is 4.04% in the taxed market, but only 2.55% in the untaxed one (p = 0.014; one-sided Wilcoxon signed ranks test; N = 12). Hence, the introduction of a Tobin Tax triggers significantly higher volatility, which is mainly due to the reduction in trading liquidity (as has been shown in the previous subsections).

As a benchmark we calculate the average absolute return of both control treatments for each third of the experiment.

$$\left|\overline{R}_{k,000}\right| = \frac{\sum_{t=1}^{T_k}\left|R_{t,k,000}\right|}{T_k}$$

t indicates tick t, k stands for one of the three parts of the experiment and T_k indicates the total number of ticks in the corresponding part of the experiment. Note that in this analysis we

calculate data from the left as well as from the right markets of the control treatments. In a next step we compare each single absolute return with the corresponding benchmark.

$$\left|R\right|_{t,k,benchmarked} = \frac{R_{t,k}}{\left|\overline{R}_{k,000}\right|}$$

In the following table we pooled all markets with no tax in the first part and a single tax in the second part of the experiment to compare the differences in volatility between the taxed and the untaxed regime within periods 7 to 12. Additionally, we want to shed some light on the question of changes in volatility from the untaxed first part of the experiment to a double taxed second part. Therefore we pooled all 02X-treatments. The first two values in the second line indicate the benchmarked average absolute return, $\left|\overline{R}\right|_{t,k,benchmarked}$, for the relevant part of the experiment. We can see that if one market is taxed and the other market is untaxed, the benchmarked average absolute return of the taxed market is around 155% higher than in the untaxed regime. A Mann-Whitney U-Test shows a significant difference at the 1% level (see the corresponding *n* in the third line of the following table). Things look different if we analyse the effect of the introduction of a tax in both markets in the second part of the experiment. We can see that the benchmarked absolute returns are not significantly different from each other.

Table 8: Market volatility

	01X – taxed [01T] vs. untaxed [01nT] market in period 7-12	02X – untaxed [02nT; period 1-6] vs. double taxed [02T; period 7-12]
	01T: 6.174	02nT: 3.195
	01nT: 2.423	02T: 3.391
	** (0.000)	(0.172)
N	01T: 247	02nT: 1205

3.5 Market efficiency and market making

As a measure of market efficiency we use the mean absolute error (MAE), i.e. the absolute difference between actual price and fundamental value of the currency.

$$MAE = \sum_{i=1}^{n} |P_i - V| * \frac{1}{n},$$

where P_i denotes the transaction price, V the fundamental value in the respective periods and n the number of transaction in the respective period.

We benchmark a market's MAE in a given part (e.g. periods 7-12, when a tax is introduced) on the same market's MAE in the previous part (e.g. in periods 1-6). Then we match again the taxed with the untaxed market and check which one has the smaller MAE. Considering the four treatments (010L, 010R, 012L, 012R) where only one market is taxed in periods 7-12, we find that the MAE is on average 1.48 in the taxed market, but only 1.08 in the untaxed market (p = 0.093; two-sided Wilcoxon signed ranks test; N = 8). Hence, the MAE rises in both markets even when a tax is introduced in only one of them, but the MAE rises significantly more strongly in the taxed market. This indicates that a Tobin Tax has a negative influence on market efficiency.

3.6 Trading Strategies / Types of Traders

One of the expected effects put forward by many proponents of a Tobin Tax is to discourage the activities of short-term speculators. In our setting we can define a measure for short-term trading as the number of times a trader switches from buying to selling. Note that there are no disincentives (other than the tax) to short-term speculation in our setting: Apart from the absence of (non-tax) transactions costs, all traders can easily and without any cost switch back and forth between LEFT and RIGHT. If the tax works as suggested, we should expect to see less switching between buy and sell transactions in the presence of the tax.

Table 9 compares four different situations: no tax in any market, no tax while the other market is taxed, tax while the other market is not taxed, tax in both markets.

Our proxy takes on its highest value for non-taxed markets, with even higher values when the other market is taxed. Interestingly enough, although the values appear to be smaller when both markets are taxed compared to the untaxed case, this difference is not statistically significant. The value is close to zero in taxed markets when there is an untaxed alternative.

Our results suggest that the tax does not induce short-term speculators to switch to a different strategy. Rather, the tax either drives short-term speculators to untaxed markets, or – if it is not possible to avoid the tax – it hardly changes speculators' behaviour.

Table 9: Short-term trading: Number of times a trader switches from buying to selling and vice versa on the market

In the first line we calculate the average number of times a trader switches from buying to selling and vice versa in the corresponding tax regimes. Below we conduct Mann-Whitney U-Tests between the different regimes (*p*-values are given in parentheses). As data points we compute the average number of switches from buying to selling for each part of the experiment. We can see that with exception of the comparison between 'double no tax regimes' and 'double tax regimes' all values are significantly different from each other on a 1% level.

	Double no tax, [0T]	No tax, other market taxed [0T']	Single tax [1T]	Double tax [2T]
Mean	5.41	8.44	0.75	4.08
No tax, other	** (0.003)			

market taxed [0T']			
Single tax [1T]	** (0.000)	** (0.000)	
Double tax [2T]	(0.125)	** (0.000)	** (0.000)

4 Conclusion

In this paper we have reported an experimental test of the consequences of a Tobin Tax. Though such a tax has gained widespread attention – both in academia as well as politics – there is practically no hard evidence on its consequences, because such a tax (on foreign exchange trading) has not been introduced on any real market. In this paper we show why the political inactivity to levy a Tobin Tax so far has to be considered fortunate.

In our experimental markets a Tobin Tax reduces trading activity (both in volume and frequency of transactions) and causes a large shift of trading from the taxed to the untaxed market (where this effect is relatively stronger when the larger market is taxed). As a result of this tax avoidance, tax revenues are very small if the tax is not levied on all markets simultaneously. The Tobin Tax also increases the volatility on the taxed markets, which runs counter to the hopes of the supporters of a Tobin Tax. Market efficiency is also lower on taxed markets than on untaxed ones, which seems to be caused by the fact that the tax reduces liquidity in a taxed market.

Our design also allowed investigating whether the unilateral introduction of a Tobin Tax has different effects from an encompassing introduction on all markets. The effects of the Tobin Tax are strongest if the tax is introduced in only one market, in particular when this market is the large (LEFT) one. Though Tobin (1996) thought that it might be sufficient to introduce a Tobin Tax on the large financial markets in the G7, our results clearly indicate that this might lead to huge distortions in market shares of different trading places. When the tax is introduced simultaneously on all markets, there are still negative effects on trading

volume and market efficiency, but these effects are relatively small compared to the ones when tax havens are available. To sum up, our results provide strong arguments against the introduction of a Tobin Tax, at least against its unilateral introduction. Should the introduction of such a tax be ultimately in the interest of politicians, they would be well-advised to take great care in introducing it simultaneously on all markets. The harm done on taxed markets cannot fully be undone, neither by lifting the tax again later on, nor by its subsequent introduction in other markets as well.

References

Aliber, Robert C., Bhagwan Chowdhry, and Shu Yan (2003), Some Evidence that a Tobin Tax on Foreign Exchange Transactions Might Increase Volatility, *European Finance Review*, Vol. 7, No. 3, 481-510.

Amihud, Y. and Mendelson, H. (1992), Transaction Taxes and Stock Values, in Modernizing U.S. Securities Regulation, Lehn, K. and Kamphuis, R. (eds.), Irwin, New York.

Bloomfield, Robert J., Maureen O'Hara, and Gideon Saar (2006), The Limits of Noise Trading: An Experimental Analysis, *Working Paper*.

Dow, J. and Rahi, R. (2000), Should Speculators be Taxed? *Journal of Business*, Vol. 73, 89-107.

Eichengreen, Barry, James Tobin, and Charles Wyplosz (1995), Two Cases for Sand in the Wheels of International Finance. *Economic Journal*, Vol. 105, 162-172.

Fischbacher, Urs (1999), Z-tree: Zurich toolbox for readymade economic experiments. University of Zurich, Institute for Empirical Research in Economics. Working Paper No. 21.

Garber, Peter and Mark P. Taylor (1995), Sand in the Wheels of Foreign Exchange Markets – a Skeptical Note, *Economic Journal*, Vol. 105, 173-180.

Habermeir, K. and Kirilenko, A. (2003), Securities Transaction Taxes and Financial Markets, in Debating the Tobin Tax, Weaver, J, Dodd, R, and Baker, J. (eds), New Rules for Global Finance Coalition, Washington D.C.

Haq, Mahbub ul, Inge Kaul, and Isabelle Grunberg (eds.), (1996) The Tobin Tax. Coping with Financial Volatility. Oxford University Press.

Kupiec, Paul H. (1995), A Securities Transactions Tax and Capital Market Efficiency, *Contemporary Economic Policy*, January 1995, 101-112.

Mannaro, Katiusca, Michele Marchesi, and Alessio Setzu (2005), The Impact of Transaction Taxes on Traders' Behaviour and Wealth: A Microsimulation. Working paper, University of Cagliari.

Palley, Thomas I. (1999), Speculation and Tobin Tax: Why Sand in the Wheels Can Increase Economic Efficiency, *Journal of Economics*, Vol. 69, 113-126.

Spahn, Paul B. (1996), The Tobin Tax and Exchange Rate Stability. *Finance and Development*, June 1996, 24-27.

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

Stotsky, Janet G. (1996), Why a Two-Tier Tobin Tax Won't Work. *Finance and Development,* June 1996, 28-29.

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

Sunder, Shyam (1995), Experimental Asset Markets: A Survey, Chapter 6 in Handbook of Experimental Economics, Roth, Al/Kagel, John (eds.), Princeton NJ, Princeton University Press (1995), 445-500.

Umlauf, Steven R. (1993), Transactions Taxes and the Behaviour of the Swedish Stock Market. *Journal of Financial Economics*, Vol. 33, 227-240.

Westerhoff, Frank H. and Roberto Dieci (2006), The effectiveness of Keynes–Tobin transaction taxes when heterogeneous agents can trade in different markets: A behavioral finance approach. Journal of Economic Dynamics and Control, Vol. 30 (2), February 2006, 293-322.