Comput Econ
DOI 10.1007/s10614-017-9649-9

The Impact of the Tobin Tax in a Heterogeneous Agent Model of the Foreign Exchange Market

Filip Stanek[1] · **Jiri Kukacka**[2,3]

Accepted: 13 January 2017

Abstract We explore possible effects of a Tobin tax on exchange rate dynamics in a heterogeneous agent model. To assess the impact of the Tobin tax in this framework, we extend the model of De Grauwe and Grimaldi (Eur Econ Rev 50(1):1–33, 2006) by including transaction costs and perform numerical simulations. Motivated by the importance of the market microstructure, we choose to model the market as being cleared by a Walrasian auctioneer. This setting could more closely resemble the two-layered structure of foreign exchanges at daily frequency than a price impact function, which is often adopted in similar studies. We find that the Tobin tax can deliver a moderate reduction of return volatility and kurtosis. In addition, simulations indicate that the Tobin tax reduces the degree of mispricing in the time series, which is primarily achieved by eliminating long-lasting deviations from fundamental value.

Keywords Tobin tax · Foreign exchange market · Agent-based modeling · Walrasian auctioneer

Support from the Czech Science Foundation (GACR) under the P402/12/G097 DYME Dynamic models in economics project and from the Grant Agency of the Charles University (GAUK) under project 588912 and project SVV 260 233 are gratefully acknowledged.

✉ Jiri Kukacka
jiri.kukacka@fsv.cuni.cz

Filip Stanek
filip.stanek@cerge-ei.cz

1 CERGE-EI, Politickych veznu 7, 110 00 Prague 1, Czech Republic

2 Institute of Economic Studies, Faculty of Social Sciences, Charles University, Smetanovo nabrezi 6, 111 01 Prague 1, Czech Republic

3 Institute of Information Theory and Automation, The Czech Academy of Sciences, Pod Vodarenskou vezi 4, 182 00 Prague 8, Czech Republic

Published online: 27 January 2017

1 Introduction

In 1972, James Tobin proposed a small, uniform tax on all foreign exchange transactions. Tobin argued that the absence of any consensus on fundamentals in foreign exchange markets in combination with low transaction costs and the limited rationality of market participants transforms the price discovery process into a "game of guessing what other traders are going to think" (Tobin 1978). However, if investors form their expectations at least partially based on the perceived expectations of other market participants, it would create positive feedback that may cause price misalignments and excessive volatility. Because most speculative transactions are made on a very short-term basis, Tobin believed that a small tax imposed on every transaction could dissuade most short-term speculators and consequently stabilize the market.

Many participants in the global economy find the promise of more predictable exchange rates appealing. Therefore, it is unsurprising that since 1972, a transaction tax levied on foreign exchange markets (henceforth referred to as the Tobin tax) has been frequently discussed. The most recent example is the Central Bank of China, which is considering adopting the Tobin tax to protect the Yuan against speculative capital flows (Li 2015, via Bloomberg). Despite the vivid public debate over the possible effects of the tax, academic scrutiny has remained relatively scant.

The most straightforward way to assess the impact of the Tobin tax is to examine some real world impositions of such a tax. Unfortunately, the time series necessary to do so are not available, as the Tobin tax has never been implemented. Aliber et al. (2003) nevertheless developed an innovative method for estimating transaction costs using foreign exchange futures and concluded that transaction costs (such as the Tobin tax) are positively associated with exchange rate volatility. In addition to the work of Aliber et al. (2003), several authors have also conducted studies addressing transaction costs on stock markets and obtained rather conflicting results (e.g., Umlauf 1993; Hau 2006; Liu and Zhu 2009). Interpreting these results is even more difficult because stock markets differ substantially from foreign exchange markets in terms of trading volumes and their microstructure. It is therefore not apparent whether these studies have any bearing on the Tobin tax.

Because the empirical evidence is sparse, researchers have employed various theoretical models to address the question of how the Tobin tax might alter exchange rate dynamics. Heterogeneous agent models have proven particularly fruitful in this regard, as such models are able to generate time series that are qualitatively similar to those observed in real markets. Ultimately, it would be non-sensical to examine the impact of the Tobin tax on trading volume or endogenous volatility using a model that is incapable of generating either. In this paper, we extend existing agent-based research regarding the Tobin tax by exploring the possible impacts of the tax in a market cleared by a Walrasian auctioneer—settings that, to the best of our knowledge, have yet to be examined in any extant study and that might be, as we will argue below, more realistic than other frequently adopted clearing mechanisms. This modification is motivated by recent research demonstrating the substantial importance of the market-clearing mechanism and liquidity provision when evaluating the impact of the Tobin tax.

There is broad consensus that the Tobin tax would reduce market depth, which may in turn increase volatility, as the price impact of a single order will be larger (e.g., Farmer et al. 2004). Proponents of the Tobin tax typically argue that because foreign exchange markets have relatively high depth, this effect will be negligible compared with the change in volatility originating from the change in the composition of the population of traders. While most studies explicitly model the structure of the trader population, few account for the taxs effect on volatility through market liquidity. This omission occurs because most studies explore the impact of the tax in a dealership market in which the market price is determined using a price impact function that is constant with respect to total trading volume. This danger of systematically overestimating the positive effect of the Tobin tax was noted by Ehrenstein et al. (2005), who demonstrated that under a more realistic price impact function[1] that decreases with respect to total trading volume,[2] the Tobin tax can have either negative or positive effects on volatility depending on how sensitive the price impact function is with respect to total trading volume. These results seriously challenge conclusions drawn from previous models.

Pellizzari and Westerhoff (2009) further supports the findings of Ehrenstein et al. (2005) by exploring the impact of the Tobin tax on a population of agents interacting in either a continuous double auction or in a dealership market. In a dealership market in which the market maker provides abundant liquidity, the Tobin tax reduces volatility. In a continuous double auction with endogenous liquidity provision, the otherwise stabilizing effect of the Tobin tax is offset by the reduction in market depth. The results of Pellizzari and Westerhoff (2009) were replicated in a laboratory experiment using human subjects by Kirchler et al. (2011), rendering the conclusions particularly sound.

In addition to studies focusing solely on the link between market microstructure and the impact of the Tobin tax, a variety of other research papers also indirectly support the claim that the models clearing mechanism and liquidity provision might be the key determinants of whether the Tobin tax can successfully stabilize the market. Because different authors naturally employ slightly different agent-based models, we can investigate the relationship between the structure of each model and its results. A majority of researchers have explored the impact of the Tobin tax using a dealership market framework with a price impact function that is constant with respect to total trading volume and, consistent with our expectations, found that the tax reduces volatility. The Tobin taxs ability to stabilize the market was reported, e.g., in Ehrenstein (2002), Westerhoff (2003, 2004b, 2008), Westerhoff and Dieci (2006), Demary (2006),[3] Bianconi et al. (2009), and Demary (2010). Most recently, Flaschel et al. (2015) confirmed the ability of the Tobin-like tax to stabilize the market in a broader macroeconomic setting.

While the impact of the Tobin tax appears to be clear when only considering dealership markets, we observe a notably different picture when examining agent-based models exploring the impact of the Tobin tax in the continuous double auction frame-

[1] For an argument regarding why this impact function is more appropriate than one constant with respect to total volume, see Farmer et al. (2004).

[2] I.e., the price impact of a single order is determined by the size of that order relative to total trading volume—the effect of liquidity is therefore included.

[3] These authors, however, reported an increase in the kurtosis of returns.

work, which allows for endogenous liquidity provision. Mannaro et al. (2005) and Mannaro et al. (2008) demonstrated that under a double auction-like mechanism, the Tobin tax increases price volatility. Further, considering a similar setting, Lavička et al. (2013) reported that the tax increases volatility while simultaneously reducing the kurtosis of returns. Results more in line with the expectations of Tobin tax proponents were obtained by Hein et al. (2006).

Motivated by the evident importance of the market microstructure, we chose to explore the impact of the tax imposed on a market that is cleared by a Walrasian auctioneer instead of a double auction or the simple price impact function. The reason that we chose to explore the Tobin tax in this particular setting is twofold.

First, actual foreign exchange markets are not organized as pure continuous double auctions or dealership markets. Dealers continuously update two-way quotes for their customers and thus act as market makers. Moreover, dealers trade with one another either via brokers (in a continuous double auction fashion) or directly using systems such as EBS or Reuters. Intra-dealer trading accounted for approximately 39% of total trading volume in 2013 (BIS 2013). While this relatively large figure[4] was the basis of criticism regarding extensive speculation on foreign exchange markets, it is not entirely true that such trading is primarily of a speculative nature. Instead, intra-dealer trading represents the tedious task of passing undesired inventories (originated by a single, possibly speculative, customer-dealer trade) along until they encounter a dealer with the opposite undesired position, such that they neutralize one another (Flood 1994). This phenomenon is conveniently termed "hot potato trading" in the literature. In addition, dealers generally change their positions dramatically throughout the trading day, but at the end of the day, a majority of them have a zero net position, just as they began the day (Cheung et al. 2004). Considering these two points, one could easily gain the impression that the intra-day trading among dealers and the adjustment of quotes for customers is some sort of intricate tâtonnement process of searching for the price under which all customers would hold desired positions with respect to their expectations and, simultaneously, none of the dealers, which together represent "the auctioneer", would be exposed to exchange rate risk.

A second, subtler point is that by using Walrasian clearing, we bypass the choice regarding the precise form for the price impact function and its relationship with liquidity. This approach is convenient because such a function is often difficult to estimate (e.g., Westerhoff 2004a) and may affect the measured impact of the Tobin tax substantially (Ehrenstein et al. 2005). Admittedly, Walrasian clearing is not a panacea, and hence this circumvention of the problematic decision regarding the price impact function comes at its own cost. In particular, the uncertainty regarding the price impact function is transferred to the utility functions of individual agents, which serve an analogous role in equilibrium pricing models—they determine how sensitive the current exchange rate is with respect to the expectations of individual agents. Thus, this approach is only as good as our knowledge of these functions. Nonetheless, it might still be considered an improvement, as there exists a relatively broad consensus regarding the utility functions used in asset-pricing models, with the majority of studies

[4] This percentage has been significantly higher in the past, see BIS (2013).

utilizing a mean-variance framework originally introduced by Sharpe (1964).[5] The research regarding price impact functions is, on the other hand, a rather novel area, and some definite conclusions regarding their form are likely still to come. For many references as well as an analysis of the possible pitfalls of such estimations, see Weber and Rosenow (2005).

Overall, to study the effects of a Tobin tax imposed on the retail market, it seems appropriate as well as convenient to model the interbank market as a Walrasian auctioneer. Considering the complexity of the interbank foreign exchange market, this approximation is somewhat crude. Nonetheless, by assessing the impact of the Tobin tax in this particular setting and by contributing to the variety of the existing research, we hope to help guide policies in markets whose complexity is beyond the reach of any single study. The remainder of the paper is structured as follows. Section 2 introduces the model and extends it by incorporating transaction costs in the optimization problem faced by retail traders. In Sect. 3, the calibration method is described, and multiple simulation runs are performed to assess the impact of the Tobin tax on selected statistics. Finally, Sect. 4 concludes.

2 Model

To study how the Tobin tax would alter price dynamics in a market cleared by a Walrasian auctioneer, we extend a heterogeneous agent model[6] of foreign exchange developed by De Grauwe and Grimaldi (2006). The original model is composed of a continuum of boundedly rational, risk-averse agents who optimize their portfolios based on anticipated exchange rate developments. They form their expectations using two distinct forecasting rules, and at the end of a trading round, they compare the accuracy of these rules and possibly switch from one rule to another. An ongoing competition between the forecasting rules creates chaotic price movements not dissimilar to those observed in real foreign exchanges, which make the model suitable for an examination of how these dynamics would be affected by the Tobin tax.

To assess the impact of the Tobin tax, we need to modify the model of De Grauwe and Grimaldi (2006). For obvious reasons, we extend the model by including a transaction tax τ in the optimization problem faced by agents. However, we also set the transaction costs in the goods market (in the original model denoted by C) to zero. This decision is motivated by the fact that transaction costs in the goods market C have a similar effect to a transaction tax τ, that is, discouraging transactions when only a small profit is expected. However, transaction costs C affect only one group of traders, which would imply that one group of traders would be predominantly affected by the newly imposed tax up to the threshold C, thus creating an interplay of effects that could obscure the analysis. Nevertheless, examining the impact of the tax for $C > 0$ might be an interesting avenue for future research. In addition, we extend the model with

[5] The strength of this framework was, among others, demonstrated by Kroll et al. (1984), who show that the mean-variance utility approximates other plausible utility functions extraordinarily well in the context of portfolio choice.

[6] For a general overview of heterogeneous agent modeling and its development, Chen et al. (2012), Hommes (2006), or LeBaron (2005) provide excellent surveys.

the possibility of imputing stochastic noise to the expectations of individual agents in order to generate realistic dynamics of traded volume. Due to technical difficulties arising from the tax τ that are described below, it is also necessary to populate the market with a finite number of agents rather than a continuum of agents. Despite these changes, with a suitable parameterization, the two models coincide, and our model can therefore be considered to be an extension of the model developed by De Grauwe and Grimaldi (2006) (conditioned on $C = 0$) that allows for financial transaction costs.

Although the description in this section is sufficient to understand the model, we recommend that the reader also consult the work of De Grauwe and Grimaldi (2006), in which the original model is presented in conjunction with numerous valuable insights concerning its dynamics. In the following paragraphs, we primarily focus on modifications to the original model.

2.1 Portfolio Optimization

The model is populated by N boundedly rational, risk-averse agents. The utility function of the i-th agent is defined as follows:

$$U\left(W_{t+1}^{i}\right) = E_{t}^{i}\left(W_{t+1}^{i}\right) - \frac{1}{2}\mu V_{t}^{i}\left(W_{t+1}^{i}\right), \quad (1)$$

where W_{t+1}^{i} represents the agent's wealth at time $t+1$. E_{t}^{i} and V_{t}^{i} are the conditional expectations and conditional variance operators, respectively. μ is a positive parameter capturing the degree of risk aversion.

Agents in the model optimize their portfolios based on anticipated exchange rate developments. In each trading round, individual agents can choose to allocate their wealth to either the domestic currency or the foreign currency. Similar to the analysis performed by De Grauwe and Grimaldi (2006), we assume that the interest rates of both countries are zero.[7] Non-zero interest rates would not alter the analysis in any way, as they merely scale the exchange rate by a constant (see De Grauwe and Grimaldi 2006, p. 4). The wealth of agent i at time $t+1$ is therefore naturally specified as follows:

$$W_{t+1}^{i} = W_{t}^{i} + d_{t}^{i}(s_{t+1} - s_{t}) - \tau\left|d_{t}^{i} - d_{t-1}^{i}\right| s_{t}, \quad (2)$$

where d_{t}^{i} are holdings of the foreign currency at time t, and s_t is the price of the foreign currency in terms of the domestic currency at time t. τ is an ad valorem transaction tax. The second term on the right-hand side of Eq. 2 represents the agents exposure to exchange rate movements. The third term is a novelty compared to the original model. It represents the additional cost of changing positions due to the transaction tax τ. The cost is proportional to the change in the position and the price, i.e., the transaction tax is paid in terms of the domestic currency and by both parties participating in the transaction.

[7] Unlike in De Grauwe and Grimaldi (2006), we do so directly when postulating the optimization problem rather than after solving it. This simplification is purely for expositional purposes.

From Eq. 2, it is immediately apparent that the model is asymmetric with respect to currencies. All agents compute their wealth in terms of the domestic currency, which is therefore considered to be safe with respect to exchange rate movements. This model could therefore represent a reasonable approximation of reality when considering, for example, two currencies, one that is widely used and another of only regional importance. It is also worth noting that the selling price of the foreign currency is equal to the buying price in Eq. 2, i.e., spreads are omitted from the model. This decision could be justified by assuming that no costs, apart from the tax τ, are connected with transactions and that the mediation of trades is a competitive endeavor with free entry. While the former assumption is admittedly unrealistic, a model with non-zero endogenously determined spreads would be difficult to design and analyze.

Maximizing utility with respect to holdings of the foreign currency d_t^i yields individual demand for the foreign currency at time t for a given $E_t^i(s_{t+1})$ and $V_t^i(s_{t+1})$:

$$d_t^i(s_t) = \begin{cases} \frac{E_t^i(s_{t+1}) - s_t(1+\tau)}{\mu V_t^i(s_{t+1})} & s_t \in \left(0, \ \frac{E_t^i(s_{t+1}) - \mu V_t^i(s_{t+1}) d_{t-1}^i}{1+\tau}\right) \\ d_{t-1}^i & s_t \in \left[\frac{E_t^i(s_{t+1}) - \mu V_t^i(s_{t+1}) d_{t-1}^i}{1+\tau}, \ \frac{E_t^i(s_{t+1}) - \mu V_t^i(s_{t+1}) d_{t-1}^i}{1-\tau}\right] \\ \frac{E_t^i(s_{t+1}) - s_t(1-\tau)}{\mu V_t^i(s_{t+1})} & s_t \in \left(\frac{E_t^i(s_{t+1}) - \mu V_t^i(s_{t+1}) d_{t-1}^i}{1-\tau}, \ \infty\right). \end{cases} \tag{3}$$

A straightforward interpretation of Eq. 3 is that for an exchange rate at time t suggested by the Walrasian auctioneer that does not sufficiently deviate from the expected exchange rate (adjusted for the current position), agent i chooses to refrain from trading because the expected realized profit from the change in positions would not exceed the transaction cost $\tau|d_t^i - d_{t-1}^i|s_t$. However, if the price offered in the tâtonnement process is sufficiently smaller (larger) than the i-th agent's expectation of the exchange rate, the agent buys (sells) the foreign currency to realize a profit (minimize loss).

2.2 Clearing Mechanism

Let Z_t be the exogenously determined supply of the foreign currency at time t. Then any $s_t > 0$ such that $\sum_{i=1}^{N} d_t^i(s_t) = Z_t$ is said to be a clearing price. For $\tau = 0$, the model is identical to De Grauwe and Grimaldi (2006) and the unique clearing price can be expressed as

$$s_t = \frac{1}{\sum_{i=1}^{N} \frac{1}{\mu V_t^i(s_{t+1})}} \left(\sum_{i=1}^{N} \frac{E_t^i(s_{t+1})}{\mu V_t^i(s_{t+1})} - Z_t \right). \tag{4}$$

However, at $\tau > 0$, one encounters several difficulties. First, although for all possible clearing prices, there exists a closed-form expression analogous to Eq. 4, this expression becomes exponentially more complex at increasing levels of time t.[8] The complexity is fortunately bounded by the number of agents N, which is why we choose to populate the model with a finite number of agents rather than with a continuum of agents as in De Grauwe and Grimaldi (2006). Second, for $\tau > 0$, the demand functions $d_t^i(s_t)$ are no longer decreasing but merely non-increasing, which implies that the clearing price may no longer be unique but instead, a whole interval of clearing prices might exist. In light of this observation, we define the actual market-clearing price s_t as the midpoint between the lowest and the highest possible clearing price:

$$s_t = \frac{\inf\left\{s_t \in R^+ : \sum_{i=1}^N d_t^i(s_t) = Z_t\right\} + \sup\left\{s_t \in R^+ : \sum_{i=1}^N d_t^i(s_t) = Z_t\right\}}{2}. \tag{5}$$

By doing so, we guarantee that the price determination rule is identical to the one in the original model in the case of $\tau = 0$, but it also naturally extends to cases where $\tau > 0$, in which possible clearing prices might constitute an interval.

2.3 Forecasting Rules and Risk Evaluation

Having selected the agent's utility functions and the clearing mechanism, we need to determine how agents form their opinions regarding the future exchange rate. In reality, traders lack the computational capabilities to solve demanding problems such as the computation of the expected exchange rate conditional on the information set available to the agent.[9] Instead, human behavior is better described as rule governed—individuals possess several simple heuristics to perform a particular task and choose between them based on their performance (see e.g., Simon 1990).

Two main types of heuristics for predicting the exchange rate were identified in foreign exchange markets: fundamental and technical analysis (e.g., Frankel and Froot 1990). Traders using technical analysis (henceforth referred to as chartists) infer the future exchange rate based on past price movements and are to a great extent the target group that Tobin intended to impair due to their supposedly destabilizing behavior. Traders using fundamental analysis (henceforth referred to as fundamentalists), by contrast, assume the convergence of the market price to some fundamental value and are believed to be stabilizing the market. Questionnaire surveys suggest that most traders in foreign exchange markets are familiar with both types of analysis and consider them to be equally important (e.g., Taylor and Allen 1992; Oberlechner 2001). In our model, we will follow De Grauwe and Grimaldi (2006) and approximate these

[8] This complexity develops because of how agents switch between different forecasting rules, as described in Sect. 2.3, and because individual demand also depends on past positions.

[9] Consider, for example, the original model of De Grauwe and Grimaldi (2006), in which fractal borders between different equilibria give rise to chaotic price movements that are unpredictable unless we possess infinitely precise estimates of the models parameters. In a stochastic and less sharply defined market more closely resembling real foreign exchange markets, this difficulty in computing the expected exchange rate would be even more severe.

two ample families of forecasting heuristics using two simple forecasting rules. Furthermore, we add an autoregressive stochastic term to the original equations to address the persistent variability of forecasting heuristics within families.

Following De Grauwe and Grimaldi (2006), we define the expectations of the i-th trader using the fundamental forecasting rule concerning the exchange rate in period $t+1$ as

$$E_t^{i,F}(s_{t+1}) = s_{t-1} - \psi\left(s_{t-1} - s_{t-1}^*\right) + e_t^{i,F} \tag{6}$$

$$e_t^{i,F} = (1-\lambda_F)\, e_{t-1}^{i,F} + \lambda_F \varepsilon \quad \varepsilon \backsim N\left(0, \sigma_F^2\right), \tag{7}$$

where the parameter $\psi \in (0, 1)$ denotes the expected speed of an adjustment toward the fundamental value s_{t-1}^*, the term $e_t^{i,F}$ captures different beliefs within the group of fundamentalists and the parameter $\lambda_F \in [0, 1]$ determines the persistence of different beliefs within the fundamentalist group.

Similarly, the expectations of the i-th trader using the chartist forecasting rule concerning the exchange rate in period $t+1$ can be expressed as follows:

$$E_t^{i,C}(s_{t+1}) = s_{t-1} + \beta \sum_{j=1}^{T_C} \alpha_j \Delta s_{t-j} + e_t^{i,C} \tag{8}$$

$$e_t^{i,C} = (1-\lambda_C)\, e_{t-1}^{i,C} + \lambda_C \varepsilon \quad \varepsilon \backsim N\left(0, \sigma_C^2\right). \tag{9}$$

An intuitive explanation for Eq. 8 is that chartists are attempting to extrapolate the future exchange rate from past exchange rate changes. The parameters a_j; $j \in \{1, 2, 3, \ldots, T_C\}$ and the overall strength of the extrapolation $\beta \in (0, 1)$ are assumed to be positive—if a chartist experienced a past price increase, he is most likely to expect the exchange rate to increase further. The term $e_t^{i,F}$ again captures different beliefs within the group of chartists.

Two aspects are worth further discussion. First, as one can see, when predicting the exchange rate in period $t+1$, agents do not use information concerning the exchange rate in period t because it is yet to be determined based on their decision at time t. This process accords with Eqs. 3 and 5. Second, all agents keep track of both their fundamental and chartist expectations, regardless of whether they are currently chartists or fundamentalists. Otherwise, they would not be able to compare the accuracy of these forecasting rules. Naturally, only the expectations matching the current type of agent are used to determine individual demand, which can be expressed by an indicator function I in the following manner:

$$E_t^i(s_{t+1}) = I(i,t)E_t^{i,F}(s_{t+1}) + (1 - I(i,t))\, E_t^{i,C}(s_{t+1}), \tag{10}$$

where

$$I(i,t) = \begin{cases} 1 & \text{iff the } i\text{-th trader at time } t \text{ is a fundamentalist} \\ 0 & \text{iff the } i\text{-th trader at time } t \text{ is a chartist} \end{cases}. \tag{11}$$

Now, we turn to the way in which individual agents evaluate the risk of their portfolios. The conditional exchange rate variance of the fundamentalist $V_t^{i,F}(s_{t+1})$ and chartist $V_t^{i,C}(s_{t+1})$ is determined by the proportion of exchange rate movements that cannot be explained by fundamentalist and chartist expectations, respectively, i.e.,

$$V_t^{i,F}(s_{t+1}) = (1-\theta)\, V_{t-1}^{i,F} + \theta \left(E_{t-2}^{i,F}(s_{t-1}) - s_{t-1} \right)^2, \tag{12}$$

$$V_t^{i,C}(s_{t+1}) = (1-\theta)\, V_{t-1}^{i,C} + \theta \left(E_{t-2}^{i,C}(s_{t-1}) - s_{t-1} \right)^2. \tag{13}$$

The parameter $\theta \in (0, 1)$ is used such that agents place greater weight on the most recent forecasting errors. Combining Eqs. 12 and 13, we obtain

$$V_t^{i}(s_{t+1}) = I(i,t) V_t^{i,F}(s_{t+1}) + (1 - I(i,t))\, V_t^{i,C}(s_{t+1}). \tag{14}$$

2.4 Fitness of the Rules

To complete the model, we specify the function $I(i, t)$, i.e., how an agent decides whether to be a chartist or a fundamentalist. In line with the models ability to generate unpredictable dynamics, we will assume that agents are incapable of computing the most profitable rule ex ante and instead employ a trial and error method. To be more precise, agents evaluate the past profitability of forecasting rules and tend to choose the one that performed better. The probability that agent i is a fundamentalist or a chartist, respectively, at time t is defined as follows:

$$P\left(I(i,t) = 1\right) = \frac{\exp\left(\gamma \pi_t^{*\,i,F}\right)}{\exp\left(\gamma \pi_t^{*\,i,F}\right) + \exp\left(\gamma \pi_t^{*\,i,C}\right)}, \tag{15}$$

$$P\left(I(i,t) = 0\right) = \frac{\exp\left(\gamma \pi_t^{*\,i,C}\right)}{\exp\left(\gamma \pi_t^{*\,i,F}\right) + \exp\left(\gamma \pi_t^{*\,i,C}\right)}, \tag{16}$$

where $\pi_{t-1}^{*\,i,F}$ and $\pi_{t-1}^{*\,i,C}$ are the risk-adjusted profits from agent i being either a fundamentalist or a chartist at time $t-1$. The parameter $\gamma \in [0, \infty)$ is often referred to as the *intensity of choice* in the literature and measures the intensity with which traders revise their forecasting rules (Brock and Hommes 1998). By setting $\gamma = 0$, agents become insensitive to past profitability, and the probability of an agent being either a fundamentalist or a chartist is constant and equal to 0.5. However, as $\gamma \to \infty$, agents choose whatever forecasting rule proved to be more profitable in the previous trading round. Negative values of γ are not economically meaningful.

We define the risk-adjusted profit of chartists and fundamentalists on a per-unit basis as in De Grauwe and Grimaldi (2006):

$$\pi_t^{*i,F} = \pi_t^{i,F} - \mu V_t^{i,F}(s_{t+1}) \tag{17}$$

$$\pi_t^{i,F} = (s_{t-1} - s_{t-2})\mathrm{sgn}\left(d_{t-2}^{i,F}\right) - \frac{\tau \left| d_{t-2}^{i,F} - d_{t-3}^{i} \right| s_{t-2}}{\left| d_{t-2}^{i,F} \right|}, \tag{18}$$

and

$$\pi_t^{*i,C} = \pi_t^{i,C} - \mu V_t^{i,C}(s_{t+1}) \tag{19}$$

$$\pi_t^{i,C} = (s_{t-1} - s_{t-2})\mathrm{sgn}\left(d_{t-2}^{i,C}\right) - \frac{\tau \left| d_{t-2}^{i,C} - d_{t-3}^{i} \right| s_{t-2}}{\left| d_{t-2}^{i,C} \right|}, \tag{20}$$

where $d_{t-1}^{i,F}$ and $d_{t-1}^{i,C}$ are the demand of agent i for the foreign currency under the assumption that the agent is a fundamentalist or a chartist, respectively, i.e., d_{t-1}^{i} provided that $I(i, t-1) = 1$ or $I(i, t-1) = 0$. The second term on the right-hand side of Eqs. 18 and 20 is a novelty relative to the original model. It represents the per-unit cost of changing positions from one period to the next. It is defined consistently with Eq. 2.

Equations 15 and 16 complete the model (provided that s_t^* and Z_t are exogenously determined) and in conjunction with Eqs. 1–14 form a non-linear system of stochastic difference equations. As is frequently experienced in heterogeneous agent modeling, the system is difficult to examine analytically. Thus, we employ numerical methods to assess the impact of the parameter τ on selected statistics of exchange rate realizations s_t.

3 Simulations

3.1 Calibration[10]

The above-presented model has 9 parameters, and the impact of the Tobin tax on the statistics generated by the model may differ depending on the exact parameter settings that we consider. To narrow the space of possible parameterizations, we calibrate[11] the foreign exchange developed by De Grauwe and Grimaldi (2006) in the model such that the time series generated by the model mimics selected statistics of real exchange rates. To do so, we focus on four statistics of returns (i.e., $r_t = \frac{s_t - s_{t-1}}{s_t}$), namely,

[10] Due to the nature of the computational methods used in this section, we are unable to describe every aspect associated with the calibration and subsequent analysis in detail. The following paragraphs are therefore intended to acquaint the reader with the methods we employed rather than to provide an exhaustive description. We will nevertheless gladly provide the interested reader with additional files such as datasets, log files, source code for the model and the random seeds employed upon request.

[11] Regarding empirical validation and calibration of agent-based models in Economics, Fagiolo et al. (2007) provide an important methodological discussion.

Table 1 Calibration statistics: the standard errors of $\widehat{sd}$ and $\widehat{kurt}$ are obtained via the bootstrap method (10,000 re-samplings)

Statistic	Estimate	SE	Average under $\mathbf{p}_0$
$\widehat{sd}$	5.956×10^{-3}	7.380×10^{-5}	6.101×10^{-3}
$\widehat{kurt}$	6.257	4.519×10^{-1}	5.242
$\widehat{\rho^r}$	-6.590×10^{-2}	1.098×10^{-2}	-4.422×10^{-2}
$\widehat{\rho^{\|r\|}}$	6.056×10^{-1}	8.788×10^{-3}	6.152×10^{-1}

The last column captures the ability of the calibrated model to reproduce the target statistics. Averages of the statistic under $\mathbf{p} = \mathbf{p}_0$ were obtained from 1000 simulation runs lasting 3400 trading rounds

- Annualized volatility: $sd = \sqrt{\frac{1}{T}\sum_{t=1}^{T}(r_t - \bar{r})^2}$
- Kurtosis of returns: $kurt = \dfrac{\frac{1}{T}\sum_{t=1}^{T}(r_t - \bar{r})^4}{\left(\frac{1}{T}\sum_{t=1}^{T}(r_t - \bar{r})^2\right)^2}$
- Autocorrelation of raw returns: $\rho^r = \dfrac{\sum_{t=2}^{T} r_t r_{t-1}}{\sum_{t=2}^{T} r_t^2}$
- Autocorrelation of absolute returns: $\rho^{|r|} = \dfrac{\sum_{t=2}^{T} |r_t r_{t-1}|}{\sum_{t=2}^{T} r_t^2}$

There are of course many interesting statistical properties of foreign exchange returns (for an extensive survey, see Sewell 2011), but the four statistics presented above capture the most fundamental characteristics. Therefore, they can be used as a rough measure for assessing the similarity of the time series generated by the model and real-world data. The calibration itself is performed by minimizing the following fitness function with respect to the vector of parameters $\mathbf{p} = (\mu, \psi, \lambda_F, \sigma_F^2, \beta, \theta, \gamma)$[12]:

$$Fitness = \frac{|sd - \widehat{sd}|}{SE_{\widehat{sd}}} + \frac{|kurt - \widehat{kurt}|}{SE_{\widehat{kurt}}} + \frac{|\rho^r - \widehat{\rho^r}|}{SE_{\widehat{\rho^r}}} + \frac{|\rho^{|r|} - \widehat{\rho^{|r|}}|}{SE_{\widehat{\rho^{|r|}}}} \tag{21}$$

The estimates $\widehat{sd}, \widehat{kurt}, \widehat{\rho^r}, \widehat{\rho^{|r|}}$ and their standard errors denoted by SE were obtained from a dataset published by the Czech National Bank (ČNB Czech National Bank 2014) containing the daily exchange rates of 34 different currencies vis-à-vis the Czech Koruna (CZK) throughout 250 trading days (2 Jan. 2013–27 Dec. 2013; 7 Nov 2013 is omitted due to an exchange rate intervention) and are depicted in Table 1.

The optimization of the fitness function with respect to the vector of parameters $\mathbf{p}$ is executed via standard genetic algorithm (see Stonedahl 2011) with Gray-Binary chromosome encoding and mutation and crossover rates of 0.01 and 0.7, respectively,

[12] To facilitate the calibration, both types of traders are assumed to be identical with respect to persistence and the level of heterogeneity in their expectations (i.e., $\lambda_F = \lambda_C$, $\sigma_F^2 = \sigma_C^2$). In addition, vector $\boldsymbol{\alpha}$ is adopted from De Grauwe and Grimaldi (2006). Parameters N, Z_t and s_t^* for $t \in \{1, 2, 3, \ldots, T\}$ are assumed to be constant and are set to 100, 0, and 1, respectively.

Table 2 Results of the calibration runs

	μ	ψ	λ_F	σ_F^2	β	θ	γ	Fitness
$\mathbf{p}_0$	64.116	0.672	0.953	0.019	0.35	0.932	55.647	8.921

and population and tournament sizes of 50 and 3, respectively. In order to find the most suitable vector of parameters, we performed 10 such independent search trials, each consisting of 10,000 simulations lasting 340 trading rounds.[13] The parameter settings $\mathbf{p}_0$, of which the individual coordinates are depicted in Table 2, performed the best in terms of the fitness function (Eq. 21) compared to all other parameter settings that we considered. Table 1 illustrates that the model under the parameter settings $\mathbf{p}_0$ is capable of mimicking the target statistic quite closely and is therefore suitable for the analysis. However, before we turn to simulations to evaluate the impact of the Tobin tax, it is necessary to stress two possible shortcomings of the calibration we performed.

The diverse results we obtained from multiple searches indicate that the parameter space was to some degree left unexplored. It is therefore more appropriate to regard $\mathbf{p}_0$ as the best parameter setting given the computational restrictions we face, rather than the best setting.

Second, the similarity in the statistics presented above is a necessary but not sufficient condition for capturing the true underlying process generating exchange rate dynamics. This problem is inherently present when examining complex systems, and aside from carefully designing the model, there is little we can do to address this issue. Despite these objections, we believe that the model under $\mathbf{p} = \mathbf{p}_0$ yields interesting insights into the possible impact of the Tobin tax in real markets.

3.2 Results

The model under $\mathbf{p} = \mathbf{p}_0$ is capable of generating rich time series that exhibit properties such as excess volatility or persistent departures from the fundamental value. One such realization is depicted in Fig. 10 in "Appendix". We choose to analyze the impact of the Tobin tax τ on four statistics of exchange rate realizations that are frequently discussed in relation to the Tobin tax:

- $sd = \sqrt{\frac{1}{T}\sum_{t=1}^{T}(r_t - \bar{r})^2}$

  The impact on volatility is a central question in the debate over the Tobin tax. While proponents of the tax argue that it could reduce volatility, the evidence appears to be mixed (see Sect. 1).

[13] All computations were made using BehaviourSearch software (Stonedahl and Wilensky 2010) integrated with the programming language NetLogo (Wilensky 1999), which we use to implement the model.

- $dist = \frac{1}{T}\sum_{t=1}^{T}|s_t - s_t^*|$

  Distortions measure the degree of mispricing in the time series. It is useful to distinguish between two distinct sources of distortions: unceasing, erratic oscillations around the fundamental value caused by the noise within the model and occasional, persistent, one-sided deviations driven by positive feedback in expectations of the chartists. The latter is characterized by a sudden drop in the popularity of the fundamentalist forecasting rule. Figure 10 in "Appendix" depicts the formation and the subsequent correction of one such persistent deviation.
- $kurt = \frac{\frac{1}{T}\sum_{t=1}^{T}(r_t - \bar{r})^4}{\left(\frac{1}{T}\sum_{t=1}^{T}(r_t - \bar{r})^2\right)^2}$

  Kurtosis is addressed here because several research papers indicate that the Tobin tax might have opposite effects on volatility and kurtosis. The results of Demary (2006) suggest that the Tobin tax reduces volatility but increases the kurtosis of returns. Conversely, Lavička et al. (2013) reported that the Tobin tax increases volatility but decreases kurtosis.
- $vol = \frac{1}{2T}\sum_{t=1}^{T}\sum_{i=1}^{N}|d_t^i - d_{t-1}^i|$

  The direction of the impact of the Tobin tax on traded volume is hardly a controversial issue. Nonetheless, the degree to which volume would be affected is often discussed, as it is crucial for estimating the possible revenues from the tax.

To assess the impact of the Tobin tax on the aforementioned statistics, we perform 1000 simulation runs lasting 340 trading rounds for each of 51 different values of τ ranging from 0 to 1% (i.e., 0.02% steps) under $\mathbf{p} = \mathbf{p}_0$. The 10th and 90th percentiles of the measured statistics for different values of τ are depicted in Figs. 1, 2, 4, and 5. The charts also contain estimates of the means of the statistics and their 95% confidence intervals. To allow for an easy comparison of effects from the tax across different statistics, we chose to present the statistics in relative terms—the scenario $\tau = 0$ is treated as the baseline.

Figure 1 indicates that the effect of the Tobin tax on volatility is not monotone. For values ranging from 0 to 0.3%, the tax negligibly (yet statistically significantly) increases the average volatility. For values between 0.3 and 1%, the tax delivers a moderate (up to 10%) decrease in volatility. This relationship between the magnitude of the tax and its qualitative effect has not, to the best of our knowledge, been reported by any agent-based model applied to the Tobin tax. Unfortunately, it is difficult to precisely determine the cause of this intriguing behavior. In practice, this result suggests that it may be necessary to impose the tax with sufficient strength in order to achieve the desired stabilizing effect. A small tax may lead to an adverse effect.

The Tobin tax also appears to be quite successful in eliminating mispricing. As Fig. 2 illustrates, the Tobin tax reduces distortions by up to 30% in the case of a 1% tax. Interestingly, the Tobin tax has a much greater effect on the second component of distortions long-lasting, one-sided deviations. Figure 3 shows the decomposition of total distortions according to the length of the individual deviations from the fundamental value for values of $\tau \in \{0, 0.005, 0.01\}$. As is apparent from the figure, the

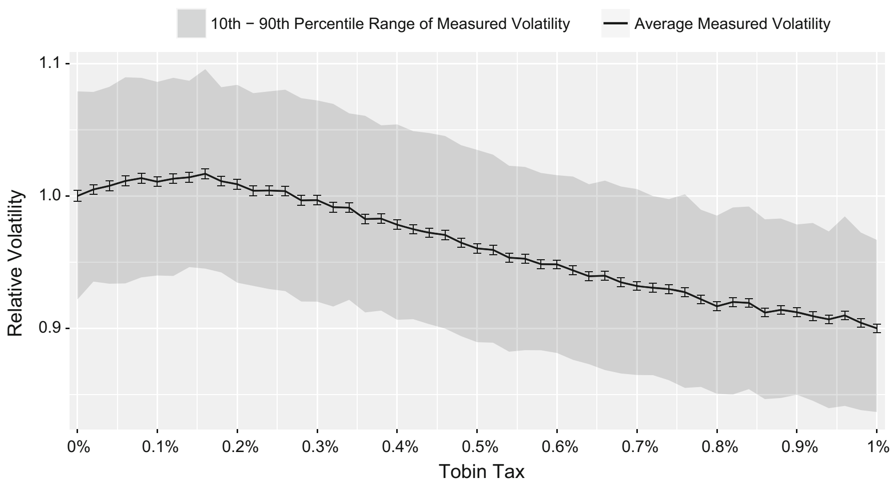

Fig. 1 Impact of τ on volatility: a statistically significant increase in mean volatility (*whiskers* represent the 90% confidence interval) for values of $\tau \in [0.02, 0.26]$ and a statistically significant decrease for values of $\tau \in [0.32, 1\%]$ (up to 10% in the case of a 1% tax). The chart is scaled such that the relative volatility in the scenario $\tau = 0$ is equal to 1. The scale factor (i.e., the average volatility under $\tau = 0$) is 6.06×10^{-3}

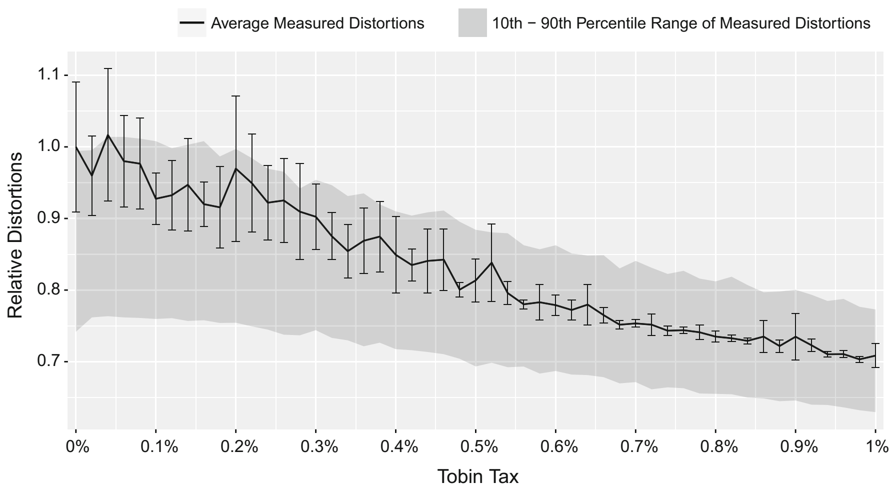

Fig. 2 Impact of τ on distortions: a statistically significant decrease in mean distortions (*whiskers* represent the 90% confidence interval) for values of $\tau \in [0.24, 1\%]$ (up to 29% in the case of a 1% tax). This result is primarily achieved by preventing the occurrence of long-lasting deviations (see Fig. 3). The chart is scaled such that the relative distortions in the scenario $\tau = 0$ are equal to 1. The scale factor (i.e., the average distortions under $\tau = 0$) is 7.81×10^{-3}

Tobin tax actually slightly increases the portion of distortions that originated from very short-term deviations lasting less than 2 periods. However, as we examine increasingly longer deviations, we see that the Tobin tax has a much stronger effect. Contributions from deviations lasting 20–30 periods are reduced by approximately half, and per-

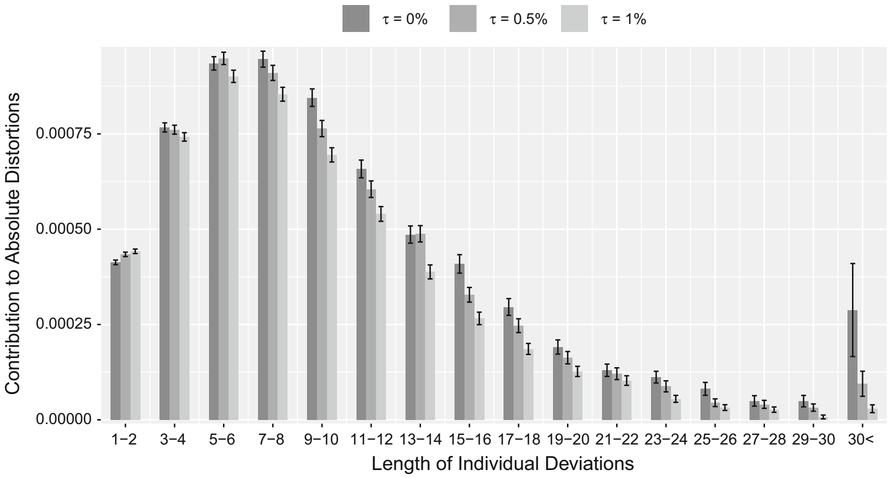

Fig. 3 Decomposition of impact of τ on distortions according to length of deviations: the reduction of total distortions by 15% in the case of a 0.5% tax and by 29% in the case of a 1% tax is not achieved by uniformly reducing all of the components of distortions. Instead, the effect on long-lasting deviations is much stronger than in the case of short-term deviations. A 1% Tobin tax appears to almost eliminate any deviations lasting more than 30 periods, while contributions from deviations lasting only 1–2 periods actually slightly increase

sistent deviations lasting more than 30 periods (such as that depicted in Fig. 10 in "Appendix") appear to be almost eliminated.

At this point, it might be worthwhile to discuss the relationship between such persistent deviations and the phenomenon of speculative bubbles. Similar to what is commonly believed to be the main feature of speculative bubbles, these persistent deviations are endogenously generated by positive feedback in the process of expectation formation. Nonetheless, because there is no universal agreement on the precise definition of a bubble, we are reluctant to denote it as such; instead, we will simply refer to it as a persistent deviation. One possible avenue for further research is to examine whether these persistent deviations indeed exhibit other properties that can be attributed to speculative bubbles, such as the steady rate of building up followed by a rapid correction.

In line with Lavička et al. (2013), we find that the Tobin tax is better at reducing the kurtosis of returns than their volatility (see Fig. 4). This effect is consistent with suppressing the incidence of the highest and the lowest returns, as lower kurtosis is by definition related to a less heavy-tailed distribution. This implies that the stabilizing effect of the tax is primarily achieved by eliminating infrequent but substantial changes in the price while leaving the magnitude of the regular day-to-day price movements mostly unaffected. A tax of only 0.1% is able to reduce the average kurtosis by 8%. This effect gradually increases for larger values of the tax.

The impact of the Tobin tax on traded volume is, as expected, negative. A 1% tax reduces the average trading volume by more than 50%. Interestingly, the effect appears to be almost linear for values of the tax larger than 0.1% (see Fig. 5). This result can be at least partially traced back to the demand of individual agents. To do

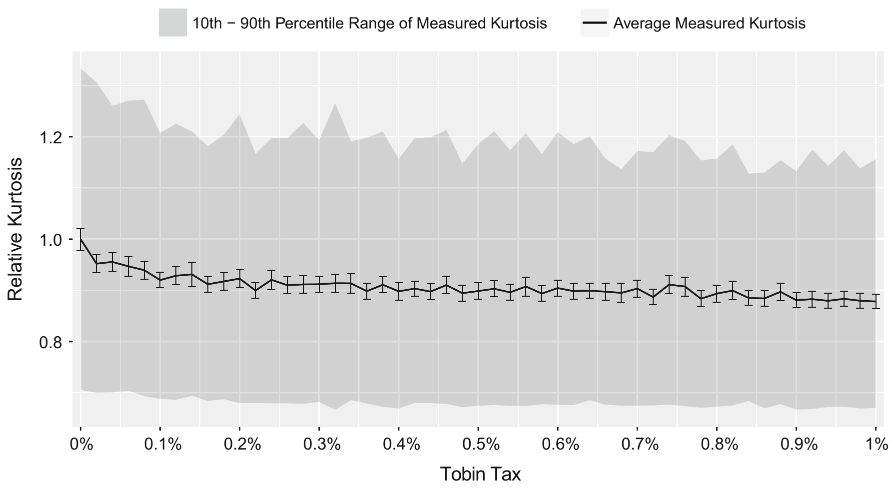

Fig. 4 Impact of τ on kurtosis: a statistically significant decrease in mean kurtosis (whiskers represent the 90% confidence interval) for values of $\tau \in [0.02, 1\%]$ (up to 12% in the case of a 1% tax). The chart is scaled such that the relative kurtosis in the scenario $\tau = 0$ is equal to 1. The scale factor (i.e., the average kurtosis under $\tau = 0$) is 5.07

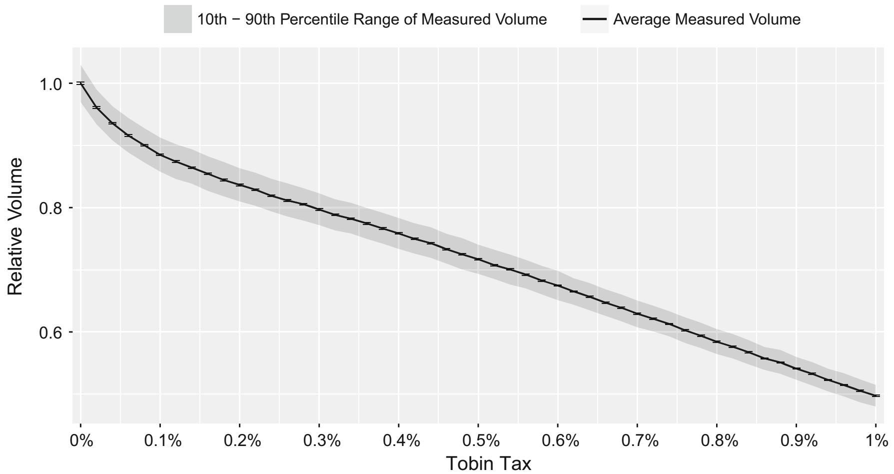

Fig. 5 Impact of τ on volume: a statistically significant decrease in mean traded volume (*whiskers* represent the 90% confidence interval) for values of $\tau \in [0.02, 1\%]$ (up to 50% in the case of a 1% tax). The chart is scaled such that the relative volume in the scenario $\tau = 0$ is equal to 1. The scale factor (i.e., the average volume under $\tau = 0$) is 659.3

so, one can consider a static example with a given clearing price s_t. The derivative of the net demand $|d_t^i - d_{t-1}^i|$ with respect to the tax τ is either $-\frac{s_t}{V_t^i(s_{t+1})}$ in the case when $E_t^i(s_{t+1})$ is deviating sufficiently from the clearing price s_t (the agent is hence actively trading, i.e., $|d_t^i - d_{t-1}^i| > 0$), or 0 in the case when it is not (the agent refrains from trading, i.e., $|d_t^i - d_{t-1}^i| = 0$). In both cases, the effect of the tax τ on the volume traded by agent i is linear.

In addition to this effect, which can be denoted as the intensive margin, the Tobin tax τ also directly affects the proportion of agents who choose to refrain from trading and whose derivatives of net demand with respect to τ are hence 0 as opposed to $-\frac{s_t}{V_t^i(s_{t+1})}$, which can be denoted as the extensive margin. For this effect, the possible linearity is difficult to judge because it is dependent on the distributions of expectations $E_t^i(s_{t+1})$, perceived variances $V_t^i(s_{t+1})$, and past demand d_{t-1}^i. Note, however, that agents who choose to refrain from trading after an increase of τ are exactly those with expectations $E_t^i(s_{t+1})$ very close to the clearing price s_t, and hence those who would otherwise have traded only very small volumes (see Eq. 3). This is likely the reason why the intensive margin effect dominates the extensive margin effect and the overall effect on volume is approximately linear.

The observed linear effect of the tax on volume has an interesting parallel in the class of disequilibrium models with a price impact function. In related literature, there is an ongoing scientific debate about its shape. While some authors suggest a simple linear form (e.g. Eisler et al. 2012; Cont et al. 2014), other authors empirically estimate various non-linear, mostly concave shapes (e.g. Kempf and Korn 1999; Lillo et al. 2002; Plerou et al. 2002), power-law or logarithmic relations (e.g. Potters and Bouchaud 2003), and also convex shapes (e.g. Weber and Rosenow 2005; Smid 2016). Our results might therefore unintendedly and indirectly contribute to this debate, slightly favoring the former linear impact-type models.

3.3 Sensitivity Analysis

3.3.1 Non-constant Fundamental Value

Thus far, we have assumed that the fundamental value is constant. We made this assumption because it is a reasonable starting point for analyzing the dynamics of the model. In reality, the fundamental value likely fluctuates, especially over longer periods of time. It is therefore necessary to determine whether our results also hold when the fundamental value follows some random process. To do so, we replicated the simulations we performed to construct Figs. 1, 2, 4, and 5, but we allow the fundamental value to follow a Gaussian random walk with zero mean and standard deviation 2×10^{-3}, 4×10^{-3}, and 6×10^{-3}.[14]

Figures 6, 7, 8, and 9 capture how the Tobin tax affects volatility, distortions, kurtosis, and volume under scenarios with different degrees of volatility in the fundamental value. The figures display the sample median rather than the sample average of the statistics because the former has been demonstrated to be less prone to random noise. The sample average would provide a similar, yet less clear story and might even be misleading because it is impossible to clearly present the additional statistics necessary to interpret it within a single chart. The directions of the effects on all of the

[14] For comparison, the estimated standard deviation of actual foreign exchange returns is approximately 6×10^{-3}.

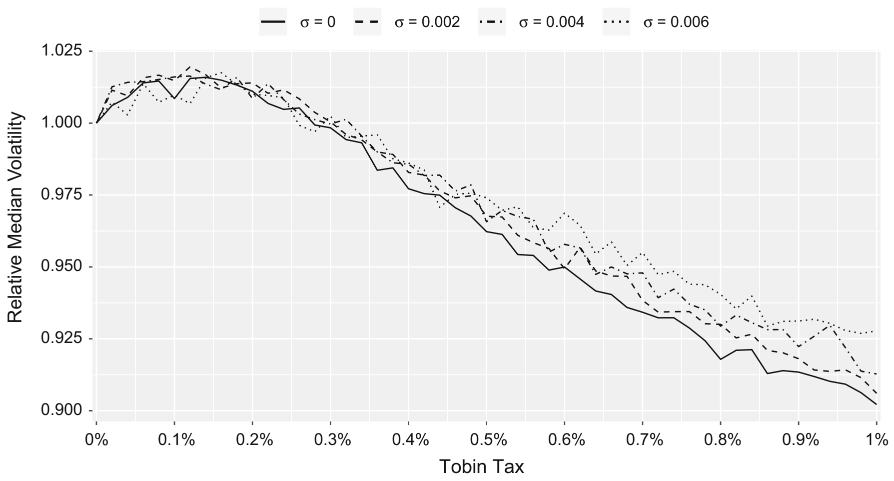

Fig. 6 Sensitivity of the effect on volatility with respect to σ: the ability of the Tobin tax to reduce the volatility of returns is mildly inversely related to the volatility of the fundamental value. The chart is scaled such that the relative volatility in scenario $\tau = 0$ is equal to 1 for each value of σ

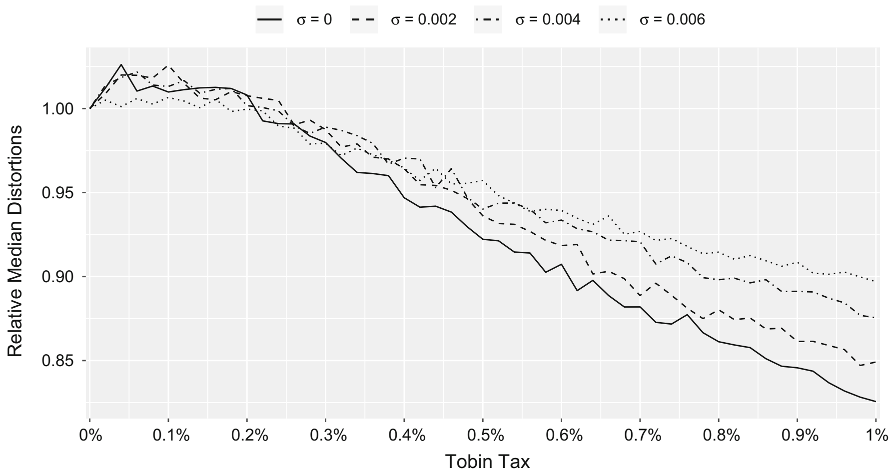

Fig. 7 Sensitivity of the effect on distortions with respect to σ: the ability of the Tobin tax to reduce distortions is inversely related to the volatility of the fundamental value. This relationship is caused by the trade-off between reducing volatility and reducing distortions in scenarios with high levels of fundamental volatility. The chart is scaled such that the relative distortion in scenario $\tau = 0$ is equal to 1 for each value of σ

examined statistics are robust to the volatility of the fundamental value. Regarding the magnitudes of the effects, we observe that the volatility of the fundamental value has no impact on the ability to reduce traded volume and a non-trivial impact on the ability to reduce the kurtosis of returns.

Interestingly, the ability of the Tobin tax to reduce both volatility and, especially, distortions appears to be inversely related to the volatility of the fundamental value.

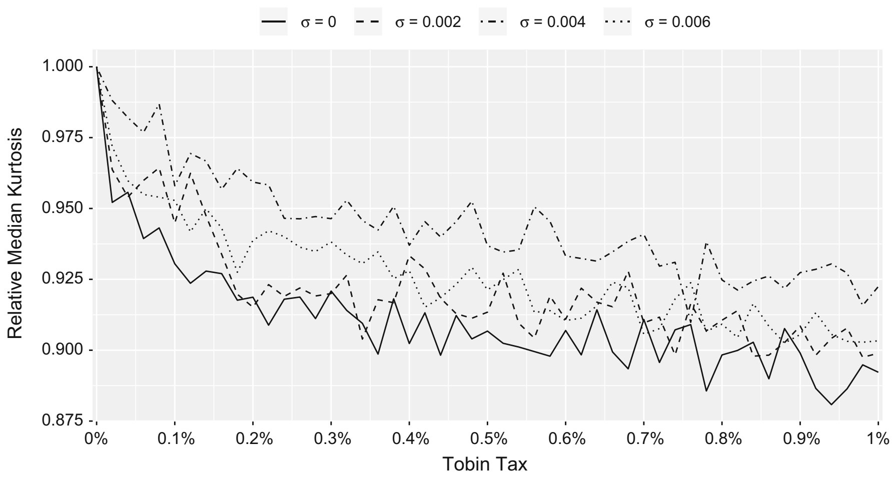

Fig. 8 Sensitivity of the effect on kurtosis with respect to σ: the ability of the Tobin tax to reduce the kurtosis of returns is non-trivially affected by the volatility of the fundamental value. The chart is scaled such that the relative kurtosis in scenario $\tau = 0$ is equal to 1 for each value of σ

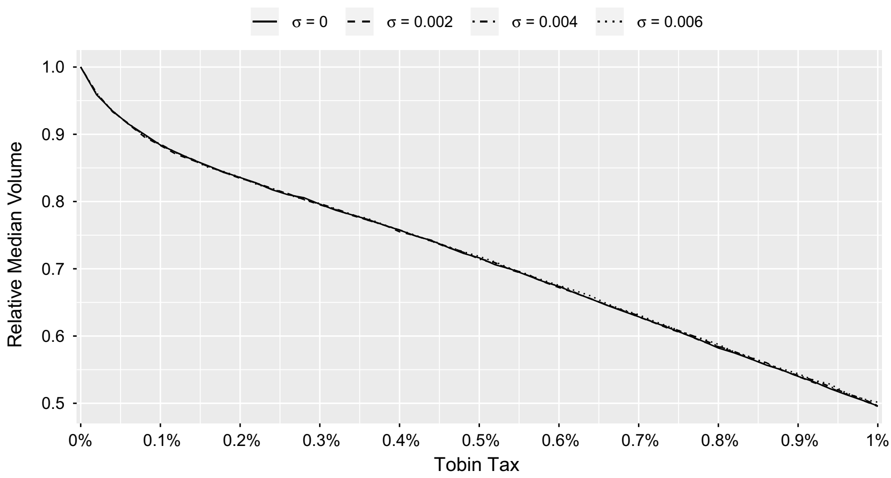

Fig. 9 Sensitivity of the effect on volume with respect to σ: the ability of the Tobin tax to reduce volume is unaffected by the volatility of the fundamental value—the individual plots almost perfectly coincide. The chart is scaled such that the relative volume in scenario $\tau = 0$ is equal to 1 for each value of σ

To some extent, this effect is simply caused by the fact that in scenarios with higher volatility of the fundamental value, there is in general higher volatility of returns and higher distortions, of which only a fixed amount is caused by the intrinsic dynamics of the model and can hence be eliminated by the tax. However, closer examination reveals that even in absolute terms, the effect of the tax is inversely related to the volatility of the fundamental value. This relationship arises because the tax affects both groups of traders. While it subdues erratic oscillations of the market exchange rate by discouraging the extrapolating behavior of chartists, it also discourages fundamentalists from trading when they observe a discrepancy between the exchange rate and the fundamental value. Consequently, the tax reduces the ability of the market exchange rate to rapidly adjust to new levels after a change of the fundamental value, which may, especially in scenarios with high volatility of the fundamental value, amplify the part of the distortions originating from short-term deviations. However, because most of the effect stems from eliminating long-lasting deviations from the fundamental value, the overall effect of the tax on distortions remains negative.

3.3.2 Number of Agents and Intensity of Stochastic Noise

Finally, we elaborate a sensitivity analysis of the two important modifications by which we extend the model, i.e., the finite number N of boundedly rational agents populating the model and the intensity of the stochastic noise σ_F^2 in their expectations. The third modification, i.e., the setting of transaction costs in the goods market to $C = 0$, does not require this type of simulation-based robustness assessment, as the rationale behind it is explained theoretically at the beginning of Sect. 2.

First, we focus on the effect of the number of agents in the model. Due to the computational burden of the analysis with larger numbers of agents, we define a grid covering an economically reasonable but computationally manageable space of possible settings to assess the effect in question. In Table 3, we present the results of the analysis with $N \in \{100, 200, 500, 1000, 2000\}$ agents combined with three important values of the Tobin tax $\tau \in \{0, 0.5, 1\%\}$ (i.e., the borderline values from our previous analysis and the median) and $\sigma_F^2 = 0.019$. For clear mutual comparability, we scale the percentage figures in Table 3 such that the scenario with $\tau = 0$ is always equal to 100%. Moreover, the first scenario with $N = 100$ for each of statistics is associated with the results of our previous analysis in Figs. 1, 2, 4, and 5.

For Volatility, we observe that the effect of delivering a moderate decrease is retained for all values of the number of agents N; however, the effect is slowly decaying. We presume that this behavior is partially caused by the fact that the number of agents N might not be neutral to the model and to some extent interplays with some other variables in the complex behavior of the system. A potential calibration under other values of N might then lead to slightly different model settings. This effect might also be connected to some form of so-called "small number of agents effect" described in Hommes (2006, p. 1156). Next, for Distortions, the ability to eliminate the mispricing is retained for all values of the number of agents N; however, the effect is again decaying. Based on our findings for Volatility, this result is not surprising because these two statistics are related. The very same explanation, therefore, applies for the effect on Distortions as well. On the other hand, for the effect of reducing the Kurtosis of returns as well as the expected negative effect on Volume, we observe robust performance over all N.

Second, we analyze the effect of the intensity of the stochastic noise in the expectations of agents. The grid covering the setting space for this sensitivity exercise is based on $\sigma_F^2 \in \{0.00475, 0.0095, 0.019, 0.0285, 0.038\}$[15] combined with the same

[15] We also test the behavior of the model for $\sigma_F^2 = 0.076$, but such a large value of the intensity of stochastic noise already causes numerical instability of the model, with many divergent runs.

Table 3 Sensitivity analysis: the effect of the number of agents N in the model

	$\tau = 0\%$	$\tau = 0.5\%$	$\tau = 1\%$
∅ effect on volatility (in %)			
$N = 100$	100	96	90
$N = 200$	100	97	92
$N = 500$	100	98	94
$N = 1000$	100	99	96
$N = 2000$	100	100	97
∅ effect on distortions (in %)			
$N = 100$	100	82	73
$N = 200$	100	91	82
$N = 500$	100	96	88
$N = 1000$	100	97	90
$N = 2000$	100	97	91
∅ effect on kurtosis (in %)			
$N = 100$	100	91	89
$N = 200$	100	90	89
$N = 500$	100	90	90
$N = 1000$	100	93	93
$N = 2000$	100	91	89
∅ effect on volume (in %)			
$N = 100$	100	72	50
$N = 200$	100	73	50
$N = 500$	100	74	51
$N = 1000$	100	74	51
$N = 2000$	100	74	51

Based on 1000 simulation runs lasting 340 trading rounds

three three important values of the Tobin tax, $\tau \in \{0, 0.5, 1\%\}$ and $N = 100$. Again, we scale the percentage figures in Table 4 such that the scenario with $\tau = 0$ is always equal to 100%, and the "middle" scenario with $\sigma_F^2 = 0.019$ is associated with our previous results in Figs. 1, 2, 4, and 5.

We must note here that the value of $\sigma_F^2 = 0.019$ is obtained as a result of the model calibration (see Sect. 3.1); therefore, we test the sensitivity of the model with respect to 25, 50, 150, and 200% of its calibrated value. By fixing one of the coefficients away from its optimal calibrated value, we in fact deflect the model into a suboptimal setting area in which some parameters might not necessarily keep their previously calibrated optimal values under a new potential calibration run. In other words, the relative effect of the Tobin tax τ might be to some extent distorted relative to the benchmark case. We therefore need to interpret the related results with caution considering this departure from calibration-based optimality of the model setting.

For Volatility, we observe that the effect of delivering a moderate decrease is violated only for the lowest value of $\sigma_F^2 = 0.00475$. This finding is a bit puzzling, but not surprisingly, the system with this minimal variation in the expectations of agents might produce distorted behavior that is potentially strongly affected by a unique

Table 4 Sensitivity analysis: the effect of the intensity of the stochastic noise in the expectations of agents

	$\tau = 0\%$	$\tau = 0.5\%$	$\tau = 1\%$
∅ effect on volatility (in %)			
$\sigma_F^2 = 0.00475$	100	87	106
$\sigma_F^2 = 0.00095$	100	89	87
$\sigma_F^2 = 0.019$	100	96	90
$\sigma_F^2 = 0.0285$	100	99	94
$\sigma_F^2 = 0.038$	100	101	97
∅ effect on distortion (in %)			
$\sigma_F^2 = 0.00475$	100	69	93
$\sigma_F^2 = 0.00095$	100	68	61
$\sigma_F^2 = 0.019$	100	86	73
$\sigma_F^2 = 0.0285$	100	92	78
$\sigma_F^2 = 0.038$	100	100	94
∅ effect on kurtosis (in %)			
$\sigma_F^2 = 0.00475$	100	81	67
$\sigma_F^2 = 0.00095$	100	90	81
$\sigma_F^2 = 0.019$	100	93	90
$\sigma_F^2 = 0.0285$	100	91	89
$\sigma_F^2 = 0.038$	100	92	92
∅ effect on volume (in %)			
$\sigma_F^2 = 0.00475$	100	18	1
$\sigma_F^2 = 0.00095$	100	52	18
$\sigma_F^2 = 0.019$	100	72	50
$\sigma_F^2 = 0.0285$	100	77	63
$\sigma_F^2 = 0.038$	100	80	69

Based on 1000 simulation runs lasting 340 trading rounds

random generation of an extreme value. The case of the highest $\sigma_F^2 = 0.038$, on the other hand, suggests that for higher intensity of the stochastic noise, the interesting deflective behavior observed in Fig. 1 for low values of τ naturally shifts with higher values of τ. The overall effect retains its direction with an increasing tax τ but slowly diminishes, which we attribute to a natural effect of a more noisy setting. In other words, we are likely to observe a comparable effect to the benchmark case with $\sigma_F^2 = 0.019$ for higher values of the Tobin tax τ in cases of higher values of σ_F^2 and vice versa—this expected pattern is clearly observable for all statistics in Table 4. For related Distortions, we observe the very same behavior as for Volatility discussed above. Kurtosis seems to be influenced in a very regular manner, with a decreasing relative effect of τ together with an increasing σ_F^2, which again indicates that a comparable effect to the benchmark case is likely to be observed for higher values of the Tobin tax τ in cases of higher values of the stochastic noise intensity σ_F^2 and vice versa. Finally, the strongest relative distinctions (but a very clear pattern) is observed for Volume. Expectedly, in a constrained system with minimal variation in expectations of agents, even a very small tax $\tau = 1\%$ almost completely hinders the trading. With increasing variation in expectations, the effect is naturally decaying. Again, we are likely to observe comparable effects to the optimally calibrated benchmark case with the Tobin tax altered accordingly.

4 Conclusion

The question of how transaction costs affect price dynamics is not only interesting from a theoretical perspective but is also of practical importance, as certain regulators are currently considering imposing such a tax. Motivated by recent research demonstrating the substantial importance of the market microstructure when imposing a Tobin tax, we explore the impact of the tax in an artificial market cleared by a Walrasian auctioneer. This setting resembles the two-layered structure of real foreign exchanges more closely than a price impact function, which is often adopted in studies on the Tobin tax. In addition, by using Walrasian clearing, we circumvent the troublesome decision concerning the precise form of the price impact function, which has a substantial influence on the results. Considering the complexity of the interbank foreign exchange market, this approximation may not be perfect. However, as the impact of the Tobin tax on exchange rate dynamics in this particular setting has, to the best of our knowledge, not been examined in any study to date, we presume that this analysis might contribute to the mosaic of existing research and hope that it further illuminates some possible effects of the tax.

To evaluate the impact of the tax in this setting, we extend the model of De Grauwe and Grimaldi (2006) by including transaction costs. The original model consists of boundedly rational agents who apply a blend of fundamental and technical analyses to predict the future exchange rate. The ongoing competition between the forecasting rules creates chaotic price movements not dissimilar to those observed in real foreign exchange markets. The calibrated model exhibits features typical of real foreign exchange markets such as excess kurtosis and volatility, autocorrelation of absolute returns and no autocorrelation in raw returns.

By performing multiple simulations, we found that the tax is capable of reducing average distortions (by up to 29% in the case of a 1% tax) and average excess kurtosis (by up to 12% in the case of a 1% tax). This result is primarily achieved by eliminating long-lasting deviations from the fundamental value—something one could loosely refer to as speculative bubbles. The effect of the tax on average volatility is more complex. We found that for small values of the tax (0–0.3%), it delivers a negligible, but statistically significant, increase in volatility, while for larger values (0.3–1%), the tax decreases average volatility by up to 10% in the case of a 1% tax. As expected, the Tobin tax notably reduces trading volume.

Our results are robust to the volatility of the fundamental value. The sensitivity analysis further indicates that the Tobin taxs ability to reduce distortions (and, to a lesser extent, the volatility of returns) is inversely related to the volatility of the fundamental value. This relationship likely arises because the Tobin tax subdues not only erratic oscillations of the market exchange rate but also the adjustment toward its fundamental value.

Acknowledgements We are grateful to an anonymous reviewer for many insightful comments and suggestions.

Appendix

See Fig. 10.

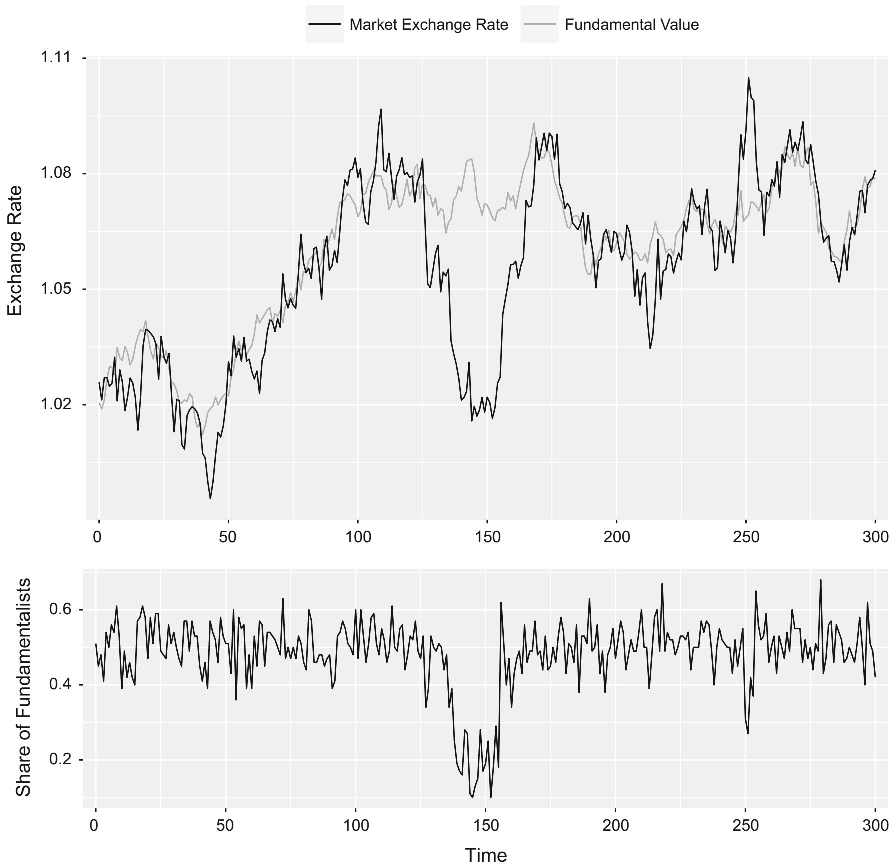

Fig. 10 Realization of time series under $\mathbf{p} = \mathbf{p}_0$ and $\tau = 0$: a realization of exchange rates capturing excess volatility and the formation of a persistent deviation during periods 120–170. The deviation emerged when several consecutive Gaussian shocks in fundamentals of negative sign happened to be sampled, which made the chartist forecasting rule more profitable. At that point, more traders adopted the chartist forecasting rule predicting a further decline of the exchange rate, which, in a self-fulfilling manner, indeed occurred. However, that speculation further increased the popularity of the chartist forecasting rule. This endogenous process continued until the market was almost exclusively populated by chartists, when, unable to fuel itself by recruiting more chartists, it finally slowed down and the exchange rate settled at an unstable equilibrium of approximately 1.02. Another Gaussian shock then occurred that triggered a similar process, except in the opposite direction

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*(3), 481–510.

Bianconi, G., Galla, T., Marsili, M., & Pin, P. (2009). Effects of Tobin taxes in minority game markets. *Journal of Economic Behavior and Organization*, *70*(1), 231–240.

BIS. (2013). *Triennial central bank survey*. Tech. rep., Bank for International Settlements, available online at: http://www.bis.org/publ/rpfx13fx.pdf. Accessed 20 October 2015.

Brock, W. A., & Hommes, C. H. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, *22*(8–9), 1235–1274.

Chen, S. H., Chang, C. L., & Du, Y. R. (2012). Agent-based economic models and econometrics. *The Knowledge Engineering Review*, *27*, 187–219. doi:10.1017/S0269888912000136, http://journals.cambridge.org/article_S0269888912000136

Cheung, Y. W., Chinn, M. D., & Marsh, I. W. (2004). How do UK-based foreign exchange dealers think their market operates? *International Journal of Finance and Economics*, *9*(4), 289–306.

Cont, R., Kukanov, A., & Stoikov, S. (2014). The price impact of order book events. *Journal of Financial Econometrics*, *12*(1), 47–88. doi:10.1093/jjfinec/nbt003, http://jfec.oxfordjournals.org/content/12/1/47.abstract, http://jfec.oxfordjournals.org/content/12/1/47.full.pdf+html

ČNB Czech National Bank. (2014). Dataset of exchange rates: 2 Jan 2013–27 Dec 2013. Available online at: https://www.cnb.cz/cs/financni_trhy/devizovy_trh/kurzy_devizoveho_trhu/rok.txt?rok=2013. Accessed 20 October 2015.

De Grauwe, P., & Grimaldi, M. (2006). Exchange rate puzzles: A tale of switching attractors. *European Economic Review*, *50*(1), 1–33.

Demary, M. (2006). *Transaction taxes, trader's behavior and exchange rate risks*. Christian-Albrechts-University of Kiel Department of Economics, Working paper 14.

Demary, M. (2010). Transaction taxes and traders with heterogeneous investment horizons in an agent-based financial market model. *Economics: The Open-Access, Open-Assessment E-Journal*, *4*(2010–8), 1–44.

Ehrenstein, G. (2002). Cont–Bouchaud percolation model including Tobin tax. *International Journal of Modern Physics C*, *13*(10), 1323–1331.

Ehrenstein, G., Westerhoff, F., & Stauffer, D. (2005). Tobin tax and market depth. *Quantitative Finance*, *5*(2), 213–218.

Eisler, Z., Bouchaud, J. P., & Kockelkoren, J. (2012). The price impact of order book events: Market orders, limit orders and cancellations. *Quantitative Finance*, *12*(9), 1395–1419. doi:10.1080/14697688.2010.528444.

Fagiolo, G., Moneta, A., & Windrum, P. (2007). A critical guide to empirical validation of agent-based models in economics: Methodologies, procedures, and open problems. *Computational Economics*, *30*(3), 195–226. doi:10.1007/s10614-007-9104-4.

Farmer, D. J., Gillemot, L., Lillo, F., Mike, S., & Sen, A. (2004). What really causes large price changes? *Quantitative Finance*, *4*(4), 383–397.

Flaschel, P., Hartmann, F., Malikane, C., & Proaño, C. R. (2015). A behavioral macroeconomic model of exchange rate fluctuations with complex market expectations formation. *Computational Economics*, *45*(4), 669–691.

Flood, M. D. (1994). Market structure and inefficiency in the foreign exchange market. *Journal of International Money and Finance*, *13*(2), 131–158.

Frankel, J. A., & Froot, K. A. (1990). Chartists, fundamentalists, and trading in the foreign exchange market. *The American Economic Review*, *80*(2), 181–185.

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

Hein, O., Schwind, M., & Spiwoks, M. (2006). A microscopic currency market model with a modified Tobin tax. In *Workshop on economic heterogeneous interacting agents (WEHIA 06)*.

Hommes, C. H. (2006). Heterogeneous agent models in economics and finance. *Handbook of Computational Economics*, *2*, 1109–1186.

Kempf, A., & Korn, O. (1999). Market depth and order size 1. *Journal of Financial Markets*, *2*(1), 29–48. doi:10.1016/S1386-4181(98)00007-X, http://www.sciencedirect.com/science/article/pii/S138641819800007X

Kirchler, M., Huber, J., & Kleinlercher, D. (2011). Market microstructure matters when imposing a Tobin tax—Evidence from the lab. *Journal of Economic Behavior and Organization*, *80*(3), 586–602.

Kroll, Y., Levy, H., & Markowitz, H. M. (1984). Mean–variance versus direct utility maximization. *The Journal of Finance*, *39*(1), 47–61. doi:10.1111/j.1540-6261.1984.tb03859.x.

Lavička, H., Lichard, T., & Novotný, J. (2013). *Sand in the wheels or the wheels in sand? Tobin taxes and market crashes*. FNSPE, CERGE-EI and Cass Business School, City University London, UK. Working paper.

LeBaron, B. (2005). *Agent-based computational finance*. Tech. rep., International Business School, Brandeis University, USA.

Li, F. (2015). *Pboc's yi proposes punitive Tobin tax to deter yuan speculators*. Available online at: http://www.bloomberg.com/news/articles/2015-10-02/pboc-s-yi-suggests-punitive-tobin-tax-to-deter-yuan-speculators. Accessed 27 November 2015.

Lillo, F., Farmer, J. D., & Mantegna, R. N. (2002). *Single curve collapse of the price impact function for the New York stock exchange*. Tech. rep., Istituto Nazionale per la Fisica della Materia and Santa Fe Institute.

Liu, S., & Zhu, Z. (2009). Transaction costs and price volatility: New evidence from the Tokyo stock exchange. *Journal of Financial Services Research*, *36*(1), 65–83.

Mannaro, K., Marchesi, M., & Setzu, A. (2005). The impact of transaction taxes on traders' behavior and wealth: A microsimulation. In *10th annual workshop on economic heterogeneous interacting agents (WEHIA 2005)*.

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*(2), 445–462.

Oberlechner, T. (2001). Importance of technical and fundamental analysis in the European foreign exchange market. *International Journal of Finance and Economics*, *6*(1), 81–93.

Pellizzari, P., & Westerhoff, F. (2009). Some effects of transaction taxes under different microstructures. *Journal of Economic Behavior and Organization*, *72*(3), 850–863.

Plerou, V., Gopikrishnan, P., Gabaix, X., & Stanley, H. E. (2002). Quantifying stock-price response to demand fluctuations. *Physical Review E*, *66*(027), 104. doi:10.1103/PhysRevE.66.027104.

Potters, M., & Bouchaud, J. P. (2003). More statistical properties of orderbooks and price impact. In *Proceedings of the International Econophysics Conference*. *Physica A: Statistical Mechanics and Its Applications*, *324*(1–2):133–140. doi:10.1016/S0378-4371(02)01896-4, http://www.sciencedirect.com/science/article/pii/S0378437102018964

Sewell, M. (2011). *Characterization of financial time series*. UCL Department of Computer Science, Research Note RN (vol. 11, no. 1, pp. 1–35).

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk*. *The Journal of Finance*, *19*(3), 425–442. doi:10.1111/j.1540-6261.1964.tb02865.x.

Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, *41*(1), 1–20.

Smid, M. (2016). Estimation of zero-intelligence models by L1 data. *Quantitative Finance*, *0*(0), 1–22. doi:10.1080/14697688.2016.1149612.

Stonedahl, F., & Wilensky, U. (2010). *Behaviorsearch (computer software)*. Evanston, IL: Center for Connected Learning and Computer Based Modeling, Northwestern University. Available online at: http://www.behaviorsearch.org

Stonedahl, F. J. (2011). *Genetic algorithms for the exploration of parameter spaces in agent-based models*. PhD thesis, Northwestern University, Evanston, IL, USA, aAI3489404.

Taylor, M. P., & Allen, H. (1992). The use of technical analysis in the foreign exchange market. *Journal of International Money and Finance*, *11*(3), 304–314.

Tobin, J. (1978). A proposal for international monetary reform. *Eastern Economic Journal*, *4*(3–4), 153–159.

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

Weber, P., & Rosenow, B. (2005). Order book approach to price impact. *Quantitative Finance*, *5*(4), 357–364. doi:10.1080/14697680500244411.

Westerhoff, F. (2003). Heterogeneous traders and the Tobin tax. *Journal of Evolutionary Economics*, *13*(1), 53–70.

Westerhoff, F. H. (2004a). Market depth and price dynamics: A note. *International Journal of Modern Physics C*, *15*(07), 1005–1012.

Westerhoff, F. H. (2004b). Speculative dynamics, feedback traders and transaction taxes: A note. *Jahrbuch für Wirtschaftswissenschaften/Review of Economics*, *55*, 190–195.

Westerhoff, F. H. (2008). The use of agent-based financial market models to test the effectiveness of regulatory policies. *Jahrbücher für Nationalökonomie & Statistik*, *228*, 195–227.

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*(2), 293–322.
Wilensky, U. (1999). *NetLogo. Center for connected learning and computer-based modeling*. Evanston, IL: Northwestern University.