

This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jeboSome effects of transaction taxes under different microstructures[☆]Paolo Pellizzari^{a,*}, Frank Westerhoff^b^a Ca' Foscari University, Dept. of Applied Mathematics, Dorsoduro 3825/E, 30123 Venice, Italy^b University of Bamberg, Department of Economics, Feldkirchenstrasse 21, 96045 Bamberg, Germany

ARTICLE INFO

Article history:

Received 8 January 2008

Received in revised form 25 August 2009

Accepted 29 August 2009

Available online 4 September 2009

Keywords:

Transaction tax

Tobin tax

Microstructures

Agent-based models

Liquidity

ABSTRACT

We show that the effectiveness of transaction taxes depends on the market microstructure. Within our model, heterogeneous traders use a blend of technical and fundamental trading strategies to determine their orders. In addition, they may become inactive if the profitability of trading decreases. We find that in a continuous double auction market the imposition of a transaction tax is not likely to stabilize financial markets since a reduction in market liquidity amplifies the average price impact of a given order. In a dealership market, however, abundant liquidity is provided by specialists, and thus a transaction tax may reduce volatility by crowding out speculative orders.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

As, for instance, illustrated by Shiller (2000), financial markets are quite volatile and may display severe bubbles and crashes. Since asset prices are determined by the orders of market participants, one may argue that speculative activity is at least at some times excessive in financial markets. Keynes (1936) and Tobin (1978) therefore suggested introducing a Transaction Tax (TT) in financial markets in order to curb speculative activity. Their basic argument rests on the assumption that there are two types of market participants: stabilizing long-term investors and destabilizing short-term speculators. A low transaction tax presumably has no impact on long-term investors, meaning that their stabilizing influence on the market should remain intact. However, even a modest transaction tax may have a strong impact on the profitability of short-term speculators. Keynes and Tobin suspect that this trader type is the main trigger for the recurrent turbulent behavior of financial markets. If destabilizing short-term speculative orders decrease, financial markets should become more efficient.

The optimistic view of Keynes and Tobin has been supported by a number of prominent economists, including Stiglitz (1989), Summers and Summers (1989) and Eichengreen et al. (1995). For a general discussion of this topic see Schwert and Seguin (1993), Ul Haq et al. (1996) and Spahn (2002). There are, however, also opponents to this proposal. For instance, a few empirical papers conclude that transaction taxes may not contribute to the stabilization of financial markets (Umlauf, 1993; Jones and Seguin, 1997; Aliber et al., 2003; Hau, 2006). However, one should note that these empirical studies are faced with some restrictions. For instance, the paper by Umlauf investigates the case where Sweden introduced a transaction tax rate of 2 percent, whereas nowadays no one is likely to recommend such a high transaction tax. Further sceptical

[☆] Part of this work was done while the first author was at the Quantitative Finance Research Center (QFRC, University of Technology Sydney), whose financial support is gratefully acknowledged.

* Corresponding author. Tel.: +39 0412346924; fax: +39 0415221756.

E-mail addresses: paolop@unive.it (P. Pellizzari), frank.westerhoff@uni-bamberg.de (F. Westerhoff).

comments concerning the empirical research may be found in [Werner \(2003\)](#). Another problematic aspect of transaction taxes may be that policy makers could have an incentive to maximize their tax revenues and thereby select a tax rate which is not efficient with respect to market stability. Note that Tobin always stressed that the key goal of transaction taxes should be the stabilization of financial markets, and that the implicit generation of tax revenues is merely a side effect.

According to survey studies (e.g. [Taylor and Allen, 1992](#)), market participants rely on both technical and fundamental trading rules to determine their orders. Guided by these observations, models have been developed that explore the impact of heterogeneous interacting agents upon the market dynamics. This approach, recently reviewed in [Hommes \(2006\)](#), [LeBaron \(2006\)](#) and [Lux \(in press\)](#), has proven to be quite successful. For instance, these models are able to replicate some important stylized facts of financial markets, such as bubbles and crashes, excess volatility and fat tails for the distribution of returns, thereby adding to our understanding of the working of financial markets. Key contributions include [Day and Huang \(1990\)](#), [Kirman \(1991\)](#), [Chiarella \(1992\)](#), [De Grauwe et al. \(1993\)](#), [Lux \(1995\)](#), [Brock and Hommes \(1998\)](#), [Farmer and Joshi \(2002\)](#) and [Rosser et al. \(2003\)](#).

Given the power of these models, it seems natural to use them as artificial laboratories to test the effectiveness of transaction taxes. An early contribution in this direction is [Frankel \(1996\)](#). This paper develops a simple exchange rate model with two types of agents: investors believe that the exchange rate will return towards its fundamental value while speculators are convinced that the exchange rate will trace out a bubble path. Frankel analytically shows that an exogenous increase in the fraction of investors leads to a reduction in the variability of the exchange rate. The opposite is true when the proportion of speculators increases. According to Frankel, a transaction tax could be expected to lower the fraction of speculators or to raise the fraction of investors. Either way, he suspects that the volatility of the exchange rate will decrease.

[Westerhoff \(2003\)](#) and [Westerhoff and Dieci \(2006\)](#) develop models in which agents may endogenously select between technical and fundamental trading rules. In addition, they may be inactive. The agents' choice process depends on the strategies' past performance. A strategy that did well in the past will be followed by more agents in the future. These two models predict that the imposition of a low transaction tax is likely to increase market stability, since it crowds out speculative activity. Only when the tax rate is set too high may market efficiency decrease.

[Ehrenstein et al. \(2005\)](#) and [Mannaro et al. \(2006\)](#) claim that transaction taxes may have a negative impact on market liquidity. This is an important observation since market liquidity is inversely related to the price responsiveness of a given order (see [Lillo et al., 2003](#); [Farmer et al., 2004](#); [Lillo and Farmer, 2005](#)). This means that the lower the liquidity, the stronger the price change is with respect to a given incoming order. Both papers find that if this effect is taken into account, the stabilizing impact of a transaction tax decreases.

So far, most models focus on the case of a (constant) uniform tax on all transactions. [Spahn \(2002\)](#) contains the early suggestion to couple the size of the tax rate to the current price volatility. For instance, during an outburst of volatility a small transaction tax may not be sufficient to prevent destabilizing speculation. A mechanism which automatically increases the tax rate in such a situation clearly appears to be worth further research.¹

The goal of our paper is to systematically reconcile the apparently contrasting results provided in the aforementioned literature. For this reason, we develop a model along the lines of [Chiarella and Iori \(2002\)](#) (CI), in which agents rely on a blend of technical, fundamental and random trading strategies. However, if past trading generates losses, a trader may also (temporarily) retreat from the market. The price adjustment is modelled in both a continuous double auction and in a dealership environment. Both settings have the potential to generate reasonable price dynamics. Therefore, our model allows a comparison of the implications of transaction taxes within different institutional market settings.

Our simulation experiments reveal that the consequences of transaction taxes depend on the total liquidity provided. When abundant exogenous liquidity is provided, the tax is stabilizing. This result, however, does not hold in other settings, as in the presence of market protocols where liquidity is endogenous and fluctuating. Most of the theoretical work, e.g. [Westerhoff and Dieci \(2006\)](#), is indeed based on a market-maker scenario in which infinite liquidity is provided. On the other hand, most empirical work collects data from more realistic market structures where liquidity imperfections are amplified by the tax, resulting in little effect or even in a deterioration of market quality. More subtly, our work suggests that finer details in the functioning of the market may ultimately decide whether the introduction of a Tobin tax has a stabilizing effect. Levying a transaction fee always reduces the volume, which, in turn, is only harmful if liquidity is affected and this triggers an increment in volatility. However, if market-makers offer liquidity widely, the tax is effective and produces less volatility. Surprisingly, the imposition of a transaction tax does not have a clear relation to the distortion, i.e. to the average absolute deviation between prices and fundamentals.

The paper is organized as follows. In Section 2, we discuss how market participants determine their orders. Our setup is related to the one presented by CI, yet we additionally assume that the market entry decision of traders is endogenous and depends on profit considerations. In Section 3, we introduce the market protocols. Order induced price adjustments take place either within a continuous double auction market or within a dealership market. In Section 4, we present our results. The last section concludes this paper and one [Appendix A](#) provides additional information on the simulated time series.

¹ We thank an anonymous referee for pointing out this aspect to us.

2. The model

We consider a market for one risky asset (stock) and cash. We assume that the interest rate $r = 0$ or, equivalently, that interest rate payments are spent elsewhere. The market is populated by N agents, whose initial endowments are S_{i0} and C_{i0} units of stock and cash, respectively. The agents' holdings are updated in the obvious way whenever there is a transaction. They cannot go short or lend/borrow money. There are multiple trading sessions (days) and every agent is selected in random order within a single day.

Agents record their trading performance and accordingly switch at the end of the day to an active or idle state. $X_i(t) \in \{A, I\}$ denotes the state of the i -th agent at time t . Inspired by the work of CI, each agent at time t , if active, submits an order to the market based on an estimate of the unknown end-of-day return. In detail, in a given instant τ of day t ($t < \tau < t + 1$), the i -th agents computes

$$r_{it}^{\text{close}} = g_1^i \frac{p^f - p_\tau}{p_\tau} + g_2^i \bar{r}_{L_i} + n_i \epsilon_{i\tau},$$

where p^f is the fundamental value, the weights $g_1^i > 0$, g_2^i , n_i represent the fundamental, chartist (trend-chasing or contrarian if g_2 is positive or negative, respectively) and random component, $\epsilon_{i\tau} \sim N(0, 1)$ and \bar{r}_{L_i} is the average past return over a time span of length $L_i \in \{1, \dots, L_{\max}\}$. As in CI, we assume that p^f is constant, known to all agents and that

$$g_1^i \sim |N(0, \sigma_1)|;$$

$$g_2^i \sim N(0, \sigma_2);$$

$$n^i \sim N(0, \sigma_n);$$

$$\bar{r}_{L_i} = \frac{1}{L_i} \sum_{j=1}^{L_i} \log p_{t-j}^{\text{close}} - \log p_{t-j-1}^{\text{close}},$$

where p_t^{close} is the closing price in session t . The weights g_1^i , g_2^i , n^i and length L_i are independently sampled only once from the respective distributions at the beginning of the simulations, and are never changed. Equipped with r_{it}^{close} , the agent can compute the expected end-of-day price as $\hat{p}_{it}^{\text{close}} = p_\tau \exp(r_{it}^{\text{close}})$.

Agents are hence heterogeneous in their own blending of different forecasting methods and in the extent of past close-to-close returns they take into account. Alternatively, we can interpret their random component as coming from liquidity shocks with no link to any strategic trading behavior.

We assume that agents can submit a unique limit order for one unit of the asset per day. A limit order is a couple, quantity-limit price (q_τ, l_τ) submitted at a randomly selected instant τ of day t . Each active trader posts an order $(q_{i\tau}, l_{i\tau})$, where

$$q_{i\tau} = \text{sgn}(\hat{p}_{it}^{\text{close}} - p_{t-1}^{\text{close}}),$$

is depending on the difference between the forecast and the last available closing price and

$$l_{i\tau} = \begin{cases} \lfloor \hat{p}_{it}^{\text{close}}(1 - \kappa_i) \rfloor & \text{if } q_{i\tau} > 0; \\ \lceil \hat{p}_{it}^{\text{close}}(1 + \kappa_i) \rceil & \text{if } q_{i\tau} < 0, \end{cases}$$

where κ_i is an individual aggressiveness parameter that reduces the bid when buying and increases the ask when selling. We independently sample individual κ_i values from a uniform distribution $U[0, K]$. Observe that all limit prices are integers, as required by the trading protocols that will be described in the sequel.

The active agents that were able to trade one unit during the day at some price $p_{i\bar{\tau}}$, $t < \bar{\tau} < t + 1$ compute their myopic profits

$$\pi(i, t) = \begin{cases} p_t^{\text{close}} - p_{i\bar{\tau}} - \text{tax} & \text{if the } i\text{-th agent is a buyer;} \\ p_{i\bar{\tau}} - p_t^{\text{close}} - \text{tax} & \text{if the } i\text{-th agent is a seller,} \end{cases}$$

where tax is a parameter denoting the fee that must be paid for each transaction. Then the same agents (that succeeded in trading) adjust an individually smoothed profit measure as

$$U_{i,t} = (1 - \eta)\pi(i, t) + \eta U_{i,t-1}.$$

With probability μ a trader considers to change strategy at the end of each day, when the closing price becomes available and evaluation of profits is possible. Agents who consider to switch strategies move to the active state with individual

probability

$$\phi_{i,t} = \frac{\exp(U_{i,t}/b)}{\exp(U_{i,t}/b) + \exp(0)}$$

Hence the future state of an agent is (conditional on switching being considered) given by

$$X_i(t+1) = \begin{cases} A & \text{with probability } \phi_{i,t}; \\ I & \text{with probability } 1 - \phi_{i,t}. \end{cases}$$

Observe that switching is based on profitability in both states: $U_{i,t}$ denotes the gains in the active state while 0 is obviously the gain in the idle state where no trade can occur. The bigger the gains in the active state, the higher is the probability to stay (or switch to) active. Here $b > 0$ is a parameter related to the intensity of switching: a high value of b makes agents insensitive to profits and prone to be idle or active with equal probability. On the contrary, a low b value will make them more likely to switch to the most profitable state at time $t + 1$.

This model of behavior assumes that the imposition of a tax pushes traders to the idle state by shrinking their gains. This is obviously an over-simplified mechanism in that agents could, say, incorporate the tax losses in their forecasts. However, this straightforward mechanism is used in a number of relevant works, see Brock and Hommes (1998), and is close to the basic intuition that (too) active agents should move to the idle state due to the tax burden.

2.1. Timing

It is useful to recap in detail how the model functions. A typical trading day (t -th day) develops as follows:

- (1) At t^- the closing price p_{t-1}^{close} of the previous session, the past returns' averages \bar{r}_{L_i} and the states of agents are known.
- (2) Beginning at time t^+ , agents trade at random times $t < \tau < t + 1$, submitting their orders to the market (with no certainty that they will be executed).
- (3) In $(t + 1)^-$ the closing price p_t^{close} for day t is known and active-successful traders can hence compute profits $\pi(i, t)$ and adjust their performance measure $U_{i,t}$. Note that if an agent is unable to trade, his/her U remains unchanged. This occurs for all idle agents.
- (4) The probability $\phi_{i,t}$ of becoming active is computed for all agents that could switch to another state to be used starting at time $(t + 1)^+$.

3. The market protocols

We consider different market architectures, namely a Continuous Double Auction (CDA) and a Dealership (Dea). Both protocols have some common features: they are organized in trading sessions (days) where agents can sequentially (in random order) submit bid and ask limit orders; at the end of the day every outstanding order is cleared. Prices and acceptable orders are quoted using a minimum tick which we assume, without loss of generality, to be 1. Hence, prices are integers.

3.1. Continuous Double Auction

The CDA is a widespread market protocol where agents place orders on separate buying and selling books. Bids (asks) are sorted in decreasing (increasing) order according to price–time priority. The largest outstanding bid is called the *best bid* and the smallest outstanding ask is called the *best ask*. If a new bid (ask) is not smaller (greater) than the best bid (ask), then the order is *marketable* and a transaction takes place at the price in the book for a unit quantity. If the incoming order is not marketable, it is inserted in the proper book for future use. Orders are canceled only when a counterpart is found or when the trading session is over.

3.2. Dealership

The dealership is a market protocol where all trades are executed by a specialist who posts at any time bids and asks (called *quotes*) valid for a unit transaction. When an agent has the chance to trade, he or she checks the dealer's quotes and if one of the two is acceptable a transaction occurs at the quoted price. If this is not the case, the agent's order is "lost" and does not become available to other agents. As the dealer is the counterpart of every trade, his inventory must be kept under control. We assume that this is done by an automated and non-strategic rule: whenever a transaction takes place, the dealer adjusts both quotes by a random integer δ , increasing prices if it was a selling transaction or decreasing quotes if he or she was a buyer. The offset δ is given by

$$\delta = \lfloor U(1, \Delta) \rfloor,$$

where $U(1, \Delta)$ denotes a uniform sample in the interval $[1, \Delta]$. Hence the dealers' behavior is completely described by two parameters, namely the fixed spread Γ between quotes and Δ . The latter value will be tuned in the sequel to obtain time

series that are somehow comparable to those we obtained in a CDA. The model of the dealership is clearly a very simple one and we are aware that in realistic markets speculators may adapt their trading behavior to the actual trading environment. However, this mechanism simply captures some features of setups where plentiful liquidity is exogenously provided.

4. Results

We consider several computational experiments. Given a market protocol, each experiment is a batch of 100 simulations (5000 trading days), where all parameters are fixed, with the exception of b and tax . This is meant to obtain 100 time series of 4999 returns, uniformly sampling b in $\{1, 2, \dots, 10\}$ and tax in $\{0, 1, 2, 3, 4\}$ in order to have representative data on a variety of settings with regard to the transaction tax imposed on agents and on their profit sensitiveness.

The parameters used in the simulations are provided in Table 1.

The values for the environmental and behavioral parameters are inspired by the aforementioned CI paper in order to obtain realistic series of returns. Hence, we vary systematically only the two parameters b , related to switching, and tax , the magnitude of the levy, together with the size N of the market. In both CDA and Dea a unique reference price is not defined, due to the presence of bids and asks. Hence, in all following statistical analyses, we have used the mid-price (midpoint of best bid and ask). Most of the time series we generated show well-known stylized facts, such as non-normal, fat-tailed and non-autocorrelated returns. This is remarkable, since the series are obtained with varying levels of tax and b . More details are provided in Appendix A.

We then consider three important indicators, namely volume, volatility (standard deviation of log returns) and distortion with respect to the fundamental value p_f . The effects of the imposition of a specific level of a TT can be estimated regressing the dependent variables volume (volatility and distortion, respectively) against the independent variables tax and b . In other words, we estimate the models

$$\begin{aligned} Volume_i &= k_1 + \alpha_1 tax_i + \beta_1 b_i, \\ Volatility_i &= k_2 + \alpha_2 tax_i + \beta_2 b_i, \\ Distortion_i &= k_3 + \alpha_3 tax_i + \beta_3 b_i, \end{aligned}$$

using the whole sample ($1 \leq i \leq 100$) and the three subsamples for which $b_i \in \{1, 2, 3\}$, $b_i \in \{4, 5, 6, 7\}$ and $b_i \in \{8, 9, 10\}$ corresponding to strong, medium, low short-term profit sensitivity, respectively.

It is useful to recall at this point the expected outcomes of a TT. The advocates of the imposition of a tax, although with differentiated nuances, claim that this should decrease volatility, presumably deterring the most speculative traders from taking part in the market. A reduction in volume should be observed together with more informative prices that might be more tightly linked to the fundamental values, due to low excess volatility.

The following two sections presents the results in a CDA and in a dealership.

4.1. Continuous Double Auction (CDA)

Table 2 exhibits the regression results for the volume in a CDA. As revealed in Panel A ($N = 1000$ traders), an increase in the transaction tax reduces volume significantly. The reduction in volume is most pronounced when b is low (-2.31 for $b \in \{1, 2, 3\}$ versus -0.87 for $b \in \{8, 9, 10\}$), a situation in which traders quickly react to the changes in the profitability of their trades. In addition, we see that volume increases with b . Hence, if b increases, more traders become active and thus volume grows. The other panels report findings for different values of N . In particular, we see that if the number of traders

Table 1
Values of the parameters used in the simulation, with a brief description.

Param.	Value	Description
T	5000	Trading days per simulation
p_0^{close}	1000	Initial price
p_f	1000	Constant fund. value
N	1000	Number of agents (base case)
C_{0i}	50,000	Initial cash endowment
S_{0i}	$\{1, \dots, 9\}$	Random initial stock endowment
l_{max}	20	Max. length of time window
σ_1	1/250	Sd of individual fund. component
σ_2	1.4/250	Sd of ind. chartist component
σ_n	3/250	Sd of ind. noisy component
K	0.2	Global aggressiveness
η	0.95	Profit smoothing parameter
μ	0.5	Probability of revising state
b	$\{1, \dots, 10\}$	Switching coefficient
tax	$\{0, 1, 2, 3, 4\}$	Tax
Γ	6	Constant spread (Dealership)
Δ	(varies with N)	Offset used by the dealer

Table 2

Volume in a CDA (units traded). The estimates are relative to different numbers of agents N (Panels A–E) and subsamples (low, medium, high, all b s). Regression t -values are in parentheses and entries are in boldface (italic) if statistically significant at the 1% (5%) confidence level.

	$1 \leq b \leq 3$	$4 \leq b \leq 7$	$8 \leq b \leq 10$	All b s
Panel A: $N = 1000$				
k_1	5.20 (10.09)	9.24 (9.15)	14.80 (4.08)	8.01 (16.73)
α_1	-2.31 (-15.88)	-1.21 (-7.50)	-0.87 (-3.85)	-1.24 (-8.97)
β_1	2.82 (12.44)	0.62 (3.66)	-0.16(-0.42)	0.74 (11.34)
Panel B: $N = 2000$				
k_1	10.73 (10.57)	22.72 (15.69)	25.16 (6.15)	16.11 (17.91)
α_1	-4.64 (-15.65)	-2.74 (-13.46)	-1.51 (-6.86)	-2.72 (-10.95)
β_1	5.61 (12.76)	0.87 (3.29)	0.27(0.60)	1.71 (13.45)
Panel C: $N = 4000$				
k_1	20.95 (9.99)	39.98 (15.54)	35.81 (4.49)	32.87 (18.12)
α_1	-8.81 (-17.66)	-4.95 (-14.12)	-3.23 (-8.21)	-5.40 (-11.16)
β_1	11.43 (14.07)	2.89 (6.27)	2.39(2.70)	3.60 (14.20)
Panel D: $N = 500$				
k_1	2.85 (5.86)	4.21 (6.71)	7.29 (5.50)	4.08 (13.94)
α_1	-1.24 (-9.82)	-0.55 (-6.40)	-0.47 (-4.13)	-0.67 (-8.59)
β_1	1.38 (6.00)	0.33 (3.19)	-0.14(-0.91)	0.33 (8.14)
Panel E: $N = 250$				
k_1	1.12 (6.48)	1.92 (4.98)	2.60(2.28)	1.47 (10.77)
α_1	-0.40 (-9.30)	-0.12(-1.83)	-0.18(-2.67)	-0.23 (-5.77)
β_1	0.48 (7.03)	0.08(1.21)	0.02(0.18)	0.17 (9.16)

increases, the total volume increases and the impact of transaction taxes on volume becomes stronger. Only if the number of traders is very low ($N = 250$, Panel E) may the impact of transaction taxes on volume become insignificant.

Table 3 shows the results for volatility in a CDA as a function of tax and b for different numbers of agents N and different subsamples with respect to b . The rightmost part of Panel A shows that a unit increment of the TT decreases volatility by 0.02% (2 basis points), while the marginal effect of b is 0.037%. The reduction due to the tax is statistically significant at the 5% but not at the 1% confidence level. More importantly, the magnitude of the decrement is rather small in relative terms: given a volatility of the order of 100 basis points, a reduction of 2 basis points may be considered negligible. Even in the most reactive subsample (when $1 \leq b \leq 3$) of simulations populated by strongly profit-sensitive traders, the decrease does not exceed 5 basis points (Panel A, leftmost part).

The same result holds in markets with many agents, see Panels B and C: the tax has a slightly greater effect and the coefficient increases from -0.020 to about -0.030 on the whole sample. Again, the leftmost part of the panels show increased efficacy for low b s but still the achievable reduction of volatility appears to be rather small in relative terms. Fig. 1 (left panel)

Table 3

Volatility in a CDA (%). The estimates are relative to different numbers of agents N (Panels A–E) and subsamples (low, medium, high, all b s). Regression t -values are in parentheses and entries are in boldface (italic) if statistically significant at the 1% (5%) confidence level.

	$1 \leq b \leq 3$	$4 \leq b \leq 7$	$8 \leq b \leq 10$	All b s
Panel A: $N = 1000$				
k_2	1.01 (17.99)	1.29 (17.23)	1.75 (5.74)	1.20 (37.76)
α_2	-0.049 (-3.08)	-0.015(-1.28)	-0.034(-1.77)	-0.020 (-2.18)
β_2	0.143 (5.75)	0.025(1.99)	-0.024(-0.75)	0.037 (8.59)
Panel B: $N = 2000$				
k_2	0.98 (31.52)	1.34 (25.42)	1.54 (12.40)	1.18 (44.12)
α_2	-0.070 (-7.72)	-0.027 (-3.66)	-0.004(-0.63)	-0.029 (-3.95)
β_2	0.165 (12.21)	0.028 (2.97)	-0.052(-0.381)	0.046 (12.03)
Panel C: $N = 4000$				
k_2	0.91 (29.96)	1.31 (28.13)	1.43 (9.60)	1.18 (40.39)
α_2	-0.068 (-9.41)	-0.014 (-2.17)	-0.019 (-2.54)	-0.030 (-3.88)
β_2	0.195 (16.48)	0.031 (3.77)	0.013(0.75)	0.050 (12.25)
Panel D: $N = 500$				
k_2	1.08 (14.27)	1.27 (11.11)	1.44 (6.10)	1.25 (30.85)
α_2	-0.043 (-2.20)	-0.030(-1.89)	-0.018(-0.87)	-0.026 (-2.38)
β_2	0.122 (3.42)	0.027(1.44)	-0.002(-0.06)	0.025 (4.50)
Panel E: $N = 250$				
k_2	1.21 (16.42)	1.40 (10.76)	1.88 (4.750)	1.34 (31.12)
α_2	-0.009(-0.51)	0.004(0.19)	-0.024(-1.02)	-0.009(-0.69)
β_2	0.067 (2.31)	-0.010(-0.45)	-0.053(-1.26)	0.003(0.44)




Fig. 1. Volatility in a CDA and in a dealership as a function of tax, with 2000 agents, for $b \in \{1, 2, 3\}, \{4, 5, 6, 7\}, \{8, 9, 10\}$ (from left to right).

depicts how the results (for $N = 2000$) are dependent on the values of parameter b , which is related to short-term profit sensitivity. In agreement with intuition, the effectiveness of the tax is visibly larger for small values of b .

Moreover, markets with few agents (Table 3, Panels D and E) are “thin” and there is little evidence of statistically significant effects, if at all. Additional comments on the importance of liquidity are deferred to the following section.

Table 4 presents the regression results obtained for the (percent) distortion, a measure of deviation from the fundamental value computed as

$$\frac{100}{T} \sum_{t=1}^T \left| \frac{p_t - p_f}{p_f} \right|.$$

Despite the common claim that a Tobin tax may help reduce distortion, we find no evidence of this effect. Fig. 2 (left panel) depicts the situation in a CDA with $N = 1000$ agents. It is apparent that a TT is unable to reduce the distortion and, in some cases, appears to mildly increase the deviation from the fundamental.

Table 4

Distortion in a CDA (%). The estimates are relative to different numbers of agents N (Panels A–E) and subsamples (low, medium, high, all b s). Regression t -values are in parentheses and entries are in boldface (italic) if statistically significant at the 1% (5%) confidence level.

	$1 \leq b \leq 3$	$4 \leq b \leq 7$	$8 \leq b \leq 10$	All b s
Panel A: $N = 1000$				
k_3	6.10 (7.71)	5.36 (6.75)	6.97 (2.00)	5.70 (17.16)
α_3	0.260 (1.165)	<i>0.301</i> (2.378)	-0.011 (-0.050)	<i>0.197</i> (2.050)
β_3	-0.228 (-0.655)	0.019 (0.145)	-0.085 (-0.231)	0.007 (0.161)
Panel B: $N = 2000$				
k_3	5.23 (9.81)	4.53 (6.99)	5.29 (3.17)	5.27 (23.62)
α_3	0.262 (1.680)	-0.169 (-1.860)	-0.022 (-0.250)	-0.023 (-0.376)
β_3	-0.203 (-0.878)	0.238 (2.022)	0.048 (0.264)	0.052 (1.648)
Panel C: $N = 4000$				
k_3	4.18 (7.69)	4.95 (8.51)	2.38 (1.25)	4.72 (22.60)
α_3	0.019 (0.149)	-0.034 (-0.435)	-0.099 (-1.055)	-0.032 (-0.579)
β_3	0.247 (1.172)	0.039 (0.375)	0.319 (1.503)	<i>0.060</i> (2.041)
Panel D: $N = 500$				
k_3	6.09 (10.15)	6.98 (7.39)	0.51 (0.12)	6.38 (15.54)
α_3	0.721 (4.616)	0.002 (0.018)	0.207 (0.570)	<i>0.229</i> (2.109)
β_3	-0.241 (-0.849)	-0.034 (-0.222)	0.715 (1.485)	0.023 (0.415)
Panel E: $N = 250$				
k_3	11.56 (8.20)	9.96 (5.95)	13.00 (2.80)	9.97 (16.27)
α_3	0.076 (0.215)	0.097 (0.344)	0.218 (0.780)	0.199 (1.132)
β_3	-0.71 (-1.28)	-0.34 (-1.19)	-0.50 (-1.00)	-0.23 (-2.83)




Fig. 2. Distortion in a CDA and in a dealership as a function of tax, with 1000 agents, for $b \in \{1, 2, 3\}, \{4, 5, 6, 7\}, \{8, 9, 10\}$ (from left to right).

4.2. Dealership

In this section we examine the effect of a TT in a dealership, according to the same measures (volume, volatility and distortion) used in a CDA. We keep the behavior of the traders fixed changing the market. The experiment is intended to test whether the results are stable across different microstructures and, in particular, if the addition of an exogenous liquidity source (the dealer) alters our findings. We fix $\Gamma = 6$ (constant bid-ask spread) and vary $\Delta = \Delta(N)$ (variation of quote) depending on the number of agents, in such a way as to get a volatility of the same order of CDA. We stress that this “calibration” exercise is hard and imperfect, as we try to align the 100 time series (one for each simulation) obtained across different parameters in institutionally different markets. By trial and error, we obtain values for Δ such that the average volatility in a dealership is similar but never exceeds the average value in the “parallel” CDA. It emerges that this volatility-based calibration also allows us to attain roughly similar volumes and distortions in the two markets. Table 5 shows descriptive statistics of the volatility in the two markets for various sizes of trader populations.

Table 5 reveals that at least 50% of the simulations (between the first and third quartile) in each panel are close in terms of volatility. The average standard deviations of returns in a CDA and in a dealership rarely differ by more than 20 basis points, despite greater differences in some extreme occasions. Read with liberty, Table 5 justifies the claim that we can reasonably

Table 5

Comparison of volatilities (%) obtained in a CDA and in a dealership with the same number of agents N . The parameter Δ varies to achieve comparable volatility with the CDA over the whole sample.

	Min.	1Qu.	Median	Mean	3Qu.	Max.
Panel A: $N = 1000, \Delta = 6$						
CDA	0.954	1.286	1.384	1.367	1.478	1.694
Dea	0.4782	1.0710	1.1840	1.1720	1.2750	1.8300
Panel B: $N = 2000, \Delta = 4$						
CDA	0.8261	1.3400	1.4340	1.3780	1.5000	1.6760
Dea	0.3434	1.0250	1.3090	1.2480	1.5290	2.2630
Panel C: $N = 4000, \Delta = 2.25$						
CDA	0.8111	1.3320	1.4380	1.3880	1.5000	1.6370
Dea	0.4153	1.0850	1.2820	1.2400	1.4590	1.9470
Panel D: $N = 500, \Delta = 11$						
CDA	0.9784	1.2310	1.3140	1.3360	1.4480	1.9100
Dea	0.4295	1.1190	1.2720	1.2710	1.4270	2.0640
Panel E: $N = 250, \Delta = 15$						
CDA	1.024	1.229	1.347	1.337	1.434	1.925
Dea	0.5837	1.0190	1.1290	1.1380	1.2430	1.8290

Table 6

Volatility in a dealership (%). The estimates are relative to different numbers of agents N (Panels A–E) and subsamples (low, medium, high, all b s). Regression t -values are in parentheses and entries are in boldface (italic) if statistically significant at the 1% (5%) confidence level.

	$1 \leq b \leq 3$	$4 \leq b \leq 7$	$8 \leq b \leq 10$	All b s
Panel A: $N = 1000$				
k_2	0.97 (7.79)	1.33 (12.88)	1.40 (6.73)	1.12 (25.34)
α_2	-0.225 (-6.992)	-0.063 (-3.438)	-0.066 (-4.516)	-0.103 (-7.381)
β_2	0.200 (3.082)	-0.003 (-0.129)	-0.003 (-0.131)	0.040 (6.502)
Panel B: $N = 2000$				
k_2	0.74 (12.87)	1.20 (7.34)	1.30 (3.73)	1.06 (18.42)
α_2	-0.198 (-14.383)	-0.146 (-6.219)	-0.065 (-2.962)	-0.149 (-9.394)
β_2	0.283 (13.397)	0.086 (3.007)	0.036 (0.945)	0.094 (12.138)
Panel C: $N = 4000$				
k_2	0.61 (6.61)	1.13 (12.58)	1.52 (6.47)	1.01 (23.36)
α_2	-0.17 (-9.10)	-0.15 (-11.92)	-0.13 (-8.22)	-0.15 (-13.14)
β_2	0.269 (6.196)	0.085 (5.322)	0.019 (0.719)	0.091 (14.644)
Panel D: $N = 500$				
k_2	1.21 (12.37)	1.27 (8.61)	0.58 (1.40)	1.36 (28.50)
α_2	-0.164 (-7.551)	-0.080 (-3.881)	-0.089 (-3.848)	-0.109 (-7.891)
β_2	0.129 (3.081)	0.043 (1.558)	0.098 (2.179)	0.024 (3.606)
Panel E: $N = 250$				
k_2	1.35 (11.58)	0.96 (9.38)	0.67 (2.47)	1.21 (26.63)
α_2	-0.199 (-5.748)	-0.072 (-5.004)	-0.003 (-0.164)	-0.069 (-5.174)
β_2	0.065 (1.271)	0.053 (2.955)	0.054 (1.867)	0.012 (1.843)

compare across different markets the effects of the introduction of a TT, given that the simulations produced are, to the best of our efforts, quite similar.²

Table 6 contains the results of the regression of the volatility in a dealership against the usual independent variables tax , b and is exactly homologous with Table 3. The marginal effect of the TT is strongly statistically significant, even in thin markets (Panels D and E). Observe that this holds despite the fact that there is “less to reduce” in a dealership whose (average) volatility is never greater than in a CDA. More fundamentally from a practical point of view, trimming volatility is amplified by a factor of 5 or more over the whole sample (see “All b s”). Virtually all subsamples show that a sizeable reduction is possible when low levels of taxation are imposed ($tax = 1, 2$ roughly equivalent to 0.1–0.2 percent proportional taxation rate), as recently hypothesized in the debate on this topic. The right panel of Fig. 1 shows the reduction of volatility when $N = 2000$, for different values of b .

Tables 7 and 8 presents the results for volume and distortion in a dealership. Similarly to the findings observed in a CDA, the volume is highly affected in a dealership. Observe that this furnishes further support to the observation that reducing volatility is not tantamount to volume cutting. Hence, if a TT visibly decreases the volume in both CDAs and dealerships, some other driver is responsible for the different outcomes in terms of volatility. A first, perhaps trivial, explanation lies in the different role taken by the bid-ask spread in the two markets. While a reduction in the traded volume is likely to widen the bid-ask spread in a CDA, one of the features of our dealership is that a constant bid-ask is provided by the dealer at any time. This liquidity provision is not affected by the number of transactions, given the dealer’s duty to quote prices without discontinuity. Second, a careful inspection of the data shows that smaller volumes produce sparse order books in a CDA. A sequence of marketable limit orders of the same type can then escalate the book, producing wide price changes. By definition, these liquidity holes are missing in dealerships, resulting in smoother price dynamics under reduced volume. Somewhat related comments on microstructures have been given, for example, in Farmer et al. (2004), Lillo and Farmer (2005) and Lillo et al. (2003).

Our simulation results confirm these interpretations. Fig. 3 shows the time-average difference of the active traders across all simulations in a dealership and in a CDA with $N = 1000$ traders. Positive values mean that there are more active agents in a dealership and the left panel of the figure shows, for different levels of sensitivity, that more agents move to the idle state in a CDA. The difference is small for large b (bottom curve) but it amounts to 1% of the number of traders in profit-sensitive markets. The right graph shows the same difference when there is no, low and high tax. The number of active agents is roughly the same in the two markets when no tax is imposed but a levy crowds out more agents in a CDA than in a dealership in the other cases. Hence, a dealership market has a larger liquidity provision by the specialist as well as more active traders.

² Observe that we calibrate a single parameter in the dealership, Δ , to “match” a single average value for volatility in a CDA. Table 5, however, shows a comparison of two distributions of values. Even forgetting the considerable differences between the two market clearing mechanisms, it is not surprising that it is difficult to obtain a closer match.

Table 7

Volume in a dealership (units traded). The estimates are relative to different numbers of agents N (Panels A–E) and subsamples (low, medium, high, all b s). Regression t -values are in parentheses and entries are in boldface (italic) if statistically significant at the 1% (5%) confidence level.

	$1 \leq b \leq 3$	$4 \leq b \leq 7$	$8 \leq b \leq 10$	All b s
Panel A: $N = 1000$				
k_1	7.63 (5.80)	12.02 (9.02)	14.55 (5.98)	9.79 (18.40)
α_1	-2.85 (-8.37)	-0.82 (-3.42)	-0.92 (-5.40)	-1.36 (-8.02)
β_1	2.76 (4.03)	0.12 (0.48)	-0.07 (-0.27)	0.59 (7.96)
Panel B: $N = 2000$				
k_1	9.06 (9.34)	17.75 (7.20)	23.06 (4.15)	14.53 (14.43)
α_1	-3.72 (-16.07)	-2.64 (-7.48)	-1.01 (-2.87)	-2.69 (-9.65)
β_1	5.26 (14.85)	1.64 (3.80)	0.31 (0.51)	1.86 (13.70)
Panel C: $N = 4000$				
k_1	11.85 (4.07)	34.14 (15.99)	43.55 (9.00)	27.59 (18.84)
α_1	-6.34 (-11.01)	-4.59 (-15.52)	-3.87 (-12.27)	-4.84 (-12.83)
β_1	10.94 (8.01)	2.63 (6.91)	0.94 (1.74)	3.33 (15.92)
Panel D: $N = 500$				
k_1	5.79 (8.62)	6.19 (5.73)	3.68 (1.40)	6.84 (20.52)
α_1	-1.22 (-8.18)	-0.64 (-4.24)	-0.57 (-3.89)	-0.80 (-8.31)
β_1	0.92 (3.20)	0.35 (1.71)	0.42 (1.48)	0.18 (3.77)
Panel E: $N = 250$				
k_1	4.68 (7.45)	3.31 (5.94)	1.07 (0.72)	4.37 (17.43)
α_1	-1.21 (-6.48)	-0.54 (-6.94)	-0.09 (-0.88)	-0.49 (-6.67)
β_1	0.60 (2.16)	0.26 (2.66)	0.32 (2.03)	0.07 (1.90)

Table 8

Distortion in a dealership (%). The estimates are relative to different numbers of agents N (Panels A–E) and subsamples (low, medium, high, all b s). Regression t -values are in parentheses and entries are in boldface (italic) if statistically significant at the 1% (5%) confidence level.

	$1 \leq b \leq 3$	$4 \leq b \leq 7$	$8 \leq b \leq 10$	All b s
Panel A: $N = 1000$				
k_3	5.56 (3.85)	5.81 (3.09)	8.52 (1.62)	7.41 (11.02)
α_3	-0.50 (-1.34)	-0.20 (-0.58)	-0.51 (-1.40)	-0.42 (-1.96)
β_3	1.553 (2.067)	0.357 (1.019)	0.077 (0.129)	0.175 (1.868)
Panel B: $N = 2000$				
k_3	7.49 (9.23)	5.81 (6.55)	4.22 (1.23)	6.28 (16.47)
α_3	0.201 (1.036)	-0.242 (-1.903)	0.110 (0.507)	0.033 (0.311)
β_3	-0.738 (-2.493)	0.122 (0.787)	0.232 (0.614)	-0.008 (-0.145)
Panel C: $N = 4000$				
k_3	6.60 (7.57)	5.11 (7.84)	5.38 (4.56)	5.26 (20.96)
α_3	0.09 (0.52)	-0.14 (-1.59)	-0.34 (-4.49)	-0.16(-2.43)
β_3	-0.938 (-2.292)	0.055 (0.475)	0.066 (0.499)	0.033 (0.920)
Panel D: $N = 500$				
k_3	9.43 (4.31)	8.68 (3.84)	17.39 (1.40)	9.38 (9.93)
α_3	-0.163 (-0.334)	-0.179 (-0.567)	0.269 (0.390)	-0.046 (-0.169)
β_3	0.028 (0.030)	0.209 (0.493)	-0.948 (-0.706)	0.001 (0.005)
Panel E: $N = 250$				
k_3	14.13 (4.60)	20.30 (3.28)	5.21 (1.39)	14.90 (10.19)
α_3	-1.62 (-1.78)	-0.79 (-0.91)	-0.36 (-1.45)	-0.86(-2.02)
β_3	0.24 (0.17)	-1.22 (-1.13)	0.48 (1.20)	-0.40(-1.99)

It is also of interest to understand which agents are induced to the idle state by the imposition of a tax. Our data show that the probability of being active for the i -th agents increases notably with κ_i , when some tax is imposed.³ This is reasonable as small mark-ups (mark-downs) are likely to result in more aggressive orders that on the one hand increase the probability to trade and the consequential tax burden and, on the other hand, reduce the gain from trading. By contrast, agents with large κ_i trade less but, conditional on trading, gain more given that their limit prices are advantageous (low bids and high asks). The increased inactivity of low κ_i agents has no effect on the bid-ask spread in a dealership but has consequences in a CDA, where exactly the most aggressive traders, that would have posted limit orders close to the best bid and ask, retreat from the market and cause a drop in the available liquidity.

³ There is no or little evidence of dependence on g_1^i, g_2^i, n_i or l_i . Of course, if there is no tax, the probability to be active is not influenced by the values of the individual parameters.




Fig. 3. Averaged difference of active traders in a dealership and in a CDA. On the left, time-averages across all simulations are computed for $b \in \{1, 2, 3\}$, $\{4, 5, 6, 7\}$, $\{8, 9, 10\}$ (from top to bottom). The right panel depicts the difference for $tax = 0$, $tax \in \{1, 2\}$ and $tax \in \{3, 4\}$ (from bottom to top).

The imposition of a tax in a Dealership has again a somewhat weak (and often null) effect on distortion, see Fig. 2, right panel. This result may look puzzling at first, as one could imagine that a reduction in trading volume and price volatility has to lead to a decrease in the distortion. However, while a decline of chartist- and noise-driven orders should push prices towards the fundamental value, it should be kept in mind that the number of stabilizing orders of fundamental traders decreases too. A constant distortion then implies that the impact of these two opposing effects roughly balance.

5. Conclusion

In this paper, we explore the effectiveness of transaction taxes in different institutional market microstructure settings. Within our model, asset prices are driven by the orders of market participants. These agents rely on technical, fundamental and random trading strategies to determine their investment positions. However, they are not forced to trade. Their decision to be active or inactive depends on the past profitability of the agents' trades. Should a trader encounter losses, he or she may decide to stop trading (and vice versa). In one of our scenarios, the orders of market participants enter a continuous double auction market. In another scenario, the orders of market participants are filled by a specialist. Both settings are generally able to mimic some stylized facts of financial markets, and thus allow a comparison of the impact of transaction taxes on the market dynamics.

Our key findings can be summarized as follows:

- In a continuous double auction market, the imposition of a transaction tax presumably has little stabilizing impact on the market dynamics. We observe that traders retreat from the market if a levy has to be paid, leading to a decrease in volume. However, liquidity also decreases so that, on average, a given order obtains a larger price impact. This, in turn, counters or even eliminates the otherwise stabilizing effect of the transaction tax. The distortion in the market remains unaltered.
- Our model predicts that in a dealership market, a transaction tax has the power to stabilize market dynamics. Also in this environment, traders retreat from the market, and volume declines. However, since liquidity is exogenously provided by a specialist, the price impact of a given order remains constant, and thus volatility declines significantly. Again, the distortion is not significantly reduced.

Concluding, we find that market microstructure details matter for the effectiveness of transaction taxes. We hope that our paper resolves some of the confusion frequently observed in the debate on the outcomes of transaction taxes.

Acknowledgements

This paper was written and revised while Paolo Pellizzari was visiting the School of Finance and Economics at UTS, whose financial support is gratefully acknowledged. We thank an anonymous referee and associate editor for useful remarks and suggestions. We are also indebted to the audiences at the seminars held at Ca' Foscari University, University of Technology Sydney (UTS), University of New South Wales (UNSW) and Royal Melbourne Institute of Technology (RMIT). Paolo Pellizzari acknowledges the support of PRIN 2007 grant 2007TKLTSR, "Computational markets design and agent-based models of trading behavior". All the errors are the authors' responsibility.




Fig. 4. Price–time series and returns density of a representative CDA simulation ($N = 1000$, $b = 9$, $tax = 4$). The dashed lines show the fundamental value (left) and a normal distribution with same mean and variance (right).

Table 9

Descriptive statistics for the kurtosis (normalized central fourth moment) of returns in a CDA and in a dealership.

	Min.	1Qu.	Median	Mean	3Qu.	Max.
CDA ($N = 1000$)	3.910	4.450	4.743	4.855	5.178	7.334
Dea ($N = 2000$)	2.574	2.897	3.022	3.059	3.121	4.564

Appendix A. Stylized facts

This appendix aims to describe the time series that are analyzed in the paper corroborating the claim that our model is able to produce reasonably realistic returns. We show in the following that the data exhibit some of the common statistical features of financial time series that are frequently dubbed “stylized facts”. For brevity, in this appendix we comment only on simulations with 1000 agents in a CDA and 2000 in a dealership.

Fig. 4 depicts a representative price trajectory in a CDA, together with the density of the log returns. The price visually displays a random-like behavior with sudden bursts and crashes. The density is clearly non-gaussian, leptokurtic and fat-tailed. Formal tests reveal that normality can be strongly rejected and lagged returns are independent (the Shapiro–Wilk test p -value is smaller than 10^{-20} ; the Box–Pierce p -value with 5 lags is 0.60, meaning that independence cannot be rejected). This example is somewhat illustrative of the whole sample obtained when the market is a CDA: normality of returns is (strongly) rejected for all simulations and the null hypothesis of linear independence of returns is rejected at the 1% confidence level by the Box–Pierce test in 8 cases out of 100.

The simulated returns in CDA show a fair amount of excess kurtosis (with respect to the gaussian value $\mu_4 = 3$), see the upper part of Table 9.

The time series obtained in a dealership are slightly less satisfactory, in that returns are closer to normality⁴ and there is a weak degree of linear predictability in some cases. The observation that some market mechanisms (like the CDA) may make the presence of stylized facts more likely is not new, see Pellizzari and Dal Forno (2007) and Maslov (2000). Fig. 5 shows a representative example of price and returns’ density. The price fluctuates quite realistically, with pronounced deviations from the fundamental. The hypothesis of normal distribution of returns cannot be rejected (the p -value is 0.075), confirming the visual proximity of the densities displayed in the right part of Fig. 5. Examining the whole sample, normality is rejected at the 5% confidence level by a Shapiro–Wilk test in 54 cases out of 100. As seen in the lower part of Table 9, there is little evidence of excess kurtosis when the market platform is a dealership. The findings are exactly the same for markets with different N .

The independence of returns is not rejected (at the 1% level) for 78 simulations. In the 22 cases where some linear structure is present in the returns, the strength of predictability is extremely low. Fig. 6 depicts the autocorrelation with the largest modulus for every simulation (35 lags are considered).

⁴ As suggested by an anonymous referee, one could assume that the fundamental value is not constant over time but follows a random process or experiences sudden shocks. Then, if a larger change materializes in the fundamental value, traders receive rather strong trading signals which may, in turn, trigger a volatility outburst. Such an outcome should also increase the tails of the distribution of the returns, implying a more pronounced deviation from normality. Although such an extension is very plausible we leave it open for future research. Within our setting one would, for instance, have to specify when the shock appears and how traders react to it.




Fig. 5. Price–time series and returns density of a typical simulation in a dealership ($N = 2000$, $b = 9$, $tax = 4$). The dashed lines show the fundamental value (left) and a normal distribution with same mean and variance (right).




Fig. 6. Autocorrelation with the largest absolute value for each of the 100 simulations of a dealership with $N = 2000$ agents.

On the one hand, the absolute magnitude of the autocorrelation rarely exceeds 0.05: this predictability is statistically significant at times but is rather low for providing trading gains. On the other hand, the prevalence of negative signs in the picture suggests that this weak autocorrelation may be due to the bid-ask bounce possibly occurring in a dealership.

References

- Aliber, R.Z., Chowdhry, B., Yan, S., 2003. Some evidence that a Tobin tax on foreign exchange transactions may increase volatility. *Review of Finance* 7, 481–510.
- Brock, W., Hommes, C., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* 22, 1235–1274.
- Chiarella, C., 1992. The dynamics of speculative behaviour. *Annals of Operations Research* 37, 101–123.
- Chiarella, C., Iori, G., 2002. A simulation analysis of the microstructure of double auction markets. *Quantitative Finance* 2, 346–353.
- Day, R., Huang, W., 1990. Bulls, bears and market sheep. *Journal of Economic Behavior and Organization* 14, 299–329.
- De Grauwe, P., Dewachter, H., Embrechts, M., 1993. *Exchange Rate Theory: Chaotic Models of Foreign Exchange Markets*. Blackwell, Oxford.
- Ehrenstein, G., Westerhoff, F., Stauffer, D., 2005. Tobin tax and market depth. *Quantitative Finance* 5, 213–218.
- Eichengreen, B., Tobin, J., Wyplosz, C., 1995. Two cases for sand in the wheels of international finance. *Economic Journal* 105, 162–172.
- Farmer, D., Gillemot, L., Lillo, F., 2004. What really causes large price changes? *Quantitative Finance* 4, 383–397.
- Farmer, J.D., Joshi, S., 2002. The price dynamics of common trading strategies. *Journal of Economic Behavior and Organization* 49, 149–171.
- Frankel, J.A., 1996. How well do foreign exchange markets function: might a Tobin tax help? NBER Working Papers 5422, National Bureau of Economic Research.
- 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.

- Hommes, C., 2006. Heterogeneous agent models in economics and finance. In: Tesfatsion, L., Judd, K. (Eds.), *Handbook of Computational Economics*, vol. 2: Agent-based Computational Economics. North-Holland, pp. 1107–1186.
- Jones, C., Seguin, P., 1997. Transaction costs and price volatility: evidence from commission deregulation. *The American Economic Review* 87, 728–737.
- Keynes, J., 1936. *The General Theory of Employment, Interest, and Money*. Harcourt Brace.
- Kirman, A., 1991. Epidemics of opinion and speculative bubbles in financial markets. In: Taylor, M. (Ed.), *Money and Financial Markets*. Blackwell, pp. 354–368.
- LeBaron, B., 2006. Agent-based computational finance. In: Tesfatsion, L., Judd, K. (Eds.), *Handbook of Computational Economics*, vol. 2: Agent-based Computational Economics. North-Holland, pp. 1187–1233.
- Lillo, F., Farmer, D., 2005. The key role of liquidity fluctuations in determining large price fluctuations. *Fluctuations and Noise Letters* 5, L209–L216.
- Lillo, F., Farmer, D., Mantegna, R.N., 2003. Econophysics: master curve for price-impact function. *Nature* 421, 129–130.
- Lux, T., 1995. Herd behaviour, bubbles and crashes. *Economic Journal* 105, 881–896.
- Lux, T., in press. Financial power laws: empirical evidence, models and mechanisms. In: Cioffi-Revilla, C. (Ed.), *Power Laws in the Social Sciences: Discovering Complexity and Non-equilibrium Dynamics in the Social Universe*.
- Mannaro, K., Marchesi, M., Setzu, A., 2006. Using an artificial financial market for assessing the impact of Tobin-like transaction taxes. *Journal of Economic Behavior and Organization* 67, 445–462.
- Maslov, S., 2000. A simple model of an order-driven market. *Physica A* 278, 571–578.
- Pellizzari, P., Dal Forno, A., 2007. A comparison of different trading protocols in an agent-based market. *Journal of Economic Interaction and Coordination* 2, 27–43.
- Rosser, J.B., Ahmed, E., Hartmann Jr., G., 2003. Volatility via social flaring. *Journal of Economic Behavior and Organization* 50, 77–87.
- Schwert, W., Seguin, P., 1993. Securities transaction taxes: an overview of costs, benefits and unresolved questions. *Financial Analysts Journal* 49, 27–35.
- Shiller, R., 2000. *Irrational Exuberance*. Princeton University Press.
- Spahn, B., 2002. On the feasibility of a tax on foreign exchange transactions. Technical Report. Federal Ministry for Economic Cooperation and Development, Bonn.
- Stiglitz, J., 1989. Using tax policy to curb speculative short-term trading. *Journal of Financial Services* 3, 101–113.
- Summers, L., Summers, V., 1989. When financial markets work too well: a cautious case for a securities transaction tax. *Journal of Financial Services* 3, 163–188.
- Taylor, M., Allen, H., 1992. The use of technical analysis in the foreign exchange market. *Journal of International Money and Finance* 11, 304–314.
- Tobin, J., 1978. A proposal for international monetary reform. *Eastern Economic Journal* 4, 153–159.
- Ul Haq, M., Kaul, I., Grunberg, I., 1996. *The Tobin Tax: Coping with Financial Volatility*. Oxford University Press.
- Umlauf, S., 1993. Transaction taxes and the behavior of the Swedish stock market. *Journal of Financial Economics* 33, 227–240.
- Werner, I., 2003. Comment on 'Some evidence that a Tobin tax on foreign exchange transactions may increase volatility'. *European Finance Review* 7, 511–514.
- Westerhoff, F., 2003. Heterogeneous traders and the Tobin tax. *Journal of Evolutionary Economics* 13, 53–70.
- Westerhoff, F., 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.