Effects of tobin taxes in minority game markets

Bianconi, Ginestra; Galla, Tobias; Marsili, Matteo; Pin, Paolo

Postprint / Postprint
Zeitschriftenartikel / journal article

Zur Verfügung gestellt in Kooperation mit / provided in cooperation with:
www.peerproject.eu

gesis
Leibniz-Institut
für Sozialwissenschaften

Mitglied der
Leibniz-Gemeinschaft

Accepted Manuscript

Title: Effects of Tobin Taxes in Minority Game markets

Authors: Ginestra Bianconi, Tobias Galla, Matteo Marsili,
Paolo Pin

Please cite this article as: Bianconi, G., Galla, T., Marsili, M., Pin, P., Effects of Tobin
Taxes in Minority Game markets, *Journal of Economic Behavior and Organization*
(2008), doi:10.1016/j.jebo.2008.10.009

Effects of Tobin Taxes in Minority Game markets [★]

Ginestra Bianconi [a], Tobias Galla [a,b], Matteo Marsili [a],
Paolo Pin [*,a,c]

[a] *The Abdus Salam International Centre for Theoretical Physics, Strada Costiera 11, 34014 Trieste, Italy*

[b] *School of Physics and Astronomy, The University of Manchester, Schuster Building, Manchester M13 9PL, UK*

[c] *Max Weber Programme, European University Institute, Via Delle Fontanelle 10, 50014 San Domenico di Fiesole (FI), Italy*

[*] Corresponding author. Tel: +39-055-4685-695; fax: +39-055-4685-804. E–mail address: pin@unive.it

Effects of Tobin Taxes in Minority Game markets

Abstract

We show that the introduction of Tobin taxes in agent-based models of currency markets can lead to a reduction of both speculative trading and the magnitude of exchange rate fluctuations at intermediate tax rates. In this regime revenues obtained from speculators are maximal for the institutions acting as market makers. We here focus on Minority Game models of markets, which are accessible by exact techniques from statistical mechanics. Results are supported by computer simulations. Our findings suggest that at finite systems sizes the effect is most pronounced in a critical region around the phase transition of the infinite system, but much weaker if the market is operating far from criticality and does not exhibit anomalous fluctuations.

Key words: Tobin tax, Minority Games.

JEL classification numbers: C63, F31, G18.

1 Introduction

In 1972 James Tobin proposed to "throw some sand in the wheels of our excessively efficient international money markets" (published in Tobin 1974) by imposing a tax of 0.05 to 0.5% on all Foreign Currency Exchange (FX). The Bretton Woods agreement (a system of fixed foreign exchange rates tied to the price of gold) at that time was gradually being dismantled, with the USA stepping out in 1971. This system had been introduced in the wake of World War II in order to rebuild global capitalism. Tobin feared the effects of countries exposed to freely fluctuating exchange rates and suggested, as a second best solution, the introduction of what is now called a 'Tobin tax' in order to suppress speculative trading, thus allowing domestic macro-economic management.

Since then, under floating exchange rates, the trading volume on international currency markets has grown sharply, especially after the introduction of electronic trading, reaching a level of 1.9 trillion US-Dollar per day in 2004 (Galati and Melvin 2004). Empirical research has shown that FX rates' evolution is 'disconnected' from the dynamics of the macro-economic fundamentals

it should depend on and, in particular, that they fluctuate much more than such a dependence would suggest (as a survey see Frankel and Rose 1995).

While the Tobin tax has never been implemented in reality, the discussion of this issue is still lively and opinions are widely divided.[1] Proponents, assuming that excess volatility is due to speculators, claim that a Tobin tax would decrease volatility, because it would make speculative trading unprofitable, hence reducing volume. In addition, the tax would improve the situation of countries damaged by international currency speculation. Opponents reject the proposal, claiming that its implementation would hardly be feasible and would be ineffective if not agreed by all countries. It could change unpredictably the market structure, eventually damaging developing countries, and moreover, through a reduction of market liquidity, Tobin taxes might indeed result in more, not less volatility.

This has called for a closer look in the micro-structure of FX markets (Eichengreen et al. 1995, Frankel 1996, Mende and Menkhoff 2003, Osler 2006). This research has shown that, as in other markets, currencies are driven by order flows, and at horizons of one day or more, these originate from two main types of traders: commercial traders (i.e. non-financial firms engaged in international trade who need currency as part of their primary business), and financial traders (i.e. institutional asset managers who care for profits generated in the trading activity). Both commercial and financial traders act on time horizons of one day or longer, and their activities are negatively correlated to each other, "meaning that at horizons of a day or longer financial demand tend to be met by commercial supply" (Osler). A third group of "liquidity or noise traders", in Osler, or banks, in Mende and Menkhoff, who trade in a manner that is unrelated to information flows, is mainly responsible for FX activity at shorter time-scales. These issues have also been addressed from the more theoretical approach to modeling financial markets. Research has departed from models based on the rational expectations paradigm such as Bacchetta and Wincoop (2000) and has focused on models that are capable of reproducing excess volatility and the typical features of fluctuations in real markets (the so-called stylized facts discussed by Cont 2001) and are therefore well-suited to shed light on the effects of introducing a Tobin tax. One strand of research has concentrated on models of interacting agents with heterogeneous (adaptive) expectations (see the survey by Hommes 2006). These consider agents with an optimizing behavior with respect to an expectation model. Two main types of expectations are considered: those of fundamentalists who expect current price to revert to the fundamental price and trend followers (or chartists), who expect exchange rates to follow a trend.

---

[1] A compendium of the different points of view and arguments is in ul Haq et al. (Eds.) 1996. The question about possible implementation is not only academic; contrasting references are Ramonet (1997) and European Banking Federation (2001).

2

DeGrauwe and Grimaldi (2006) have shown that this approach can be adapted to modeling FX markets, obtaining three main empirical facts: the 'disconnect' puzzle, excess volatility, and non-normality. Westerhoff (2003) analyzed Tobin tax in this framework, showing that its introduction reduces volatility. Westerhoff and Dieci (2006) find moreover that if agents can trade on multiple markets and a Tobin tax is introduced in one of them, that markets stabilize and become attractive for risk–averse investors. If market regulators are competing on investors, the Tobin tax introduction would follow in the other markets as well. [2]

A completely different approach, has been taken by Ehrenstein et al. (2005) in zero intelligence atomistic models based on percolation theory (Cont and Bouchaud 2000). These models relate the emergence of excess volatility and fat tails in the distribution of returns to herding behavior in the population of traders. Ehrenstein et al. show that generally the introduction of a Tobin tax brings about a reduction in volatility, as long as the tax rate is not too high to cause liquidity problems. This second class of models, though neglecting any notion of optimizing behavior, has the virtue of taking into account the discrete nature of traders, whereas in the heterogeneous agent approach only the effect of the aggregate demand of different types of traders matters.

The present paper addresses similar general questions but from a different approach, that of Minority Games (MG) (see Challet et al. 2005, Coolen 2005 and Johnson et al. 2003). These are stylized models that depict a financial market as an ecology of different types of agents interacting along an 'information food chain' where speculators 'predate' on market inefficiencies (arbitrage opportunities) created by other investors. [3] This is particularly suited for FX markets as it captures the interplay between commercial traders and financial speculators. In addition, at odds with the heterogeneous agent approach, the demands on the market are not just the aggregated demands of a few trader types, but in principle every agent differs from the others and full heterogeneity is considered. [4]

---

[2] As we discussed above, one argument against the Tobin tax is that it would be difficult to coordinate all markets for its adoption. The result of Westerhoff and Dieci is a good counterargument.

[3] In the MG setup the speculators try to play against the market, so they act as contrarians. This assumption seems incompatible with the usual one of speculators as trend followers. What these agents, however, try to do is to anticipate a reversion of the market trend (e.g. ride a bubble up to the very last minute). This is compatible with most of the technical trading, from moving averages to 'scalping', so that contrarian speculators are actually trend followers who try to anticipate the market (see e.g. the model proposed in Chiarella et al. 2006).

[4] Empirical works such as Frankel, and Osler, find the negative correlation discussed above between the orders of commercial and financial traders. The micro–founded models (Westerhoff, DeGrauwe and Grimaldi, and Westerhoff and Dieci) mimic this negative correlation with the interplay between fundamentalists and chartists: when real prices' drift departs from the fundamental prices, the orders of these two types

3

The MG highlights the tradeoff between volatility and market efficiency in a vivid though admittedly simplified and stylized way. Indeed the analysis of the MG has revealed that within this model framework excess volatility and market efficiencies are identified as two sides of the same coin, both resulting as consequences of speculative trading. The MG exhibits two different regimes, one in which the market is fully efficient and another in which arbitrage opportunities are not entirely eliminated by the dynamics of the agents. These regimes are separated sharply in the parameter space of the model, and it turns out that the boundary at which the market becomes efficient coincides precisely with the locus of a phase transition in the language of statistical physics. At the same time non-trivial fluctuations in the returns, very similar to the stylized facts observed in real market data, emerge in the vicinity of this transition but not further away (Challet and Marsili 2003).[5] Hence, at odds with previous models, the MG allows us to investigate how the introduction of transactions taxes affects the information ecology and market efficiency.

A further advantage of the MG over the other heterogeneous agent–based approaches discussed above lies in the fact that, despite its stylized nature, it exhibits a remarkably rich phenomenology that can be understood fully analytically with tools of statistical physics of disordered systems (as discussed in Challet et al. 2005 and Coolen). Analytical tractability provides an understanding that is much more complete than that based on numerical simulations. In particular it is possible to derive closed analytical expressions for key observable variables such as the market volatility, the trading activity of the agents and the revenue for the market maker (as will be defined).

In brief, our main result is that within the picture of the MG model, a small tax decreases volatility whenever a market with a finite number of traders exhibits anomalous fluctuations. Indeed the fundamental effect of the tax is to draw the market away from the critical point discussed above. As the occurrence of anomalous fluctuations and their size depend inversely on the size (i.e. number of agents) of the market, the introduction of a Tobin tax has a weaker impact as the size of the market increases. At the same time, the tax introduces an information inefficiency, and thus too high a tax might not be advisable. The total revenue from the tax exhibits a maximum for intermediate rates similar to what was found in earlier works on different models (as

---

of traders are anticorrelated. Also the two types of agents considered in the MG submit orders that are on average negatively correlated, but in a more heterogeneous way.

[5] The two regimes of the model are characterized respectively by high market efficiency and many speculators at one side, market inefficiency and few speculators at the other. If speculators were allowed to enter the market when expected profits are high and leave it when they are low, the system would balance itself exactly in–between the two regimes (i.e. close to the transition point). This stylized model is therefore explicative on why we observe only the phase–transition dynamics in real markets, as discussed in Challet and Marsili.

4

in Ehrenstein 2002). Furthermore, the effects of imposing a tax materialize in the market behavior only after times that scale inversely with the tax rate. Extremely small tax rates may thus need a long time to stabilize turbulent markets. This can be quite relevant if this time scale becomes comparable to that over which market's composition changes because then speculators may fail to 'learn' how to coordinate on low volatility states. Introducing a Tobin tax reduces the time needed to reach coordination, thus reducing the volatility significantly even in an *infinite* system.

In the following we shall first introduce the grand-canonical MG (GCMG) and re-iterate its known main features, while technical details are in the Appendix (available on the JEBO website). In Sections 3 and 4 we comment on how a tax on transactions can be introduced and discuss the effects on the market within the present model. In Section 5 we turn to a brief discussion on how to relate these outcomes to real markets, summarize our results, and give some final concluding remarks.

2 The Grand-canonical Minority Game

2.1 The model

The so-called grand-canonical MG (GCMG) describes a simple market of N agents $i = 1, \ldots, N$ who at each round of the game make a binary trading decision (to buy or to sell) or who each may decide to refrain from trading. They thus each submit bids $b_i(t) \in \{-1, 0, 1\}$ in every trading period $t = 0, 1, 2, \ldots$, resulting in a total excess demand of $A(t) = \sum_{i=1}^{N} b_i(t)$.

These trading decisions are taken to be based on a stream of information available to the agents. This common information on the state of the market (or on other factors relevant to the market) is encoded in an integer variable $\mu(t)$ at time t taking values in $\{1, 2, \ldots, P\}$.[6] Here we assume that $\mu(t)$ models an exogenous news arrival process and that the $\{\mu(t)\}$ are drawn at random from the set $\{1, \ldots, P\}$, independently and with equal probabilities at each time.[7] The objective of each agent is to be in the minority at each time-step

---

[6] This is an assumption on the cognitive limitation of agents because the game itself will generate much richer information than just the sequence $\mu(t)$.

[7] This contrasts with the definition of $\mu(t)$ in the original MG, where the information was endogenously generated by the market, with $\mu(t)$ encoding the sign of the past $M = \log_2 P$ price changes. Most collective properties of the model have been shown to be affected only weakly by the origin of information. We focus here on the simpler case of exogenous information, which is analytically more convenient. Note however that analytical approaches to MGs with endogenous information are also feasible, but involve much more intricate mathematics (as in Coolen).

5

(i.e. to place a bid $b_i(t)$ that has the opposite sign of the total bid $A(t)$.)
This minority setup, discussed in previous section, corresponds to contrarian
behavior and can be derived from a market mechanism taking into account the
expectations of the traders on the future behavior of the market (Marsili 2001).
In order to do so, each agent has a 'trading strategy' at his disposal. Trader i's
strategy is labeled by $\mathbf{a}_i = (a_i^\mu)_{\mu=1,...,P} \in \{-1,1\}^P$ and provides a map from all
values of the information μ onto the binary set $\{-1,1\}$ of actions (buy/sell).
Upon receiving information μ the trading strategy of agent i thus prescribes to
take the trading action $a_i^\mu \in \{-1,1\}$. These strategies are assigned at random
and with no correlations at the beginning of the game, and then they remain
fixed. [8] Agents in the GCMG are adaptive and may decide not to trade if they
do not consider their strategy adequate. More precisely, each agent keeps a
score $u_i(t)$ measuring the performance of his strategy vector. He then trades
at a given time-step t only if his strategy has a positive score $u_i(t) > 0$ at
that time. Therefore, the bids of agents take the form $b_i(t) = n_i(t)a_i^{\mu(t)}$ with
$n_i(t) = 1$ if $u_i(t) > 0$ and $n_i(t) = 0$ otherwise. Accordingly the excess demand
is given by

$$A(t) = \sum_{i=1}^{N} n_i(t)a_i^{\mu(t)}. \tag{1}$$

Agent i keeps a record of the past performance of his strategy a_i^μ by updating
the score $u_i(t)$ as follows:

$$u_i(t+1) = u_i(t) - a_i^{\mu(t)}A(t) - \varepsilon_i, \tag{2}$$

at each step, with constant ε_i. The first term $-a_i^{\mu(t)}A(t)$ is the Minority Game
payoff; it is positive whenever the trading action $a_i^{\mu(t)}$ proposed by i's strategy
vector and the aggregate bid $A(t)$ are of opposite signs, and negative whenever
i joins the majority decision. The idea of Equation (2) is that whenever the
payoff $-a_i^{\mu(t)}A(t)$ is larger than ε_i, the score of player i's strategy is increased;
otherwise it is decreased. The constant ε_i in (2) thus captures the inclination of
agent i to trade in the market. This inclination will in general be heterogeneous
across the population of agents, with agents with high values of ε_i being more
cautious to trade than agents with low ε_i. In our simplified model we consider
only two types of agents. First we assume that there are $N_s \leq N$ financial
speculators who trade only if their perceived market profit obtained by using
their strategy exceeds a given threshold, and hence we set $\varepsilon_i = \varepsilon \geq 0$ for such
agents. Here ε can be considered as the speculative margin of gain in a single

---

[8] We here assume that each agent holds only one trading strategy. Generalizations
to multiple strategies per player are straightforward (Shayeghi and Coolen 2006 and
De Martino et al 2007) and can be seen not to alter the qualitative behavior of the
model. Hence we here restrict to the simplest case.

6

transaction.

The remaining $N_c = N - N_s$ agents (the commercial traders) are assumed to trade no matter what, so a tax would not affect their behavior. Mathematically this is implemented by setting $\varepsilon_i = -\infty$ for this second group of agents. They have $n_i(t) = 1$ for all times t. For convenience we will order the agents such that speculators carry the indices $i = 1, \ldots, N_s$ and commercial traders the labels $i = N_s + 1, \ldots, N$.

2.2 Price process, volatility and predictability

Within the MG setup the market volatility is then given by

$$\sigma^2 = \frac{\langle A(t)^2 \rangle}{N}, \tag{3}$$

where $\langle \ldots \rangle$ will stand for a time-average in the stationary state of the model from now on. The normalization to the number of agents N is here introduced to guarantee a finite value of σ^2 in the infinite-system limit, with which the statistical mechanics analysis of the model is concerned.

The information variable $\mu(t)$ allows one to quantify information-efficiency of the model market by computing the predictability

$$H = \frac{1}{PN} \sum_{\mu=1}^{P} \langle A | \mu \rangle^2 \tag{4}$$

where $\langle \ldots | \mu \rangle$ denotes an average conditional on the occurrence of information pattern $\mu(t) = \mu$. A value $H \neq 0$ indicates that for some μ the minority payoff is statistically predictable $\langle A(t) | \mu \rangle \neq 0$, whereas the market is unpredictable and fully informationally efficient when $H = 0$.

The simplest way to relate this picture to a financial market is to postulate a simple linear impact of $A(t)$ on the (logarithmic) price (or the exchange rate); that is, assume that

$$p(t + 1) = p(t) + \frac{A(t)}{\lambda}, \tag{5}$$

where $\{p(t)\}$ denotes a price (exchange rate) process and where λ is the liquidity. Equation (5) can be justified, as usual, assuming a market maker with a high currency availability who satisfies orders and balances the prices. [9] The liquidity typically increases with the volume (i.e. with the number of traders in

---

[9] The market maker hypothesis was first adopted in heterogeneous agent–based models for financial markets by Beja and Goldman (1980), then by Day and Huang

this context). For convenience, [10] we make the simple assumption $\lambda = \sqrt{N}$, so that σ^2 becomes the volatility of the price process, $\sigma^2 = \langle (p(t+1) - p(t))^2 \rangle$. Notice that the notion of information efficiency provided by the value of H is somewhat more primitive than the usual one, which relates to the deviation of prices from fundamentals as it directly relates to the possibility of predicting the market. In a naïve view, one might regard the action of commercial traders as mimicking the contribution of fundamental trading (i.e. trading related solely on the stream of information $\mu(t)$ arriving to the market). This contribution arrives through the behavior of financial speculators (i.e. of adaptive agents).

2.3 The behavior of the GCMG

The GCMG has been studied in great detail in Challet and Marsili and Challet et al. (2006), with methods well established in statistical mechanics. The analysis is here generally concerned with the stationary states of the system (i.e. with the behavior reached after running the learning dynamics of the agents for some sufficiently long transient equilibration time).

The statistical mechanics approach provides exact results for the model in the limit of infinite market sizes, where one takes the number of agents $N = N_s + N_c$ and the number P of possible different information states to infinity, while at the same time keeping the ratios $n_s = N_s/P$ and $n_c = N_c/P$ fixed and finite. n_s and n_c along with ε are thus control parameters of the model. This approach makes it possible to derive exact expressions for several quantities, including the predictability H, and upon neglecting time-dependent correlations accurate approximations for the volatility σ^2 can be found (see also Marsili and Challet 2001, and Heimel and Coolen 2001). We will not enter here the detailed mathematics of the calculations, based on the replica method

---

(1990), and became a standard assumption thereafter. Assuming the presence of a market maker in real foreign exchange markets is a strong simplification. This is because of the less rigorous institutionalization of foreign exchange markets as compared to financial markets. What we mean by 'market maker' is the set of actively trading banks in the market and of the public institutions that may benefit from taxation.

[10] The dependence of liquidity on trading volume has indeed been advocated as a possible source of problems when a tax on transaction is levied. The reason is that a reduction in volume may make the market less liquid and hence increase volatility. This argument does not apply to our model, because the volume of commercial traders is not affected by the tax, and therefore the reduction in the volume is only moderate. We explicitly checked that, even in the case where λ is proportional to the number of active traders, as suggested in Marsili on the basis of a different market mechanism, the qualitative behavior of numerical simulations is the same as that discussed below for the simpler choice, $\lambda = \sqrt{N}$.

8

of statistical physics (Mezard et al. 1987). The resulting equations for the key quantities in the stationary states as well as a sketch of their derivation are found in the Appendix (available on JEBO website). Further details regarding the statistical mechanics analysis of MGs and GCMGs can be found in Challet and Marsili, Challet et al. (2005), Coolen, and in references therein.

[Figure 1 approximately here]

The overall picture that emerges is the following: at fixed n_c, the statistical behavior of the model is characterized by a critical line at $\varepsilon = 0$ that extends beyond some critical value $n_s \geq n_s^\star(n_c)$. This is illustrated in Figure 1. As this line (segment) is approached in parameter space the market becomes more and more efficient (i.e. $H \to 0$ as $\varepsilon \to 0$ for $n_s \geq n_s^\star$). On the critical line ($\varepsilon = 0, n_s \geq n_s^\star$) itself the market is fully efficient, and one finds $H = 0$ exactly in the limit of infinite system size. In addition, numerical simulations of the model at finite sizes close to the critical line reveal fluctuation properties similar to those observed in real markets (and discussed in Cont). In particular $A(t)$ has a fat tailed distribution, and one observes volatility clustering. These effects become weaker as the system size is increased at constant values of n_s, n_c and ε, and similarly they disappear gradually when one moves away from the critical line at fixed system size.

The case $\varepsilon = 0$ and $n_s > n_s^\star(n_c)$ is peculiar because it turns out that the stationary state is here not unique, but rather it depends on the initial conditions from which simulations are started. In the following we will not consider the case of a strictly vanishing ε, but will assume instead that speculators have a positive profit margin $\varepsilon > 0$, even if the latter may be small. All simulations on which we report are started from zero initial conditions $u_i(t = 0) = 0$ for $i = 1, \ldots, N_s$.

We finally remark that a detailed analysis of the transient dynamics demonstrates that for small ε the stationary state is reached after a characteristic equilibration time that scales as $1/\varepsilon$. For $\varepsilon \ll 1$, such a long equilibration becomes relevant if the market composition changes in time, as discussed in Section 4.

3 Tobin tax in the GCMG

Within the model setup the introduction of a tax τ on each transaction can be accounted for by a change $\varepsilon_i \to \varepsilon_i + \tau$ for all $i = 1, \ldots, N$. Indeed by raising ε by an amount τ, an additional cost τ incurs every time for any given agent who trades and no costs for agents who refrain from trading. Note that the trading volume of any fixed (active) agent is one unit in our simple setup so that τ indeed corresponds to a transaction tax per unit traded. Hence we will

9

assume

$$\varepsilon_i = \begin{cases} \varepsilon + \tau \ i \le N_s \\ -\infty \ i > N_s \end{cases} \tag{6}$$

in the following. While this will discourage speculators from trading (via the reduction of their strategy score), such a tax will have no effect on the participation of commercial traders. They will trade at every time step as before ($n_i(t) = 1 \ \forall i > N_s$ at all times t).
The total revenue from this tax τ is then given by

$$R = \frac{\tau}{N} \sum_i \langle n_i(t) \rangle = \frac{\tau}{N} \sum_{i=1}^{N_s} \langle n_i(t) \rangle + \tau \frac{n_c}{n_s + n_c} \equiv R_s + R_c, \tag{7}$$

where the first term corresponds to the revenue R_s from speculators and the second (R_c) to that obtained from the commercial traders.
An evaluation of the effects of levying a tax τ on the GCMG then amounts to studying the behavior of the model as a function of ε at fixed model parameters n_s and n_c. It turns out that here one has to distinguish between two different regimes, namely close and far away from the phase transition.

[Figure 2 approximately here]

Figure 2 reports the effects of introducing a tax on markets whose parameters are far from the critical line. We here consider $n_c = 1$ and $n_s = 1 < n_s^\star(n_c = 1) \approx 4.15$, so that one operates sufficiently far to the left of the critical region depicted in Figure 1. For such parameter values the results of numerical simulations follow the curves predicted by the analytical theory perfectly, and no anomalous fluctuations are present in the corresponding price time-series. The tax has very mild effect both on volatility and on the information efficiency, as long as $\varepsilon + \tau \ll 1$. Figure 2 also shows that the contribution of speculators to the tax revenue has a peak at intermediate tax rates, but that at the same time the revenue R_s obtained from speculators is smaller than that from institutional investors.

[Figure 3 approximately here]

Figure 3 on the other hand illustrates the response of the market at the other extreme where $n_s \gg n_s^\star(n_c)$. For such values of the parameters one is within the critical region (for small enough $\varepsilon + \tau$), and the model exhibits strong anomalous price fluctuations for market sizes of a few thousand agents. The deviations from the analytical theory (which is valid only in the limit of infinite systems) mark strong finite-size effects in the critical region. As illustrated in the lower panel of Figure 3, imposing a sufficiently large tax may in such

10

markets have a pronounced effect on the volatility, whereas smaller transaction fees may influence the time-series of the market only marginally.

[Figure 4 approximately here]

Figure 4 presents a systematic account of these effects and shows the dependence of the volatility, the predictability, and the revenue from the tax on the system size and the tax rate τ at $n_s \gg n_s^\star(n_c)$. In particular, a significant reduction of the market volatility can be obtained while still keeping the market relatively information efficient. Furthermore, the contribution to the tax revenue of speculators largely outweighs that of producers, and it is peaked at a value close to that where the volatility is minimal. The effect of a tax, as shown in Figure 4, also depends on the size of the market. The volatility at low $\varepsilon + \tau$ indeed decreases with the size of the system and approaches the theoretical line, making a tax more effective in small than in large markets.

[Figure 5 approximately here]

Figure 5 relates these two extremes and discusses the dependence of the volatility on n_s at intermediate number of speculators for fixed $n_c = 1$. We here fix the (effective) system size by keeping $L = PN_s$ constant. For small values of n_s one is then well outside the critical region, and the numerical results follow the analytical predictions (solid lines in Figure 5). As discussed in Challet and Marsili, the simulations then deviate systematically from the theory at large n_s when the system has entered the zone near the phase transition line. More precisely, one finds a threshold value $\overline{n}_s(L) > n_s^\star$ so that numerical simulations agree with the theoretical lines for $n_s < \overline{n}_s(L)$, but deviations and anomalous fluctuations are observed for $n_s > \overline{n}_s(L)$. As L is increased $\overline{n}_s(L)$ is found to grow as well in simulations (not shown here), and in particular one has $\lim_{L \to \infty} \overline{n}_s(L) = \infty$ (at $\varepsilon + \tau \neq 0$) so that the critical region vanishes in the limit of infinite systems.

4   Markets with evolving composition of agents

In the previous sections we have assumed that all traders stay in the market for an infinite amount of time and that their trading strategies remain fixed forever. Individual agents have the option to abstain from trading at intermediate times and to join the market again at a later stage, but no agent in the setup considered thus far can actually modify his strategy vector $\{a_i^\mu\}$. Thus the composition of the population of traders does not change over time. In real-market situations however it would appear more sensible to expect some fluctuation in the population of traders and to assume that the market composition will evolve and/or that strategies get replaced after some time. In the

11

latter case, one would expect poorly performing strategies to be removed from the market and replaced by new ones.

In this section we consider the simplest case of an evolving composition of the market, namely a situation in which agents (or equivalently their strategies) are replaced randomly, irrespective of their performance. More precisely, at each time step each speculator is removed with a probability $1/(\theta N_s P)$ and replaced by a new one with randomly drawn strategy and zero initial score. Here θ is a constant, independent of the system size. This choice $\theta = \mathcal{O}(L^0)$ guarantees that the expected survival time of any individual agent scales as $N_s P$ so that one exit/entry event occurs in the entire population on average over a period of θP transaction time steps. Relaxation times in Minority Games are known to be of the order of P, so the above scaling of θ results in the composition of the market changing slowly on times comparable with those on which the system relaxes. Indeed, extensive numerical simulations show that the behavior of the volatility on θ as well as that of other quantities characterizing the collective behavior of the system is independent of the system size (see Figure 6). This is in sharp contrast with the strong finite size effects observed for $n_s \gg n_s^\star$ at a fixed composition of the population of agents (Figure 4).

[Figure 6 approximately here]

The main feature of the MG market with an evolving population of agents is a pronounced minimum of the volatility as a function of $\varepsilon + \tau$ in the crowded regime $n_s > n_s^\star$. In particular the volatility increases as $\varepsilon + \tau$ is decreased, even in the limit of large system sizes (in which a corresponding system with fixed agent population would equilibrate to the flat theoretical line as shown in Figure 4). This behavior of the system with changing agent structure can be related to the fact that relaxation time of the GCMG scales as $1/(\varepsilon + \tau)$. Thus, when $\epsilon + \tau$ is very small, the time it would take a fixed population of agents to equilibrate collectively can be much larger than the time scale θP over which the market composition changes. In this case, the market remains in a highly volatile state indefinitely because the agents do not have sufficient time to 'coordinate' and to adjust their respective behavior since the strategy pool represented in the evolving population of agents changes too quickly. Figure 6 demonstrates that introducing a tax in such markets with dynamically evolving trader structure can reduce the volatility considerably.

5 Conclusions

We have shown how the theoretical picture derived for the GCMG (Challet and Marsili, Challet et al. 2005, Coolen) can be used to characterize fully the impact of a Tobin tax on this toy model of a currency market. The main re-

12

sult of our study (see Figure 4) is that within the GCMG the introduction of a Tobin Tax reduces the market volatility whenever the market is operating close to the critical line ($\varepsilon = 0$, $n_s > n_s^\star$) of informationally efficient markets ($H = 0$). This region of parameter space is the one relevant for real markets *1)* because real markets are believed to be nearly efficient and *2)* because only in this region does the model exhibit stylized facts, such as a fat tailed return distributions and volatility clustering, qualitatively similar to those of real markets.

Within the simplified picture of the GCMG, we find that if the market is far from the critical line, the introduction of a tax has a weak effect on the volatility. If instead the market operates close to the critical line, then a tax τ draws the market away from the critical region in parameter space, and thus reducing volatility. Furthermore, the size of anomalous fluctuations and of the region where they occur depend inversely on the system size, suggesting that the introduction of a tax might be particularly effective in small markets.

In all cases we find that the contribution of speculators to the revenue for the market maker from the tax attains a maximal value at intermediate tax rates. In the case of a market near criticality, this occurs approximately at a tax rate at which the reduction in volatility is largest. In both shown examples (Figures 2 and 4) the revenue for the market maker appears to weight more on commercial traders than on speculators. Only for unrealistic cases when the number of commercial traders is extremely small or when $\varepsilon < 0$ does one observe instances in which the contribution of speculators is largest.

Finally our findings demonstrate that imposing a tax can also reduce the market volatility in cases where the composition of the population of traders changes slowly over time. In this case, the tax allows the agents to reach a coordinated state faster so that the market can reach a stationary state of relatively low volatility.

Given the stylized nature of the MG, it is hard to make a connection between the model parameter τ and an actual tax rate in a real-world market. At any rate it seems reasonable to assume that a realistic tax rate should be of the order or smaller than the margin of profit of speculators, which is gauged by ε in the GCMG. The optimal tax rate τ might be unrealistically large compared to ε. For example in Figure 4 if $\varepsilon = 0.01$ volatility can be substantially reduced only for tax rates τ that are more than ten times larger.

The GCMG can at best be seen as a minimalistic, simplified version of a real market. Hence our conclusions on the behavior of the MG are at most suggestive with respect to what might happen in the real world. Still, thanks to its analytic tractability, the Minority Game provides a coherent picture of the interplay between stochastic fluctuations and information efficiency going far beyond the insights of zero–intelligence or agent based models. The picture developed here can be extended in a number of directions toward more realistic market modes without giving up analytical tractability. One of the most interesting future directions might be to endow agents with individual wealth variables that evolve according to their relative success when trading in the

model market.

Appendix

We here sketch the theoretical analysis of the model with $\tau = 0$ and general values of ε. The introduction of a tax τ can be accounted for by replacing $\varepsilon \to \varepsilon + \tau$ in all equations below. The starting point of the statistical mechanics approach is the function

$$H_\varepsilon[\{\phi_i\}] = \frac{1}{P} \sum_{\mu=1}^{P} \left[\sum_{i=1}^{N} a_i^\mu \phi_i \right]^2 + \frac{2\varepsilon}{P} \sum_{i=1}^{N_s} \phi_i \tag{8}$$

of the mean activities $\{\phi_i = \langle n_i(t) \rangle\}$ of the speculators $i = 1, \ldots, N_s$. The ϕ_i, $i = 1, \ldots, N_s$ are continuous variables within the interval $[0, 1]$; recall that commercial traders are always active and have $\phi_i = 1$, $i = N_s + 1, \ldots, N = N_s + N_c$. Note that this function depends explicitly on the strategy assignments $\{a_i^\mu\}$, so H_ε is a stochastic quantity. The strategy vectors, which are fixed at the beginning of the game, correspond to what is known as 'quenched disorder' in statistical mechanics. It turns out that the learning dynamics (2) minimizes the function H_ε in terms of the $\{\phi_i\}$ for any fixed choice of the strategy vectors. Computing the stationary states of the model thus reduces to identifying the minima of H_ε. It is here possible to characterize these minima using the so-called replica method of statistical physics (Mezard et al.). A different statistical mechanics approach is based on so-called generating functionals and deals directly with the update dynamics (2); see Coolen. Both methods ultimately lead to the same equations describing the stationary states of the model, so we here restrict the discussion to the former approach.

In the following we give a brief sketch of the so-called replica analysis of the model, allowing us to compute the minima of the random function H_ε. To this end one first introduces the partition function

$$Z_\varepsilon(\beta) = \int_0^1 d\phi_1 \cdots \int_0^1 d\phi_{N_s} \exp\left(-\beta H_\varepsilon[\phi_1, \ldots, \phi_{N_s}]\right) \tag{9}$$

at an 'annealing temperature' $T = 1/\beta$. In the limit $\beta \to \infty$ these integrals are dominated by configurations $\{\phi_1, \ldots, \phi_{N_s}\}$ that minimize H_ε so that the evaluation of $\lim_{\beta \to \infty} Z_\varepsilon(\beta)$ allows one to characterize the minima of H_ε.

This procedure is in general not feasible for individual realizations of the random strategy assignments as the dependence of H_ε on the $\{a_i^\mu\}$ is quite intricate. Instead we will compute 'typical' quantities in the limit of infinite

14

systems (i.e averages over the space of all strategy assignments).
The key quantity to compute here is the free energy density

$$f_\varepsilon(\beta) = -\lim_{N \to \infty} \frac{1}{\beta N} \ln Z_\varepsilon(\beta). \tag{10}$$

The limit $N \to \infty$ is here taken at fixed ratios $n_s = N_s/P$ and $n_c = N_c/P$ (so that N_s, N_c, P are taken to infinity as well). All relevant properties of the typical minima of H_ε can be read off from the disorder-average of $\lim_{\beta \to \infty} f_\varepsilon(\beta)$. Using the identity $\ln Z = \lim_{n \to 0} \frac{Z^n - 1}{n}$ this problem can be reduced to computing averages of Z_ε^n that corresponds to an n-fold replicated system with no interactions between the individual copies. This is referred to as the replica method in statistical physics and is a standard tool for the analysis of problems involving quenched disorder (Mezard et al., Challet et al. 2005, Coolen). The averaging procedure leads to an effective interaction between the replicas and requires an assumption regarding the symmetry of the solution with respect to permutations of the replicas. In principle this symmetry may be broken, as different replica copies may end up in different minima of H_ε. In the non-efficient phase of this model, instead, the so-called 'replica symmetric' ansatz is exact, simplifying the analysis considerably. We will here not report the detailed intermediate steps of the calculation, but will only quote the final outcome, namely a set of closed equations for the variables characterizing the minima of H_ε (and hence the stationary states of the GCMG). Further details of the replica analysis are found in Challet and Marsili, Challet et al. (2005) and Coolen.
The minima of H_ε turn out to be described by two independent variables K and ζ, uniquely determined from the following two relations:

$$\zeta = \frac{1}{\sqrt{n_s(Q(\zeta, K) + n_c/n_s)}} \ ,$$
$$K = \varepsilon \left[1 - \frac{n_s}{2} \left(\mathrm{erf}[(1 + K)\zeta/\sqrt{2}] - \mathrm{erf}(K\zeta/\sqrt{2})\right)\right]^{-1} \ , \tag{11}$$

with

$$Q(\zeta, K) = \frac{1}{2} \ \mathrm{erfc}[(1 + K)\zeta/\sqrt{2}] + \frac{1}{\zeta\sqrt{2\pi}} \left[(K - 1)e^{-(1+K)^2\zeta^2/2} - Ke^{-K^2\zeta^2/2}\right]$$
$$+ \frac{1}{2} \left(K^2 + \frac{1}{\zeta^2}\right) \left(\mathrm{erf}[(1 + K)\zeta/\sqrt{2} - \mathrm{erf}(K\zeta/\sqrt{2})\right) \ , \tag{12}$$

where erf and erfc are, respectively, the error function encountered when integrating the normal distribution ($\mathrm{erf}(x) \equiv \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$) and its complementary ($\mathrm{erfc}(x) \equiv 1 - \mathrm{erf}(x)$). These equations are easily solved numerically and one

15

obtains K and ζ as functions of the model parameters $\{n_s, n_c, \varepsilon\}$. The disorder-average of quantities such as the predictability H or the mean activity of the speculators $\phi = \lim_{N \to \infty} N_s^{-1} \sum_{i=1}^{N_s} \phi_i$ can then be expressed in terms of K and ζ. One finds

$$
\begin{aligned}
H &= \varepsilon^2 \frac{n_c + n_s Q(\zeta, K)}{(n_s + n_c)K^2} \; , \\
\phi &= \frac{1}{2}\mathrm{erfc}[(1+K)\zeta/\sqrt{2}] + \frac{1}{\zeta\sqrt{2\pi}}\left(e^{-K^2\zeta^2/2} - e^{-\zeta^2(1+K)^2/2}\right) \\
&\quad + \frac{K}{2}(\mathrm{erf}(K\zeta/\sqrt{2}) - \mathrm{erf}[(1+K)\zeta/\sqrt{2}]) \; .
\end{aligned}
\tag{13}
$$

These results are fully exact in the thermodynamic limit, with no approximations (except for the replica-symmetric ansatz) made at any stage. Finally, neglecting certain dynamical correlations between agents, the volatility can be approximated as

$$
\sigma^2 = \varepsilon^2 \frac{n_c + n_s Q(\zeta, K)}{(n_s + n_c)K^2} + n_s \frac{\phi - Q(\zeta, K)}{n_s + n_c}.
\tag{14}
$$

As shown in the main text of the paper this approximation is in excellent agreement with numerical simulations (up to finite-size and equilibration effects).

References

[1] Bacchetta, P., van Wincoop, E., 2000. Does Exchange-Rate Stability Increase Trade and Welfare? The American Economic Review 90, 1093–1109.

[2] Beja, A., Goldman, M. B., 1980. On the dynamic behavior of prices in disequilibrium. Journal of Finance 35, 235–248.

[3] Challet, D., De Martino, A., Marsili, M., Perez Castillo, I., 2006. Minority games with finite score memory. Journal of Statistical Mechanics: Theory and Experiments, P03004.

[4] Challet, D., Marsili, M., 2003. Criticality and market efficiency in a simple realistic model of the stock market. Physical Review E 68, 036132.

[5] Challet, D., Marsili, M., Zhang, Y. C., 2005. Minority Games. Oxford: Oxford University Press.

[6] Chiarella, C., He, X. Z., Hommes, C., 2006. A dynamic analysis of moving average rules. Journal of Economic Dynamics and Control 30, 1729-1753.

[7] Cont, R., 2001. Empirical properties of asset returns: stylized facts and statistical issues. Quantitative Finance 1, 223–236.

16

[8] Cont, R., Bouchaud, J. P., 2000. Herd behaviour and aggregate fluctuations in financial markets. Macroeconomic Dynamics 4, 170–196

[9] Coolen, A. C. C., 2005. The Mathematical Theory of Minority Games. Oxford: Oxford University Press.

[10] Day, R.H., Huang, W., 1990. Bulls, bears and market sheep. Journal of Economic Behavior and Organization 14, 299–329.

[11] De Grauwe, P., Grimaldi, M., 2006. Exchange rate puzzles: A tale of switching attractors. European Economic Review 50, 1–33.

[12] De Martino, A., Castillo, P., Sherrington, D., 2007. On the strategy frequency problem in batch Minority Games. Journal of Statistical Mechanics: Theory and Experiments, P01006.

[13] Ehrenstein, G., 2002. Cont-Bouchaud Percolation Model including Tobin tax. International Journal of Modern Physics 13, 1323–1331.

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

[15] Eichengreen, B., Tobin, J., Wyplosz, C., 1995. Two cases for throwing sand in the wheels of international finance. Economic Journal 105, 162-172.

[16] European Banking Federation, 2001. FBE Letter. A tax on foreign exchange transactions: A false solution to the challenges posed by financial markets, March 2001.

[17] Frankel, J. A., 1996. How well do markets work: might a Tobin tax help? In: ul Haq, M., Kaul, I., Grunberg, I., (Eds). The Tobin tax, Oxford: Oxford University Press, 41–82.

[18] Frankel, J. A., Rose, A. K., 1995. A Survey of Empirical Research on Nominal Exchange Rates. In: Grossman, G., Rogoff, K. (Eds). Handbook of International Economics. Amsterdam: North Holland, 1689–1729.

[19] Galati, G., Melvin, M., 2004. Why has FX trading surged? Explaining the 2004 triennial survey. Bank of International Settlements Quarterly Review, Vol. 4, 67–74.

[20] ul Haq, M., Kaul, I., Grunberg, I., (Eds). The tobin tax. Oxford: Oxford University Press.

[21] Heimel, J.A.F., Coolen, A.C.C., 2001. Generating functional analysis of the dynamics of the batch minority game with random external information. Physical Review E 63, 056121

[22] Hommes, C. H., 2006. Heterogeneous Agent Models in Economics and Finance. In: Tesfatsion, L., Judd, K. L. (Eds.). Handbook of Computational Economics, Vol. 2. Amsterdam: Elsevier 1109–1186.

[23] Johnson, N. F., Jefferies, P., Hui, P. M., 2003. Financial market complexity. Oxford: Oxford University Press.

[24] Marsili, M., 2001. Market mechanism and expectations in minority and majority games. Physica A 299, 93–103.

[25] Marsili, M., Challet, D., 2001. Continuum time limit and stationary states of the minority game. Physical Review E 64, 056138.

[26] Mende, A., Menkhoff, L., 2003. Tobin tax effects seen from the foreign

17

exchange market's microstructure. International Finance 6, 227–247.

[27] Mezard, M., Parisi, G., Virasoro, M., 1987. Spin glass theory and beyond. Singapore: World Scientific Publishing.

[28] Osler, C. L., 2006. Macro lessons from microstructure. International Journal of Finance and Economics 11, 55-80.

[29] Ramonet I., 1997. Disarming the markets. Le Monde Diplomatique, Vol. 12.

[30] Shayeghi, N., Coolen, A. C. C., 2006. Generating functional analysis of batch minority games with arbitrary strategy numbers. Journal of Physics A: Mathematical and General 39, 1392113950.

[31] Tobin, J., 1974. The New Economics One Decade Older. Princeton: Princeton University Press.

[32] Westerhoff, F., 2003. Heterogeneous traders and the Tobin tax. International Journal of Modern Physics 13, 1323–1331.

[33] Westerhoff, F., Dieci, R. 2006. The effectiveness of Keynes-Tobin transaction taxes when heterogeneous agents can trade in different markets: A behavioral finance approach. Journal of Economic Dynamics and Control, 30, 293-322.
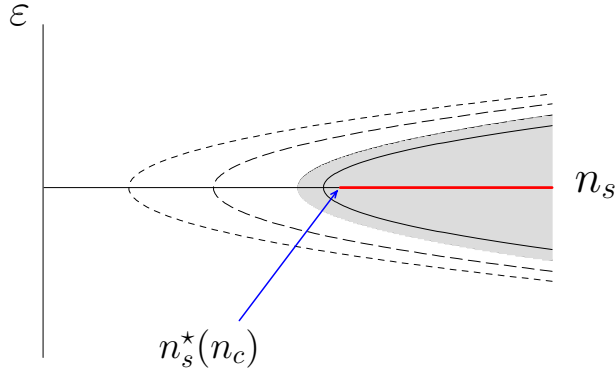
18

Figures

Fig. 1. Phase diagram of the GCMG in the (n_s, ε)-plane at fixed n_c. The red line segment at $\varepsilon = 0$ and $n_s \geq n_s^{\star}(n_c)$ marks the phase transition in the limit of infinite system size. At finite size anomalous fluctuations and stylized facts are found in a region around this critical line, as indicated by the shaded area. This so-called 'critical region' indicates regions with strong dynamical and finite size effects and is large for small systems and shrinks towards the critical line segment in the infinite-size limit.
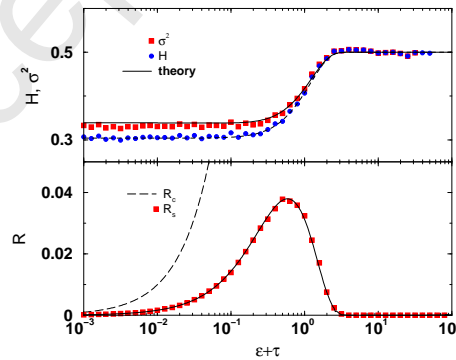
Fig. 2. Volatility, predictability (top) and revenue from speculators and producers (bottom) for $n_s = n_c = 1$. Symbols are data obtained from numerical simulations with $PN_s = 6000$, every data point is an average over at least 1000 samples, simulations are run for $1000P$ steps, with measurements taken in the second half of this interval. Lines are the corresponding predictions for infinite systems obtained from the analytical theory.
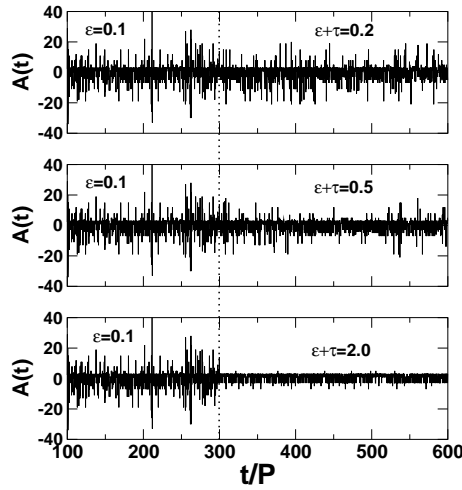
19

Fig. 3. Effect of introducing a tax on the exchange rate fluctuations $A(t)$. Each panel corresponds to a single run with parameters $n_s = 60$, $n_c = 1$, $PN_s = 1600$, $\varepsilon = 0.1$. In the initial period up to $t = 30P/\varepsilon$ no tax is imposed ($\tau = 0$). At $t\varepsilon/P = 30$ a tax rate ($\tau > 0$) is introduced and then kept fixed for the rest of the simulation. The tax rate increases from top ($\varepsilon + \tau = 0.2$) to bottom ($\varepsilon + \tau = 2$). .
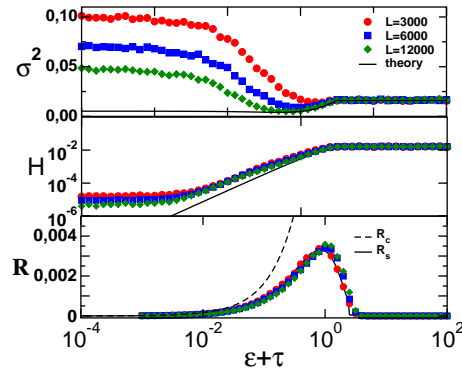

Fig. 4. Volatility (top), predictability (middle) and revenue (bottom) from speculators and producers for $n_s = 60$ and $n_c = 1$. Markers are data obtained from numerical simulations of systems with different (effective) size $L = PN_s$ with circles, squares and diamonds corresponding to $L = 3000, 6000, 12000$ respectively. Every data point represents an average over at least 1000 different strategy assignments. Simulations are run for $400P + 20Ns/(\varepsilon + \tau)$ steps, with measurements taken in the second half of this interval. Lines are the predictions of the analytical theory for the infinite system.
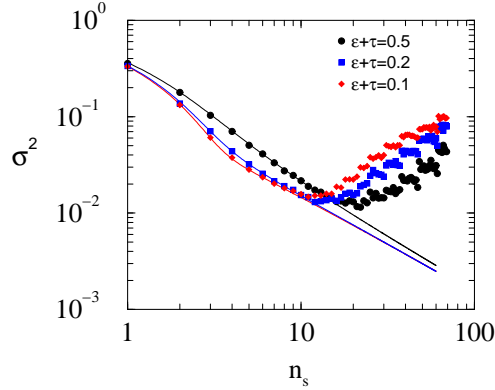
20

Fig. 5. Volatility as a function of n_s at fixed $n_c = 1$ for different tax rates ($\varepsilon + \tau = 0.1$, 0.2 and 0.5). Markers represent data obtained from numerical simulations with $PN_s = 3000$, run for $200P + 200P/(\varepsilon + \tau)$ steps (with measurements in the second half of this interval). An average over 3000 samples is taken. Lines are the corresponding predictions obtained from the analytical theory.
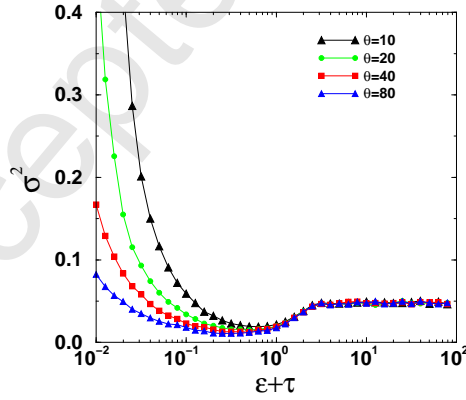

Fig. 6. Volatility as a function of $\varepsilon + \tau$ in markets with slowly changing composition. The parameter θ indicates the rate at which agents are replaced; on average one replacement event occurs in the entire population every θP time-steps (see main text for further details). Simulations are at $n_c = 1, n_s = 20$ and $L = PN_s = 5000$ with an equilibration time $T = 400P/(\varepsilon + \tau)$ and different values of θ as indicated.

21