The Economic Consequences of a Tobin Tax – An Experimental Analysis*

Michael Hanke, Jürgen Huber, Michael Kirchler and Matthias Sutter†

University of Innsbruck

February 8, 2010

Abstract

The effects of a Tobin tax on foreign exchange markets have long been disputed. We present an experiment with currency trading on two markets, where either none, one, or both markets are taxed. Our results confirm the hitherto undisputed issues: a tax reduces trading volume, shifts market share to untaxed markets, and leads to negligible tax revenues if tax havens exist. Concerning the controversial issues we find that (i) volatility effects depend on the existence of tax havens and on market size, (ii) market efficiency decreases in taxed markets when tax havens exist, and (iii) short-term speculation is reduced.

JEL classification: C91, E62

Keywords: Tobin tax, Experiment, Foreign exchange, Market efficiency, Trading volume, Volatility

*We thank two referees, a co-editor and participants of the 2006 ESA World Meeting (Atlanta), the 2006 ESHIA conference (Bologna), and the 2007 ESA World Meeting (Rome) for helpful comments and discussions. Financial support from the Center of Experimental Economics (sponsored by *Raiffeisen-Landesbank Tirol*), and the Austrian Science Fund (grant P-20609). Sutter thanks the German Science Foundation (by the Leibniz-Award to Axel Ockenfels) for financial support.

†Kirchler is corresponding author. Department of Finance, University of Innsbruck, Universitaetsstr. 15, 6020 Innsbruck, Austria and University of Gothenburg, Center for Finance, 40530 Gothenburg, Sweden. Phone: +43 512 507 7587, Fax: +43 512 507 2846. E-mail: Michael.Kirchler@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 markets was advocated by James Tobin in the early 1970s, and it has been controversial among economists and politicians ever since.[1] Of course, the actual implementation of a Tobin tax on real-world foreign exchange markets would resolve the controversies over its alleged consequences on volatility, efficiency, and short-term speculation, to name but a few of the disputed issues. Since a Tobin tax has not been implemented on any real foreign exchange market so far, however, we use the method of experimental economics to assess the effects of a Tobin tax.

In the political debate, the Tobin tax has gained popularity as a candidate instrument to fight speculation and stabilize foreign exchange markets. Its intended effects (according to many of its proponents) include a decrease in volatility and an increase in market efficiency. These expected benefits of a Tobin tax have been the reason for the Canadian House of Commons to speak out for a Tobin tax in recent years and for several political proposals in the U.S. to introduce a securities transaction tax (Bloomfield et al., 2009). Although the tax revenues are often downplayed as "side-effects", expected fiscal benefits obviously also increase the political appeal of a Tobin tax. For instance, when taking over the EU-presidency in 2006, the Austrian Federal Chancellor Wolfgang Schüssel proposed the introduction of a Tobin tax to provide a stable revenue basis for the EU budget.

In the academic debate, the Tobin tax has often been linked to the more general issue of how a transaction tax affects financial markets (e.g., Stiglitz, 1989; Summers and Summers, 1989; Schwert and Seguin, 1993; Jones and Seguin, 1997; Subrahmanyam, 1998; Dow and Rahi, 2000). The economics literature has reached a consensus on several issues such as the negative effects of a Tobin tax on trading volume or market shares (see, e.g., the contributions in Haq et al., 1996; Weaver et al., 2003). However, some other issues are still disputed,

[1] For collections of articles on various aspects of a Tobin tax see Haq et al. (1996), Habermeier and Kirilenko (2003) or Weaver et al. (2003).

e.g. the impact of a Tobin tax on market efficiency and volatility.

Parts of the controversy concerning the latter issues are probably due to different modeling approaches concerning the coverage of the tax, either uniformly across all markets or applying only to a subset of markets. Assuming full coverage of the tax across all markets, Kupiec (1995) relies in his analysis partly on the empirical evidence concerning a transaction tax on stocks in Sweden (Umlauf, 1993). Kupiec then argues that a Tobin tax would increase mispricing, i.e. decrease informational efficiency, and lead to lower liquidity. The latter result is also established in a model with only one market by Subrahmanyam (1998). Palley (1999) presents a microeconomic model with two groups of risk-neutral traders (fundamentalists and noise traders). He shows that noise traders (speculators) cause inefficiencies and higher costs for fundamentalists. Therefore, anything that reduces the volume of noise trading without harming fundamentalists would be considered positive. Palley 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, a Tobin tax would reduce noise trading and, so he claims, increase market efficiency, contrary to the conclusions by Kupiec (1995). More recent models by Ehrenstein (2002) and Westerhoff (2003) also predict that a Tobin tax will increase informational efficiency by reducing the degree of mispricing (i.e., the difference between market prices and fundamental values). Cipriani and Guarino (2008) focus on the effects of a transaction tax on informational cascades, and hence market efficiency in incorporating information in market prices, in a laboratory financial market. While theory would predict the transaction tax to reduce informational efficiency, they find in the experiment no significant effect, which is due to less irrational behavior in the presence of transaction costs. Summing up the evidence on transaction taxes and market efficiency, it seems fair to say that the literature is still inconclusive.

Turning to the effects of a Tobin tax on price volatility we start by noting that Kupiec (1995) does not arrive at a clear-cut prediction for the influence of a Tobin tax on volatility, because a possible reduction in volatility might be wiped out by an increase in liquidity premia. Many other papers (see, e.g., Frankel, 1996; Westerhoff, 2003; Ehrenstein et al., 2005) expect a decrease in price volatility. However, an empirical study by Aliber et al. (2003) provides conflicting evidence. They consider the Tobin tax as a particular type of transactions costs on currency markets. Therefore, they investigate the impact of the size of transactions costs on trading volume and volatility. Using an innovative approach to derive transactions costs from futures prices, they show that higher transactions costs are associated with *higher* volatility and lower trading volume on foreign exchange markets.[2] Similar results are presented in Hau (2006). Hence, there is no general agreement on the consequences of a Tobin tax on price volatility, although two recent contributions may be able to resolve the contradictions. Haberer (2006) presents a model with a U-shaped relationship between volatility and market volume. The reduction of market volume due to the introduction of a Tobin tax can then have different consequences for volatility, depending on the relative market volume. Taxing relatively large markets may decrease volatility, whereas a tax on relatively small markets may increase volatility. This will be one of the key findings of our experiment. Pellizzari and Westerhoff (2009) have investigated in computer simulations the impact of market microstructure on the effects of a transaction tax on market volatility. They have found that different trading institutions – either a continuous double auction or a dealership market – yield different effects of a transaction tax on market volatility. While there is no significant effect in a continuous double auction (where the tax reduces liquidity), the introduction of a transaction tax reduces volatility in a dealership market (where abundant liquidity is provided by specialists and the tax crowds out speculative orders).

Interestingly, the implications of tax havens have only recently been explicitly modeled. Mannaro et al. (2008) and Westerhoff and Dieci (2006) analyze models with two markets where traders can choose on which market to trade and where a Tobin tax is either implemented on both markets or on just one of

[2]Werner (2003) raises the important question in which way causality is running, however. It could be from trading costs to volatility, or the other way round.

them, leaving the other market as a tax haven. Both papers show that introducing the tax on only one market leads to a strong decrease in trading volume on the taxed market. Whereas Mannaro et al. (2008) expect an increase in volatility on the taxed market, Westerhoff and Dieci (2006) claim that volatility decreases on the taxed market, but increases on the untaxed market. The latter paper stresses that the interplay between liquidity and volatility (via the price impact of orders) is difficult to assess in practice, so Westerhoff and Dieci (2006) explicitly call for an experimental analysis of the question.

Bloomfield et al. (2009) run a controlled laboratory experiment to study trading behavior on markets when a securities transaction tax (STT) is introduced. They are particularly interested in the effects of a STT on three different types of traders whom they call informed traders, liquidity traders, and noise traders. Their experimental results suggest that a STT leads to less noise trading, which then increases informational efficiency. Market volume is driven down by the tax, whereas market volatility is hardly affected. A limitation of the setting used in Bloomfield et al. (2009) is its restriction to only a single market. In such a setting, it is impossible to examine how one market is affected by a Tobin tax if there are other markets that remain untaxed, i.e. if there are tax havens.[3]

In our experiment we let subjects trade currencies on two distinct markets. Initially, there is no tax on any of these markets, but then a transactions tax is either introduced on one of the two markets or on both. In order to study whether some effects of the tax persist even after its abolishment – an aspect which has not been explored in the literature so far – we consider also a scenario where the tax is abolished again after its introduction.

We let 480 participants trade in a continuous double auction for 18 trading periods. We had 7 different treatments, defined by the sequence of taxing none, one or both markets and by the prevalent tax rate.

[3]A recent paper by Kaiser et al. (2007) considers also only a single market that is taxed, therefore also missing the opportunity to examine the effects of tax havens. Since they allow for several tax rates, but run only one market per tax rate, Kaiser et al. (2007) consider their findings as preliminary.

Our key findings for the case of a unilateral introduction of the tax are that (i) the tax causes a dramatic shift in trading volume to the untaxed market, (ii) tax revenues are negligible, (iii) volatility on the taxed market may decrease or increase, depending on market size, while (iv) volatility on the untaxed market is reduced significantly as a consequence of an increase in liquidity, and (v) market efficiency decreases in the taxed market.

If a Tobin tax is introduced simultaneously on both markets, we find that (i) overall trading volume is reduced, (ii) price volatility remains unchanged, and (iii) market efficiency remains unchanged as well.

Through an analysis of individual trading patterns we can examine the microfoundations of these aggregate effects on the market level. Taking two different measures of speculation, we find that a Tobin tax reduces speculative trading. Although this was presumably one of the motivations for James Tobin's proposal (Tobin, 1978; Eichengreen et al., 1995), it has to be stressed, though, that the effect of this reduction in short-term speculation on volatility can go in either direction. Hence, individual trading patterns may be a misleading indicator for aggregate market effects.

The rest of the paper is organized as follows: In Section 2, we present our market model and the experimental design. Sections 3 and 4 report the experimental results. Section 5 concludes the paper by relating our results to previous findings and by discussing the practical implications of our results.

2 Market model and experimental design

2.1 Model description

There are two markets (denoted LEFT and RIGHT) on which a foreign currency (Taler) can be traded for the home currency (Gulden). Both markets are implemented as continuous double-auction markets with open order books. Traders can be active on both markets simultaneously and both markets are displayed on the screen at the same time. 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.

Apart from prohibiting short sales, there are no limitations to trading, meaning that traders are allowed to freely place limit and market orders. Limit orders consist of the number of Talers a participant wants to trade and the amount of Gulden offered or asked for each Taler. Price is given priority over time for the execution of limit orders. Market orders are executed immediately. For the sake of simplicity we introduce a symmetric information structure where traders know the current fundamental Gulden value of the Taler, but of course not its future values. We implement the development of 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 Experimental design

We set up groups of 20 subjects each who can trade currencies on two markets, LEFT and RIGHT. Each trader is initially endowed with 200 Taler and 8,000 Gulden. The experiment consists of 18 trading periods, each of them lasting 100 seconds.[4] At the start of each trading period, order books are empty. Subjects receive information about the fundamental value of the Taler at the start of each period. During a trading period subjects are continuously informed about all open orders, their own holdings of both currencies, the transaction prices on both markets, their individual transaction prices, and their wealth. The latter is calculated as the sum of the Gulden holdings and the Gulden value of their Taler holdings (number of Talers held multiplied by the current Taler price in Gulden).[5] When a trading period 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, the total trading volume, the

[4]To avoid strategic behavior towards the end of the experiment, we told participants that the experiment would be terminated between periods 15 and 25. The 18 actual periods were preceded by 5 unpaid trial periods to accustom subjects with the trading environment. See the experimental instructions with some screenshots in the online Supplementary Material.

[5]If the current Taler prices on the two markets deviate, the price on the market with the higher volume of the last transaction is used to value the Taler.

amount of taxes paid (only if applicable), and the trading volume of the subject on the relevant market. Moreover, the current holdings of Taler and Gulden are displayed, as well as a 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 the transaction value (price multiplied with Talers traded) 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 phases in our experiment (periods 1-6, 7-12, and 13-18). In the first phase all treatments are identical, starting with identical instructions and without any tax. We consider the absence of a tax as the most realistic starting condition. Only after period 6 are subjects informed about the introduction of a tax in the respective treatments. This is announced by the experimenter and indicated on a separate screen. Yet, it is not revealed at this stage that the taxation of markets will change again after period 12. The same holds for the introduction or abolishment of a tax after period 12.

Insert Table 1 about here

The treatment abbreviations in Table 1 are to be read as follows. The numbers "0" and "2" specify whether no market ("0") or both markets ("2") are taxed. If one market is taxed, we use the letters "L" and "R" to indicate whether the taxed market is the LEFT or the RIGHT one. For instance, in treatment 02R the tax is introduced on both markets from period 7 to period 12, but from period 13 onwards it is sustained only on the RIGHT market, while it is abolished on the LEFT market.

Table 1 shows that in the second phase of the experiment the tax is either introduced on one market (0L0, 0R0, 0L2, 0R2) or on both markets (02L, 02R). In the third phase the tax is either abolished from one previously taxed market (0L0, 0R0, 02L, 02R), or it is introduced on one hitherto untaxed market (0L2, 0R2). This design allows us to study both the effects of applying as well as abolishing a Tobin tax, and to explore the effects from the existence of tax havens.

Our treatments are also balanced in the following way. When a tax is introduced on one market from periods 7-12, we set up an equal number of sessions where the tax applies only to the LEFT or the RIGHT market, resp. When the tax is abolished in periods 13-18, we also do that in a symmetric way either in LEFT or in RIGHT.

In most of the treatments, we use a tax rate of 0.5%. As a robustness check on the influence of the tax rate, we also consider one treatment (with sequence of taxation 0L2) with a lower tax rate of 0.1% in order to see whether our main findings are sensitive to the tax rate (see the last row in Table 1).

The number of sessions per treatment is provided in the last column of Table 1. For each treatment except the first two (with two sessions each), we ran four sessions with 20 traders. This corresponds to 48 markets in total (one LEFT and one RIGHT market per session). The 480 participants in our 24 sessions were business students at the University of Innsbruck. Sessions were computerized (using zTree by Fischbacher, 2007) and lasted about 75 minutes.

Since foreign currency traders in reality typically do not trade on their own accounts and are compensated contingent on their relative performance in comparison with the market, we also used a benchmarking system in our experiment. More precisely, incentives for trading were such that a trader's compensation per period was based on his performance benchmarked by the performance of all other traders in his market. The earnings per period were accumulated over all trading periods, and subjects earned on average 17 Euros.

3 Descriptive overview of results

In this section, we present descriptive data on several key market variables. In Section 4 we provide econometric estimations on how a Tobin tax affects the trading volume, exchange rate volatility, market efficiency, and speculative behavior.

3.1 Trading volume and number of transactions

Table 2 shows the development of trading volume by reporting the average trading volume in Gulden for each period of the three phases of the experiment. The percentages below the figures refer to the change in the average trading volume per period of the current phase in relation to the average trading volume per period of the previous phase. Note that we do not distinguish between the tax being levied on the LEFT or RIGHT market in Table 2, but rather pool the treatments that have otherwise the same sequence of taxation (and the same tax rate). For example, we pool 0L0 and 0R0 into 010.

Insert Table 2 about here

The figures in Table 2 reveal that taxing both markets leads to a strong reduction in the trading volume (of 24% on average) in relation to the preceding phase, whereas the effect of taxing only a single market is markedly smaller.

Across all treatments, the average number of transactions per period and market is 25.2 when both markets are untaxed. When both markets are taxed simultaneously, the average number of transactions drops to 21.3. In case of taxing only one market, but not the other, the average number of transactions on the taxed markets falls sharply to 5.9, while it rises to 39.8 in the untaxed market. Hence, taxing only one market has strong repercussions on both markets.

Figure 1 shows on the left-hand side the average number of transactions in the different treatments. The dashed vertical lines indicate changes in taxation (by either introducing a tax on one or both markets or by abolishing it in a hitherto taxed market).

Insert Figure 1 about here

When the tax is levied in periods 7-12 only on the LEFT market in treatments 0L0 and 0L2 (panel A in Figure 1), there is a very strong shift in the number of transactions from the LEFT market to the RIGHT market from period 7 onwards. When the tax is either abolished after period 12 (in 0L0) or the RIGHT market is also taxed (in 0L2) much of the trading activity floats back to the LEFT market, without reaching pre-tax levels, though.

Panel B of Figure 1 refers to the treatments where only the RIGHT market is taxed in periods 7-12. We see that the number of transactions on RIGHT is less affected by the tax than if the tax is levied on LEFT (compare the sharp kinks after periods 6 and 12 in the previously discussed graph with the much smoother transitions in this graph). This is a first indication that introducing a tax on LEFT or RIGHT has different consequences. Section 3.2 will get back to this issue in more detail.

Panel C of Figure 1 presents the number of transactions in treatments 02L and 02R. By looking at the transition from period 12 to period 13, we see that abolishing the tax in either LEFT or RIGHT leads to a very strong increase in transactions on the respective market.

3.2 Market shares of LEFT and RIGHT

Figure 1 shows on the right-hand side the development of the market share of LEFT, i.e. the ratio of the trading volume on the LEFT market to the trading volume of both markets combined. A first notable fact is that in periods 1-6 the LEFT market always has a considerably larger market share than the RIGHT market. On average, 70% of trading in periods 1-6 takes place on the LEFT market, but only 30% on the RIGHT market, even though both markets are set up identically. Hence, the visual positioning of markets on subjects' screens creates one relatively big and one relatively small market.[6]

Panel D of Figure 1 shows that the introduction of the tax in 0L0 causes a drop in the market share of the LEFT market from almost 80% in period 6 to less than 20% in period 7 and less than 10% in periods 9-12. Hence, the shift in trading volume as a consequence of introducing the tax is very rapid and very strong. Abolishing the tax after period 12 leads to an increase in LEFT's market share to about 46% in periods 13-18, but the pre-tax levels (of, on average, 67%

[6] This may be the outcome of a coordination game between market participants. If subjects prefer more liquid markets and if they expect other market participants to trade on the RIGHT (LEFT) market, then it may be reasonable for them to enter the RIGHT (LEFT) market. If more subjects expect the others to be active on the RIGHT than the LEFT market, the RIGHT market ends up as the thicker one.

in periods 1-6) are not reached anymore. Considering treatment 0R0, we note that LEFT gains in market share when the tax is levied on the RIGHT market in periods 7-12. The RIGHT market's share drops from 21% in periods 1-6 to 7% in periods 7-12. Abolishing the tax from the small RIGHT market brings back the market share (22%) almost exactly to pre-tax levels.

Panel E of Figure 1 conveys a similar message as panel D. When the tax is introduced on the large LEFT market in treatment 0L2, it causes a huge drop in market share (of 79%) when comparing periods 7-12 to periods 1-6). Trading does not shift back completely when the RIGHT market is also taxed from period 13 onwards. Yet, if the small RIGHT market is taxed in treatment 0R2 there are only minor effects. The RIGHT market loses only about 47% 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. In sum, panels D and E of Figure 1 imply that introducing a Tobin tax on a relatively larger market leads to a relatively stronger drop in market share than when the tax is introduced on a smaller market.

Panel F of Figure 1 presents the two treatments where the Tobin tax is first levied on both markets, and afterwards abolished on one of them. The introduction of the tax in period 7 leads to a (small) shift in market shares from the large to the small market. When the tax is abolished on one of the markets after period 12, this market captures almost the whole trading activity. For instance, when the tax is maintained on LEFT, but is abolished on RIGHT (02L, see solid line in panel F), the market share of LEFT drops from 78% in period 12 to 10% in period 13.

3.3 Tax revenues

Naive estimates of the revenues from a Tobin tax would multiply total turnover before the introduction of the tax by the intended tax rate. Such an approach could be highly misleading, though, as Table 3 shows.

Insert Table 3 about here

The first column in Table 3 reports the average "hypothetical" tax revenues per period and market in phase 1, assuming that the whole turnover in these periods could have been taxed. Considering the treatments with the tax rate of 0.5% we see that these hypothetical revenues range from 237 Taler to 402 Taler. When both markets are actually taxed (see periods 7-12 in 02L or 02R, and periods 13-18 in 0L2 or 0R2) the tax revenues range only from 144 Taler to 248 Taler, and they are always smaller than the hypothetical revenues from periods 1-6 in the same treatment. If only one market is taxed, however, the tax revenues per period are at most 31 Taler, and on average only 21 Taler. This indicates massive tax avoidance through a shift in trading activity from the taxed to the untaxed market, as has been documented already in Figure 1.

3.4 Speculative behavior

Table 4 provides summary statistics for two different measures – based on individual trading behavior – of speculative behavior. Given our information structure where new information arrives only at the beginning of trading periods, we can define a first measure for short-term speculation based on the number of times a trader switches from buying to selling within a given trading period. Since a currency's fundamental value remains constant within a trading period, the frequency of switching between buying and selling within a period indicates the extent to which a trader speculates on short-term price movements which are not driven by fundamentals. The validity of this measure may be adversely affected in periods when the price oscillates around the fundamental value. Therefore, we can construct a second proxy for short-term speculation from observing that short-term speculators typically favor a quick execution of orders, hence prefer market orders over limit orders.

Insert Table 4 about here

For the first measure, we calculate the ratio of the absolute frequency of switching to the total number of trades minus one (for each trader). For example, consider a trader who first buys, then sells, then buys again and finally sells again within a given period. Hence, this trader switches three times, which is the maximum number of switches possible with four trades, yielding 3/(4–

1)=100%. The second proxy is calculated as the ratio of market orders to limit orders and termed "acceptance ratio".

In the first phase, no market is taxed, so this serves as a benchmark in the following. When taxes prevail, we distinguish between taxed and untaxed markets. Both measures of speculation yield qualitatively identical results: If one of the markets is taxed, speculative behavior on this market is reduced sharply, while it increases on the untaxed market. In line with the literature, speculators, who trade often and are therefore very sensitive to this type of tax, move very quickly to tax havens if they exist. When both markets are taxed, speculation is reduced to about 80% of the benchmark level from phase 1.

3.5 Impact of the tax rate level

To investigate the impact of the level of the tax rate, we replicated the 0L2 treatments with a reduced tax rate of 0.1%. The last line of Table 2 (0L2_0.1) contains the results for the trading volume, which are roughly in line with the corresponding 0.5% tax treatments. Table 3 shows that tax revenues in the 0.1% treatments are roughly one fifth of revenues in the corresponding 0.5% treatments. If a lower tax led to a smaller reduction in trading volume, we would see higher tax revenues in the bottom line of this table.

4 Econometric analysis

The analysis in this section is based on the following panel regression equation (where y is a generic placeholder for the dependent variables to be considered in the following subsections):

$$y_{j,k} = \beta_{\mathrm{AR}}\mathrm{AR}(1) + \sum_{i \in \{b,t,u\}} \beta_i T^i_{j,k} + \epsilon_{j,k}. \tag{1}$$

The market index[7] is denoted by j, and $k = 1, \ldots, 18$ refers to the trading period index. $\mathrm{AR}(1) = y_{j,k-1}$ denotes the dependent variable lagged by one period. $T^i_{j,k}$ are tax dummies, where T^b equals 1 if both markets are taxed, T^u

[7]Remember that we have 48 markets in total, cf. Section 2.3.

is set to 1 if the respective market is untaxed, but the other one taxed, and T^t takes on the value 1 for a taxed market when the other market is untaxed.

For each dependent variable we test for fixed or random unobserved effects both in the time domain and in the cross-section. Whenever such effects are detected, they are accounted for in the estimation procedure. We account for any remaining (time-) heteroskedasticity as well as autocorrelation within cross-sections by using a period SUR (PCSE) method to compute robust covariances (Beck and Katz, 1995).

We start our analyses by focusing on the treatments with a 0.5% tax rate ($j = 1, \ldots, 40$). The final subsection considers also the markets with the lower tax rate of 0.1% ($j = 41, \ldots, 48$).

4.1 Market volume

If only one market is taxed and the other one remains as a tax haven, the literature is unanimous about the likely consequences: Tax avoidance will lead to a reduction in volume on the taxed market and to an almost equivalent increase on the untaxed market. When both markets are taxed, a reduction in volume can be expected if the tax is successful in reducing short-term speculation (see, e.g., the contributions in Haq et al., 1996; Weaver et al., 2003).

To examine the effects of a Tobin tax on the volume of trading, we use equation (1) with the level of trading volume as the dependent variable. Trading volume is defined as the amount of Taler traded within each period. The first column in Table 5 shows that the trading volume on the taxed market is significantly reduced when only one market is taxed (see the negative coefficient of T^t), which is mirrored by an increase in volume on the untaxed market (see T^u). However, the net effect is significantly negative, indicating that taxation (even on one market only) reduces overall trading volume. This finding is perfectly in line with the literature and can be explained by tax avoidance. When both markets are taxed, the volume decreases (see T^b), and the decrease is larger than if only a single market is taxed. However, the decrease in volume is not significantly higher in the former than in the latter case (i.e., the sum of coefficients for T^t and T^u is not significantly different from the coefficient for T^b).

Insert Table 5 about here

4.2 Volatility

One key issue in the academic and political debate about the Tobin tax is its effect on market volatility. While some papers predict an increase in volatility (see, e.g., Aliber et al., 2003; Hau, 2006), others expect volatility to decrease (see, e.g., Westerhoff, 2003; Ehrenstein et al., 2005). We proxy volatility by the average of absolute returns across each trading period (|Ret|):[8]

$$y_{j,k} = \frac{\sum_{\theta=1}^{\Theta} |\mathrm{Ret}_{j,k,\theta}|}{\Theta}, \tag{2}$$

$$|\mathrm{Ret}_{j,k,\theta}| = |\ln(P_{j,k,\theta}) - \ln(P_{j,k,\theta-1})|. \tag{3}$$

Here θ stands for each transaction and Θ measures the total number of transactions within a certain period k and market j. Previous experimental studies of continuous double auction markets have found that the volatility of transaction prices is generally decreasing across trading periods (for a survey see Sunder, 1995). Since the Tobin tax has only been introduced in period 7 or later in our experiment, it would be inadequate to compare the volatility of prices before and after introducing the tax without controlling for a time trend. Therefore, we include a linear trend term that absorbs any linear dependence on the number of trading periods.

The results are presented in the second column of Table 5, showing that volatility remains unchanged if a market is taxed. This holds regardless of whether the tax is encompassing (see T^b) or introduced only on a single market (see T^t). On the untaxed market, however, volatility *decreases* significantly (see T^u), supposedly due to higher liquidity as a consequence of an increase in trades and trading volume.

It seems important to note already at this point, though, that the insignificant coefficient of T^t is a consequence of two opposing interaction effects of the

[8] As a robustness check, we repeated our estimation using the standard deviation of returns as the dependent variable. Reassuringly, our results did not change qualitatively. Details are available upon request.

tax with market size. Section 4.5 will show in detail that volatility increases when the small market is taxed, but decreases when the large market is taxed. On aggregate, both effects balance each other, leading to the seeming null-effect on volatility when only one market is taxed.

4.3 Short-term speculation

We continue our analysis by looking at the effects of a Tobin tax on individual trading patterns, because one of the desired effects advocated by many proponents of a Tobin tax is to discourage the activities of short-term speculators. Bloomfield et al. (2009) find a decrease in short-term speculation in their single-market setting. With our two-market setup, we can examine the effects on speculation also in the presence of tax havens.

In subsection 3.4 we have introduced two measures of speculation. The dependent variable "switching frequency" is constructed as the overall average of switching per period across all 20 traders. The dependent variable "acceptance ratio" is calculated analogously for the ratio of market orders to limit orders. The results from the regressions using these dependent variables can be found in columns 3 and 4 of Table 5. Both proxies yield very similar results.

When only one market is taxed, short-term speculation is reduced on the taxed market (see T^t), as advocates of a Tobin tax would expect. However, the decrease in short-term speculation on the taxed market is accompanied by an increase in short-term speculation on the untaxed market (see T^u), which partly – but not fully – offsets the reduction in speculative behavior caused by the tax. Again, we can show that the Tobin tax has repercussions also on markets where the tax is *not* levied. When both markets are taxed, we see that short-term speculation is significantly reduced (see T^b), confirming that the tax has the desired effects if no tax haven exists.

As regards our second proxy for short-term speculation, we find that the acceptance ratio decreases when a market is taxed. Yet, this decrease is not caused by a reduction in the number of limit orders (which are unaffected by a tax) but by a reduction in market orders, i.e., in the willingness to accept limit orders.

4.4 Market efficiency

Another issue on which the literature is ambiguous is market efficiency. While Ehrenstein (2002) and Westerhoff (2003) predict an increase in efficiency due to a Tobin tax, others expect the exact opposite (Kupiec, 1995; Subrahmanyam, 1998). The results of Bloomfield et al. (2009) are somewhat in between, reporting no change in efficiency following the introduction of the tax, which is a result similar to Cipriani and Guarino (2008).

We measure informational efficiency by the absolute deviation between the average price within a trading period ($\overline{P}$) and the fundamental value (V), standardized over its average per market and period:

$$y_{j,k} = \frac{|\overline{P_{j,k}} - V_{j,k}|}{\frac{\sum_{k=1}^{18} |\overline{P_{j,k}} - V_{j,k}|}{18}}. \tag{4}$$

A higher value for this measure indicates a less efficient market. The results are shown in the last column of Table 5. If only one market is taxed, inefficiency on the taxed market increases significantly. The most plausible explanation for this observation is the reduction in liquidity due to decreased trading volume. In the case of an encompassing tax, we do not find any evidence of a change in market efficiency due to the presence of a Tobin tax.

4.5 Influence of market size

Figure 1 has already suggested that there might be an interaction between market size and the effects of a Tobin tax. In order to substantiate this conjecture, we split our sample into two subsamples: The first (second) subsample contains those markets that had the larger (smaller) market share in the first six trading periods.[9] In the following we discuss the interaction effects separately for the different tax regimes. The estimation results are summarized in Table 6.

Insert Table 6 about here

[9]We prefer splitting the sample to the alternative approach via dummy variables because the fixed effects transformation is not applicable for dummies that are constant for any cross-section.

4.5.1 Tax on both markets

When the tax covers both markets (i.e., when there are no tax havens), the volume and market share of the large market decrease significantly as trading activity shifts to the smaller market. The market share of the small market increases correspondingly, while the volume on the small market remains unchanged (see row T^b in Table 6). Hence, an encompassing tax reduces the overall trading volume. Speculative behavior also decreases on the large market. On the small market, one of the two measures (acceptance ratio) decreases, while the other remains unchanged. Volatility and efficiency are not affected on either market, though (see the headings "Av. abs. return" and "Efficiency").

4.5.2 Tax on the large market only

When the tax is levied on the large market only (see row T^t in columns "large"), we observe a massive shift in trading activity towards the small market (see row T^u in columns "small"). Both volume and market share decrease significantly on the large market, while simultaneously increasing (volume to a smaller extent, though) on the small market. Speculators are driven away from the large market by the tax, increasing speculative activity on the small market. However, the total effect (viewing both markets together) is a reduction in speculation on the taxed market, which is accompanied by an increase in speculative behavior on the untaxed market. While volatility is not affected on either market, efficiency on the taxed market decreases.

4.5.3 Tax on the small market only

When the tax is levied on the small market only, half of its volume is lost (see row T^t in columns "small"). About 90% of this loss in volume is due to a shift of market volume to the untaxed large market (T^u), implying that the overall reduction in trading volume across both markets is hardly noticeable. Short-term speculative behavior in the small market is significantly reduced and speculation does not shift to the large market, but simply vanishes. Volatility even *increases* on the small market when it is taxed, and decreases on the larger market. Efficiency on the taxed market decreases significantly.

Taken together, the results in subsections 4.5.2 and 4.5.3 are consistent with a recent model that postulates a U-shaped relationship of volatility and market volume (Haberer, 2004, 2006). While on the highly liquid large market a tax-driven reduction in volume (taken as a rough proxy for liquidity) leads to a decrease in volatility, the opposite effect is observed when a comparatively illiquid small market is taxed, thus driving out liquidity and increasing volatility.

4.6 Impact of the tax rate level

To investigate a possible impact of the level of the tax rate, we finally analyze the markets with a tax rate of 0.1%. We start with regressions described in equation (1) for the markets with the 0.1% tax rate ($j = 41, \ldots, 48$). This allows us to examine whether the qualitative results established under a tax rate of 0.5% (see Table 5) also prevail under a smaller tax rate. The estimation results for the markets with the tax rate of 0.1% are shown in Table 7.

As far as market volume is concerned, the signs of the coefficients are the same as in Table 5, which means that our results remain qualitatively unchanged by the lower tax rate. The same holds for the switching frequency and acceptance ratio, our measures of short-term speculation. With respect to volatility and market efficiency we note some differences, though, when we only consider the markets with the tax rate of 0.1%: Volatility on the untaxed market is unaffected, and efficiency decreases not only for single-taxed markets, but also when both markets are taxed. When we consider all 012-markets, however, these differences do not prevail. Table 8 combines all 012-markets (both with 0.5% and 0.1% tax rates). While the tax dummies T^b, T^t and T^u are set for all the corresponding treatments regardless of the level of the tax rate, the dummies with subscript 0.1% are only set for the treatments with the lower tax rate to capture its differential effects. We find that the coefficients for all dummies of the latter type are insignificant. This suggests that the tax rate itself has no significant impact on our results.

5 Conclusion

James Tobin has triggered a lively debate about the pros and cons of a transaction tax on foreign exchange markets. In this paper, we have examined in a controlled experiment many of the disputed issues. While one may consider it a big leap from the laboratory to the real world, the experimental approach seems justified because of the lack of empirical evidence from real foreign exchange markets. Of course, the actual implementation of a Tobin tax on real-world foreign exchange markets would ultimately resolve the controversies over the tax's alleged consequences. In the meantime, experimental studies on a Tobin or transactions tax on foreign exchange markets can provide insights into the consequences of such taxes.

Our experimental results confirm many of the theoretically expected effects of a Tobin tax, while at the same time questioning some of its alleged effects. The results on issues where there is broad consensus could have been easily anticipated, of course: Trading volume is negatively affected when all markets are subjected to a Tobin tax, and the tax reduces the number of transactions. The large degree of the shift of trading and transactions from the taxed to the untaxed market – if a tax haven exists – may also be regarded as self-evident, as it is the outcome of massive tax avoidance, leading to almost negligible tax revenues in the presence of tax havens. If these experimental results applied to the real world, this would clearly question many politicians' expectation of using the Tobin tax as a stable basis for tax revenues.

The more interesting issues are those where the results would have been difficult to anticipate. For instance, the interaction of market size and Tobin tax on market activity seems less straightforward. In fact, trading volume and trading activity are much more affected if the Tobin tax is levied on the larger of the two markets. The different positioning of both markets on subjects' screens has turned out to yield strong differences in market size (with an approximate ratio of 2:1 for the LEFT market when there is no tax), which has made it possible to detect these intricate interaction effects of market size and Tobin tax. The stronger influence of the tax on the larger market seems to be driven by drying up the hitherto very liquid large market. When the tax is introduced on the small market, this relatively illiquid market recovers its rather low level of liquidity when the tax is abolished again. The latter effects has not been dependent on the tax rates applied on our experimental financial markets.

One disputed key issue in the debate on the Tobin tax has been market volatility. We find in our experiment no reduction in volatility due to the introduction of a Tobin tax if the tax is encompassing. This result is clearly in conflict with the hopes of the supporters of a Tobin tax. When the tax is introduced unilaterally on the larger (smaller) of the two markets, volatility decreases (increases). Noting that market size in our experiments is closely linked to liquidity, this result confirms theoretical results by Haberer (2004). Furthermore, it is important to note that a Tobin tax on one market has been found to decrease volatility on the untaxed market, which is mainly caused by a shift in trading volume. This effect has not been documented in the literature so far. However, as the results of Pellizzari and Westerhoff (2009) suggest, this finding may depend on the trading institution prevalent on a market. One avenue for future research would be to examine this conjecture by running experimental markets with different trading institutions.

Another key issue besides volatility is the question of how a Tobin tax affects market efficiency. Confirming earlier experimental findings of Bloomfield et al. (2009) and Cipriani and Guarino (2008) we observe no impact of the tax on informational efficiency when both markets are taxed. If only one market is taxed, inefficiency on the taxed market increases significantly.

We have also been able to document the microeffects of a Tobin tax on individual trading patterns. As intended by the supporters of a Tobin tax, it affects in particular the trading activities of those traders who might be classified as short-term speculators. The frequency of switching between buying and selling is adversely affected by the introduction of a tax, as is the traders' willingness to issue market orders. This result is clearly in line with earlier findings of Bloomfield et al. (2009) who have shown that a securities transaction tax limits the activities of noise traders. In contrast to Bloomfield et al. (2009)

we have been able to document these effects in a broader range of settings that includes both the uniform taxation of all foreign exchange markets and the parallel existence of taxed and untaxed markets. The latter case seems to be the one that is more likely – provided some politicians deem the (economic) consequences of a Tobin tax desirable and implement such a tax, whereas others abstain from it in order to benefit from shifts in trading volume towards tax havens. Our results show that speculators will mostly evade the tax, hence short-term speculation will shift to tax havens rather than vanish. If our results – in particular those in Section 4.5.2 – would hold on real financial markets, then the most important political implication would be that the distortions caused by introducing a Tobin tax not worldwide, but on a subset of markets, may not be undone completely by later on abolishing the tax again. This would suggest that politicians should think twice before they use the financial markets in their countries for a real-time field experiment on the economic consequences of a Tobin tax. Of course, our paper can only be considered a first step into analyzing the effects of a Tobin tax using experimental economics. More research is clearly needed. A next step – also worthy of future investigation – would be to examine the behavioral reactions of actual foreign-currency traders on the introduction of a Tobin tax on a controlled experimental financial market. This would contribute to the issue of the external validity of our experiment. Interestingly, a paper by Haigh and List (2005) shows that professional traders do not better (and partly worse) than university students in an investment task that examines myopic loss aversion. It remains to be seen whether the same holds true for trading behavior in foreign exchange markets. If it does, this would be no good news for the (political) supporters of a Tobin tax.

References

Aliber, R.C., Chowdhry, B., Yan, S., 2003. Some evidence that a Tobin tax on foreign exchange transactions might increase volatility. European Finance Review 7, 481–510.

Beck, N., Katz, J.N., 1995. What to do (and not to do) with time-series cross-section data. American Political Science Review 89, 634–647.

Bloomfield, R.J., O'Hara, M., Saar, G., 2009. How noise trading affects markets: An experimental analysis Review of Financial Studies 22, 2275–2302 .

Cipriani, M., Guarino, A., 2008 Transaction costs and informational cascades in financial markets. Journal of Economic Behavior and Organization 68, 581–592.

Dow, J., Rahi, R., 2000. Should speculators be taxed? Journal of Business 73, 89–107.

Ehrenstein, G., 2002. Cont-Bouchaud percolation model including Tobin tax. International Journal of Modern Physics C 13, 1323–1331.

Ehrenstein, G., Westerhoff, F. Stauffer, D., 2005. Tobin tax and market depth. Quantitative Finance 5, 213–218.

Eichengreen, B., Tobin, T., Wyplosz, C., 1995. Two cases for sand in the wheels of international finance. Economic Journal 105, 162–172.

Fischbacher, U., 2007. z-tree: Zurich toolbox for readymade economic experiments. Experimental Economics 10, 171–178.

Frankel, J., 1996. How well do foreign exchange markets work: might a Tobin tax help? Haq, M. ul, Kaul, I., Grunberg, I. (Eds.). The Tobin Tax: Coping with Financial Volatility, Oxford: Oxford University Press, 41–81.

Haberer, M., 2004. Might a securities transactions tax mitigate excess volatility? Some evidence from the literature. CoFE Discussion Paper 04-06, University of Konstanz.

Haberer, M., 2006. Regulierung internationaler Finanzmärkte durch Transaktionssteuern: Die Wirkung einer Tobin-Steuer auf Handelsvolumen und Wechselkursvolatilität. Ph.D. thesis, University of Konstanz.

Habermeier, K., Kirilenko, A., 2003. Securities transaction taxes and financial markets. IMF Staff Papers 50, 181–186.

Haigh, M.S., List, J.A., 2005 Do professional traders exhibit myopic loss aversion? An experimental analysis Journal of Finance 60, 523–534.

Haq, M. ul, Kaul, I., Grunberg, I., 1996. The Tobin Tax. Coping with Financial Volatility. Oxford: Oxford University Press.

Hau, H., 2006. The role of transaction costs for financial volatility: Evidence from the Paris bourse. Journal of the European Economic Association 4, 862–890.

Jones, C.M., Seguin, P.J., 1997. Transaction costs and price volatility: Evidence from commission deregulation. American Economic Review 87, 728–737.

Kaiser, J., Chmura, T., Pitz, T., 2007. The Tobin tax - a game-theoretical and an experimental approach Working Paper.

Kupiec, P.H., 1995. A securities transactions tax and capital market efficiency. Contemporary Economic Policy 113, 101–112.

Mannaro, K., Marchesi, M., Setzu, A., 2008. Using an artificial financial market for assessing the impact of Tobin-like transaction taxes. Journal of Economic Behavior and Organization 67, 445–462.

Palley, T.I., 1999. Speculation and Tobin tax: Why sand in the wheels can increase economic efficiency. Journal of Economics 69, 113–126.

Pellizzari, P., Westerhoff, F., 2009. Some effects of transaction taxes under different microstructures. Journal of Economic Behavior and Organization 72, 850–863.

Schwert, W.G., Seguin, P.J., 1993. Securities transaction taxes: An overview of costs, benefits, and unresolved questions. Financial Analysts Journal 27–35.

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

Subrahmanyam, A., 1998. Transaction taxes and financial market equilibrium. Journal of Business 71, 81–118.

Summers, L., Summers, V.P., 1989. When financial markets work too well: A cautious case for a securities transaction tax. Journal of Financial Services Research 3, 261–286.

Sunder, S., 1995. Experimental asset markets: A survey. Roth, A., Kagel, J., (Eds.). Handbook of Experimental Economics, chap. 6: Princeton University Press, 445–500.

Tobin, J.T., 1978. A proposal for international monetary reform. Eastern Economic Journal 4, 153–159.

Umlauf, S.R., 1993. Transactions taxes and the behaviour of the Swedish stock market. Journal of Financial Economics 33, 227–240.

Weaver, J., Dott, R., Baker, J.E., 2003. Debating The Tobin Tax. New Rules for Global Finance Coalition.

Werner, I.M., 2003. Comment on some evidence that a Tobin tax on roreign exchange transactions may increase volatility. European Finance Review 7, 511–514.

Westerhoff, F.H., 2003. Heterogeneous traders and the Tobin tax. Journal of Evolutionary Economics 13, 53–70.

Westerhoff, F.H., Dieci, R., 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 30, 293–322.

Table 1: Experimental treatments.

	Phase 1		Phase 2		Phase 3		
Treatment	LEFT	RIGHT	LEFT	RIGHT	LEFT	RIGHT	No. of sessions
0L0	-	-	0.5%	-	-	-	2
0R0	-	-	-	0.5%	-	-	2
0L2	-	-	0.5%	-	0.5%	0.5%	4
0R2	-	-	-	0.5%	0.5%	0.5%	4
02L	-	-	0.5%	0.5%	0.5%	-	4
02R	-	-	0.5%	0.5%	-	0.5%	4
0L2_0.1	-	-	0.1%	-	0.1%	0.1%	4

Entries show the two-way tax rate for taxed markets (LEFT and/or RIGHT), dashes indicate the absence of taxes.

Table 2: Trading volume in Gulden per phase, aggregated treatments.

Tax regime	Phase 1	Phase 2	Phase 3
010 (0L0 + 0R0)	31940	28766	30340
		–9.9%	+5.5%
012 (0L2 + 0R2)	21837	20346	18681
		–6.8%	–8.2%
021 (02L + 02R)	26981	20486	22795
		–24.1%	+11.3%
0L2_0.1	19734	18963	18304
		–3.9%	–3.5%

Percentage numbers represent changes in trading volume relative to the previous phase.

Table 3: Tax revenues in Gulden per phase and treatment.

Treatment	Phase 1	Phase 2	Phase 3
0L0	*237*	18	*294*
0R0	*402*	16	*320*
0L2	*197*	24	178
0R2	*240*	31	195
02L	*239*	144	16
02R	*301*	248	23
0L2_0.1	*39*	5	37

Italic figures represent hypothetical tax revenues when both markets are untaxed, assuming that the actually observed turnover could have been taxed.

Table 4: Speculative behavior per phase and treatment.

	Switching frequency					Acceptance ratio				
	Phase 1	Phase 2		Phase 3		Phase 1	Phase 2		Phase 3	
	untaxed	taxed	untaxed	taxed	untaxed	untaxed	taxed	untaxed	taxed	untaxed
0L0	9.7	0.8	16.0		9.8	91.6	55.9	3.4		97.0
0R0	14.6	1.0	22.0		13.8	84.9	33.1	90.3		74.6
0L2	11.7	3.4	18.5	9.6		86.0	43.4	85.3	70.9	
0R2	12.0	2.3	17.3	9.6		90.9	54.6	110.7	79.5	
02L	12.7	10.8		1.5	19.2	81.4	63.4		33.1	84.1
02R	12.9	8.8		1.4	16.8	82.6	64.7		33.7	91.1
0L2_01	11.2	3.3	17.4	10.3		99.9	66.6	114.3	88.5	

Table 5: Panel regressions for treatments with 0.5% tax rate (based on equation (1)).

	Volume	Average absolute return in %	Switching frequency in %	Acceptance ratio in %	Inefficiency in %
intercept	322.231***	4.581***	12.037***	85.503***	97.683***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
T^b	−67.448***	−0.060	−2.555***	−15.365***	−19.056
	(0.005)	(0.917)	(0.003)	(0.000)	(0.219)
T^t	−224.704***	0.542	−8.978***	−40.638***	52.843***
	(0.000)	(0.398)	(0.000)	(0.000)	(0.003)
T^u	167.917***	−1.309**	5.306***	5.110	−4.459
	(0.000)	(0.032)	(0.000)	(0.161)	(0.802)
AR(1)	0.326***	7.955*	28.677***	-	-
	(0.000)	(0.062)	(0.000)		
lin. trend		−0.150***			
		(0.001)			
Fixed effects	CS	CS	CS	CS	P
R^2 (in %)	79.1	25.0	66.8	53.0	16.5
n	680	660	680	720	707

p-values are given in parentheses. T^b ... both markets taxed, T^u ... this market untaxed, but other market taxed, T^t ... this market taxed, but other market untaxed. CS... cross-section fixed effects, P... period fixed effects. *, ** and *** represent the 10%, 5% and the 1% significance levels.

Table 6: Panel regressions conditional on market size in phase 1 (periods 1-6) for treatments with 0.5% tax rate.

	Volume		Market share		Av. abs. return in %		Accept. ratio in %		Switching freq. in %		Inefficiency in %	
market size	large	small	large	small	large	small	large	small	large	small	large	small
intercept	469.884^{***}	201.710^{***}	68.500^{***}	31.500^{***}	4.517^{***}	4.785^{***}	96.412^{***}	74.644^{***}	17.780^{***}	7.942^{***}	102.864^{***}	92.413^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
T^b	-105.243^{***}	-15.231	-5.301^{*}	5.301^{*}	-0.226	0.352	-19.962^{***}	-10.890^{***}	-3.066^{***}	-0.987	-27.463	-10.319
	(0.000)	(0.605)	(0.096)	(0.096)	(0.756)	(0.678)	(0.000)	(0.008)	(0.006)	(0.324)	(0.205)	(0.635)
T^t	-309.057^{***}	-110.458^{***}	-53.119^{***}	-15.560^{***}	-0.528	1.523^{*}	-53.073^{***}	-25.999^{***}	-11.735^{***}	-5.007^{***}	40.704^{*}	67.094^{***}
	(0.000)	(0.003)	(0.000)	(0.000)	(0.507)	(0.098)	(0.000)	(0.000)	(0.000)	(0.000)	(0.100)	(0.008)
T^u	95.220^{***}	239.838^{***}	15.560^{***}	53.119^{***}	-1.778^{**}	-0.600	-0.894	10.390^{**}	1.095	9.412^{***}	-17.599	7.308
	(0.001)	(0.000)	(0.000)	(0.000)	(0.034)	(0.494)	(0.851)	(0.025)	(0.366)	(0.000)	(0.497)	(0.763)
AR(1)	-0.128^{**}	0.376^{***}	31.737^{***}	31.737^{***}	16.408^{***}	-	-	-	-	24.292^{***}	-	-
	(0.032)	(0.000)	(0.000)	(0.000)	(0.004)					(0.000)		
lin. trend	-	-	-	-	-0.129^{**}	-0.190^{***}	-	-	-0.206^{**}	-	-	-
					(0.034)	(0.005)			(0.044)			
Fixed effects	CS	CS	CS	CS	CS	CS	CS	CS	CS	CS	P	P
$R2$ (in %)	83.0	69.8	84.9	84.9	34.6	21.9	63.1	43.1	71.5	63.0	21.7	18.5
n	360	340	340	340	335	350	360	360	360	340	357	350

p-values are given in parentheses. T^b . . . both markets taxed, T^u . . . this market untaxed, but other market taxed, T^t . . . this market taxed, but other market untaxed. CS. . . cross-section fixed effects, P. . . period fixed effects. *, ** and *** represent the 10%, 5% and the 1% significance levels.

Table 7: Panel regressions for the low-tax markets (based on equation (1)).

	Volume	Average absolute return in %	Switching frequency in %	Acceptance ratio in %	Inefficiency in %
intercept	248.440***	1.183***	11.164***	99.934***	67.943***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
T^b	−36.428	−0.099	−0.867	−11.685***	27.618*
	(0.400)	(0.711)	(0.537)	(0.000)	(0.075)
T^t	−168.592***	0.441	−8.883***	−36.090**	102.488**
	(0.006)	(0.174)	(0.000)	(0.011)	(0.009)
T^u	167.793***	0.070	7.218***	17.088	37.552
	(0.006)	(0.825)	(0.002)	(0.224)	(0.313)
AR(1)	-	-	-	-	-
lin. trend	-	−0.150***	-	-	
		(0.001)			
Fixed effects	CS	CS	CS	CS	-
R^2 (in %)	71.6	20.4	58.2	37.1	6.4
n	136	143	144	144	143

p-values are given in parentheses. T^b ... both markets taxed, T^u ... this market untaxed, but other market taxed, T^t ... this market taxed, but other market untaxed. CS... cross-section fixed effects. *, ** and *** represent the 10%, 5% and the 1% significance levels.

Table 8: Panel regressions – comparison between high and low tax regimes in the 012-markets.

	Volume	Average absolute return in %	Switching frequency in %	Acceptance ratio in %	Inefficiency in %
intercept	263.290***	3.310***	11.642***	92.282***	77.923***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
T^b	−45.509	1.184	−2.173*	−13.245***	30.741*
	(0.150)	(0.291)	(0.077)	(0.000)	(0.036)
T^t	−154.897***	1.177**	−6.937***	−32.695***	71.029***
	(0.000)	(0.023)	(0.000)	(0.000)	(0.001)
T^u	111.884***	−0.533	4.795***	2.809	−5.766
	(0.007)	(0.495)	(0.003)	(0.725)	(0.788)
$T^b_{0.1\%}$	16.356	0.900	0.826	1.561	−13.103
	(0.765)	(0.241)	(0.697)	(0.780)	(0.488)
$T^t_{0.1\%}$	−15.307	−0.313	−2.324	−3.395	21.479
	(0.830)	(0.768)	(0.410)	(0.806)	(0.509)
$T^u_{0.1\%}$	60.045	1.638	1.403	14.279	33.339
	(0.402)	(0.115)	(0.618)	(0.303)	(0.283)
AR(1)	0.361***	-	26.888***	-	-
	(0.000)		(0.000)		
lin. trend	-	−0.193**	-	-	
		(0.042)			
Fixed effects	CS	CS	CS	CS	-
R^2 (in %)	73.5	32.8	61.9	51.1	5.1
n	408	398	408	432	426

p-values are given in parentheses. T^b … both markets taxed, T^u … this market untaxed, but other market taxed, T^t … this market taxed, but other market untaxed; these dummies are set for all treatments regardless of the tax rate. The corresponding dummies with subscript 0.1% flag the treatments with the 0.1% tax rate to capture the differential effect of the lower tax rate. CS… cross-section fixed effects. *, ** and *** represent the 10%, 5% and the 1% significance levels.

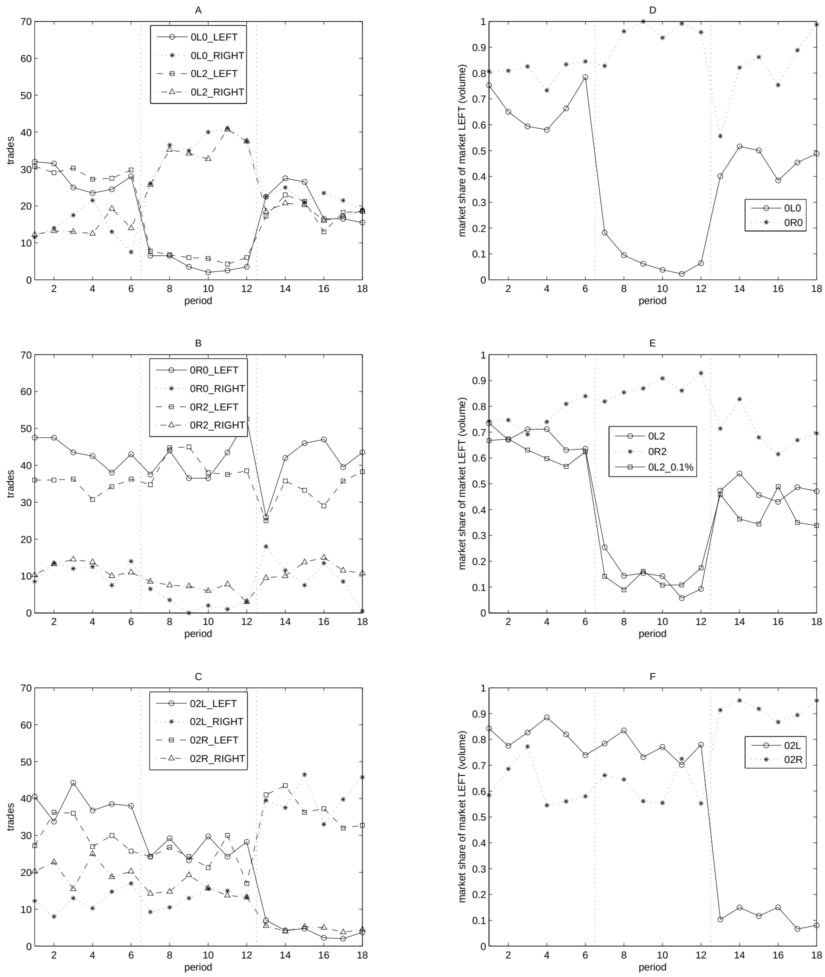

Figure 1: Left column, from top to bottom: Transactions per period when only LEFT is taxed in periods 7-12, transactions per period when only RIGHT is taxed in periods 7-12, transactions per period when the tax is abolished on one market after period 12. Right column, from top to bottom: Market share of LEFT in 0L0 and 0R0, market share of LEFT in 0L2 and 0R2, market share of LEFT in 02L and 02R.

Supplementary Material

Instructions for the Experiment

Background of the experiment

In this experiment 20 traders whose home currency is Gulden can trade Taler in two independent markets for 15-25 consecutive periods.

Market Mechanism

Each participant receives an initial endowment of 8,000 Gulden (G) and 200 Taler (T). Trading will occur through double auctions in two independent markets (called LEFT and RIGHT). In both markets the initial price of Taler is 40 Gulden. This implies that each trader holds 50 percent of his/her total wealth in Taler and 50 percent in Gulden. The price of the Taler is determined by your and the other traders' actions in the market. As orders will not always be identical on both markets, there may be two different prices for Taler. No interest is paid for Taler or Gulden. The value of the Taler is determined by a set of economic factors (not modeled here) and follows a random-walk process without drift:

$$P_k = P_{k-1} + \epsilon_k,$$

with P_k being the value of Taler in period k and ϵ_k following a standard normal distribution. This implies that this period's value is the best estimator for next period's value. Each trader is free to buy or sell Taler at any time on any market. You can trade on both markets at the same time. You may buy in one market and sell in the other, buy in both, etc. All traders always receive the same information on the value of the Taler and everybody sees all transaction prices of the respective period.

Trading strategies

You are free to follow any trading strategy you want.

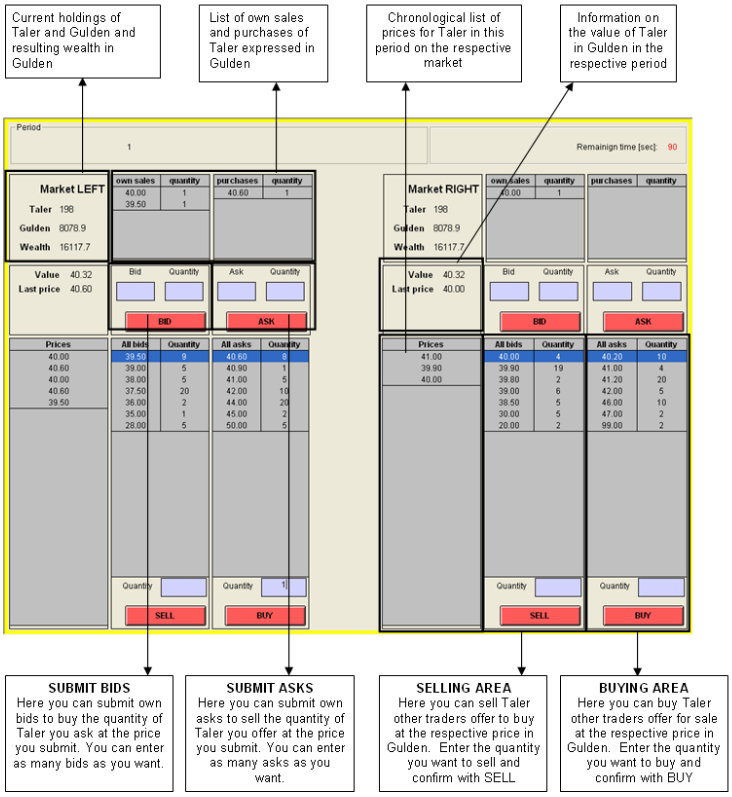

Trading

The trading mechanism is a double auction. This means that each trader can act as seller and as buyer. You are free to submit as many bids and asks (in the range of 10 to 500 with up to two decimal places) as you wish. For each order you have to enter a price and the number of Taler you want to trade. Your holdings of Taler and Gulden can not fall below zero.

Total wealth

Your wealth in Gulden is the sum of your Gulden holdings plus the Gulden value of your Taler (the number of Taler you hold multiplied with the current price; of the two markets the price with the higher volume is used). If you buy Taler your Gulden holdings are immediately reduced and vice versa. Your wealth will change during a period as the market price changes, even if you do not trade; the most recent trading price will be used to value your Taler.

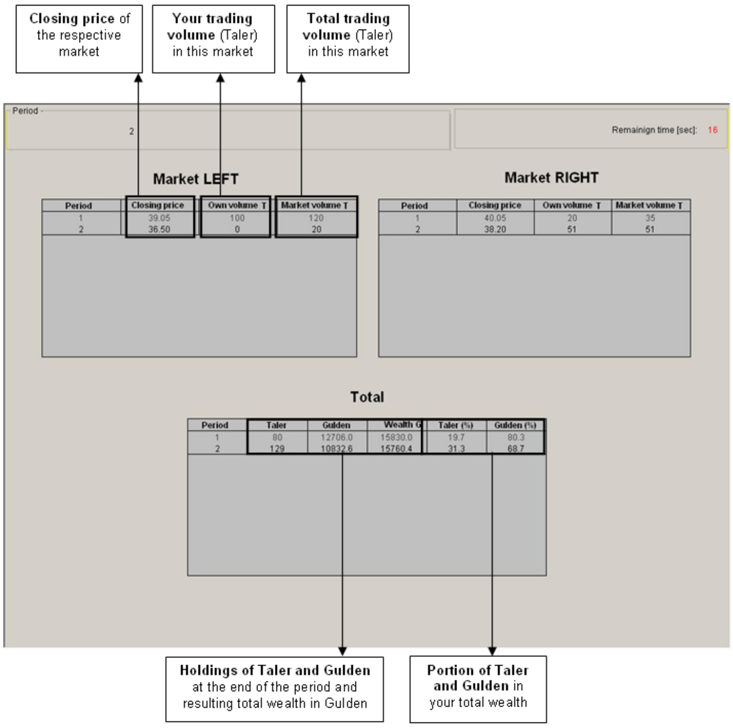

Important details

- Each period lasts 100 seconds.
- The experiment will be terminated randomly between periods 15 and 25.
- Your payment at the end of the experiment depends on your trading success relative to all other participants. This incentive structure is usually used for professional funds managers. Specifically: at the end of each period your wealth will be divided by the average wealth of all participants. The resulting numbers (smaller than 1 if your wealth is below average, equal 1 if it is exactly the average and larger than 1 otherwise) for all periods are added up, yielding a final score. The higher this score, the higher your payment will be. In total we will distribute 340 Euros for this session. To determine your share of this amount the scores of all traders will be added up and 340 will be divided by the total score. Multiplying the result by your score gives your final payment.