

Fricke, Daniel; Lux, Thomas

Working Paper

The effects of a financial transaction tax in an artificial financial market

Kiel Working Paper, No. 1868

Provided in Cooperation with:

Kiel Institute for the World Economy – Leibniz Center for Research on Global Economic Challenges

Suggested Citation: Fricke, Daniel; Lux, Thomas (2013) : The effects of a financial transaction tax in an artificial financial market, Kiel Working Paper, No. 1868, Kiel Institute for the World Economy (IfW), Kiel

This Version is available at:

<https://hdl.handle.net/10419/80349>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Working Papers

Kiel Institute
for the World Economy



The Effects of a Financial Transaction Tax in an Artificial Financial Market

by Daniel Fricke and Thomas Lux

No. | August 2013

Kiel Working Paper No. | August 2013

Title: The Effects of a Financial Transaction Tax in an Artificial Financial Market

Author: Daniel Fricke, Thomas Lux

Abstract:

We investigate the effects of a Financial Transaction Tax (FTT) in an order-driven artificial financial market. FTTs are meant to limit short-term speculative behavior by reducing the amount of excess liquidity in the system. To quantify these effects, adjustments in trading strategies and their effects on liquidity need to be taken into account. We model an agent-based continuous double-auction, allowing for a continuum of investment strategies within the chartist/fundamentalist framework. For certain parameter combinations, our model is able to reproduce certain stylized facts of financial time-series. We find largely positive effects of the FTT for small tax rates. Additionally, for large tax rates we find the effects not to be as negative as previously found.

Keywords: Transaction Tax, Tobin Tax, Market Microstructure, Agent-Based Models, Speculative Bubbles

JEL classification: H20, C63, D44

Daniel Fricke

Institute for New Economic Thinking,
Oxford Martin School;
CABDyN Complexity Centre,
Saïd Business School,
University of Oxford,
Park End St, Oxford OX1 1HP.
E-mail: daniel.fricke@sbs.ox.ac.uk

Kiel Institute for the World Economy
24100 Kiel, Germany
Telephone: +49 431 8814 229

Thomas Lux

Kiel Institute for the World Economy
24100 Kiel, Germany
Telephone: +49 431 8814 278
E-mail: Thomas.lux@ifw-kiel.de

Christian-Albrechts-University Kiel
Department of Economics
Chair of Monetary Economics and
International Finance
24098 Kiel, Germany

Bank of Spain chair of Computational Finance
Department of Economics, University Jaume I
Castellón, Spain

The responsibility for the contents of the working papers rests with the author, not the Institute. Since working papers are of a preliminary nature, it may be useful to contact the author of a particular working paper about results or caveats before referring to, or quoting, a paper. Any comments on working papers should be sent directly to the author.

Coverphoto: uni_com on photocase.com

The Effects of a Financial Transaction Tax in an Artificial Financial Market.[†]

Daniel Fricke^{‡§} Thomas Lux^{§¶||}

This version: August 2013

Abstract

We investigate the effects of a Financial Transaction Tax (FTT) in an order-driven artificial financial market. FTTs are meant to limit short-term speculative behavior by reducing the amount of excess liquidity in the system. To quantify these effects, adjustments in trading strategies and their effects on liquidity need to be taken into account. We model an agent-based continuous double-auction, allowing for a continuum of investment strategies within the chartist/fundamentalist framework. For certain parameter combinations, our model is able to reproduce certain stylized facts of financial time-series. We find largely positive effects of the FTT for small tax rates. Additionally, for large tax rates we find the effects not to be as negative as previously found.

Keywords: Transaction Tax, Tobin Tax, Market Microstructure, Agent-based models, Speculative Bubbles

JEL-Codes: H20, C63, D44.

[†]Funding for an earlier version of this paper by the Paul Woolley Centre for the Study of Capital Market Dysfunctionalities at the University of Technology Sydney is acknowledged. We are grateful for helpful comments by Karl Finger, Reiner Franke, Tony He, and Daniel Ladley. Correspondence: daniel.fricke@sbs.ox.ac.uk

[‡]Institute for New Economic Thinking, Oxford Martin School; CABDyN Complexity Centre, Saïd Business School, University of Oxford, Park End St, Oxford OX1 1HP.

[§]Kiel Institute for the World Economy, Hindenburgufer 66, 24105 Kiel.

[¶]Department of Economics, University of Kiel, Olshausenstr. 40, 24118 Kiel

^{||}Banco de España Chair in Computational Economics, University Jaume I, Campus del Riu Sec, 12071 Castellón.

1 Introduction and Existing Literature

For several decades the financial transaction tax (FTT) has been discussed as an instrument to curb financial market volatility, cf. Keynes (1936), chapter 12, and Tobin (1978). Only recently -given the surging government deficits from responses to the global financial crisis- the focus has shifted to the FTT's large potential monetary revenues.¹ In this paper we investigate the effects of a FTT in an agent-based artificial financial market.

The FTT's appeal stems from its potential to limit short-term speculative behavior, and thus transaction volumes, on financial markets. This seems a reasonable aim given the divergence of financial market and 'real' activity during the last decades, when increases in financial market transaction volumes continuously exceeded those of the real economy. The exponential growth of financial transaction volumes was fueled by a continuous fall in transaction costs for many assets due to the technological progress in computer-based trading and an increased competition between stock exchanges. One result of this development is the increased presence of so-called high-frequency trading (HFT), which is predominantly employed by large hedge funds.² Indeed, higher liquidity³ seems to have come along with higher fragility in the sense that financial crises, i.e. the build-up and bursting of speculative bubbles, became more frequent.⁴ In this way, a FTT that favors longer-term investments could have the effect of reducing the decoupling of financial markets from real activity and could additionally free resources from the financial sector for more productive uses.⁵

Critics of the FTT, most importantly from the financial industry, usually bring forward the following arguments: (1) market liquidity will dry up, (2) volatility may thereby in fact increase, (3) banks will pass on the tax burden to firms and other bank customers, raising capital costs in general, and (4)

¹These revenues are estimated to range between 1 and 3% of national GDPs. See, e.g. Pollin *et al.* (2003).

²We should note, however, that many economists actually favor HFT, arguing that more liquid markets should be much more resilient, cf. Brogaard (2010). This might be justified if HFT activities were largely equivalent to market making. However, insofar as many of these strategies might have a destabilizing tendency, their 'net effect' on market efficiency and volatility might be ambiguous.

³Liquidity is the ability to trade large size quickly, at low costs, see Harris (2003).

⁴Bordo *et al.* (2001) find that the frequency of crises since 1973 has been twice that of the Bretton Woods and classical gold standard periods. Two important explanatory factors are financial globalization and expectations of bail-outs encouraging financial institutions to take on higher risks.

⁵The different activity patterns of financial markets and goods markets are also emphasized by Aoki and Yoshikawa (2007), chapter 10.

there is a danger of capital flights from a taxed market towards untaxed markets. In this paper we are concerned with the first two (interrelated) points.⁶

High liquidity, i.e. small transaction costs, fuels excess volatility (compared to ‘fundamentals’) as it makes round-trips relatively cheap, cf. Shiller (1981). Empirical evidence suggests that FTTs, despite applying for all market participants, harm short-term speculators disproportionately more. For example, it has been found that an increase of a transaction tax has increased asset holding periods, while transaction volumes have decreased.⁷ However, this does not imply that volatility will decrease as well. In theory, there could be a U-shaped relationship: for small tax rates volatility should decrease, since (destabilizing) short-term oriented speculation becomes unprofitable. However, larger tax rates will affect (stabilizing) longer-term strategies as well, thereby reducing liquidity and potentially increasing volatility. Empirical evidence on point (2) is therefore rather mixed: some studies find that volatility decreases, increases or does not react at all in response to a tax increase.⁸

Given these contradicting results, simulations of artificial financial markets are a promising way to non-invasively evaluate the effects of regulatory measures in general, see Westerhoff (2008) for a discussion. More detailed (realistic) models are usually hard to tackle analytically, so numerical simulations are needed. Agent-based models are such computerized simulations, containing a number of components (agents) interacting with each other through prescribed rules, thus taking all the necessary ingredients for modelling complex systems into account, cf. Aoki (2002). Numerous agent-based models, usually within the chartist-fundamentalist framework, are able to replicate many of the stylized facts characterizing financial market data.⁹ When used to evaluate regulatory policies, however, using overly simplified models could affect the conclusions. For example, many authors assume that a market-maker provides infinite liquidity, in which case FTTs are potentially stabilizing for small tax rates. For a single asset market, see Ehrenstein

⁶We are currently dealing with point (4) as well. The basic idea is to have two asset markets where only one of them is being taxed.

⁷See Jackson and O'Donnell (1985) and Baltagi *et al.* (2006). For example Baltagi *et al.* (2006) find that a tax increase from 0.3 to 0.5% reduced trading volume in China by roughly 1/3.

⁸See, e.g. Jones and Seguin (1997), Hau (2006) and Roll (1989), respectively.

⁹An early example is Beja and Goldman (1980). See also LeBaron *et al.* (1999), Challet and Zhang (1997), Chiarella and Iori (2002), Lux and Marchesi (1999, 2000), Lux and Schornstein (2005), Raberto *et al.* (2003) and Chiarella *et al.* (2009). Among others, Allen and Taylor (1990) and Menkhoff (1998) provide empirical evidence on the use of chartist and fundamentalist strategies.

(2002) and Westerhoff (2003), and Westerhoff (2004a). However, Giardina and Bouchaud (2003) find that only substantial trading costs will actually stabilize the market, while a small tax (of the order of a few basis points) would have no real effect.¹⁰ However, since liquidity is a major determinant of volatility in real markets, cf. Mike and Farmer (2008), it is crucial when discussing the effects of FTTs. In fact, the findings from both laboratory experiments and simulation studies indicate that the effects of a FTT may depend on the structure of the market, see Kirchler *et al.* (2012) and Pellizzari and Westerhoff (2009). Therefore we explicitly take the microstructure of real markets and provision of liquidity into account by simulating an order-driven continuous double-auction (CDA).

While the models based on the marker-maker setting typically incorporate important psychological factors that drive the system's properties, e.g. through herding and imitation, CDA models work at shorter time-scales where psychological factors are either hard to model or simply assumed to be absent. For example, Bak *et al.* (1997) treat the limit order book (LOB) as a system of particles with each particle (order) having a mass (order size) and a price (spatial position). Price variations stem from diffusion and annihilation of particles, well-known processes in physics, which allows to obtain analytical results. Even though many important insights can be gained from such approaches, the usual criticism is that these models operate within a zero intelligence framework.¹¹ This is the exact opposite to 'homo oeconomicus' in mainstream economics, but traders are unlikely to be either fully rational or plainly stupid. Another problem is that, by avoiding detailed behavioral assumptions, these models typically ignore budget constraints and wealth dynamics. Nevertheless, since these models are able to replicate certain stylized facts of LOB data, the structure of the trading protocol is likely to have a significant effect on the data-generating process. To date, few attempts have been made to model the LOB based on detailed strategic interactions between many boundedly rational agents, while incorporating economic constraints.¹² Our model aims at bridging the gap between models with short

¹⁰For two ex-ante identical markets, with one country unilaterally introducing the tax, Westerhoff and Dieci (2006) finds that the taxed market is stabilized while volatility in the tax haven strongly increases. Using laboratory experiments with markets of different size, Hanke *et al.* (2010) find that volatility decreases (increases), when the tax is introduced in the large (small) market.

¹¹See Cliff and Bruten (1997). Gode and Sunder (1993) were the first to introduce the zero intelligence framework in a trading setup. The authors show that the double auction mechanism ensures allocative efficiency irrespective of the level of rationality of the agents.

¹²To our knowledge, Chiarella and Iori (2002) were the first to incorporate trading strategies into a CDA setup. The authors state that without these strategies, it is impossible to generate realistic time-series. See also Chiarella *et al.* (2009).

and long time-scales.

To our knowledge, only few studies have been dealing with FTTs in detailed order-driven markets. Two examples are Mannaro *et al.* (2008) and Pellizzari and Westerhoff (2009). Mannaro *et al.* (2008) use a zero intelligence framework combined with a once-a-day supply-demand based market-clearing rule, deleting all orders not executed during the clearing session and thus strongly limiting their impact. In this setting the FTT is found to be destabilizing. Pellizzari and Westerhoff (2009) compare the effects of the FTT in different market settings. The main finding is that the FTT destabilizes a CDA market, while it stabilizes a dealership market where specialists provide abundant liquidity. One important assumption in the dealership setup is that the dealer (or market-maker) is exempt of the tax, which is hardly the case for a general FTT that should apply to all market participants. Moreover, these studies suffer from the assumption that all agents act with equal probability, i.e. they neglect the importance of heterogeneous investment horizons.¹³ In this way, the FTT's effect of more severe taxation of short-term speculation is missed.¹⁴ Another novelty compared to the existing literature is that the limit orders in our model emerge from a rule-based decision process, rather than from a pure zero intelligence framework.

In our model two groups of agents compete in the market: Noise traders act as liquidity providers, by posting random orders. Informed traders use information about past prices and the fundamental value when forming their price expectations. As in Youssefmir *et al.* (1998), their price expectations depend on three different time horizons (Figure 1): the investment horizon (denoted by H^w) basically defines how often a particular agent acts and how long his planning horizon is when making investment decisions. Two different trend horizons model the trend chasing behavior of agents: the backward trend horizon (H^b) defines how many past price observations are relevant when calculating the trend. The forward trend horizon (H^t) defines how long the agent expects his calculated trend to last before the price will start returning to the fundamental value. This setting is very flexible concerning the strategies and we essentially allow for all combinations of time horizons within a certain set. Most importantly, the relative size of forward trend and investment horizon defines whether an agent is a chartist, a fundamentalist or something in between.¹⁵

¹³The importance of time-scales is a relatively recent research topic. See for example Zumbach and Lynch (2001) and Borland and Bouchaud (2005).

¹⁴See Anufriev and Bottazzi (2004) for the importance of investment horizons. In an infinite-liquidity model, Demary (2010) incorporates investment horizons and finds that investment horizons increase for small tax rates.

¹⁵Below, we will impose symmetry between the trend horizons, i.e. set $H^b = H^t$.




Figure 1: Time horizons in the model. In order to reduce the complexity of the model, we will set $H^b = H^t$ in the following. More details can be found below.

Our main conclusions can be summarized as follows: First, the model is able to replicate certain stylized facts of real financial time-series for several parameter combinations, e.g. the model replicates the building up and bursting of price bubbles. Second, we find the usual trade-off between monetary revenues (a kind of Laffer curve) and stability, as higher tax revenues come along with higher volatility. This finding is in line with the results from the existing literature. However, we find somewhat different results for very small and large tax rates, indicating that the effects of the tax may not be entirely negative. In any case, the tax allows to generate substantial tax revenues, which could be used for a number of more productive purposes.

The remainder of this paper is structured as follows: Section 2 introduces the structure of the model. Section 3 presents pseudo-empirical results and Section 4 concludes.

2 Model

2.1 CDA and Information

The basic model structure is as follows: the financial market consists of N heterogeneous agents trading one asset (which pays no dividend and has fixed supply) against cash. Cash earns zero interest, so there are no interest payments (or they are spent elsewhere). In order to avoid cash being sucked out of the system due to the FTT, tax revenues are regularly redistributed equally among all agents. The market is order-driven and the quoted price p_t (midprice) is the average of the best ask (a_1) and best bid (b_1) in the limit-order book (LOB), while $a_1 - b_1 > 0$ is the bid-ask spread. In case there are no orders in the LOB, the quoted price is simply the last quoted price. Prices are discrete and can only be submitted on a specified grid, defined by

the tick-size Δ .

We simulate a CDA market, where two types of orders exist: a market order specifies to buy or sell a certain amount of the asset at the best available price. A limit order additionally specifies a limit price at which the agent is still willing to trade. In general, market orders are guaranteed execution but not price, since with a market order a trader is assured that it will be executed against the best price in the LOB within a short amount of time. Limit orders, on the other hand, are guaranteed price but not execution as they will only be executed at, or below (above) for buy (sell) orders, the specified price which may never happen if no matching order is found. Each transaction involves a market order transacting against an existing limit order. Generally speaking, patient investors are more likely to place limit orders, while impatient investors place market orders. The price of immediacy is simply the bid-ask spread. Thus, choosing a limit price is a strategic decision that induces a trade-off between patience and (expected) profit, cf. Harris and Hasbrouck (1996). The price dynamics within the LOB are therefore driven by three forces: limit order arrivals, market order arrivals (i.e. trades) and cancellations of limit orders.¹⁶

Table 1 illustrates the structure of the CDA: Buy orders are stored on the bid side (left), while sell orders are on the ask side (right). The two relevant features are price-priority and time-priority. Price-priority means that the best orders are placed on top of the book, i.e. the order with the highest bid price (best bid) and the order with the lowest ask price (best ask). Obviously, orders stored in the LOB cannot be executed immediately: in the example, the best bid (100.50) is smaller than the best ask (101.50) such that currently no trade is possible. Time-priority means that, after providing price-priority, orders with the same limit price are sorted according to submission date. Therefore the best bid is placed above the second best bid (with the same limit price), since it was submitted earlier to the LOB. In the example the quoted price would be $\frac{a_1+b_1}{2} = 101.00$. Note however, that this quoted price is just a proxy for the price of an immediate transaction: For example, assume there arrives a new sell order with a quantity of 25 and limit price 100.00. In this case the order is marketable, such that the offered 25 assets are sold at a price of 100.00, which differs from the quoted price of 101.00.¹⁷

Despite disregarding dividends,¹⁸ we assume a constant and positive fun-

¹⁶There is a growing literature on the stylized facts of LOB data, see e.g. Bouchaud *et al.* (2002), Farmer *et al.* (2004), and Mike and Farmer (2008).

¹⁷Note how time- and price-priority favor the buying agent, i.e. the trade initiator, in the Example: He submitted a limit price of 100.50 but only pays 100.00.

¹⁸Dividend payments are negligible on a short-term basis, since they are only paid once a year and usually only have a small effect on wealth. Ignoring dividend payments simplifies

Bid			Ask		
Price	Quantity	Time	Price	Quantity	Time
100.50	20	12:38:39	101.50	10	12:15:01
100.50	10	12:42:08	105.00	5	12:28:40
95.60	8	12:10:52	110.50	10	09:01:05
87.90	5	10:15:23	125.50	8	12:40:18

Table 1: Example for a LOB at a certain point in time.

damental value p^f of the asset.¹⁹ Only informed agents know the fundamental value, whereas the state of the LOB and the history of the quoted prices are public information.²⁰ Each agent is initially endowed with $S_0^i = N_s$ assets and $C_0^i = N_s p_0$ units of cash. We impose short-selling and capital constraints, such that $S_t^i, C_t^i \geq 0$ for all t, i .

2.2 Trader Types

Two groups of traders, differing in the way they form their price expectations and choose their limit prices, compete in the market: there is a fraction $\theta \in [0, 1]$ of informed traders and a fraction $(1 - \theta)$ of noise traders. Hence there are $N\theta = N_\theta$ informed and $(N - N_\theta)$ noise traders.²¹

One general point worth mentioning is that all agents act strategically in terms of their order submission. This means that all agents create two orders, where the first order (with corresponding limit price $p^{(1)}$) is always being sent to the LOB, while the submission of the second order (limit price $p^{(2)}$) depends on whether the first order was fully executed. In real markets the second order corresponds to a take-profit order. For example, an agent buying the asset today will try to sell it again at a higher price later on. Quite interestingly, while this ‘buy-low/sell-high’ framework is straightfor-

the analysis, since (without the FTT and with cash earning zero interest) the total amount of stocks and cash is constant over time.

¹⁹We model trading dynamics on very short time-scales where the fundamental value is unlikely to change significantly. Furthermore, this assumption makes it possible to ignore adverse selection problems due to news arrival. As long as the fundamental volatility is relatively small (compared to the volatility of noise traders expectations), this does not affect many of the qualitative results.

²⁰Of course, ‘true’ fundamentals are unobservable in reality. Another interesting feature would be to model costly acquisition of the fundamental value.

²¹Liquidity providers would be another possible label for the group of noise traders. The term noise traders however, emphasizes the random nature of their random limit price determination.

ward and incorporates the aim of wealth/utility maximization, we do not know of a single study dealing with these conditional (take-profit) orders within a simulation model.

2.2.1 Noise Traders

Noise traders are typically modelled as simple-minded investors. In our setting, their limit prices are chosen randomly around the current best prices according to

$$p^{(k)} = p_t e(\epsilon_t^k), \quad (1)$$

with $k = 1, 2$ and ϵ denoting $\text{iin}(0, \sigma_\epsilon)$ random numbers. With this approach, we directly obtain two limit prices, $p^{(1)}$ for the first (unconditional) order and $p^{(2)}$ for the second (conditional) take-profit order. Based on these limit prices, we identify the market side that the trader acts upon by comparing the limit prices, as we impose that their orders should not create a sure loss, i.e. the agent buys first and sells later if $p^{(2)} > p^{(1)}$, or the agent sells first and buys later if $p^{(2)} < p^{(1)}$. In everything that follows, we treat market orders as limit orders with limit prices equal to the best opposite price (marketable limit orders). For example, an agent will submit a buy limit price at most equal to the best ask. Note that we could make the distribution of ϵ explicitly dependent on other variables, for example positively related to the historical volatility of returns. This type of volatility feedback, i.e. having noise traders choose their limit prices from a broader distribution when historical volatility is large, would then for example generate volatility clustering mechanically.²² However, given that we are more interested in the behavior of informed traders and their effect on the system's properties, we are reluctant to impose this feedback and hold σ_ϵ constant in the following.²³

We would like to stress here that we incorporate noise traders as liquidity providers in our model, since, as will become clear in the next section, it is possible that many of the informed traders appear on the same market side. Thus, noise traders provide liquidity when the informed agents are not willing to do so or at least not sufficiently to generate trades (and hence price

²²For example, the noise traders in Raberto *et al.* (2003) are constructed exactly in such a way, i.e. in their model informed traders are not necessary to reproduce volatility clustering and excess kurtosis. However, this is a very 'direct' way to guarantee volatility clustering in a model. It is not clear why agents should behave like this and there is in fact some evidence that past price volatility tends to lead the arrival of limit orders, see Zovko and Farmer (2002).

²³Note that since the width of the distribution is fixed, noise traders are more likely to submit market orders when the spread is small, while they are more likely to submit limit orders when the spread is large. This is in line with empirical findings, e.g. Biais *et al.* (1995), Bae *et al.* (2003), and Foucault *et al.* (2005).

changes). Given the random structure of their limit prices, noise traders tend to lose money to the informed traders on average, in particular when there are pronounced trends. A relatively small number of noise traders is already sufficient for the model to work. However, with such a small number of noise traders the generated bubbles appear relatively smooth, such that prices and returns would be autocorrelated. Therefore, θ will not be too large during the simulations.

2.2.2 Informed Traders: Chartists and Fundamentalists

Price Expectations. Whether an informed agent buys or sells the asset depends on his expectation of the asset's future price at the end of his investment horizon. When forming price expectations, informed traders use information about past prices and fundamental values. Expectations evolve, following Youssefmir *et al.* (1998), as

$$\frac{d\hat{p}_{t+\tau}^i}{d\tau} = -\frac{\hat{p}_{t+\tau-1}^i - p^f}{H^{t,i}} + \left(T_t^i + \frac{p_t - p^f}{H^{t,i}}\right) e\left(-\frac{\tau}{H^{t,i}}\right), \quad (2)$$

where \hat{p}_t^i is agent i 's expected price at t , T_t^i is the calculated trend and $H^{t,i}$ is the forward trend horizon over which agent i expects the trend to last. The trend itself is an weighted average rate of price changes over a backward horizon $H^{b,i}$ of the form

$$T_t^i = \frac{1}{H^{b,i}} \int_{t_0}^t \frac{dp}{d\tau} e\left(-\frac{t-\tau}{H^{b,i}}\right) d\tau, \quad (3)$$

where dp is the price change between $(t-\tau)$ and t and $H^{b,i}$ is the backward trend horizon of i . As noted by Youssefmir *et al.* (1998), Eq. (3) can be integrated as $T_t^i = \frac{p_t - \langle p_t \rangle_{H^{b,i}}}{H^{b,i}}$, where $\langle p_t \rangle_{H^{b,i}}$ is the exponential average price over the horizon $H^{b,i}$. Thus the trend measures the deviation from the moving average of prices, which is a popular approach among technical analysts.

The evolution of trends can be obtained by taking the time derivative of Eq. (3) which yields

$$\frac{dT_t^i}{dt} = \frac{1}{H^{b,i}} \left(\frac{dp}{dt} - T_t^i \right). \quad (4)$$

Subject to the boundary condition $\hat{p}_t^i = p_t$, each agent formulates his expected price development over the next $H^{w,i}$ time-steps via Eq. (2) using the calculated trend from Eq. (4). This system incorporates, depending on the corresponding horizons, chartist and fundamentalist components. In principle, all agents are fundamentalists in the sense that for $H^{w,i} \rightarrow \infty$ (given

$H^{t,i}$) the expected price will collapse towards the fundamental value. As agents do not have infinite investment horizons in general, the relative magnitude of $H^{t,i}$ to $H^{w,i}$ matters: agents with a small value of $H^{t,i}/H^{w,i}$ can be considered as fundamentalists, a large value indicates a more chartist strategy, and intermediate values are a combination of both. We should stress already here that we found heterogeneity in H^b to be of minor importance compared to H^w and H^t . Therefore, in everything that follows, we simply take $H^{b,i} = H^{t,i}$ for all agents. Note that this substantially reduces the total number of strategies.

Technically, the nonlinear price expectations are influenced by three terms: first, agents expect the observed trend and the difference between price and fundamental value to continue in the near term (this is the second part on the right-hand side of Eq. (2)). However, the influence of this term decreases for increasing τ (and hence for large $H^{w,i}$). Second, via the decreasing impact of the calculated trend, the agent expects the price to eventually relax towards the fundamental value at a rate of $-\frac{\hat{p}_t^i - p^f}{H^{t,i}}$. Note that fundamentalism is defined in terms of the expected price at the end of the investment horizon, but a fundamentalist may nevertheless try to make a profit based on short-term trends.




Figure 2: Example: Development of expected price for different forward trend horizons, $T_t^i = 1$, $H^{w,i} = 500$, $p_t = 110$ and $p_t^f = 100$.

In the simulations, we discretize Eqs. (2) and (4), with agents forming their price expectations over the next $H^{w,i}$ days and updating the trends at the end of each day. We will discuss the different time-scales of our model in more detail in section 2.5. As an illustration, Figure 2 shows, for different

forward trend horizons, the expected price development of an agent with $T_t^i = 1$ and an investment horizon of 500. The current price is 110 and the fundamental value equals 100. Obviously, for relatively small trend horizons the agent expects the price to revert towards the fundamental value soon. For larger trend horizons, the agent expects the trend to last in the near term but the price to revert towards the fundamental value at the end of his planning period. For very large trend horizons, this no longer holds. Consequently there is a low level of speculation for small trend horizons, in which case the dynamics are dominated by the first term in Eq. (2).

Note that this way of modelling informed agents implies that all of them are trend-followers, at least to some extent. This clearly affects the autocorrelation of prices and returns and makes the (in-)efficiency of the market a serious issue, at least when the fraction of informed traders is not too small.

Limit Prices. The limit price determination of informed agents can be split into three parts: In the first part the agent uses Eqs. (2) and (4) to forecast the evolution of the midprice between t and $t + H^{w,i}$. In case the agent expects the price to be higher (lower) than the current price he will submit a buy (sell) order. The limit prices of this order and the corresponding conditional order depend on the expected development of the price between t and $t + H^{w,i}$.

In the second step, the agent uses information about the midprice and a proxy for the expected price volatility to form more detailed expectations about the best bid and ask over time. As proxy for the expected price volatility, we use the average distance between the current best and second best prices on the two market sides.²⁴ This is calculated as

$$\hat{\sigma}_t^i = \frac{(a_2 - a_1) + (b_1 - b_2)}{2}, \quad (5)$$

where a_2 and b_2 denote the second best ask and bid prices, respectively. In this view, the average price change due to immediate market orders wiping out the best prices on either market side is being calculated. Note that this is the only channel in our model where (informed) traders use higher-order information on limit prices. We leave it to future research to model the information usage of LOB data in more detail.

Equipped with expectations about the midprice, the corresponding price volatility, and the decision to buy or sell from the first step, the agent then

²⁴The motivation for using the average gap is based on the finding that large price changes are in fact due to gaps in the LOB, see e.g. Farmer *et al.* (2004) and Farmer and Lillo (2004).

chooses his limit prices as follows: Defining

$$\begin{aligned}\hat{p}_{\max}^i &= \max\{\hat{p}_{t:\tau:t+H^{w,i}}\} \\ \hat{p}_{\min}^i &= \min\{\hat{p}_{t:\tau:t+H^{w,i}}\}\end{aligned}\tag{6}$$

gives the maximum and minimum of the expected midprice over the investment horizon of agent i . For an initial buy order, the agent then has to decide between buying the asset right away using a market order with price a_1 or setting a limit order at the minimum expected ask, i.e. $\hat{p}_{\min}^i + 0.5\hat{\sigma}_t^i$. Quite intuitively, he will choose the minimum of the two, in order to take favorable future developments into account. Therefore, if he expects the best ask to drop significantly below the current value in the near future, he will submit a limit order with a price below the current best ask. For the limit price of the conditional sell order the agent has to decide between the maximum expected bid, i.e. $\hat{p}_{\max}^i - 0.5\hat{\sigma}_t^i$ and the expected bid at the end of his investment horizon, and naturally takes the maximum of the two. Similar arguments can be used for the case of an initial sell order. More formally, the strategies of finding limit prices can be written as follows:

Definition 1. *Buy limit prices:* $\hat{p}_{t+H^{w,i}} > p_t$.

$$\begin{aligned}p^{(1)} &= b^i = \min(a_1, \hat{p}_{\min}^i + 0.5\hat{\sigma}_t^i) \\ p^{(2)} &= a^i = \max(\hat{p}_{\max}^i - 0.5\hat{\sigma}_t^i, \hat{p}_{t+H^{w,i}} - 0.5\hat{\sigma}_t^i).\end{aligned}\tag{7}$$

Definition 2. *Sell limit prices:* $\hat{p}_{t+H^{w,i}} < p_t$.

$$\begin{aligned}p^{(1)} &= a^i = \max(b_1, \hat{p}_{\max}^i - 0.5\hat{\sigma}_t^i) \\ p^{(2)} &= b^i = \min(\hat{p}_{\min}^i + 0.5\hat{\sigma}_t^i, \hat{p}_{t+H^{w,i}} + 0.5\hat{\sigma}_t^i).\end{aligned}\tag{8}$$

With this definition it may happen, in particular for large spreads, small trends and/or small deviations from the fundamental value, that the two limit prices are not in line with the agents' price expectations. For example, an agent with an expected price increase might end up with limit prices $p^{(1)} > p^{(2)}$. To ensure consistency, such orders will not be submitted to the LOB.

Figure 3 illustrates the concept for a sell order: Again the current price is equal to 110 and the fundamental value equals 100. The investment horizon is $H^{w,i} = 500$. The expected price at the end of the investment horizon is below the current price, therefore the agent will first sell the asset and try to buy it back at a lower price. Since the agent expects a positive trend to continue in the near term, \hat{p}_{\max}^i exceeds the current price, while \hat{p}_{\min}^i coincides with $\hat{p}_{t+H^{w,i}}$. She will, therefore, place a sell order with limit price $a^i = \hat{p}_{\max}^i - 0.5\hat{\sigma}_t^i$

and a conditional buy order with limit price $b^i = \hat{p}_{t+H^w,i} + 0.5\sigma_t^i$. The expected return (after tax) of the agent equals $r^e = |\ln(p^{(2)}/p^{(1)})| - \chi$, with χ denoting the two-sided tax rate. Thus, the tax drives a wedge between the agents' two limit prices.




Figure 3: Example: Limit price determination for a sell order.

Price Bubbles and FTT. At this point, a brief explanation on the building up and bursting of price bubbles, and their relation to the FTT, is in order:²⁵ At the beginning of the simulation, random strategies are assigned to the informed traders. Thus, approximately half of the population is willing to buy the asset while the other half is willing to sell, so the price will randomly move upwards or downwards. Suppose it moves upwards, then the trend signals of informed agents with high-frequency strategies will turn positive. This can be the beginning of the bubble, where exactly those traders induce additional positive price changes through their positive demand, which in turn affects the calculated trend variables of lower-frequency traders. Thus, the buy pressure increases even more, where the sellers of the asset are informed agents using more fundamentalist strategies or noise traders. The larger the distortion between price and fundamental value, the more fragile the bubble becomes: first, the buying power of potential buyers decreases due to the higher price and lower amount of cash available. This implies

²⁵See Giardina and Bouchaud (2003) for a similar explanation.

that the trend signals of high-frequency traders become smaller, possibly even turning negative. A random negative shock, induced by a fundamentalist trader or a noise trader, raises the possibility of a sudden downturn, where again the trend signals of lower-frequency traders follow those of the high-frequency traders. In this case a negative bubble may appear and the described process starts again, however now with negative signs.

This description of a bubble process also clarifies the arguments favoring a FTT: High-frequency oriented chartists have an incentive to create bubbles, since trend-following strategies are only profitable when there is some trend to follow. These traders are also the first to leave the sinking ship when the bubble bursts, again amplifying the negative trend. This is particularly true for high-frequency strategies which would be very costly when the buy/sell signal is constantly wrong. Thus, with a large proportion of fundamental traders in the population, the initial bubble might not even build up. The FTT should be one way to increase the proportion of fundamentalists, as it reduces the profitability of marginally profitable trades in the first place. In this way, it should lower the frequency of ‘false’ signals appearing in the system. In Section 3 we will discuss the effects of such a tax within our model.

2.3 Asset Demand

Any (limit) order is a commitment to trade (at most) a certain quantity at the specified limit price. In this part, we focus on the determination of the order sizes. While defining a strategy that maps price expectations into order sizes appears to be a trivial task, our imposed short-selling and budget constraints complicate things considerably.²⁶ In our setting all agents are initially endowed with the same level and composition of wealth. The wealth of agent i at time t is simply $W_t^i = p_t S_t^i + C_t^i$ and his wealth at the next date is $W_{t+1}^i = p_{t+1} S_t^i + C_t^i = W_t^i + dp_t S_t^i$. Therefore, given the price expectations of the agent, the trading behavior reduces to an optimization problem with respect to the asset holdings S_t^i .

As future price developments are uncertain, we assume agents to be risk-averse. To some extent, this risk-aversion is reflected in the determination of the informed agents’ limit prices in Eqs. (7-8), but should also be present in the choice of the order sizes. The usual approach in the literature is to either use some form of utility maximization (often CARA or CRRA utility functions, see e.g. Chiarella *et al.*, 2009, and Bottazzi *et al.*, 2005), to use random order sizes (see e.g. Mannaro *et al.*, 2008), or use rules-of-thumb, most

²⁶See Franke and Asada (2008).

importantly unit orders with constant size equal to one (see e.g. Pellizzari and Westerhoff, 2009). There are problems with all these approaches: First, much of the literature favors the CARA approach, mainly due to the fact that under this approach the desired asset holdings are independent of wealth (for Gaussian returns). However, when taking accumulated asset positions into account, the actual order size (desired minus actual holdings) is by definition not independent of wealth, cf. Franke (2008). Imposing short-selling and credit constraints will then yield order sizes that are not in line with economic principles. From an economic viewpoint, random and unit order sizes are not very appealing as well, as they imply that order sizes are independent of current wealth and behavioral parameters. We overcome these problems by combining economic variables with rule-of thumb behavior.²⁷

The order size depends on three crucial variables: the agent's aggressiveness, his available resources, and, in case of a market order, the liquidity available at the best opposite price. We specify the demand function as

$$d_t^i = \begin{cases} \left\lceil \alpha^i(\cdot) \frac{C_t^i}{(1+\chi/2)b^i} \right\rceil & \text{if buy order,} \\ \left\lfloor \alpha^i(\cdot) S_t^i \right\rfloor & \text{if sell order,} \end{cases} \quad (9)$$

where $\lceil x \rceil$ denotes rounding towards minus infinity and $\alpha^i(\cdot) \in (0, 1]$ is an aggressiveness parameter determining the proportion of cash/assets an agent actually wants to use for investment. Note that only half the tax rate is taken into account since buyer and seller will share the tax burden equally.

In principle, $\alpha^i(\cdot)$ could take any functional form, which is why we left the arguments unspecified. Following Giardina and Bouchaud (2003) and Martinez-Jaramillo (2007), we simply set

$$\alpha^i(\cdot) = \bar{\alpha}, \quad (10)$$

i.e. a fixed parameter identical for all agents.²⁸ Note that, by construction, the agent's budget constraint is never binding since he willingly only uses a fraction of his wealth to invest in the risky asset.²⁹ Thus, agents are reluctant to submit very large orders which are likely to have strong market impact.³⁰

²⁷Note that rule-of-thumb behavior, although having the weakness of being 'ad-hoc', is more realistic in terms of how actual people make decisions, see Gigerenzer (2008).

²⁸As an alternative, we could make agents' aggressiveness explicitly dependent on economic variables (such as volatility) or on the relative weight of chartism and fundamentalism. In such a setting, the aggressiveness of chartists would be higher, since chartists are usually found to be less risk-averse than fundamentalists, see e.g. Menkhoff and Schmidt (2005). In order to reduce the complexity of the current model, we leave that for future research.

²⁹This ensures that agents do not run out of assets/cash.

³⁰See Harris (2003), Ch. 15.

Additionally, when submitting market orders, agents will at most trade the amount available at the opposite best price, denoted as d_1 , i.e.

$$\bar{d}_t^i = \min(d_1, d_t^i). \quad (11)$$

This takes into account the empirical fact that orders removing more than the volume available at the opposite best quote are rare.³¹

A brief note on the effects of the FTT is in order here. The tax affects transaction volumes negatively in two ways: first via Eq. (9) single order sizes necessarily become smaller due to the negative impact of the tax. Second, agents will only post orders with an expected return larger than the tax rate. Without the tax, the only requirement to post an order is that the two limit prices differ by at least one tick. Depending on the tax rate, orders with rather small expected returns will not be posted anymore. In this way, possible liquidity reductions for higher tax rates may be both due to noise and informed traders. Thus, the tax drives a wedge between the agents' two limit prices. Note that for larger tax rates, price changes should become less frequent, but larger on average. For example, in case of an initial market order, the limit price of the second order needs to be substantially far away from the current price in order to have a positive expected return. The probability of this second order to be executed is quite small, but if so the resulting price change will be quite large.

2.4 Cancellation

The price dynamics of the LOB are driven by the non-trivial interplay between liquidity takers and liquidity providers.³² Prices may change due to the arrival of market and limit orders, and the cancellation of existing limit orders. Limit orders can disappear from the LOB in four different ways: 1) a newly arriving market order is executed against an existing limit order, 2) a limit order can remain at most H^w time-steps (days) in the LOB, afterwards it will be deleted automatically, 3) an agent being chosen to act again will cancel any outstanding orders, 4) an agent cancels his outstanding orders autonomously (possibly even at random), thus the order 'evaporates'.

While the first three channels are obvious, the fourth channel has usually attracted not as much attention in agent-based modelling. However, Farmer

³¹Farmer and Lillo (2004) have shown that roughly 87% of the market orders creating an immediate price change have a volume equal to the volume at the opposite best, while 97% of the market orders creating an immediate price change have a volume at most of the sum of volumes available at the two best opposite prices.

³²See Bouchaud *et al.* (2003), Bouchaud *et al.* (2004) and Toth *et al.* (2011).

et al. (2004) find that cancellation occurs roughly 32% of the time at the best price and 68% of the time inside the book. Challet and Stinchcombe (2001) find that typically 80% (20%) of all orders are cancelled (executed). The main argument for this large number is that placing and cancelling limit orders is usually free of charge and therefore a strategic opportunity for all types of traders.³³ Thus, when modelling the LOB, cancellation of orders cannot be neglected.³⁴

There may well be an important link between investment horizons and the average order lifetime: It is widely believed that a power-law in the distribution of investment horizons may be the driving force behind the power-law tail of price changes.³⁵ However, when investment horizons are indeed fat tailed, the same should be true for order lifetimes. And indeed, the lifetime of an order increases as one moves away from the best bid/ask. Patient investors are therefore less likely to cancel their orders, as found in Potters and Bouchaud (2003).

For the sake of simplicity, we assume a Poisson process of order cancellation for the noise traders as in Daniels *et al.* (2001). At the beginning of each time-step, each noise trader cancels his outstanding orders with probability π_{canc} . For informed traders, we neglect this channel of order cancellation, since this would inject a significant amount of (additional) randomness to their strategies.

2.5 Trading Process and Time

This section contains more details on the trading process and the issue of timing. While we model the LOB at the highest possible frequency, our basic analysis is concerned with the daily frequency. The agents' time horizons are therefore measured in terms of days, such that, for example, each agent approaches the market every $H^{w,i}$ days on average (Poisson waiting times).³⁶ At the beginning of each day t , we randomly reshuffle the list of agents willing

³³See e.g. Cao *et al.* (2008). Note that fundamental traders will post limit orders with prices far away from the best quotes. If the agent is not patient enough, he will cancel his order prematurely.

³⁴See also Challet and Stinchcombe (2001).

³⁵There is some indirect evidence of a power-law distribution in time-scales, see e.g. Lynch and Zumbach (2003). For a theoretical argument, see Lillo (2007). There, heterogeneity in the time horizons is identified as the most likely explanation of the fat-tailed distribution of limit-order prices.

³⁶Thus the probability of a particular agent being chosen equals $(H^{w,i})^{-1}$. Agents with relatively small investment horizons are thus acting more frequently than those with longer horizons.

to approach the market on this given day (active agents).³⁷ Within the day, each active agent approaches the market and potentially submits new orders. Informed traders calculate their price expectations over the next $H^{w,i}$ days, using the most recent input variables in Eq. (2) and the last observation for the trend in Eq. (4), which is being updated only at the end of the day. Orders will remain in the LOB for at most $H^{w,i}$ days. In everything that follows, we restrict ourselves to analyzing daily data only.

As a summary, the algorithm does the following on each day:

1. **Start of the day:** Construct the list of active agents and randomly reshuffle it. Agents then sequentially approach the market.
2. **Intradaily activity:**
 - The active agent i deletes his outstanding orders and (possibly) generates his two orders. The first order is submitted to the LOB.
 - Execute all possible trades (sequentially) taking into account price-time priority and send conditional orders to the LOB. Repeat until no trades possible anymore. Update wealth continuously.
3. **End of the day:** Save the closing midprice. Update trends and wealth accordingly. Update outstanding orders (reduce order lifetime) and deleted expired ones. Cancel outstanding orders of noise traders with probability π_{canc} . At times, redistribute the tax revenues equally across all agents. Go back to step 1 until the desired number of time-steps (days) has been reached.

At the start of the simulations, we need to choose values for the different time horizons. For simplicity, we fix H^t for all informed traders using the same value (equal to the average value in the admissible range) and only incorporate heterogeneity in H^w .³⁸ In the following, we define the admissible range of H^w as all values between 20 and 640, in steps of 20, i.e. $H^w \in \{20, 40, 60, \dots, 640\}$. By default, investment horizons are uniformly chosen from this set and kept constant throughout the simulation. Somewhat surprisingly, we found learning to have no discernable effect within our model. We will briefly comment on this issue in section 3.2.2.

3 Pseudo-Empirical Results

In this section we present pseudo-empirical results from the model simulations. If not stated otherwise, the reported results are the outcome of Monte-Carlo simulations of 22,500 days, disregarding an initial period of

³⁷Hence, we ignore strategic considerations on behalf of the agents on the exact (intra-day) time of approaching the market.

³⁸Recall that we also replaced the backward trend horizon by the forward trend horizon.

Parameter		Value	Description
H^w	\in	$\{20, 40, 60, \dots, 640\}$	Admissible range, H^w
N	$=$	500	Total number of agents
N_s	$=$	100	Parameter for initial endowment
p_0, p^f	$=$	100	Starting value: price/fundamental value
$\bar{\alpha}$	$=$	0.10	Order aggressiveness
Δ	$=$	10^{-3}	Tick size for the price
θ	$=$	0.20	Fraction of informed traders
π_{canc}	$=$	0.01	Cancellation probability (noise traders)
σ_ϵ	$=$	10^{-3}	Volatility for noise traders' expectations
χ	$=$	0	Tax rate

Table 2: Baseline parameter setting for the simulations.

2,500 days, each of which are repeated 20 times with different random seeds.³⁹ In order to get a feeling for the model's properties, we will first present time-series of single simulation runs. For such single runs, we will always present time-series for the most interesting variables (e.g. price and log-returns) and comment on certain statistical properties. Afterwards, we will investigate the effects of a FTT based on the Monte-Carlo approach. The basic parameter values used in our simulations and brief descriptions for all parameters are given in Table 2.

At this point, we should mention one of the main drawbacks of agent-based modelling, namely the large number of degrees of freedom in the choice of the parameter values. This holds even more, when modelling very complex decision-making of agents, as in our case. While one should employ empirical estimates whenever possible, in case there is no (and perhaps never will be) empirical estimate, the modeler has to decide about this value. This is often denoted as 'calibration', which is a neat description for something which can be dangerously misleading. Obviously, a model cannot be robust to changes in all parameters, but should be considered relevant only if (1) it is able to produce realistic dynamics for (economically) plausible parameters values, and (2) is robust with respect to changing certain parameters.

To stress this point, consider for example the parameter θ : What would be a reasonable value for the fraction of informed traders? A priori we should expect a relatively large number of agents to use information about past prices and the fundamental value when forming price expectations in real markets. If so, how many of those agents will be chartists and fundamentalists, respectively? Within our model, we found θ to be a very important parameter

³⁹We will see below that the variability across simulations is typically quite small, so this small number of runs is indeed already sufficient.

for the time-series properties. Consequently, while there is always scope for fine-tuning of the parameters in order to obtain more ‘realistic’ time-series we found the qualitative results to be rather robust with respect to parameter changes as compared to our baseline scenario in Table 2.⁴⁰ For the effects of the FTT, we found the results to be quite robust as well.

3.1 Baseline Scenario and Dependence on θ

One obvious question is whether our model is able to replicate some of the stylized facts of empirical financial time-series. Without going into the details, the most basic stylized facts of asset prices and returns can be summarized as follows:⁴¹

- Martingale property (unit root) of prices: Price dynamics close to a random walk. Zero expected return, with only the very first lags positively autocorrelated (at least for high-frequency data).
- Fat-tailed return distribution: Positive excess kurtosis and power-law tails both imply more probability mass in the center and the tails of the return distribution (compared to Gaussian). Tail exponent around 3.
- Volatility clustering: Autoregressive dependence with very slow (hyperbolic) decay in various measures of volatility.

Recently, many more stylized facts of order-book data have been identified which we will not comment on in the following.⁴²

Here, we do not aim to test quantitatively whether all of these stylized facts are present in the model. Rather we present several basic properties, which are illustrated by individual representative time-series from the model. For the sake of brevity, we restrict ourselves to explaining the baseline scenario in detail. Since we found that the parameter θ plays an important role for the model properties, we briefly illustrate the effects of changing this parameter. Additionally, we investigated the model properties with respect to certain parameter variations in more detail (unreported results, available upon request from the authors). In the next section we then move on to the introduction of a FTT. Again, we should stress that agents’ strategies are kept constant throughout the simulations, see section 3.2.2 below.

⁴⁰In principle, we would be happy to use an approach similar to Franke and Westerhoff (2012), where the ‘optimal’ parameters (with respect to the stylized facts) are estimated via moment-matching criteria. However, given the complexity of our model, and the related high-number of degrees-of-freedom, parameter estimation would be prohibitive.

⁴¹See e.g. Lux (2008).

⁴²See Bouchaud *et al.* (2008) for an extensive overview.

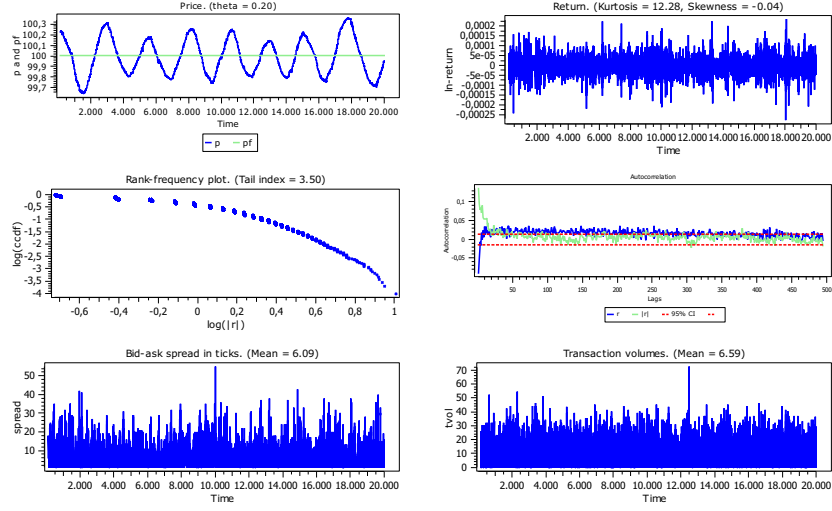


Figure 4: Simulation results: Single run, $\theta = 0.2$, baseline scenario. Top left: price (blue) and fundamental value (green). Top right: log-returns. Center left: rank-frequency plot (ccdf, log-log-scale) of log-returns. Center right: Autocorrelation of raw (blue), absolute returns (green), and 95% confidence interval for absence of autocorrelations. Bottom left: bid-ask spread (in ticks). Bottom right: transaction volumes.




Figure 5: Simulation results: Single run, baseline scenario. Dynamics of relative wealth of average informed trader vs. average noise trader.

Figure 6 shows the results for a single simulation run of the baseline scenario. The top left panel shows the price and the fundamental value. We see that the price fluctuates around the fundamental value over time, and there is a continuous building up and bursting of bubbles. On the top right panel we see the corresponding log-returns, with kurtosis and skewness being 12.28 and -.04, respectively. Thus, the return distribution is highly non-normal but roughly symmetric around zero. In this regard, the center left panel further quantifies the fat tail of the return distribution by means of the complementary cumulative distribution (ccdf) of the return time-series on a log-log scale. The tail region shows a somewhat linear decay, with an estimated tail exponent of 3.50.⁴³ This value lies within the range observed in real markets. The center right panel shows the autocorrelation function (ACF) of the raw and absolute returns. For the raw returns (blue), the bid-ask bounce is the major source for the first few negative lags. However, due to the trend-chasing behavior of informed traders, larger values of θ lead to more autocorrelated prices and returns (see below), which wash-out the bid-ask bounce. For higher lags, the autocorrelations are marginally insignificant. For the absolute returns, the first lags are significantly positively autocorrelated, i.e. there is a small level of volatility clustering. However, the decay of the autocorrelation function is much faster than empirically observed. As we will see below, the trend-chasing behavior of informed traders tends to induce a somewhat larger level of autocorrelation in the raw returns, as compared to the absolute returns. The bottom left panel shows the bid-ask spread (in ticks). We should first note that the average spread is rather large, here with an average value of 6.09 ticks. We see some persistence in the spread, so large (small) spreads tend to be followed by large (small) spreads. The bottom right panel shows the transaction volumes, i.e. the number of stocks traded per time step. Quite interestingly, the Figure shows a smaller (if any) level of persistence in the transaction volumes, such that large price changes appear to be more or less unrelated to large volumes traded.

We checked that large price changes are in fact driven by gaps in the LOB (unreported), as argued by Farmer and Lillo (2004). Furthermore we checked that, as long as there are some trends to follow, informed traders tend to gain, while noise traders tend to lose on average. Figure 5 illustrates this for the simulation run in Figure 4, where we show the average wealth difference between informed and noise traders over time, which steadily increases over time. Even though the interpretation is different, this finding is consistent

⁴³For the estimation of the tail parameter, we used the usual Hill (1975) estimator based on the top 15% observations of the absolute returns, i.e. ignoring signs. The values are not affected by focusing on positive or negative returns only.

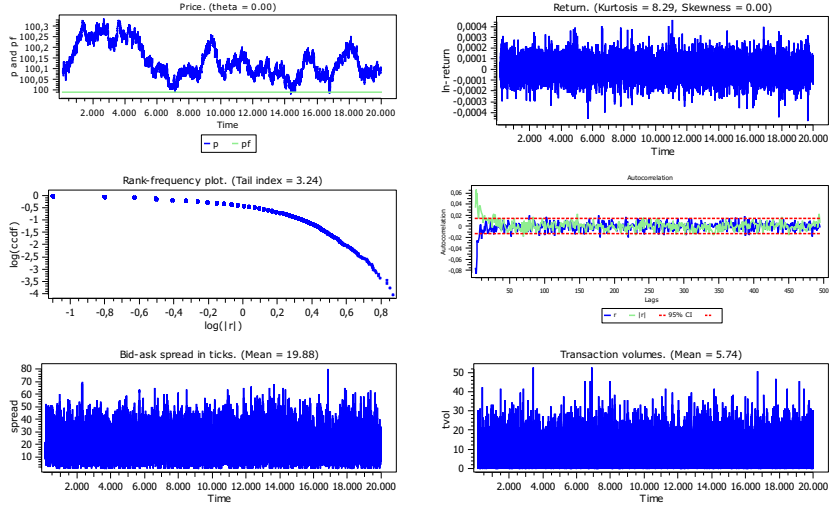


Figure 6: Simulation results: Single run, $\theta = 0$. Top left: price (blue) and fundamental value (green). Top right: log-returns. Center left: rank-frequency plot (ccdf, log-log-scale) of log-returns. Center right: Autocorrelation of raw (blue), absolute returns (green), and 95% confidence interval for absence of autocorrelations. Bottom left: bid-ask spread (in ticks). Bottom right: transaction volumes.

with the results of Kyle (1985), where more informed traders tend to gain from noise traders. However, we should stress that these are only average values. There are also informed traders that lose substantial amounts of their wealth, in part to other informed traders, but also to some of the noise traders. In section 3.2.1 below we look at this issue in more detail, with a particular focus on the effects of the FTT on different subsets (in terms of trading strategies) of the population of informed traders.

As a next step, we briefly compare the findings for different values of θ . Figures 6 and 7 show the results for a single simulation run for $\theta = 0$ and $.5$, respectively. For $\theta = 0$, the price is completely unrelated to the fundamental value: here it continuously exceeds the fundamental value. Without informed traders there is no force that would push the price towards the fundamental value. Quite interestingly, the return distribution has a fat tail already in the case with noise traders only (kurtosis 8.29, tail parameter 3.24) and the first lags of the ACFs are again significant. In contrast, there is no fundamental persistence in the returns, the bid-ask spread, and the transaction volumes, whatsoever. Also the bid-ask spread is substantially larger (19.88 ticks on average), while transaction volumes are smaller. This changes for large values of θ , where we see very smooth bubbles, which translate into

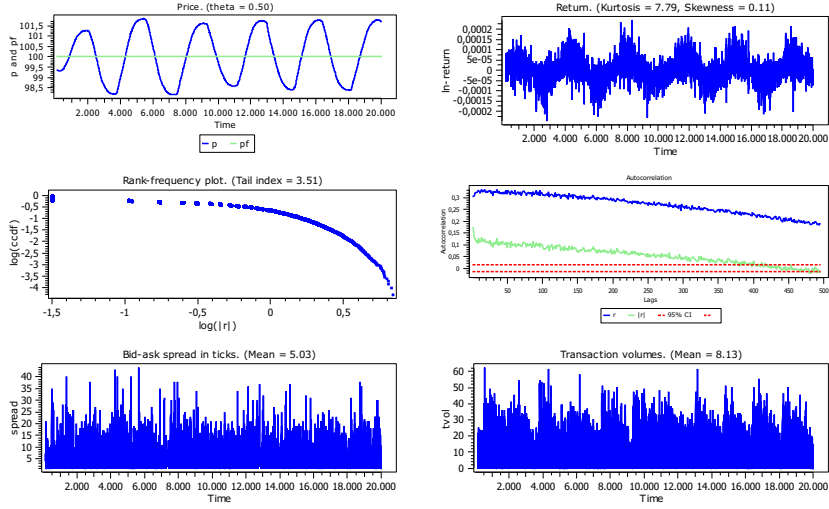


Figure 7: Simulation results: Single run, $\theta = 0.5$. Top left: price (blue) and fundamental value (green). Top right: log-returns. Center left: rank-frequency plot (ccdf, log-log-scale) of log-returns. Center right: Autocorrelation of raw (blue), absolute returns (green), and 95% confidence interval for absence of autocorrelations. Bottom left: bid-ask spread (in ticks). Bottom right: transaction volumes.

highly autocorrelated raw and absolute returns, with the raw returns having a substantially larger positive autocorrelation at all lags, as compared to the absolute returns.⁴⁴ Choosing θ such that the ACF of the absolute returns decays hyperbolically results in a very inefficient market, with the ACF of the raw returns decaying similarly slowly. In part, we already saw this for the simulation of the baseline scenario. Still, our choice of θ in the baseline scenario takes this trade-off into account.

An important reason for the small level of volatility clustering is probably the absence of the leverage effect in our model.⁴⁵ The basic idea is that leverage would increase the necessity of portfolio adjustments after large negative price changes, i.e. large downward spirals.⁴⁶ We leave this model extension for future research.

⁴⁴The simulated prices are qualitatively very similar to the results for the unstable case in Youssefmir *et al.* (1998). There, the analysis is based on a much simpler version of the model, solved based on a mean-field approximation.

⁴⁵In simple terms, the leverage effect corresponds to a negative correlation between past returns and future volatility, see e.g. Bouchaud *et al.* (2001).

⁴⁶See e.g. Thurner *et al.* (2012).

3.2 Effects of the FTT

Now we turn to the effects of a (two-sided) FTT. We will present aggregate results from Monte-Carlo simulations for 18 different tax rates.⁴⁷ The parameters used are summarized in Table 2.

Figure 8 summarizes the most important results from Monte-Carlo simulations for varying tax rates. The top left panel shows the tax revenues in dependence on the tax rate. We find the usual Laffer curve relationship, i.e. the tax revenues tend to increase for small tax rates until they start decreasing again for larger ones. The top right panel shows that the tails become fatter for larger tax rates. This indicates that the variance of the most extreme observations becomes substantially larger. Note that the tail exponent is usually larger than 2, except for very large tax rates where we might end up in the Levy-stable regime. The center left panel shows that the tax leads to a significant increase of the bid-ask spread. Thus, as expected, the liquidity is significantly reduced.⁴⁸ The center right panel shows that transaction volumes tend to decrease in the tax rate, which is not surprising given that all trades become less profitable. Most interestingly, the bottom panels show the distortion (left) and volatility (right), respectively. Quite surprisingly, the distortion tends to decrease for very small tax rates, increases later on and reaches the initial level without tax only for relatively large tax rates. In contrast, the volatility tends to increase for very small to intermediate tax rates, but decreases later on.⁴⁹ We will comment on the ‘kinks’, i.e. the strong changes in some of the variables for very small tax rates in section 3.2.1. The findings for the distortion and volatility can be explained by the wedge that the tax drives between agents’ two limit prices. Only those limit prices corresponding to a positive expected return (post-tax) are submitted to the LOB. However, these orders are unlikely to be executed. For example, in case of an initial market order, the limit price of the second order needs to be substantially far away from the current price. However, this

⁴⁷We used the following tax rates (in percent): 0 (baseline scenario), 0.01, 0.02, 0.03, ..., 0.18%. We use 0.18% as the maximum value, since the LOB may be empty at times for larger tax rates and we require at least one order to be on each side of the book to be existent at any point in time.

⁴⁸Note the relatively large change for very small tax rates (also present in the transaction volumes and the distortion). This effect is mostly driven by liquidity reductions from very short-term oriented informed traders, whose trades become unprofitable even for these very small tax rates. For larger tax rates, also longer-term oriented strategies are affected and the distortion increases again.

⁴⁹For larger tax rates, the volatility increases substantially. We do not show the results, since the LOB might become very sparse, with only few or no orders present at certain points in time.

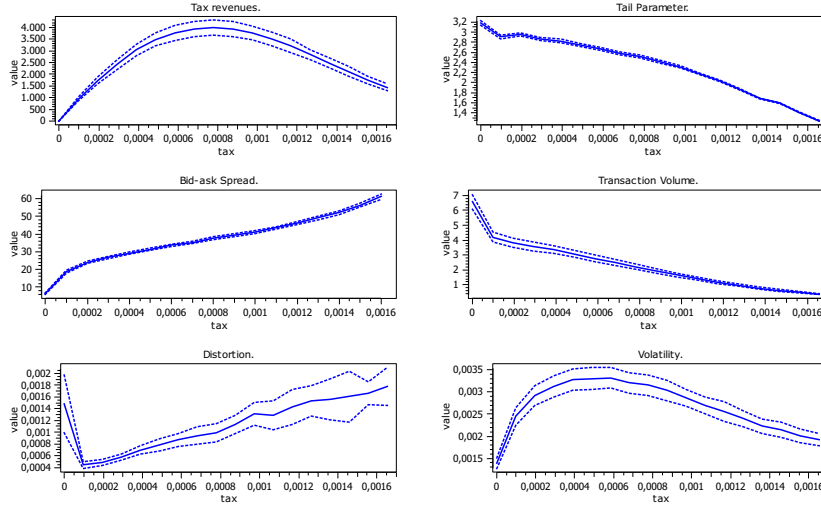


Figure 8: Simulation results: Monte-Carlo simulations. Dependence on χ . Top left: tax revenues. Top right: tail exponent (top 15% observations). Center left: bid-ask spread. Center right: transaction volumes. Bottom left: distortion ($|\ln(p/p^f)|$). Bottom right: volatility (absolute price change). Plotted are mean values (solid lines), plus and minus one standard deviation (dashed lines).

second order will probably never be executed. In this way, price fluctuations are relatively rare (volatility decreases), but if the price changes it does so substantially (tail index smaller). Additionally, the informed traders make sure that prices do not depart too much from the fundamental value. We should also note that the tax rate with the maximum volatility level is close to the tax rate that maximizes tax revenues, representing the usual trade-off between stability and tax revenues.

Summing up, the results for small tax rates are roughly in line with those from the literature, except for the strong decrease in distortion for very small tax rates.⁵⁰ This suggests that a small tax rate should reduce liquidity and transaction volumes, but might actually bring prices closer to the fundamental values and would only marginally increase volatility. Additionally, our findings also show that larger tax rates may not create entirely negative effects. In this case, comparable values for distortion and volatility as in the no-tax case come along with substantial tax revenues.

To some extent, our findings might be driven by the high level of stability in the baseline scenario. As one of the referees of this paper rightly pointed out, the stabilizing potential of the FTT is limited in our setting since the

⁵⁰See Mannaro *et al.* (2008) and Pellizzari and Westerhoff (2009) for the CDA case.

average levels of both distortion and volatility are arguably quite small compared to real markets in the baseline scenario. Thus, it might not be too surprising that the attempt to stabilize an already relatively stable financial market may not yield the desired outcome. Nevertheless, it seems even more remarkable that the average level of distortion is being reduced for small tax rates in our setup. We would expect similar effect for the distortion in real markets, whereas the negative impact on volatility might be less severe.

3.2.1 Investment Horizons, Performance, and FTT

In the previous section, we analyzed the effects of the FTT in terms of aggregate system properties. Here, given that we expect the tax to affect the behavior of certain groups of traders in different ways, we take a closer look at different subgroups of the population of the informed traders. In everything that follows, we use the parameters from the baseline scenario and a set of rather small tax rates compared to the previous sections.⁵¹ The main reason for using these smaller rates is the relatively high sensitivity of informed traders' trading decisions, in particular for very small investment horizons, to minor changes in the tax rate. Here we mainly aim at explaining the 'kinks', for example in distortion, in Figure 8 for very small tax rates, but focusing on the behavior of certain groups of informed traders. Here we do not randomly assign investment horizons to the informed traders within the admissible range, but rather divide them into three discrete groups: All agents in group 1 use the same small investment horizon, those in group 2 an intermediate value, and those in group 3 a large value. To be precise, for the three groups H^w corresponds to the minimum, the midpoint, and the maximum investment horizons in the admissible range, respectively. Thus, group 1 uses more of a chartist strategy, group 3 more of a fundamentalist strategy, and group 2 something in between.

We see that informed traders tend to gain, while noise traders tend to lose on average (cf. Figure 5). In the following, by analyzing the wealth dynamics of the three groups (relative to the performance of noise traders), we will show that the distribution of gains and losses of informed traders is far from uniform. Most importantly, we find that the increase in investment horizons shown before is mainly due to the very poor performance of group 1, i.e. the chartists. This can be seen from Figures 9-11, where we plot the average wealth difference (in absolute terms) between the three groups and the group of noise traders over time for single simulation runs and different tax rates.

⁵¹The maximum tax rate used here is only 0.0035%.




Figure 9: Single run, baseline scenario with three groups of informed traders with small, intermediate, and large H^w . Average wealth differences between the three groups and noise traders over time. Group 1 (3) are the chartists (fundamentalists), while group 2 uses an intermediate strategy.




Figure 10: Single run, baseline scenario with small tax rate and three groups of informed traders with small, intermediate, and large H^w . Average wealth differences between the three groups and noise traders over time. Group 1 (3) are the chartists (fundamentalists), while group 2 uses an intermediate strategy.




Figure 11: Single run, baseline scenario with larger tax rate and three groups of informed traders with small, intermediate, and large H^w . Average wealth differences between the three groups and noise traders over time. Group 1 (3) are the chartists (fundamentalists), while group 2 uses an intermediate strategy.

Figure 9 shows that, without a tax group 2 (green line) performs best, i.e. the intermediate strategy is the most profitable one, as it tends to strongly gain in wealth over time. Similarly, group 3 tends to gain as well, but on a smaller scale, so the fundamental strategy is also profitable. Somewhat surprisingly, the worst performers are in group 1, whose average wealth tends to be comparable to the wealth of noise traders (at times even smaller as the wealth difference may be negative and highly cyclical). Thus, chartists tend to perform very poorly on average. An explanation for the good performance of group 2 is that group 1 is the major source of predictable bubbles (high-frequency traders), while group 3 tries to drive prices back towards the fundamental level (low-frequency traders). The intermediate strategy works on a higher frequency than the fundamental traders, thereby leading them to follow the trend at times or expecting reversal towards the fundamental value at other times.

Figure 10 shows the wealth dynamics for a very small tax rate. We see that the ordering is conserved, since group 2 still performs better on average than the other groups. As before, chartists tend to perform very poorly, in this example even losing money over time, while fundamentalists and those with an intermediate strategy tend to gain. Thus, for very small




Figure 12: Simulation results: Monte-Carlo simulations. Fraction of order submissions relative to the activity frequency for noise traders (black), and group 1-3 in dependence of the tax rate.

tax rates, chartists continue posting very unprofitable orders. This changes for very large tax rates, as we can see from Figure 11: here the average wealth of chartists now exceeds that of the noise traders. Large tax rates reduce the amount of strong chartist orders, so chartists post fewer orders (i.e. effectively work on lower frequencies) and perform better. Also we see that their wealth dynamics are significantly less cyclical, so reducing their posting of high-frequency orders tends to reduce the occurrence of bubbles and bursts. Interestingly, we see that the fundamental strategy becomes most profitable for larger tax rates. Note that in this case there are practically no cyclical fluctuations, so the calculated trends (necessary for calculating the expected prices) tend to be very small. This implies that the trending component for all strategies becomes negligible, so practically all traders tend to expect the price to revert towards the fundamental value. From this viewpoint, it is clear that fundamental traders tend to be the most successful, since their orders have the longest lifetime, i.e. the highest probability that their two limit prices are being hit by the noise traders. The longer investment horizon also implies that the difference between the current price and the expected price can be larger, such that the majority of orders with expected profits (after tax) are posted by the group of fundamental traders.

Figure 12 illustrates the effects of different tax rates on the order submission process of the three groups. There we show the average fraction of

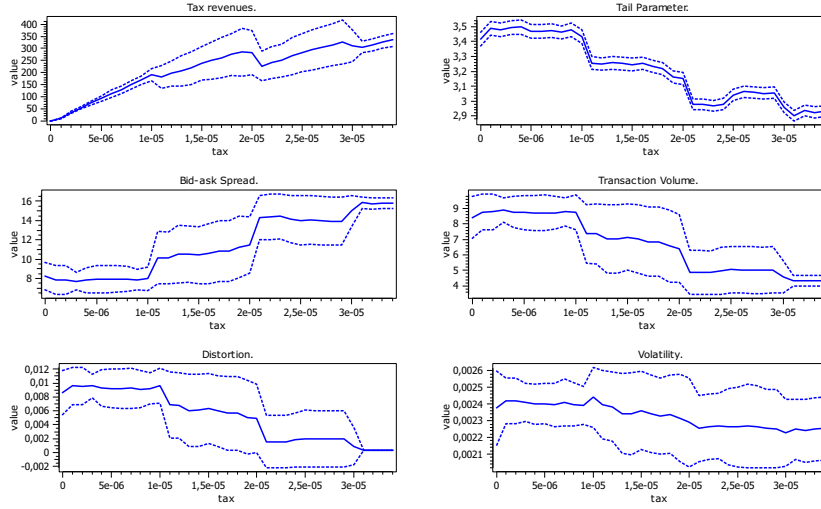


Figure 13: Simulation results: Monte-Carlo simulations. Dependence on χ for relatively small tax rates with three groups of informed traders. Top left: tax revenues. Top right: tail exponent (top 15% observations). Center left: bid-ask spread. Center right: transaction volumes. Bottom left: distortion ($|\ln(p/p^f)|$). Bottom right: volatility (absolute price change). Plotted are mean values (solid lines), plus and minus one standard deviation (dashed lines).

order submissions relative to the frequency of being active for all groups. First, we see that noise traders (black) are practically unaffected by these relatively small tax rates, as they post orders practically every time when they become active (i.e. the fraction of order submissions is close to 1). In contrast, we see that group 1 posts orders roughly 60% of the time without a tax, but this fraction approaches zero relatively quickly. We see a step-wise relationship, which is due to the fact that prices (and thus the taxes to be paid) are not continuous, but depend on the tick size. For groups 2 and 3 we observe similar relationships, but the values are on a significantly higher level for these relatively small tax rates. Quite interestingly, at the highest tax rate in this analysis, group 3 continues posting orders 80% of the time. The kinks in the previous section are therefore mainly driven by the inactivity of chartists already for relatively small tax rates (cf. Figure 13). As expected, the tax affects strategies with higher trading frequency significantly more. Note that for such small tax rates, the other groups continue posting orders, so the absence of strong speculative bubbles reduces the average distortion. The fraction of order submissions for group 2 and 3

also tend to approach zero for larger tax rates, with group 3 being active also for larger tax rates (unreported). For very large tax rates, i.e. those used in the previous sections, even the fundamental traders act rarely, such that the average distortion starts increasing again.

3.2.2 FTT and Learning

An important drawback of our analysis is the absence of learning. Since the FTT is actually meant to harm short-term oriented strategies more than proportionally, we would expect adjustments in trading strategies in real markets. In our setting, agents can only respond to the FTT by deciding not to trade, but, as usual in most related models, we keep their underlying strategies fixed throughout the simulation. Hence, adjustments in trading strategies in response to the FTT might affect our results.

Quite surprisingly, our findings are quite robust with respect to the presence of learning, cf. Fricke (2013) for details. This strengthens our confidence in the results presented here, but we should stress that the performance of the learning algorithms is typically quite poor.⁵² Our simulation experiments indicate that the market development is apparently too complex and erratic to warrant systematic learning at least over the time horizons that we have considered. Moreover, we found that defining an adequate fitness function for our setting is not as straightforward as it seems. Based on the results from the previous section, one might want to use relative changes in wealth (over some past horizon) as an indicator of agents' trading performances. However, this fitness function comes with certain problems as it is difficult to distinguish profits coming from trading actively and those due to changes in the book value of assets.

4 Conclusions

In this paper we have presented a detailed artificial financial market, where agents compete against each other within a CDA mechanism. While incorporating the usual chartist/fundamentalist/noise trader framework, our model has the advantage of explicitly accounting for the importance of time horizons in financial markets. We showed evidence that the model is able to replicate certain well-known stylized fact of financial time-series, among them the martingale property, fat-tailed return distributions, and, albeit to a lesser extent, volatility clustering. Moreover, for certain parameter combinations

⁵²We tested a variety of (social) learning algorithms, using genetic algorithms, where successful strategies tend to replicate and spread through the population.

the model is able to generate the building-up and bursting of asset price bubbles.

The main focus of this paper was on the effects of a FTT in the artificial financial market. In this regard, we find the usual trade-off between monetary revenues (Laffer curve) and stability, as higher tax revenues come along with higher volatility. The results for small tax rates are roughly in line with those from the literature, except for the strong decrease in distortion for very small tax rates.⁵³ This suggests that a small tax rate should reduce liquidity and transaction volumes, but might actually bring prices closer to the fundamental values and would only marginally increase volatility. Additionally, we also show that larger tax rates may not create entirely negative effects. In this case, comparable values for distortion and volatility as in the no-tax case come along with substantial tax revenues. These revenues could be used for a number of productive purposes. Additionally, the reduced market activity also frees-up resources, both in terms of financial and human capital, that could be directed to other parts of the economy.

We have also discussed a major drawback of our approach, namely the absence of learning. Moreover, in reality we would need to weigh the different effects in order to come up with a welfare-optimizing solution. For example, here we simply redistributed the tax revenues among the traders to keep total wealth constant. In political discussions it has often been proposed to use these revenues for investment in developing countries; more recently the motivation has been to compensate for the costs of the financial crisis. Given the extraordinary high transaction volumes in real markets, partly driven by the arrival of HFTs, it appears promising to introduce a small tax to reduce the possibly distorting effects of their activities and generate large tax revenues at the same time. In part, our results suggest that imposing a very small tax would make HFT strategies highly unprofitable. However, since their trading algorithms are usually not meant to follow trends or drive prices towards some fundamental value, their effects on the macro-properties of the system are still under debate. In the end, even a tiny FTT would lead to a shrinking of the financial sector, allowing to extract highly productive resources (e.g. human capital) for other purposes.

The presented model is very flexible and serves as an illustration of the complexity of the optimization task in real markets. The self-referential nature of this task makes the extraction of valuable information quite difficult. It is worth noting that the order generation process of individuals is still poorly understood, in contrast to the aggregate order-book dynamics for

⁵³See Mannaro *et al.* (2008) and Pellizzari and Westerhoff (2009) for the CDA case.

which a number of scaling-laws have been identified.⁵⁴ We hope that future research, for example by means of laboratory experiments, may help us in deciphering these processes. It is crucial to understand the agents' individual behavior at the micro-level to generate more realistic dynamics at the macro-level. In the end, we believe that our model is an ambitious first step towards more realistic 'wind-channels' for testing regulatory policies. In future research, we plan to tackle several of the issues mentioned throughout the text.

⁵⁴See e.g. Zovko and Farmer (2002).

Bibliography

- ALLEN, H., AND M. P. TAYLOR (1990): “Charts, Noise and Fundamentals in the London Foreign Exchange Market,” *The Economic Journal*, 100(400), 49–59.
- ANUFRIEV, M., AND G. BOTTAZZI (2004): “Asset Pricing Model with Heterogeneous Investment Horizons,” LEM Papers Series 2004/22, Laboratory of Economics and Management (LEM), Sant’Anna School of Advanced Studies, Pisa, Italy.
- AOKI, M. (2002): “Open Models of Share Markets with Two Dominant Types of Participants,” *Journal of Economic Behavior & Organization*, 49(2), 199–216.
- AOKI, M., AND H. YOSHIKAWA (2007): *Reconstructing Macroeconomics: A Perspective from Statistical Physics and Combinatorial Stochastic Processes*. Cambridge University Press.
- BAE, K.-H., H. JANG, AND K. S. PARK (2003): “Traders’ choice between limit and market orders: evidence from NYSE stocks,” *Journal of Financial Markets*, 6(4), 517–538.
- BAK, P., M. PACZUSKI, AND M. SHUBIK (1997): “Price Variations in a Stock Market with many Agents,” *Physica A: Statistical and Theoretical Physics*, 246(3-4), 430 – 453.
- BALTAGI, B., D. LI, AND Q. LI (2006): “Transaction Tax and Stock Market Behavior: Evidence from an Emerging Market,” *Empirical Economics*, 31(2), 393–408.
- BEJA, A., AND M. B. GOLDMAN (1980): “On the Dynamic Behavior of Prices in Disequilibrium,” *Journal of Finance*, 35(2), 235–48.
- BIAIS, B., P. HILLION, AND C. SPATT (1995): “An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse,” *Journal of Finance*, 50(5), 1655–89.
- BORDO, M., B. EICHENGREEN, D. KLINGEBIEL, AND M. S. MARTINEZ-PERIA (2001): “Is the Crisis Problem Growing More Severe?,” *Economic Policy*, 16(32), 51–82.
- BORLAND, L., AND J.-P. BOUCHAUD (2005): “On a Multi-Timescale Statistical Feedback Model for Volatility Fluctuations,” *Science & Finance*

- (CFM) working paper archive 500059, Science & Finance, Capital Fund Management.
- BOTTAZZI, G., G. DOSI, AND I. REBESCO (2005): “Institutional Architectures and Behavioral Ecologies in the Dynamics of Financial Markets,” *Journal of Mathematical Economics*, 41(1-2), 197–228.
- BOUCHAUD, J.-P., J. D. FARMER, AND F. LILLO (2008): *How Markets Slowly Digest Changes in Supply and Demand*, In: *Handbook of Financial Markets* (Eds.: T. Hens and K. R. Schenk-Hoppe), Ch. 2. North Holland.
- BOUCHAUD, J.-P., Y. GEFEN, M. POTTERS, AND M. WYART (2003): “Fluctuations and Response in Financial Markets: The Subtle Nature of ‘Random’ Price Changes,” Science & Finance (CFM) working paper archive 0307332, Science & Finance, Capital Fund Management.
- BOUCHAUD, J.-P., J. KOCKELKOREN, AND M. POTTERS (2004): “Random Walks, Liquidity Molasses and Critical Response in Financial Markets,” Science & Finance (CFM) working paper archive 500063, Science & Finance, Capital Fund Management.
- BOUCHAUD, J.-P., A. MATA CZ, AND M. POTTERS (2001): “The Leverage Effect in Financial Markets: Retarded Volatility and Market Panic,” Science & Finance (CFM) working paper archive 0101120, Science & Finance, Capital Fund Management.
- BOUCHAUD, J.-P., M. MÉZARD, AND M. POTTERS (2002): “Statistical Properties of Stock Order Books: Empirical Results and Models,” *Quantitative Finance*, 2(4), 251–256.
- BROGAARD, J. (2010): “High Frequency Trading and its Impact on Market Quality,” *Social Science Research Network Working Paper Series*.
- CAO, C., O. HANSCH, AND X. WANG (2008): “Order Placement Strategies In A Pure Limit Order Book Market,” *Journal of Financial Research*, 31(2), 113–140.
- CHALLET, D., AND R. STINCHCOMBE (2001): “Analyzing and Modelling 1+1d Markets,” *Physica A*, 300(1-2), 285–299.
- CHALLET, D., AND Y. C. ZHANG (1997): “Emergence of Cooperation and Organization in an Evolutionary Game,” *Physica A: Statistical and Theoretical Physics*, 246(3-4), 407 – 418.

- CHIARELLA, C., AND G. IORI (2002): “A Simulation Analysis of the Microstructure of Double Auction Markets,” *Quantitative Finance*, 2(5), 346–353.
- CHIARELLA, C., G. IORI, AND J. PERELLO (2009): “The Impact of Heterogeneous Trading Rules on the Limit Order Book and Order Flows,” *Journal of Economic Dynamics and Control*, 33(3), 525–537.
- CLIFF, D., AND J. BRUTEN (1997): “Zero is Not Enough: On the Lower Limit of Agent Intelligence for Continuous Double Auction Markets,” Technical Report HPL-97-141, Hewlett-Packard Laboratories.
- DANIELS, M. G., J. DOYNE FARMER, L. GILLEMOT, G. IORI, AND E. SMITH (2001): “A Quantitative Model of Trading and Price Formation in Financial Markets,” *ArXiv Condensed Matter e-prints*.
- 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).
- EHRENSTEIN, G. (2002): “Cont-Bouchaud Percolation Model Including Tobin tax,” *International Journal of Modern Physics*, 13, 1323–1331.
- FARMER, J. D., L. GILLEMOT, F. LILLO, S. MIKE, AND A. SEN (2004): “What Really Causes Large Price Changes?,” *Quantitative Finance*, 4(4), 383–397.
- FARMER, J. D., AND F. LILLO (2004): “On the Origin of Power Law Tails in Price Fluctuations,” *Quantitative Finance*, 4(1), 7–11.
- FOUCAULT, T., O. KADAN, AND E. KANDEL (2005): “Limit Order Book as a Market for Liquidity,” *Review of Financial Studies*, 18(4), 1171–1217.
- FRANKE, R. (2008): “On the Interpretation of Price Adjustments and Demand in Asset Pricing Models with Mean-Variance Optimization,” Economics Working Papers 2008,13, Christian-Albrechts-University of Kiel, Department of Economics.
- FRANKE, R., AND T. ASADA (2008): “Incorporating Positions Into Asset Pricing Models With Order-Based Strategies,” *Journal of Economic Interaction and Coordination*, 3(2), 201–227.
- FRANKE, R., AND F. WESTERHOFF (2012): “Structural Stochastic Volatility in Asset Pricing Dynamics: Estimation and Model Contest,” *Journal of Economic Dynamics and Control*, 36(8), 1193 – 1211.

- FRICKE, D. (2013): “Coping with the Complexity of Financial Markets,” Phd thesis, Christian-Albrechts-Universität Kiel.
- GIARDINA, I., AND J.-P. BOUCHAUD (2003): “Volatility Clustering in Agent Based Market Models,” *Physica A: Statistical Mechanics and its Applications*, 324(1-2), 6 – 16.
- GIGERENZER, G. (2008): *Gut Feelings: The Intelligence of the Unconscious*. Penguin (Non-Classics).
- GODE, D. K., AND S. SUNDER (1993): “Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality,” *Journal of Political Economy*, 101(1), 119–37.
- HANKE, M., J. HUBER, M. KIRCHLER, AND M. SUTTER (2010): “The Economic Consequences of a Tobin Tax - An Experimental Analysis,” *Journal of Economic Behavior and Organization*, 74, 58–71.
- HARRIS, L. (2003): *Trading and Exchanges - Market Microstructure for Practitioners*. Oxford University Press.
- HARRIS, L., AND J. HASBROUCK (1996): “Market vs. Limit Orders: The SuperDOT Evidence on Order Submission Strategy,” *Journal of Financial and Quantitative Analysis*, 31(02), 213–231.
- 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.
- HILL, B. M. (1975): “A Simple General Approach to Inference About the Tail of a Distribution,” *The Annals of Statistics*, 3(5), 1163–1174.
- JACKSON, P., AND A. O’DONNELL (1985): “The Effects of Stamp Duty on Equity Transactions and Prices in the UK Stock Exchange,” Discussion Paper 25, Bank of England.
- JONES, C. M., AND P. J. SEGUIN (1997): “Transaction Costs and Price Volatility: Evidence from Commission Deregulation,” *American Economic Review*, 87(4), 728–37.
- KEYNES, J. M. (1936): *The General Theory of Employment, Interest and Money*. Macmillan, London.

- KIRCHLER, M., J. HUBER, AND D. KLEINLERCHER (2012): “Market Microstructure Matters When Imposing a Tobin Tax. Evidence from the Laboratory Experiments,” *Journal of Economic Behavior and Organization*, 80, 586–602.
- KYLE, A. S. (1985): “Continuous Auctions and Insider Trading,” *Econometrica*, 53(6), 1315–35.
- LEBARON, B., W. B. ARTHUR, AND R. PALMER (1999): “Time Series Properties of an Artificial Stock Market,” *Journal of Economic Dynamics and Control*, 23(9-10), 1487–1516.
- LILLO, F. (2007): “Limit order Placement as an Utility Maximization Problem and the Origin of Power Law Distribution of Limit Order Prices,” *The European Physical Journal B*, 55(4), 453–459.
- LUX, T. (2008): *Stochastic Behavioral Asset Pricing Models and the Stylized Facts*, In: *Handbook of Financial Markets* (Eds.: T. Hens and K. R. Schenk-Hoppe), Ch. 3. North Holland.
- LUX, T., AND M. MARCHESI (1999): “Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market,” *Nature*, (397), 498–500.
- (2000): “Volatility Clustering in Financial Markets: a Microsimulation of Interacting Agents,” *International Journal of Theoretical and Applied Finance*, 3(4), 169–196.
- LUX, T., AND S. SCHORNSTEIN (2005): “Genetic Learning as an Explanation of Stylized Facts of Foreign Exchange Markets,” *Journal of Mathematical Economics*, 41(1-2), 169–196.
- LYNCH, P., AND G. ZUMBACH (2003): “Market Heterogeneities and the Causal Structure of Volatility,” *Quantitative Finance*, 3(4), 320–331.
- MANNARO, K., M. MARCHESI, AND A. SETZU (2008): “Using an Artificial Financial Market for Assessing the Impact of Tobin-like Transaction Taxes,” *Journal of Economic Behavior & Organization*, 67(2), 445–462.
- MARTINEZ-JARAMILLO, S. (2007): “Artificial Financial Markets: an Agent Based Approach to Reproduce Stylized Facts and to Study the Red Queen Effect,” *PhD Thesis, Centre for Computational Finance and Economic Agents (CCFEA), University of Essex*.

- MENKHOFF, L. (1998): “The Noise Trading Approach – Questionnaire Evidence From Foreign Exchange,” *Journal of International Money and Finance*, 17(3), 547–564.
- MENKHOFF, L., AND U. SCHMIDT (2005): “The Use of Trading Strategies by Fund Managers: Some First Survey Evidence,” *Applied Economics*, 37(15), 1719–1730.
- MIKE, S., AND J. D. FARMER (2008): “An Empirical Behavioral Model of Liquidity and Volatility,” *Journal of Economic Dynamics and Control*, 32(1), 200–234.
- PELLIZZARI, P., AND F. WESTERHOFF (2009): “Some Effects of Transaction Taxes Under Different Microstructures,” *Journal of Economic Behavior & Organization*, 72(3), 850–863.
- POLLIN, R., D. BAKER, AND M. SCHABERG (2003): “Securities Transaction Taxes for U.S. Financial Markets,” *Eastern Economic Journal*, 29(4), 527–558.
- POTTERS, M., AND J.-P. BOUCHAUD (2003): “More Statistical Properties of Order Books and Price Impact,” *Physica A: Statistical Mechanics and its Applications*, 324(1-2), 133 – 140, Proceedings of the International Econophysics Conference.
- RABERTO, M., S. CINCOTTI, S. FOCARDI, AND M. MARCHESI (2003): “Traders’ Long-Run Wealth in an Artificial Financial Market,” *Computational Economics*, 22(2), 255–272.
- ROLL, R. (1989): “Price Volatility, International Market Links, and Their Implications for Regulatory Policies,” *Journal of Financial Service Research*, 3, 211–246.
- SHILLER, R. J. (1981): “Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?,” *American Economic Review*, 71(3), 421–36.
- THURNER, S., J. D. FARMER, AND J. GEANAKOPLOS (2012): “Leverage Causes Fat Tails and Clustered Volatility,” *Quantitative Finance*, 12(5), 695–707.
- TOBIN, J. (1978): “A Proposal for International Monetary Reform,” *Eastern Economic Journal*, 4(3-4), 153–159.

- TOTH, B., Z. EISLER, F. LILLO, J. BOUCHAUD, J. KOCKELKOREN, AND J. DOYNE FARMER (2011): “How Does the Market React to Your Order Flow?,” *ArXiv e-prints*.
- WESTERHOFF, F. (2003): “Heterogeneous Traders and The Tobin Tax,” *Journal of Evolutionary Economics*, 13(1), 53–70.
- (2004a): “Speculative Dynamics, Feedback Traders and Transaction Taxes: a Note,” *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,” *Journal of Economics and Statistics (Jahrbuecher fuer Nationaloekonomie und Statistik)*, 228(2+3), 195–227.
- WESTERHOFF, F. H., AND R. DIECI (2006): “The Effectiveness of Keynes-Tobin Transaction Taxes when Heterogeneous Agents Can Trade in Different Markets: A behavioral finance approach,” *Journal of Economic Dynamics and Control*, 30(2), 293–322.
- YOUSSEFMIR, M., B. A. HUBERMAN, AND T. HOGG (1998): “Bubbles and Market Crashes,” *Computational Economics*, 12(2), 97–114.
- ZOVKO, I., AND D. FARMER (2002): “The Power of Patience: a Behavioural Regularity in Limit-Order Placement,” *Quantitative Finance*, 2(5), 387–392.
- ZUMBACH, G., AND P. LYNCH (2001): “Heterogeneous Volatility Cascade in Financial Markets,” *Physica A: Statistical Mechanics and its Applications*, 298(3-4), 521 – 529.