Forthcoming in the *Journal of Economic Dynamics & Control*

Asset Prices, Traders' Behavior and Market Design*

Mikhail Anufriev [a] Valentyn Panchenko [b,†]

[a] *CeNDEF, Department of Economics, University of Amsterdam, Roetersstraat 11, NL-1018 WB Amsterdam, Netherlands*

[b] *School of Economics, University of New South Wales, Sydney, NSW 2052, Australia*

Abstract

The dynamics of a financial market with heterogeneous agents are analyzed under different market architectures. We start with a tractable behavioral model under Walrasian market clearing and simulate it under different trading protocols. The key *behavioral* feature of the model is the switching by agents between simple forecasting rules on the basis of a fitness measure. By analyzing the dynamics under order-driven protocols we show that the behavioral and structural assumptions of the model are closely intertwined.

*An earlier version of this paper was prepared while Mikhail Anufriev was visiting UNSW, whose hospitality he gratefully acknowledges. We wish to thank all the participants of the Symposium on Agent-Based Computational Methods in Economics and Finance in Aalborg, of the Workshop on "Complexity in Economics and Finance" in Leiden, and of the seminars in Amsterdam, Pisa, Sydney, and Trieste for useful comments and stimulating questions. We especially thank Giulio Bottazzi, Buz Brock, Cars Hommes, and Florian Wagener for providing very detailed comments during different stages of the work. We also thank the editor and two anonymous referees. Financial support for this research was provided by the Complex Markets E.U. STREP project 516446 under FP6-2003-NEST-PATH-1 and the Research Grant of the Australian School of Business. The usual disclaimers apply.

†Corresponding author. Tel.: +61 2 9385 3363; fax: +61 2 9313 6337; e-mail: v.panchenko@unsw.edu.au.

The high responsiveness of agents to a fitness measure causes excess volatility, but the frictions of the order-driven markets may stabilize the dynamics. We also analyze and compare allocative efficiency and time-series properties under different protocols.

JEL codes: G12, D44, D61, C62.

Keywords: Asset pricing model, Heterogeneous beliefs, Learning, Trading protocols, Market architecture

1 Introduction

Models of financial markets often assume a simplistic mechanism for market clearing: the Walrasian scenario. This observation also applies to an innovative research area of heterogeneous agent models, wherein the heterogeneity of traders' expectations is a key to explaining the properties of markets. In reality, however, markets function in a different way. Agents are allowed to transmit only a finite amount of information in the form of orders to buy or sell. Furthermore, many markets employ continuous trade in the form of sequential orders. In this paper we study the impact of market organization on the dynamic properties of the asset pricing model populated by adaptive, boundedly rational agents with heterogeneous forecasting rules. We demonstrate that the adaptive abilities of the agents can be impaired by frictions inherent in the order-driven mechanisms. Surprisingly, this may stabilize the price dynamics. We also analyze how the market efficiency and statistical properties of prices are affected by the interplay of behavioral and institutional assumptions.

Statistical properties of real financial data have been thoroughly investigated in the past, see for example Fama (1970), Pagan (1996), Brock (1997), and Cont (2002). This line of research established a number of regularities in financial data, so-called "stylized facts," many of which are observed universally in all time periods and on different stock exchanges. Some of these regularities, for example the absence of significant autocorrelations in price returns, are well in agreement with the prevailing theory, called the Efficient Market Hypothesis, which suggests that markets are informationally efficient, i.e. new information is immediately reflected in asset prices. At the same time, such regularities as large and persistent trading volume, significant positive autocorrelations in variance of returns (volatility clustering), and heavier-than-normal tails of the return distribution are left unexplained within the classical paradigm. A seminal paper by Shiller (1981) detected that asset prices are more volatile than underlying fundamentals. The discovered excess volatility undermined the completeness of the Efficient Market Hypothesis.

Explaining these empirical properties by means of a simple model is an important but difficult task and there are several different directions, which deviate from the classical paradigm with a rational, representative agent (see, e.g., Lucas, 1978) leading to this goal. One way is to acknowledge that the assumption of full rationality is too demanding in the complex environment of financial markets. Models with heterogeneous agents using a bounded rational procedure as proposed in Sargent (1993) and Evans and Honkapohja (2001) may be more appropriate. A number of agent-based simulations of markets and rigorous analytical Heterogeneous Agent Models (HAMs) have been developed, which allow agents with different expectations to coexist in one market.[1] If one group of agents, called *fundamentalists*, believes that price typically reflects a fundamental value, and another group, *chartists*, extrapolates price trends, then the prices in a market can deviate from the fundamental value when chartists are in a majority. In Brock and Hommes (1998), this simple story is augmented by the evolutionary dynamics of relative fractions of fundamentalists and chartists. In this Adaptive Belief System agents not only update their forecasts as new data become available but also switch from one forecasting technique to another depending on techniques' past performances. Gaunersdorfer et al. (2008) show that even a simple version of such an adaptive model can generate dynamics with some realistic properties. Since the extrapolative expectations of chartists can be self-confirming, prices can deviate from the fundamental level and exhibit excess volatility. Furthermore, for certain parameter values the underlying deterministic system possesses two attractors, the fundamental steady state and a cycle around it, with small volatility for the former and high volatility for the latter. When dynamic noise is added to the system, price trajectory can interchangeably visit the basins of these two attractors producing volatility clustering. Gaunersdorfer and Hommes (2007) show that with a sufficiently large level of noise, this model indeed generates the dynamics that are qualitatively similar to a real market.

[1] The Santa Fe artificial market introduced in Arthur et al. (1997) and the model of microscopic simulations in Levy et al. (2000) are two known examples of a computational approach focused on bounded rationality in the formation of expectations. They are accompanied by parsimonious models in Day and Huang (1990), Lux (1995), Brock and Hommes (1998), Chiarella and He (2001), Farmer and Joshi (2002), Diks and van der Weide (2005), Anufriev et al. (2006), and Anufriev and Dindo (2007). See LeBaron (2006) and Hommes (2006) for recent reviews focused, respectively, on computational and analytical models with heterogeneous agents.

Alternatively, one can focus on the market design as a possible origin of stylized facts. Many classical models and all the HAMs quoted above use a Walrasian market clearing. It may be the case, however, that specific design features of the real markets bring some structure into the data. Simulations in Cohen et al. (1978) imply that the limit order book causes significant return autocorrelations. LiCalzi and Pellizzari (2003) show that an artificial market with realistic architecture, namely an order-driven market under electronic book protocol, is capable of generating satisfactory statistical properties of price series (e.g., leptokurtosis of the returns distribution) even with minimal behavioral assumptions. Furthermore, simulations in Bak et al. (1997) and Maslov (2000) suggest that desirable distributional properties can arise in the order-driven market even in the absence of any behavioral assumptions on the side of the agents.

These two streams of literature model either behavioral or structural features, but not both and, therefore, may provide only a partial explanation for the statistical regularities of financial markets. As opposed to those studies, recent agent-based models in Chiarella and Iori (2002), LeBaron and Yamamoto (2006), and Chiarella et al. (2007) incorporate the agents' heterogeneity in the order-driven markets. However, the interplay between behavioral and structural assumptions is far from trivial in these models, so that it often becomes difficult to understand how the two sets of assumptions contribute to the models' results. Consequently, our approach in this paper will be to start with a parsimonious model, which is analytically tractable under a Walrasian market clearing, and then increase the complexity by adding price- and order-driven trading protocols. The latter versions of the model are investigated through computer simulations.

Our research strategy is largely inspired by the work of Bottazzi et al. (2005). Motivated by empirical evidence from the world's stock exchanges that market micro-structure does influence statistical properties of returns, they compare dynamics under different trading protocols in the market populated by two types of traders: *chartists* and *noise traders*. The proportions of both types are *fixed*. Bottazzi et al. find that market architecture plays a larger role in shaping the time series properties than the behavioral aspects of the model. The authors also analyze the allocative efficiency of the market and show that, as opposed to the time series properties, the allocative efficiency depends mainly on the traders' behavior.

This paper follows a similar research strategy. However, in contrast to Bottazzi et al. (2005), our model is based on the Adaptive Belief System of Brock and Hommes (1998). In our model the market is populated by *fundamentalists* and *trend-followers* whose proportions are *evolving* on the basis of differences in past profits. A key behavioral parameter of the model is the intensity of choice, measuring the sensitivity of agents to this difference. If the market clears in the Walrasian way and the number of agents approach infinity, our model can be approximated by the deterministic model similar to the one analyzed in Gaunersdorfer et al. (2008). With our choice of forecasting rules, there exist two regimes in the market: tranquil and volatile. When the intensity of choice is low, i.e., smaller than a certain critical value, there is no excess volatility and prices remain tranquil on the fundamental level in the absence of dividend payments. When the intensity of choice is high, i.e., larger than this critical value, the volatile regime occurs with persistent deviations of prices from the fundamental level and there is excess volatility.

Our simulations reveal that similar regimes are also displayed under price- and order-driven trading protocols. Interestingly, the critical value of the intensity of choice is higher in the order-driven markets, implying a larger parameter's range of market tranquility. Given the noisy nature of the order-driven trading protocols, this result is surprising. We explain this result by the interplay of our behavioral assumptions and the market design. We also compare the properties of market dynamics over different market mechanisms, and show, in particular, that the order-driven trading protocols bring volatility clustering to the model dynamics.

The paper is organized as follows. In the next section we briefly describe different market mechanisms and introduce the behavioral part of our model. In Section 3 we analyze the model for a simple case of Walrasian market with a large number of agents and explain how two different market regimes arise. We then proceed by detailing our implementation of different market mechanisms in Section 4. Results of the simulations are discussed in Section 5 and Section 6 provides some final remarks.

2 The Model

We consider a standard asset-pricing model with two assets. The numéraire of the economy is the elastically supplied riskless asset which yields constant gross return $R = 1 + r$ per period. The risky asset pays a random dividend y_t at the beginning of period t. Realizations of dividend are independently drawn from a distribution with positive support and mean $\bar{y}$. The fundamental price of the risky asset is defined as a discounted value of the expected dividends and equal to $p^f = \bar{y}/r$. The risky asset is traded in the market, and its actual price dynamics are influenced both by evolution in the demand/supply of traders and by the precise mechanism for price determination. The following four trading protocols will be compared.

Under *Walrasian market-clearing* (WA), agents submit complete demand and supply schedules, and the price of the risky asset p_t in period t is defined as an intersection of a sum of the individual demand curves with a sum of the individual supply curves. The second protocol assumes the existence of a special agent, so-called *market-maker* (MM), who quotes price p_t before the start of trading session. Agents trade the quantities determined from their demand and supply schedules with the market-maker, who uses his own inventory to clear the market. Before the next trading session the market-maker adjusts the price on the basis of his accumulated inventory. In a market organized as a *batch auction* (BA), agents simultaneously post the buy or sell orders.[2] Cumulative demand and supply curves are then derived, and the price of the risky asset p_t is an intersection of these curves. In the fourth market type, the agents submit orders sequentially during the trading session, and the matching is accommodated by an electronic *order book* (OB) which stores unsatisfied orders. If a submitted order finds a matching order of the opposite type in the book, it is satisfied (completely or partially). An unsatisfied part of the order is stored in the book. In such a market there is no unique price during period t, and notation p_t is used to denote the closing price in this market, i.e., the price of the last transaction.

These four mechanisms are interesting because they range from the settings preferred in theoretical literature to the protocols used in real markets. Moreover, they differ on a number

[2] This mechanism is sometimes referred to as a *call auction* or a *sealed-bid double auction* in the literature.

of dimensions, such as information required from the traders and timing of order submission. The WA is a standard theoretical tool to model the market clearing process. However, it requires infinite information from the agents and therefore is hardly implementable in practice.[3] This problem is solved under the MM, a price-driven protocol, by announcing a single trading price in the beginning of the trading session. Under order-driven protocols, such as the BA and OB, the trading price is not known in advance, but the agents have to submit only a finite number of orders. The BA is used in a number of stock exchanges, typically to define a starting price of a trading session. Nowadays, however, most stock exchanges use the OB mechanism as it is more efficient for continuous trading.

For easy comparison, agents' behavior will be modeled in a similar way under the four institutional market settings. We populate the market by N myopic, expected utility maximizers whose demand functions depend on their expectations of next period price. The demands of agents are not homogeneous because agents can form their expectations according to one of two forecasting rules. Fundamentalists, i.e. the adopters of the fundamental rule, compute the fundamental value of the risky asset and expect that the price will move towards this value. Trend-followers simply extrapolate from past price changes. Relative fractions of fundamentalists and trend-followers affect, of course, the price determination in a given trading session. These fractions, in turn, change between trading sessions and depend on the relative past performances of the two groups using different rules. As a performance measure, we take an average return earned by fundamentalists (trend-followers) during the previous trading session.

In the rest of this section we explain how the demand of agents is defined and then introduce the evolutionary dynamics in the model.

2.1 Agents' demand

Agents are risk-averse, expected utility maximizers with common risk aversion coefficient a. Let $A_{i,t}$ and $B_{i,t}$ denote, respectively, the number of shares of the risky and the riskless asset

[3]The English clock auction with inter-period bids can provide a close approximation to the WA.

possessed by agent i at time t. In order to obtain the optimal portfolio composition, agent i maximizes at time t the conditional expectation of negative exponential utility of wealth in the next period, $W_{i,t+1}$. The wealth is uncertain both because the market price of the risky asset may change and because the random dividend is paid. Agent i solves the following problem

$$\max_{A_{i,t},B_{i,t}} \left\{ E_{i,t}[-\exp(-a\,W_{i,t+1})] \right\} \tag{1}$$

subject to

$$\begin{aligned} W_{i,t+1} &= A_{i,t}\big(p_{t+1} + y_{t+1}\big) + B_{i,t}(1+r)\,, \\ W_{i,t} &= A_{i,t}\,p + B_{i,t} \end{aligned} \tag{2}$$

The notation $E_{i,t}$ in (1) stresses the fact that the expectation is conditional on the information available at the beginning of time t and that the expectation is agent-specific. From the constraints (2), the evolution of wealth is derived

$$W_{i,t+1} = W_{i,t}(1+r) + A_{i,t}\big(p_{t+1} + y_{t+1} - (1+r)p\big)\,.$$

Assuming conditional normality of the wealth at time $t+1$, the optimization problem above is equivalent to the mean-variance optimization

$$\max_{A_{i,t}} \left\{ A_{i,t}E_{i,t}\big[\big(p_{t+1} + y_{t+1} - (1+r)p\big)\big] - \frac{a}{2}A^2_{i,t}\,V_{i,t}[p_{t+1} + y_{t+1}] \right\},$$

where $V_{i,t}[p_{t+1} + y_{t+1}]$ stands for the conditional expectations of trader i about the variance of price cum dividend at time $t+1$. From the first-order condition we derive *the desired position* of trader i in the risky asset as a function of price p:

$$A_{i,t}(p) = \frac{E_{i,t}[p_{t+1} + y_{t+1}] - (1+r)\,p}{a\,V_{i,t}[p_{t+1} + y_{t+1}]}\,. \tag{3}$$

The difference, $q_{i,t}(p) = A_{i,t}(p) - A_{i,t-1}$, between the desired position and holdings at the end of round $t-1$ constitutes agent's demand (supply) function at time t, when it is positive (negative). As in Brock and Hommes (1998), we assume that traders have homogeneous and time-invariant expectations about conditional variance $V_{i,t}[p_{t+1} + y_{t+1}] = \sigma^2$ and share correct expectations about dividend, $E_{i,t}[y_t] = \bar{y}$. It simplifies the model and allows us to concentrate on the heterogeneity of traders' expectations.

At any trading session, each trader chooses one of two possible forecasting rules, reflecting two trading attitudes commonly observed in the real markets. The fundamental forecasting rule

$$E_t^1[p_{t+1}] = p^f + v\,(p_{t-1} - p^f)\,, \qquad v \in [0,1]\,, \tag{4}$$

predicts that any price deviation from the fundamental level will be corrected. In one limiting case, $v = 0$, immediate correction is expected, while in another limiting case, $v = 1$, agents rely on the market, expecting that the last observed price is the best predictor. The trend-following forecasting rule

$$E_t^2[p_{t+1}] = p_{t-1} + g\,(p_{t-1} - p_{t-2})\,, \qquad g > 0\,, \tag{5}$$

predicts that past trends in the price will hold and uses coefficient g for extrapolation.

Notice that the fundamental forecasting rule (as opposed to the trend-following forecasting rule) requires a knowledge of fundamental value. Consequently, we assume that to use the fundamental rule the agent has to pay cost $C > 0$ per period, whereas the second rule is available for free.

2.2 Evolutionary updating of expectations

At the end of every trading round, agents update their forecasting strategy, i.e., choose which of the two rules, (4) or (5), will be used during the next session. The choice of forecasting rule for the next period is based upon the commonly observed deterministic component reflecting the past performances of two rules and stochastic error component reflecting the measurement error or imperfect computations of agents. The choice is modeled as follows.

At the end of trading round t, first, an individual *realized excess profit* is computed as a product of holdings of the risky asset at the end of round $t-1$ and its excess return, that is

$$A_{i,t-1}\left(p_t + y_t - (1+r)\,p_{t-1}\right). \tag{6}$$

Notice that in the case of continuous trading, the excess return is evaluated on the basis of closing prices. We stress also that under order-driven protocols (BA and OB), the realized position of an agent, $A_{i,t-1}$, can differ from the agents' desired position, $A_{i,t-1}(p_{t-1})$, due to possible rationing and/or difference between quoted and transacted prices. More details will be provided, when we discuss the market protocols.

Once individual profits have been computed, the performances of fundamental and trend-following forecasting rules U_t^1 and U_t^2 are defined as the *averages* of the individual realized excess profits over all fundamentalists and all trend-followers, respectively. From (6) it is clear that performance of the rule is the average holdings of the followers of this rule times the excess return. Thus, if the risky asset has earned a positive (negative) return, then the performance of the group with larger average possessions of the asset is higher (smaller).

Finally, agent i chooses the predictor for which the following maximum is realized

$$\max\left(U_t^1 - C + \xi_{i,t}; U_t^2 + \zeta_{i,t}\right), \tag{7}$$

where C is the cost of fundamental predictor, and $\xi_{i,t}$ and $\zeta_{i,t}$ are random variables independent over time and of agents. The choice can be rewritten in terms of probabilities for the special case of a Gumbel distribution of error terms. In this case, individual i chooses the predictors with *discrete choice* probabilities[4]

$$n_{t+1}^1 = \frac{\exp\left(\beta(U_t^1 - C)\right)}{\exp\left(\beta(U_t^1 - C)\right) + \exp\left(\beta U_t^2\right)}, \qquad n_{t+1}^2 = 1 - n_{t+1}^1, \tag{8}$$

with a subscript indicating that these probabilities shape the population which trades at period $t+1$. Parameter $\beta \geq 0$ is the *intensity of choice*, measuring the sensitivity of agents with respect to the difference in past performances of the two strategies. If the intensity of choice is infinite, the traders always switch to the historically most successful strategy. At the opposite extreme, $\beta = 0$, agents are equally distributed between different forecasting types independent of their past performances. The intensity of choice β is inversely related to the variance of the noise terms $\xi_{i,t}$ and $\zeta_{i,t}$.

[4]Our specification of the error terms is common in the literature on random utility models; see Anderson et al. (1992). Implied probabilities are used to model a choice in a number of theoretical models with a different range of applications, see, e.g., Brock (1993), Brock and Hommes (1997), Camerer and Ho (1999), and Weisbuch et al. (2000).

The timing of our model is as follows.[5] At the end of period t, the average profit earned by fundamentalists from their holdings between periods $t-1$ and t is computed and learned by all the traders. The average profit of trend-followers is learned, analogously. On the basis of these two performances, at the beginning of period $t+1$, agents independently choose their new forecasting types according to the probabilities defined in (8).[6]

At the same time, the desired positions for every forecasting type are computed on the basis of past prices by plugging the expectation rules (4) and (5) into equation (3). By subtracting the current holdings from the desired positions, demand/supply functions $q_{i,t+1}(p)$ for every agent are found, and trading session $t+1$ starts. Then, under the MM, BA and OB every agent is allowed to submit only one order per period, which is a point from the demand/supply curve. The definition of price p_{t+1} depends on the trading mechanism assumed. At the end of the trading session agents receive the profits earned on the holdings between period t and $t+1$ and update their portfolios.

3 Walrasian market clearing and large market limit

Our first mechanism, *Walrasian protocol* (WA), assumes that at every period the market is in temporary Walrasian equilibrium with demand equal to supply. At time t every agent submits the demand/supply function $q_{i,t}(p)$. The price of the risky asset is determined by the market clearing condition $\sum_i q_{i,t}(p) = 0$. Since the demand/supply function is strictly decreasing, there exists a unique equilibrium price, which we denote as p_t.

In this paper we concentrate on a special case of zero outside supply of shares of the risky asset.[7] The pricing equation becomes $\sum_i A_{i,t}(p_t) = 0$, which we solve in *deviations* from the

[5] One may think of one period as a trading day.

[6] In our simulations, we always assure that every forecasting type has at least one representative in any trading round. If an independent random draw does not produce a fundamentalist/trend-follower, we simply repeat the procedure until the population contains at least one fundamentalist/trend-follower.

[7] The model in Brock and Hommes (1998) is solved under the same assumption. As they show, this assumption can be made without loss of generality, since a positive supply case is equivalent to a re-definition of the dividends. Hommes et al. (2005) consider the model with positive supply.

fundamental price, $x_t = p_t - p^f$. Furthermore, we normalize the risk aversion coefficient so that $a\sigma^2 = 1$. Using (3), (4), and (5), the solution is given by

$$x_t = \frac{1}{(1+r)\,N} \sum_{i=1}^{N} \Big(v\, x_{t-1}\, \mathcal{I}^1_{i,t} + \big(x_{t-1} + g(x_{t-1} - x_{t-2}) \big)\, \mathcal{I}^2_{i,t} \Big), \tag{9}$$

where for $h \in \{1, 2\}$, index function $\mathcal{I}^h_{i,t}$ is equal to 1 if agent i forms at time t expectation $E^h_t[p_{t+1}]$; if not, it is equal to 0. Eq. (9) shows that the market price is a discounted average of individual expectations. Therefore, it depends on the relative number of fundamentalists and trend-followers. Under the WA clearing the agents' demands/supplies are always satisfied. Therefore, at any given period the positions of all the agents of the same forecasting type are the same. Their next period excess profits defined in (6) are also the same. The performance measures of fundamentalists and trend-followers are then given by

$$U^1_t = \big(v\, x_{t-2} - R\, x_{t-1} \big) \big(x_t - R\, x_{t-1} + \delta_t \big), \tag{10}$$

and

$$U^2_t = \big(x_{t-2} + g\,(x_{t-2} - x_{t-3}) - R\, x_{t-1} \big) \big(x_t - R\, x_{t-1} + \delta_t \big), \tag{11}$$

respectively, where $R = 1 + r$ is the gross return for the riskless asset. The random term $\delta_t = y_t - \bar{y}$ represents the shock to the excess return due to the stochastic dividend realization.

3.1 Large market limit

An important feature of our setting is that the dynamics under the WA can be studied by means of the dynamical system theory in a special case where the number of agents, N, becomes large. Indeed, for $N \to \infty$, the Law of Large Numbers guarantees a convergence of the actual fractions of fundamentalists and trend-followers to the probabilities defined in (8). We refer to this situation as the large market limit, shortly LML. The model is described then by one equation of the fourth order (or, equivalently by the 4-dimensional system) consisting of the market clearing equation coupled with an update of the fractions of fundamentalists

(which we will denote simply as n_t omitting the superscript)

$$\begin{cases} x_{t+1} = \dfrac{1}{R}\Big(vx_t\, n_{t+1} + \big(x_t + g(x_t - x_{t-1})\big)\,(1 - n_{t+1})\Big) + \varepsilon_{t+1} \\ n_{t+1} = \exp\Big(\beta\Big[\big(v\, x_{t-2} - R\, x_{t-1}\big)\,\big(x_t - R\, x_{t-1} + \delta_t\big) - C\Big]\Big)\Big/ Z_{t+1} \end{cases}, \tag{12}$$

where normalization factor $Z_{t+1} = \exp\big(\beta(U_t^1 - C)\big) + \exp\big(\beta U_t^2\big)$ is introduced with the performance measure of fundamentalists (10) and the performance measure of trend-followers (11). The dynamics in (12) are stochastic and there are two sources of noise. The first source, δ_t, is the stochastic component of the dividend realization entering the performance measure. The second source, ε_{t+1}, is added to the pricing equation to represent the dynamic noise. When both of these terms are zero, the corresponding system is deterministic and the following result takes place.

Proposition 3.1. *Consider system* (12) *with $\delta_t = 0$, i.e., when $y_t = \bar{y}$, and with $\varepsilon_t = 0$. This system has a unique steady state with $x^* = 0$ and $n^* = e^{-\beta C}/(1 + e^{-\beta C})$.*

(i) For $g \leq R$, this steady state is locally stable.

(ii) For $g > R$, this steady state is locally stable for $\beta < \beta^* = -\ln\big((g - R)/R\big)/C$. When $\beta = \beta^*$ the steady state exhibits the Neimark-Sacker bifurcation, and for $\beta > \beta^*$ it is locally unstable.

Proof. See Appendix A. □

The only steady state of system (12) is "fundamental," with price staying on the level p^f, implying deviation $x^* = 0$. In this situation both forecasting rules give correct predictions, but fundamentalists have to pay positive cost C. Consequently, they have a smaller relative share than the trend-followers: $n^* < 0.5$. With growing β, the equilibrium fraction of trend-followers increases. When trend-followers extrapolate weakly ($0 < g < 1+r$), the fundamental steady state is stable; when they extrapolate strongly ($g > 1 + r$), the fundamental steady state is stable for small β and unstable for high β. When β crosses its critical value, the stable quasi-cyclic attractor is created through the Neimark-Sacker bifurcation. Notice that

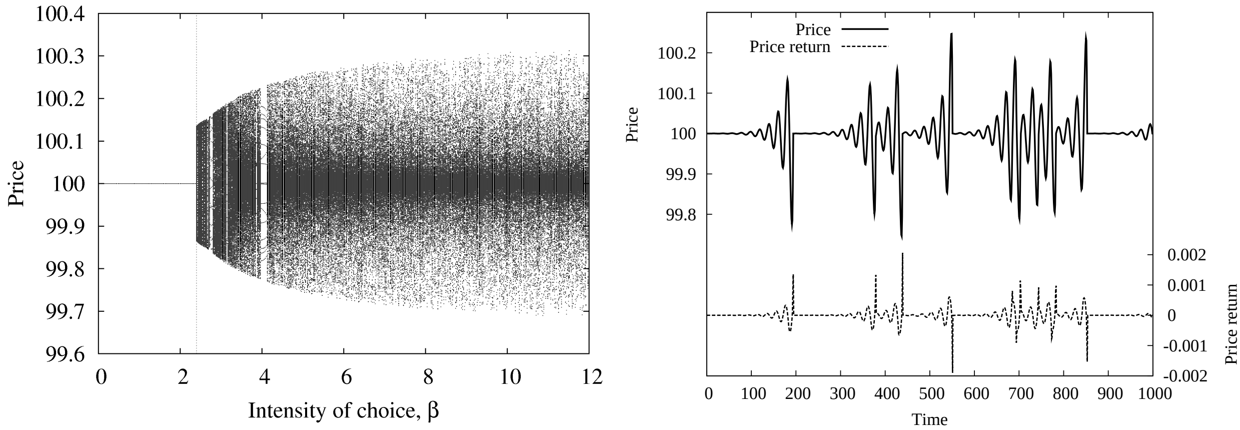

Figure 1: Bifurcation diagrams for the WA under the LML. For each $\beta \in (0, 12)$, 300 points after 10000 transitory periods are shown. **Left panel:** Neimark-Sacker bifurcation of fundamental steady state. Parameters are: $r = 0.1$, $\bar{y} = 10$, $v = 0.1$, $g = 1.2$, and $C = 1$. **Right panel:** Time series of the price and return with random dividend and small dynamic noise. Both noise terms are i.i.d. normal with standard deviation equal to 0.005.

the bifurcation value of β does not depend on the value of parameter v in the fundamentalists' forecasting rule.

Qualitatively, Proposition 3.1 implies that, for $g > 1 + r$, depending on the intensity of choice, market dynamics can be in one of two regimes: *tranquil*, with price staying at the fundamental level, and *volatile*, with large systematic deviations from the fundamental level. In the first regime there is no excess volatility in price and the trading volume is zero. In the second regime periods of overvaluation of the asset are followed by periods of its undervaluation, and the price exhibits bubbles and crashes with positive trading volume and excess volatility. The coexistence of two regimes and their dependence on a simple behavioral parameter, makes the case of strong extrapolation especially interesting.[8] Benchmark parameters $r = 0.1$ and $g = 1.2$ for our simulation are chosen to guarantee the coexistence of the two regimes, since we are interested in the dependence of "bifurcation scenario" on the market architecture. Other parameters are set to $\bar{y} = 10$, implying fundamental price $p^f = 100$, and $C = 1$.[9] Finally, we

[8]Similar regimes were identified in simulations of a Santa Fe artificial market model by Arthur et al. (1997), and in an analytical treatment by Brock and Hommes (1998).

[9]Note that these parameter values are not fully consistent with daily trading. Similar parameter values are

choose $v = 0.1$, implying that fundamentalists expect very small deviation from fundamental value. Again, the precise value of v is not important. However, as our simulations show, parameter v must be small enough in order for the price dynamics to be bounded.

The results of Proposition 3.1 are illustrated for these benchmark parameters in Fig. 1. The left panel shows a bifurcation diagram, where for each $\beta \in (0, 12)$ we simulate the long-run behavior of price. In accordance with our result for any $\beta < \beta^* \approx 2.398$, the price converges to the fundamental price p^f. When the intensity of choice increases to β^*, the fundamental equilibrium loses stability and a stable quasi-periodic cycle around p^f is created.[10] In the volatile regime, the price fluctuates around p^f. The bifurcation diagram shows that the amplitude of these fluctuations slightly increases with the intensity of choice.

Typical patterns of price and return dynamics for the volatile regime of the model ($\beta = 5$ in this example) are shown in the right panel of Fig. 1. In this regime, characterized by persistent deviations from the fundamental level, dynamics repeatedly go through qualitatively similar phases. At the beginning of each phase, price fluctuates only slightly around the fundamental value. With time, however, fluctuations get wilder, but at a certain point the price stabilizes and exhibits only small oscillations. These price fluctuations are endogenously determined by the proportion of the fundamentalists and trend-followers operating at the market.

We have demonstrated that in the special case of the Walrasian market with a large number of agents the model generates two different regimes, tranquil and volatile. The second regime is consistent with excess volatility. Will these two regimes be observed in a market with an alternative trading mechanism? How does the critical value of intensity of choice depend on the market trading rules? And how are the properties of price dynamics affected by the

used in Brock and Hommes (1998) and Gaunersdorfer et al. (2008). For easier comparison we stay close to their values. One could carefully calibrate the parameters, but it would not change qualitative behavior of the model.

[10] The model exhibits coexistence of attactors that is not captured in the Proposition 3.1. Numerically we find that the fundamental steady state is *globally* stable for $\beta < \beta^{**} \approx 2.23$. When $\beta \in [\beta^{**}, \beta^*]$, a locally stable fundamental steady state coexists with a quasi-periodic attractor. This coexistence seems to be a consequence of a so-called Chenciner bifurcation, thoroughly discussed in Gaunersdorfer et al. (2008) in a similar model.

mechanisms? These are the questions which we analyze below.

4 Market Mechanisms

A market mechanism is a well-defined procedure which transforms input from agents into an output of price and quantity traded. There are numerous market mechanisms in the theoretical literature and in the trading practice, among which we select four stylized procedures. The behavioral model is simulated separately for every mechanism and the results are compared.

4.1 Walrasian Auction

The implementation of Walrasian auction (WA) was described in the beginning of the previous section. The dynamics of the model is described by the pricing equation (9), while the forecasting type of a trader is determined by probabilities (8). Recall that the demands/supplies of agents are always satisfied under the WA, implying that the performance measures are given by (10) and (11). The only difference between the LML analyzed in Section 3.1 and the WA simulations reported below is the number of agents, which is infinite in the former case and finite in the latter.

A simple but informative illustration of the WA can be found in the left panel of Fig. 2. In this example there are five fundamentalists and five trend-followers arriving at the market at time t with zero shares of the risky asset (i.e., the demand and desired position coincide for these agents). Let us assume that $p_{t-2} < p_{t-1} = p^f$. Since fundamentalists forecast price p^f for the next period, they demand the risky asset when $p_t < p^f$, and supply it otherwise. The long-dashed lines show individual demand (supply) schedules of fundamentalists. Trend-followers expect that prices will rise to the level of $p^c > p^f$, so that they demand the risky asset when $p_t < p^c$ and supply it when $p_t > p^c$. The short-dashed lines show individual demand (supply) schedules for trend-followers. All individual demand and supply have the same slope, R, in absolute value. Horizontal summation of five corresponding curves gives the aggregate demand and aggregate supply both shown by thick solid lines. When price is inside the interval

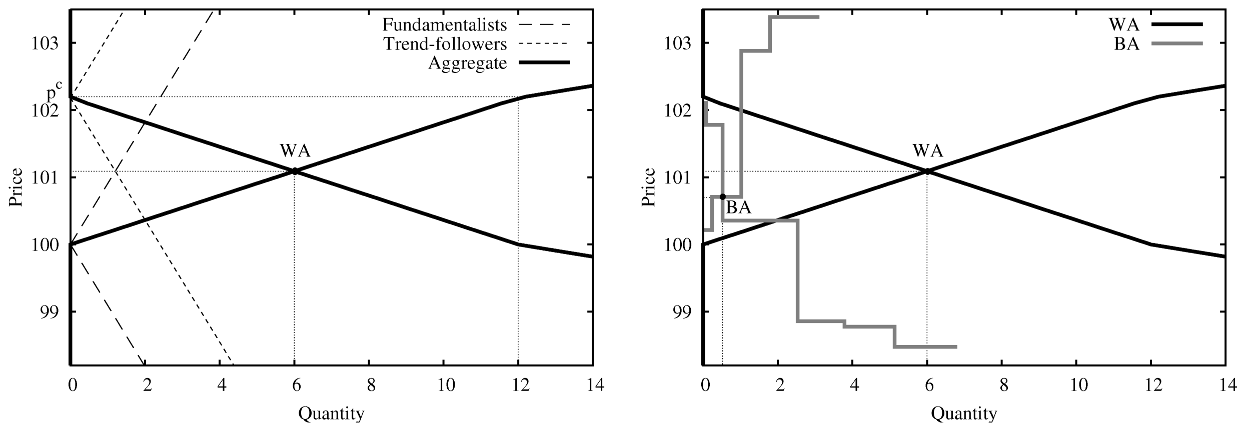

Figure 2: Comparison of different market-clearing mechanisms. **Left panel:** Walrasian price and quantity are found as intersections of the aggregate demand and supply schedules, built starting from the individual curves. **Right panel:** Batch auction compared with Walrasian auction.

$[p^f, p^c]$ all the demand is generated by the trend-followers and all the supply is generated by the fundamentalists. When price is below p^f, all ten agents want to buy. Thus, the aggregate supply is zero, while the aggregate demand curve has a kink at price p^f. Analogously, the aggregate demand is zero above p^c and the aggregate supply curve has a kink at price p^c. The aggregate demand and the aggregate supply curves intersect at the point labeled WA, whose ordinate is the equilibrium price and abscissa is the equilibrium quantity under a Walrasian market-clearing. All the agents trade their quantities at the equilibrium price. Notice that in this example, the equilibrium price is half way between p^f and p^c, which are "no-trade" prices for fundamentalists and trend-followers respectively. Another distribution of agents between the two forecasting types would lead to a different outcome.

4.2 Market maker

Price p_t under a market maker (MM) clearing is set by a special agent, a market maker, before trade starts. His goal is to provide liquidity. Observing p_t, other traders compute their desired positions $A_{i,t}(p_t)$ as in (3). Notice that since the current price can now be used in forecasting, predictors (4) and (5) are changed in the MM case to $E_t^1[p_{t+1}] = p^f + v\,(p_t - p^f)$ and $E_t^2[p_{t+1}] = p_t + g\,(p_t - p_{t-1})$, respectively. The difference between the desired positions and previous holdings gives agents' demand/supply $q_{i,t}(p_t)$. Then the fundamentalists and trend-followers trade with the market maker at price p_t their quantities $q_{i,t}(p_t)$, buying from him if these quantities are positive and selling to him if they are negative. As under the WA, the demands/supplies of all N traders are satisfied after the trade. However, to clear the market, the market maker keeps his own inventories. Under the assumption of zero outside supply of shares, the inventories of the market maker after the trading session t are opposite to the aggregate possessions of all other traders and equal to $-\sum_i A_{i,t}(p_t)$.

Before the next session the market maker changes the quoted price. If he is short in the risky asset, he adjusts the price up in order to induce more traders to sell some of their holdings. Otherwise, the specialist adjusts the price down, inducing traders to buy some of his inventories. We assume the following linear adjustment rule

$$p_{t+1} = p_t + \mu \sum_{i=1}^{N} A_{i,t}(p_t)\,, \tag{13}$$

where positive coefficient μ measures the speed of adjustment.[11] As for the other traders, their behavior is the same as under the WA. Namely, they compute the performance measures and define their predictor probabilistically according to (8). When the number of agents N is infinitely large the probabilities may be substituted by the actual fractions and the model becomes analytically tractable (see Appendix B). It turns out that choosing $\mu = 1/R$ leads to the same local stability conditions as under the WA.

4.3 Batch auction

Under the batch auction (BA), agent i submits at time t one *limit order*, which is a price/quantity combination $(p_{i,t}, q_{i,t})$. When the ordered quantity is positive (negative), the order is of a buy (sell) type. The price in the limit order defines the largest (smallest) price accepted by the submitter for execution of the buy (sell) order.

To make a comparison between the BA and the WA meaningful, we require agents to submit those price/quantity combinations which belong to their demand/supply schedules

[11] Recently Hommes et al. (2005) has studied an analytical model with the same stylized price adjustment rule.

$q_{i,t}(p_{i,t})$. We use simple strategic considerations to determine the price in the order generation process. Namely, we assume that the price of the limit order $p_{i,t}$ of agent i is determined as a random draw from a normal distribution with mean p_{t-1}, the price of the previous trading session, and standard deviation σ_o. The realizations are independent over time and of agents. Under this price selection rule, an agent reasonably believes that there is a high chance that his order will be executed at a price which is close to the last closing price p_{t-1}. The larger is the deviation from this price, the higher may be the potential gains from the trade, but the lower is the likelihood of such an order being executed. Therefore, only on rare occasions will an agent experiment with an order priced considerably far from the previous closing price p_{t-1}.[12]

All N orders are submitted simultaneously. Then, all the limit buy orders are sorted such that their price sequence is decreasing. It gives us a step-level, market demand curve. Analogously, the limit sell orders are sorted so that their price is increasing to define a step-level, market supply curve. The price p_t is determined as an intersection of constructed demand and supply curves (or the average of the lowest and the highest clearing prices, if the intersection is an interval rather than a point). In those cases when demand and supply curves do not intersect, price p_t is set to the price of the previous period, p_{t-1}. The corresponding quantities are traded on this price between those agents who submitted bids (asks) no lower (no higher) than p_t. Traders who submitted orders exactly at p_t may be rationed accordingly, while no other traders make any trades keeping their previous portfolios.

In the right panel of Fig. 2 we use the previous example to construct the schedules for the BA. For every agent one price was generated randomly around price $p_{t-1} = p^f$ and then the corresponding individual demanded or supplied quantity was found. Sum of the quantities gave two step curves whose intersection, labeled "BA", determines equilibrium price and quantity. Of course, the precise schedules depend on the particular random draw, but certain general tendencies can be seen from this example. Obviously, the quantity traded under the BA is

[12] In their empirical study of the order-driven market with an electronic book on the London Stock Exchange, Mike and Farmer (2008) find that the shape of the actual distribution of prices of submitted orders can be well approximated by the Student's t-distribution around the best price.

always smaller than the quantity traded under the WA, while the equilibrium price under the BA can be both above or below the price under the WA.

In computation of time-t performance measures, notice that a number of traders did not trade during time $t-1$, so that their positions were left unchanged. Thus, their performances are evaluated as $A_{i,t-2}\big(p_t+y_t-(1+r)p_{t-1}\big)$, which is an earned excess profit on the old holdings. However, those traders who did trade (and were not rationed) changed their positions on the basis of their submitted quantities. Their performance is then equal to $A_{i,t-1}(p_{i,t-1})\big(p_t + y_t - (1+r)p_{t-1}\big)$, and it does not necessarily coincide with the performance under the WA since the transaction price, p_{t-1}, almost always differs from the submitted price of bid, $p_{i,t-1}$. Thus, the BA mechanism, by its nature, modifies the individual agents' performances.

4.4 Order book

In the order-book market, there are many transactions during one trading session at time t. Each agent can place only one buy or one sell order during the session. To make a reasonable comparison with the BA, the order generation process is identical to the one described in Section 4.3, while the sequence in which agents place their orders is determined randomly and varies between different sessions.

During the session the market operates according to the following mechanism. There is an electronic book containing unsatisfied agents' buy and sell orders placed during the current trading session. When a new buy or sell order arrives into the market, it is checked against the counter-side of the book. The order is partially or completely executed if it finds a *match*, i.e., a counter-side order at requested or better price, starting from the best available price. An unsatisfied order or its part is placed in the book. At the end of the session all unsatisfied orders are removed from the book. The mechanisms for determining the type of the order, its price and its quantity are equivalent to those described for the batch auction. Price p_t is the closing price of the session, i.e., the price of the last transaction. Agents' realized excess profit is computed as in (6), but similarly to the BA mechanism at the end of $t-2$ the agents' actual holdings $A_{i,t-2}$ may be different from the agents' desired position $A_{i,t-2}(p_{t-2})$. Contrary to

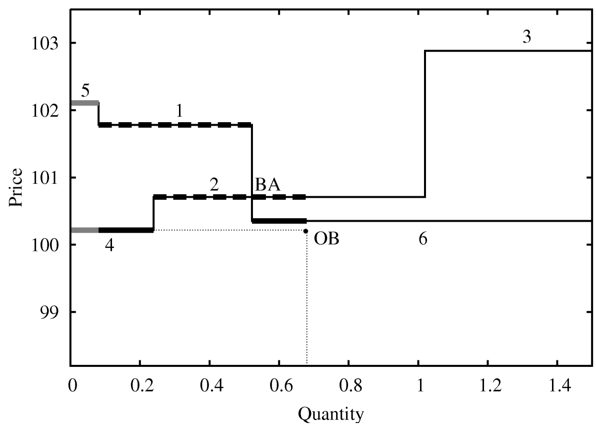

Figure 3: Order-book market compared with BA.

the BA case, however, the order in which agents arrive at the market may influence whether an agent is rationed or not.

In Fig. 3 we show a possible order book realization for the same limit orders as were generated for the BA illustration in Fig. 2. The integer labels show the sequence in which the buy/sell orders arrive, while different types of horizontal lines show which part of the order is satisfied. When the first order to buy (labeled by 1) arrives, it is automatically saved in the book. Order 2 to sell partially matches order 1 and after the corresponding transaction one sell order (part of order 2) remains in the book. Then sell order 3 arrives, which is written in the book without affecting the best ask price. The next sell order 4 is also stored in the book, since the buy side of the book is empty. However, order 4 has price priority over orders 2 and 3. Thus, when buy order 5 arrives it is executed at the price of order 4. The remaining part of the order 4 still occupies the best position on the sell part of the book. Finally, buy order 6 arrives, which can be partially matched with order 4. Point OB shows the outcome of the trade with the OB mechanism. The price p_t is the price of order 4, while the volume for the session is the total traded number of shares. Notice that the flexibility of this trading mechanism allows a larger traded volume than under the BA. At the same time, the OB price p_t is typically different from the BA price and may significantly depend on the order in which transactions happen.

Parameter	Symbol	Value (Range)
Intensity of choice	β	$[0, 12]$
Interest rate	r	0.1
Mean dividend	$\bar{y}$	10
Normalized risk-aversion	$a\sigma^2$	1.0
Trend-followers' extrapolation	g	1.2
Fundamentalists' reversion	v	0.1
Fundamentalists' costs	C	1.0
MM price adjustment coefficient	μ	0.9
Stand. deviation of limit order price	σ_o	3.0
Number of agents	N	1000
Transient period	Tr	2000

Table 1: Parameter values used in simulations.

5 Simulations and Results

We simulate the system with a finite number of agents under different market mechanisms, namely the WA, MM, BA, and OB mechanism.[13] In all cases we keep the dividend constant and do not add any dynamic noise. While under these two assumptions in the LML with a Walrasian market clearing the system is deterministic, in the simulations with a finite number of agents and under different trading mechanisms the amount of randomness increases from one mechanism to another. Under the WA and MM the system becomes stochastic since the realized fractions of fundamentalists and trend-followers are no longer equal to their analytic probabilities. In the BA the amount of stochasticity is higher due to the fact that agents choose a random point on their individual demand/supply schedules. In the OB mechanism the level of stochasticity is further increased by random sequencing of order submissions.

The parameters that we use for the simulations are summarized in Table 1. The model

[13] The software for the simulation is written in C++ and is a modification of the YAFiMM package created for Bottazzi et al. (2005). The YAFiMM package is publicly available at `http://www.sssup.it/~bottazzi`.

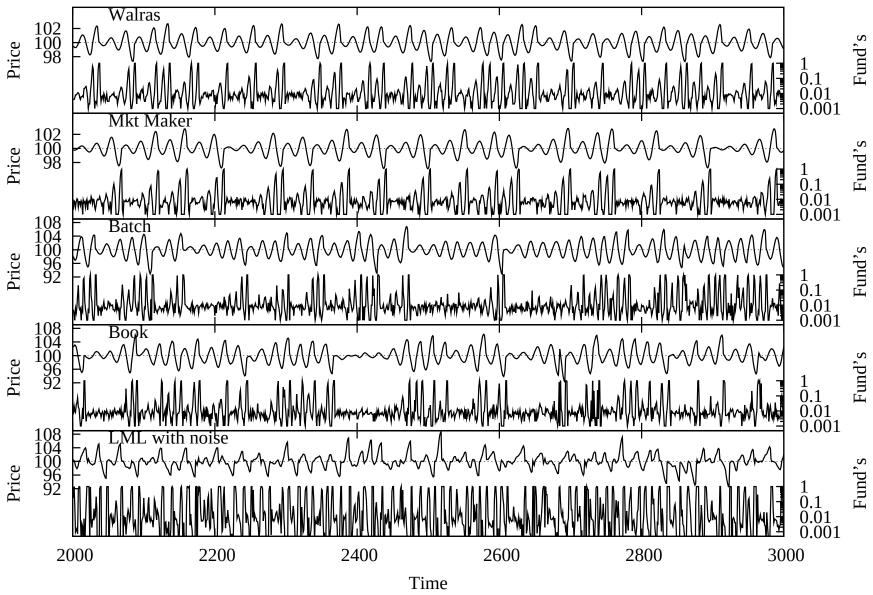

Figure 4: Time series of prices (left axis) and share of fundamentalists (right axis in the log-scale) for $\beta = 5$ in WA, MM, BA, OB, and LML with model approximation error. In the latter case the error is independently drawn from normal distribution with zero mean and standard deviation $\sigma_\varepsilon = 0.3$.

parameters are the same as those used in the analysis of the LML, when non-trivial and non-divergent dynamics are generated. The behavioral parameter we mainly focus on is the intensity of choice β. In all our simulations we set the number of agents to $N = 1000$. This number is high enough to obtain dynamics close to the deterministic LML under the WA. The transient period is set to 2000 to avoid any transitory effects.

Before starting a detailed analysis of the stochastic system let us have a look at the time series from the model simulated under different market mechanisms. In Fig. 4 we show time series of the price (upper part, left axis) and the share of fundamentalists (lower part, right axis) for $\beta = 5$. Recall from Fig. 1 that for this level of the intensity of choice, the system in the LML does not converge to the fundamental steady state, but oscillates around it. The top panel shows that under the WA the system behaves similarly to the deterministic LML with some distortions caused by the finite number of agents. Under the MM, BA and OB (the second, third and fourth panel from the top, respectively) the dynamics are similar but certain differences can be observed even by the naked eye. First, under the BA and OB the price deviations from the fundamental level have larger amplitude than under the WA. Second, looking at the dynamics of the fraction of the fundamentalists in all four markets, one can distinguish the two alternating phases in the dynamics, but these phases are much more visible under the MM, BA and OB protocols than under the Walrasian market. These phases are (1) stable ecology, when the fractions of fundamentalists/trend-followers in the market exhibit only moderate changes and (2) turbulent ecology, when in every period a large fraction of agents switch from one forecasting type to another. Third, under the BA and OB these populational dynamics translate into price dynamics and the observed phases of stable and turbulent ecology correspond to the phases of small and large price oscillations, respectively. These phases can be linked to the phenomenon of volatility clustering which is observed in real financial data. Under the MM, the changes between small and large price fluctuations are not so abrupt.

Finally, we want to verify whether similar price behavior can be obtained by adding a dynamic noise to the LML.[14] In the bottom panel we add the Gaussian noise to the LML dynamics given by the system (12). The standard deviation 0.3 is chosen to match the amplitude of the price fluctuations under the order-driven mechanisms. Contrary to the last two mechanisms, we no longer observe clear-cut phases in the time series. Therefore, the randomness inherent in the order-driven market mechanisms distorts the price evolution differently from simple dynamic noise.

5.1 Change of market regimes

Next, we investigate in detail the interplay of agents' behavior and market mechanisms on price dynamics.

Fig. 5 depicts the dependence of the stable distribution of prices on the intensity of choice parameter β. The parameter β ranges from 0 to 12 with a linear step of 0.05, which gives 240 points in total. The distribution for each level of beta is represented by a gray-shade

[14]For instance, Gaunersdorfer and Hommes (2007) try to reproduce the "stylized facts" by adding the dynamic noise to a similar deterministic model.

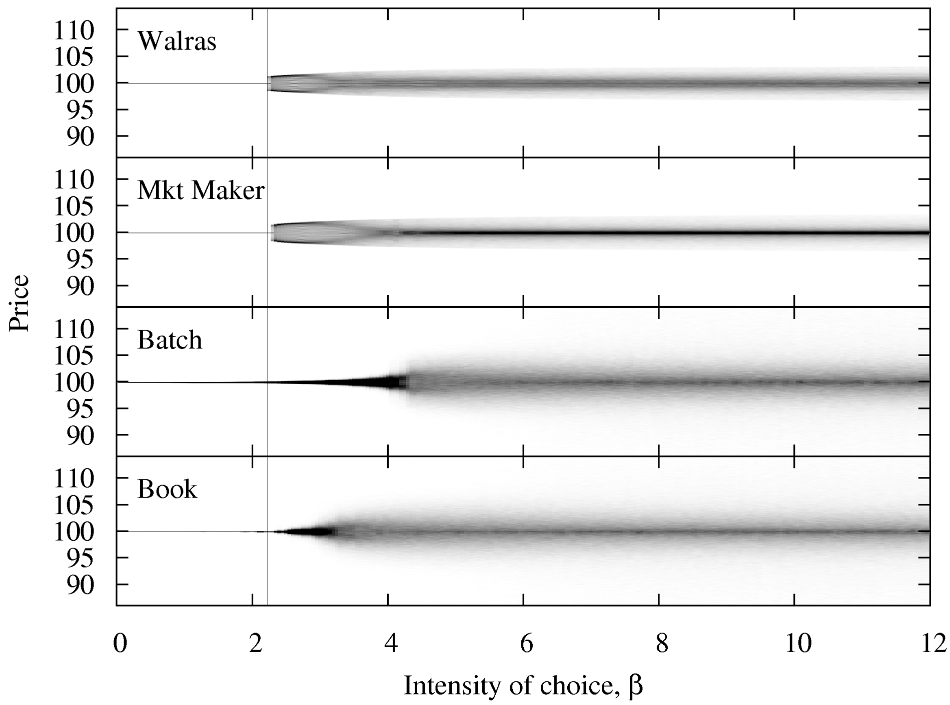

Figure 5: Stable distribution of prices for different market mechanisms as a function of the intensity of choice.

coded histogram. Darker shades correspond to areas of higher density. The histogram is computed using price levels from 10000 periods after 2000 transient periods. A bifurcation in the stochastic system corresponds to the qualitative changes in the stable distribution which we attempt to identify graphically. Under the WA and MM, around the point $\beta = 2.23$ we observe a dramatic increase in the variance of distribution and an emergence of bimodal distribution. Around this point of bifurcation, the system transits from the tranquil regime to the volatile regime discussed in details in Section 3.1. The bifurcation is delayed for the BA and occurs around $\beta = 4$. Under the OB mechanism, the bifurcation value of β is higher then under the WA and lower then under the BA, and is around the point $\beta = 3$. For additional evidence about the point of bifurcation under four different mechanisms, see Fig. 6 (left panel) which shows standard deviation of the price.

The delay in bifurcation observed for order-driven mechanisms can be explained by the interplay between agents' behavior and the characteristics of the mechanism. As we pointed out in Section 2.2, the intensity of choice parameter β is inversely proportional to the level of noise in the performance of a forecasting rule (see Eq. 7). Thus, the higher level of noise in the performance measure would correspond to the lower level of β. Recall also that the performance measure of a forecasting rule is the average individual performance of all agents using the same forecasting rule, and that individual performances are computed on the basis of the agents' holdings of the risky asset and the excess return of this asset as in (6). When describing trading mechanisms in Section 4, we emphasized that the agents' demands/supplies are fully satisfied only under the WA and MM. Under the BA and OB mechanisms, the agents' demands/supplies are translated into orders, some of which may be rationed, while others may be executed at a price different from the one specified in the order. Thus, the agents' holdings of the risky asset may be inconsistent with their chosen predictors, resulting in distorted performance measures. This in turn results in a higher level of noise in the average performance measure of the strategy, lowering the "effective" value of the intensity of choice parameter β. The larger the amount of rationed orders, the higher the level of noise. The fact that the amount of rationed orders is larger under the BA than under the OB (see Fig. 3) suggests that the actual "bifurcation" level of β is higher under the BA mechanism than under the OB mechanism, which is exactly what we observe in Fig. 5. Additional evidence of the amount of rationed orders under the order-driven mechanisms is provided in Fig. 6 (right panel) where we report an average traded volume of the risky asset.

Thus, in the Adaptive Belief Scheme where heterogeneous agents choose their forecasting rules on the basis of past performances the order rationing inefficiencies introduced by the order-driven mechanisms lead to market stabilization in the sense of a wider interval of the intensity of choice parameter β for which the market is in the tranquil regime. It is, however, important to keep in mind that in the volatile regime, when the intensity of choice is high enough, the market fluctuations are larger in amplitude under the order-driven markets (see Figs. 4 and 5).

5.2 Informational efficiency

The Efficient Market Hypothesis postulates that in the efficient market the price should reflect all available information about the asset value. In our setting with random i.i.d. dividend, the

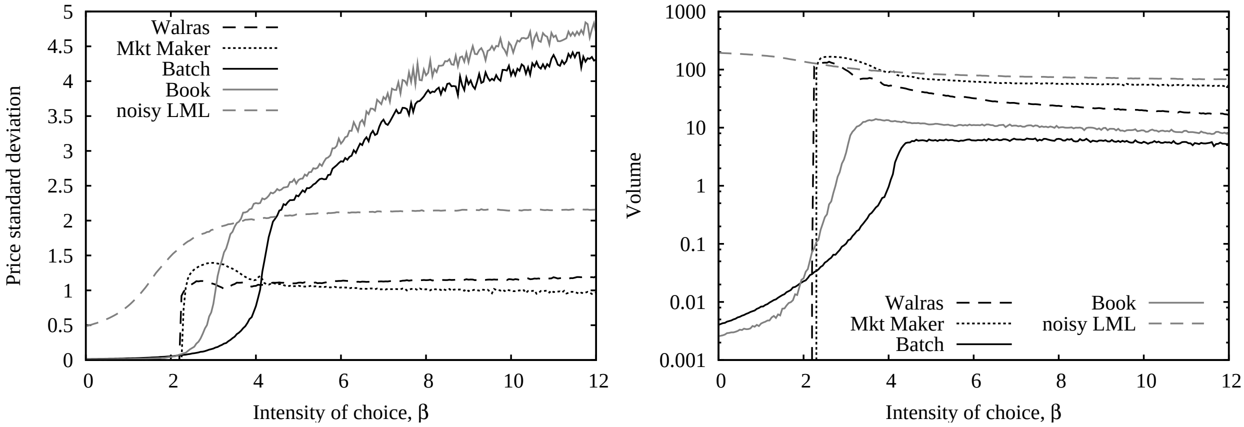

Figure 6: Two measures of market information inefficiency for different market mechanisms as a function of the intensity of choice β. **Left panel:** Standard deviation of price. **Right panel:** Traded volume.

fundamental value of the asset is simply the discounted sum of all future expected dividends, i.e., fundamental price p^f. The informational efficiency is often measured by comparing the volatility of the observed price with the volatility of the fundamental dividend process (see Shiller, 1981). In order to abstract from an effect of time-varying dividend in our model, we keep the dividend process constant. Under this assumption, the Efficient Market Hypothesis would predict constant price over time and zero trading volume. Therefore, price volatility and trading volume can be used as measures of information efficiency.

In Fig. 6 we compare the standard deviation of the price (left panel) and the average traded volume (right panel) under different values of β for four market mechanisms and the LML with the Gaussian noise ($\sigma_\varepsilon = 0.3$). We ignore the first 2000 transitory steps and compute the standard deviation of the price and the average traded volume over the next 1000 periods. To eliminate the dependence of our results on a particular realization of random seed, we repeat this process for 100 random seeds and report an average of the statistics of interest. Averaging over different random seeds accounts for possible dependence of our results on initial conditions. We also compute 95% confidence bounds for the reported averages. Given the length of the series and the large number of random seeds, the confidence bounds are very tight. For clarity, we do not plot them on the figures but they can be easily inferred from statistical variations for neighboring values of β.

The standard deviation of the price (Fig. 6, left panel) depends on both the intensity of choice parameter β and the market mechanism. For $\beta < 2$, all the mechanisms without added dynamic noise have the standard deviation close to zero. The standard deviation rises rapidly at the point of bifurcation and quickly converges to the level of around 1.0 for the WA and MM, while for the BA and the OB, it continues to grow with β and shows some signs of stabilization at 4.0 and 4.5 respectively when $\beta > 11$. The standard deviation for the OB is always higher then for the BA because of earlier bifurcation and an extra layer of stochasticity (order sequencing) specific to the OB mechanism. The dynamic noise added to the LML is magnified from the initial level of $\sigma_\varepsilon = 0.3$ to the level increasing from 0.5 to 2.0 and stabilizing at 2.0 when $\beta > 4$. The observed volatility pattern for the LML with dynamic noise is different from the pattern produced by the order-driven mechanisms, which confirms that the time-series produced under the latter could not be produced by adding dynamic noise to the analytic LML. Based on the volatility measure, we conclude that information efficiency depends on both behavioral and institutional assumptions. For $\beta > 4$, the WA and MM provided the most informationally efficient outcomes followed by the BA and the OB in that order. However, for smaller β the order-driven mechanisms can be superior even to the WA due to the delayed bifurcation.

The average traded volume (Fig. 6, right panel) also depends on the intensity of choice. For $\beta < 2$, when the price is very close to the fundamental, the average traded volume is almost 0 for all mechanisms except the LML with dynamic noise. The dynamic noise added to the LML creates a very high level of volume which slowly levels off with β increasing, which as before, is in sharp contrast with the patterns produced by the order-driven mechanisms. For $\beta > 2.3$, the traded volume is always higher for the MM, which is followed by the WA, OB and BA. The latter two show lower volume because of the order rationing which is higher for the BA than for the OB (see Fig. 3). Interestingly, the average traded volume decreases in β, when $\beta > 5$. For large values of β, the fraction of one forecasting type is much larger than the fraction of the other type, which leads to the lower volume.

In this section we showed that in the volatile regime, i.e., when the intensity of choice

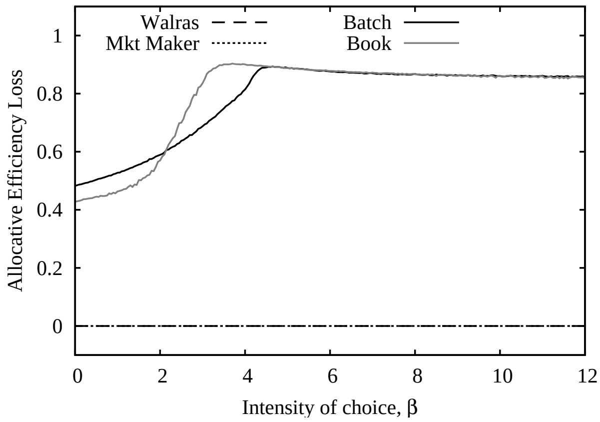

Figure 7: Allocative efficiency loss for different market mechanisms as a function of the intensity of choice β.

is large enough, the order-driven markets are less informationally efficient from the point of view of price volatility. A smaller trading volume in these markets does not suggest higher informational efficiency, but indicates lower liquidity. The latter is connected to allocative efficiency, which we investigate next.

5.3 Allocative efficiency

The main purpose of any trading mechanism is an efficient allocation of resources, that is an allocation which fully satisfies agents' demands/supplies at a realized price. By way of its nature, the WA and MM always achieve an efficient allocation. However, for more realistic trading mechanisms such as the BA and the OB, an efficient allocation is not necessarily guaranteed. Following Bottazzi et al. (2005) we define a measure of allocative efficiency loss, $L_{i,t}$:

$$L_{i,t} = 1 - \frac{1}{1 + |A_{i,t}(p_t) - A_{i,t}|\, p_t}, \tag{14}$$

where $A_{i,t}(p_t)$ is a desired position in risky asset of agent i at closing price p_t of period t, and $A_{i,t}$ is the realized holding of the risky asset of agent i at the end of period t. Note that the measure is always between 0 and 1. Obviously under the WA and MM the demands of the agents are fully satisfied and $L_{i,t} = 0$.

Fig. 7 shows the measure of allocative efficiency loss averaged across 1000 agents as a function of the intensity of choice β for the BA and the OB mechanisms. As before, we take an average over 1000 time periods after the 2000 transitory periods which is averaged again over 100 random seeds. Consistent with our previous result, we observe that the allocative efficiency loss depends on β. Before the bifurcation point the inefficiency is lower since we are close to the fundamental price and agents' desired positions in risky assets is relatively small. After the bifurcation the price amplitude increases, which translates into the larger desired positions, and the allocative efficiency loss increases. For $\beta > 5$, the allocative efficiency loss stabilizes close to 0.85 for both order-driven mechanisms and then slowly levels off. This small increase in efficiency is again explained by the higher concentration of price distribution around the fundamental price for larger values of the intensity of choice.

It is remarkable that the allocative efficiency of the two order-driven mechanisms is exactly the same for a given $\beta > 5$. While the effect of the order rationing is more pronounced under the BA, the OB produces higher price deviations. Apparently, both effects are of the same magnitude in terms of an influence on the allocative efficiency loss. Similar to Bottazzi et al. (2005), in our model the precise implementation of the clearing system on the order-driven market does not affect the allocative efficiency.

5.4 Time-series properties

We compare the times series of the price returns generated under different market mechanisms through the prism of "stylized facts" established in the literature which was shortly discussed in the Introduction. The returns are defined as $r_t = (p_t - p_{t-1})/p_{t-1}$, i.e., as relative price changes. All the statistics below were computed over 1000 periods after 2000 transient periods, and averaged over 100 random seeds.

The averages of the returns over time are close to zero for all four mechanisms and for all considered values of β. Similarly, the skewness of return, which measures the asymmetry of distribution, is close to zero for all mechanisms and all β. Both these statistics are in agreement with real data. In discussing other statistics, notice that in the tranquil regime under the WA and MM, the price converges to the constant fundamental level, and thus the

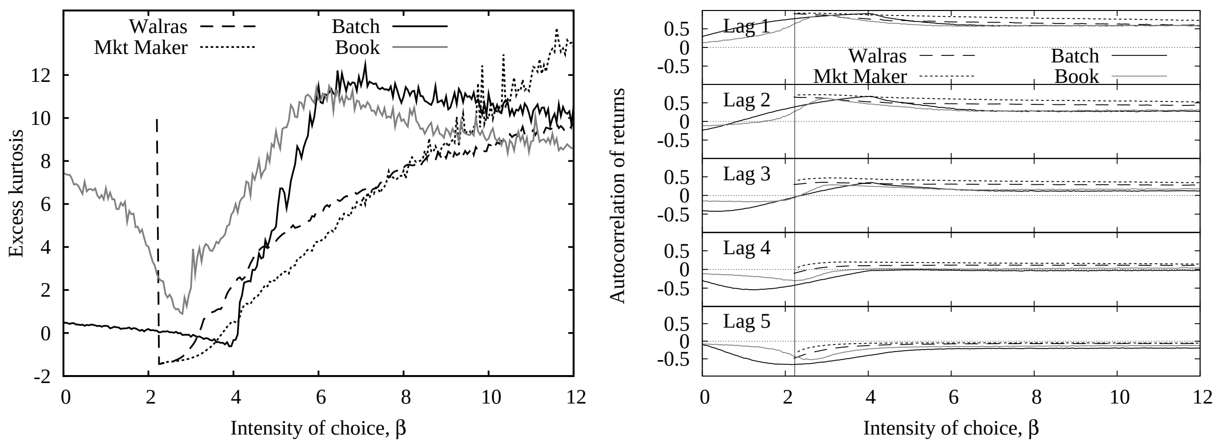

Figure 8: Time-series properties as a function of the intensity of choice β. **Left panel:** Excess kurtosis of the return distribution. **Right panel:** Autocorrelation of returns.

higher order statistics are not defined.

Comparison of the empirical return distributions with the normal distribution reveals that in the real markets, returns exhibit a higher concentration around the mean and also the fatter tails. These properties of distribution can be measured by *excess kurtosis* with respect to the kurtosis of the normal distribution, which is equal to 3. We find that the excess kurtosis of returns (see the left panel of Fig. 8) depends both on the value of the intensity of choice and on the market clearing mechanism. As β increases and reaches the critical value of the market regime change, the kurtosis first drops sharply, then grows monotonically and finally converges to a relatively stable level. Close to the point of the regime change, the kurtosis is the highest under the OB and the lowest under the WA and MM. For higher values of the intensity of choice, under all four mechanisms the kurtosis converges to a value close to the one observed for the S&P 500 Index, which is 8.5 according to Gaunersdorfer and Hommes (2007).[15] The dependence of kurtosis on behavioral parameters is in sharp contrast to the conclusions of Bottazzi et al. (2005) who find that skewness and kurtosis values depend only on the market mechanism.

Linear unpredictability of the stock returns is another well-established regularity. It is

[15]Note that estimates of kurtosis depend on the particular financial series and the time period used for estimation.

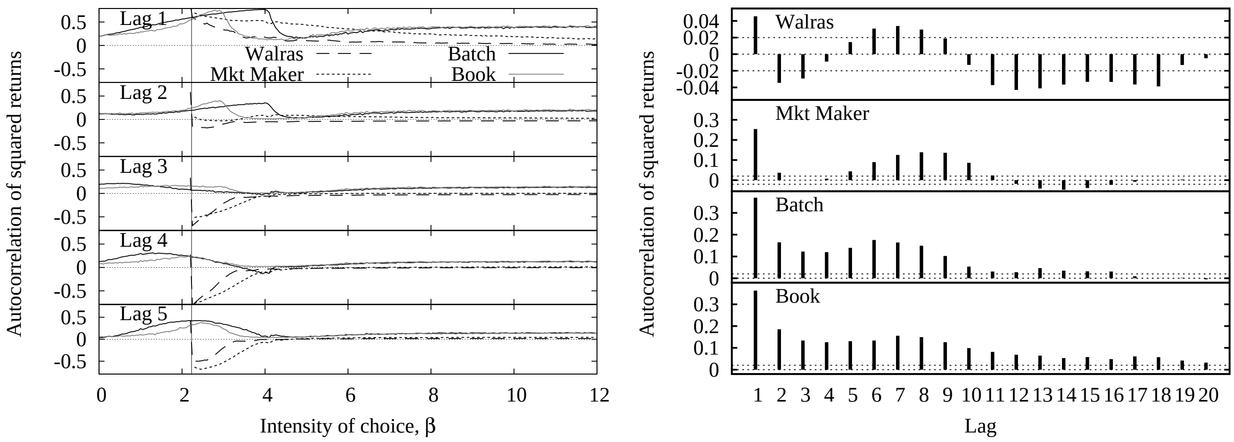

Figure 9: Autocorrelations of squared returns. **Left panel:** Autocorrelations at lags $1-5$ (top to bottom) as a function of the intensity of choice β. **Right panel:** Autocorrelation function at lags $1-20$ for fixed $\beta = 8.0$.

usually verified by computing the autocorrelations of returns. For real financial data, the autocorrelations of returns are insignificant already at the first lag. In the right panel of Fig. 8, we show the autocorrelations of returns for the first five lags as a function of the intensity of choice for all four mechanisms. In all cases we observe relatively large autocorrelations for the lags $1-3$ which is in contrast to the stylized facts. Relatively large autocorrelations of the returns produced by the model are the consequence of our behavioral assumptions, and, in particular, of the dominating trend-following behavior. Even if the modeling of a more realistic market architecture cannot "kill" the autocorrelations completely, it contributes to a certain improvement in the statistics. Indeed, when that the market is in the volatile regime, the autocorrelations are the largest under the MM and WA.[16] Bottazzi et al. (2005) do not find conclusive evidence about the sign and magnitude of the autocorrelations of returns produced by their model, but indicate that they are significant only at the first few lags.

Finally, we turn to the volatility clustering universally observed in the data. Volatility clustering can be identified by the presence of significant autocorrelations in the squared returns for a number of lags. The presence of volatility clustering suggests that despite the

[16]The returns autocorrelations can be lowered to the zero level by adding a sufficiently large amount of dynamic noise, see, e.g., Gaunersdorfer and Hommes (2007). Since it comes in a cost of understanding the dynamics of the model, we do not follow this path.

linear unpredictability of returns, the returns are not statistically independent. In the left panel of Fig. 9 we show the autocorrelations of squared returns at lags $1-5$ (top to bottom) as a function of β. One immediately observes not only an expected dependence on the intensity of choice parameter β but also a significant dependence on the market architecture. Namely, under both order-driven protocols, the autocorrelations of the squared returns are always positive and relatively large. They decay slowly, which is consistent with the volatility clustering. In contrast, under the MM and WA, the autocorrelations of squared returns are generally close to zero or even negative.

To verify whether the squared returns generated by our model exhibit long memory, in the right panel of Fig. 9 we plot the autocorrelation function for the fixed value of $\beta = 8$ for the four different mechanisms. The thin lines indicate 0.95 confidence limits. Under the BA and OB, the squared returns show positive slow decaying correlation, while under the MM and WA the auto-correlations are small and their pattern is atypical for a financial series. We conclude that in our model the realistic patterns of volatility clustering can be attributed to the realistic order-driven mechanisms. Similarly, Bottazzi et al. (2005) also find volatility clustering and long memory of squared returns under the BA and OB mechanisms.

6 Conclusion

Simulations presented in this paper contribute to the analysis of the interplay between behavioral ecologies of markets with heterogeneous traders and institutional market settings. Our work is motivated by many regularities observed in financial markets which still need a structural explanation and by different approaches which economists exploit to explain these regularities. Since the dynamics of financial markets are an outcome of a complicated interrelation between behavioral patterns and the underlying market mechanism, we offer a route in between, starting with a simple, analytically tractable model based on simple behavioral assumptions and simulating it in a more realistic market setting.

Our gradual approach of introducing different market mechanisms in the market with heterogeneous agents was inspired by the work of Bottazzi et al. (2005). While in their model agents did not change their strategies over time, in our setup agents were able to do so, which seems to be more realistic. The results of Bottazzi et al. (2005) suggest that the time-series properties are largely driven by market architecture. We, however, clearly see that certain behavioral features are also important. In our model, no matter which type of market clearing is used, two different regimes with completely different dynamical properties occur depending on the value of the intensity of choice. On the other hand, the trading protocol strongly affects the critical value of the intensity of choice, playing the role of the border line between the two regimes. Namely, we find that the frictions introduced by the order-driven protocols increase the range of the intensity of choice parameter for which the system is in the tranquil regime. Furthermore, when the market is in the volatile regime, the trading protocol also dictates the time-series property.

We also investigate an allocative efficiency of the market. A seminal paper by Gode and Sunder (1993) suggests that the continuous double auction leads to an allocatively efficient outcome even when agents trade at random. LiCalzi and Pellizzari (2007) explore this line of research and compare performances of four market protocols in terms of different criteria, such as the time needed to converge to the equilibrium, traded volume, and price volatility generated during this convergence. Agents' valuations, or the so-called environment, are fixed in both papers. We consider a dynamic model with an ever-changing environment and observed in this setting that there is simply not enough time to converge to an allocative efficiency outcome under the order-driven mechanisms. We find that the allocative efficiency loss is comparable for the BA and OB mechanisms.

Appendix A. Proof of Proposition 3.1

From the first equation of (12), one gets that in an equilibrium, $R\,x^* = vx^*\,n^* + x^*\,(1-n^*)$. Since $v < 1$ and the fraction of fundamentalists n^* should belong to the interval $[0,1]$, this equation has a unique solution $x^* = 0$. Substituting zero deviation into the performance measure, we derive $n^* = e^{-\beta C}/(1+e^{-\beta C})$.

To derive the stability conditions, the Jacobian matrix of the system should be computed. Substituting the second equation of (12) into the first and introducing the lagged variables, the Jacobian matrix in the fundamental equilibrium reads:

$$\boldsymbol{J} = \left\| \begin{array}{cccc} \big((v-1-g)n^*+1+g\big)/R & g(n^*-1)/R & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{array} \right\| .$$

This matrix has two zero eigenvalues, while two others are derived from a 2×2 matrix with trace $T = \big((v-1-g)n^*+1+g\big)/R$ and determinant $D = g(1-n^*)/R$. Standard conditions for eigenvalues of this matrix to be inside the unit circle are $D < 1$, $T < D+1$, and $T > -D-1$. The last two conditions are always satisfied, while the first is simplified to $n^* > 1 - R/g$. When $g \leq R$, this is also satisfied and the fundamental steady state is locally stable. If $g > R$, the bifurcation value is a solution of $e^{-\beta C}/(1+e^{-\beta C}) = 1 - R/g$, which gives us the result of Proposition 3.1*(ii)*.

Appendix B. Market Maker under the Large Market Limit

Similarly to the WA case, as $N \to \infty$, the Law of Large Numbers guarantees a convergence of the actual fractions of fundamentalists and trend-followers to the probabilities defined in (8). The model can then be represented by the 4-dimensional system consisting of the market clearing equation and equation of the evolution of the fractions of fundamentalists.

The system in deviations from fundamental price reads

$$\begin{cases} x_{t+1} = x_t + \mu\Big(n_t\, v\, x_t + (1-n_t)\big(x_t + g\,(x_t - x_{t-1})\big) - R\, x_t\Big) + \varepsilon_{t+1} \\ n_{t+1} = \exp\Big(\beta\Big[\big(v\, x_{t-1} - R\, x_{t-1}\big)\,\big(x_t - R\, x_{t-1} + \delta_t\big) - C\Big]\Big)\Big/ Z_{t+1}\,, \end{cases}$$

where, similar to (12), n_t is the fraction of fundamentalists, δ_t and ε_{t+1} are two stochastic components representing dividend realization and dynamic noise, respectively, and Z_{t+1} is the normalization factor.

As in the case of the WA LML, in an equilibrium, $R\,x^* = vx^*\,n^* + x^*\,(1-n^*)$ that has a unique solution $x^* = 0$ leading to $n^* = e^{-\beta C}/(1+e^{-\beta C})$.

To derive the stability conditions, we compute the Jacobian matrix of the system:

$$\boldsymbol{J} = \begin{Vmatrix} 1+\mu(n^*v+(1-n^*)(1+g)-R) & \mu g(n^*-1) & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{Vmatrix}.$$

Two eigenvalues of this matrix are equal to zero, and two others are derived from a 2×2 matrix with trace $T = 1+\mu((v-1-g)n^*+1+g-R)$ and determinant $D = \mu g(1-n^*)$. The eigenvalues of this matrix are inside of the unit circle when $D < 1$, $T < D+1$, and $T > -D-1$. The second condition is always satisfied, while the first is simplified to $n^* > 1-1/(\mu g)$. The third condititon can be written as $\mu(1-n^*)(1+2g) + \big(2+\mu(n^*v-R)\big) > 0$ The first term of the LHS of the condition is always positive, while the second is increasing in n^* and takes its minimum value when $n^* = 0$. Therefore, $\mu R < 2$ is sufficient for the satisfaction of the third condition. To satisfy this condition we need to choose small values of μ.

Note that setting $\mu = 1/R$ gives us the same bifurcation point β^* as in the WA LML.

References

Anderson, S.P., de Palma, A. and Thisse, J.-F., 1992. Discrete Choice Theory of Product Differentiation. MIT Press, Cambridge, MA.

Anufriev, M., Bottazzi, G. and Pancotto, F., 2006. Equilibria, stability and asymptotic dominance in a speculative market with heterogeneous agents. Journal of Economic Dynamics and Control, 30, 1787 – 1835.

Anufriev, M. and Dindo, P., 2007. Wealth selection in a financial market with heterogeneous agents. CeNDEF Working Paper 2007-10. University of Amsterdam.

Arthur, W.B., Holland, J.H., LeBaron, B., Palmer, R. and Tayler, P., 1997. Asset pricing under endogenous expectations in an artificial stock market, in: Arthur, W.B., Durlauf, S.N. and Lane, D.A. (Eds.), The Economy as an Evolving Complex System II. Addison-Wesley, pp. 15–44.

Bak, P., Paczuski, M. and Shubik, M., 1997. Price variations in a stock market with many agents. Physica A, 246, 430–453.

Bottazzi, G., Dosi, G. and Rebesco, I., 2005. Institutional achitectures and behavioural ecologies in the dynamics of financial markets: a preliminary investigation. Journal of Mathematical Economics, 41, 197–228.

Brock, W.A., 1993. Pathways to randomness in the economy: Emergent nonlinearity and chaos in economics and finance. Estudios Económicos, 8, 3–55.

Brock, W.A., 1997. Asset price behavior in complex environment, in: Arthur, W.B., Durlauf, S.N. and Lane, D.A. (Eds.), The Economy as an Evolving Complex System II. Addison-Wesley, pp. 385–423.

Brock, W.A. and Hommes, C.H., 1997. A rational route to randomness. Econometrica, 65, 1059–1095.

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

Camerer, C.F. and Ho, T.-H., 1999. Experienced-weighted attraction learning in normal form games. Econometrica, 67, 827–874.

Chiarella, C. and He, X., 2001. Asset price and wealth dynamics under heterogeneous expectations. Quantitative Finance, 1, 509–526.

Chiarella, C. and Iori, G., 2002. A simulation analysis of the microstructure of double auction markets. Quantitative Finance, 2, 346–353.

Chiarella, C., Iori, G. and Perelló, J., 2007. The impact of heterogeneous trading rules on the limit order book and order flows. Working Paper. Available at `http://xxx.lanl.gov/pdf/0711.3581/`.

Cohen, K.J., Maier, S.F., Schwartz, R.A. and Whitcomb, D.K., 1978. Limit Orders, Market Structure, and the Returns Generation Process. The Journal of Finance, 33, 723–736.

Cont, R., 2002. Empirical properties of asset returns: stylized facts and statistical issues. Quantitative Finance, 1, 45–60.

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

Diks, C. and van der Weide, R, 2005. Herding, a-synchronous updating and heterogeneity in memory in a CBS. Journal of Economic Dynamics and Control, 29, 741–763.

Evans, G.W. and Honkapohja, S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press.

Fama, E.F., 1970. Efficient capital markets: a review of theory and empirical work. Journal of Finance, 25, 383–422.

Farmer, J.D. and Joshi, S., 2002. The price dynamics of common trading strategies. Journal of Economic Behaviour and Organization, 49, 149–171.

Gaunersdorfer, A. and Hommes, C.H., 2007. A nonlinear structural model for volatility clustering, in: Teyssière, G. and Kirman, A. (Eds.), Long Memory in Economics. Springer, pp. 265–288.

Gaunersdorfer, A., Hommes, C.H. and Wagener, F.O.O., 2008. Bifurcation routes to volatility clustering under evolutionary learning. Journal of Economic Behavior and Organization; forthcoming.

Gode, D.K. and Sunder, S., 1993. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. Journal of Political Economy, 101, 119–137.

Hommes, C.H., 2006. Heterogeneous agent models in economics and finance, in: Judd, K.L. and Tesfatsion, L. (Eds.), Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics. North-Holland, Amsterdam, pp. 1109–1186.

Hommes, C.H., Huang, H. and Wang, D., 2005. A robust rational route to randomness in a simple asset pricing model. Journal of Economic Dynamics and Control, 29, 1043–1072.

LeBaron, B., 2006. Agent-based computational finance, in: Judd, K.L. and Tesfatsion, L. (Eds.), Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics. North-Holland, Amsterdam, pp. 1187–1234.

LeBaron, B. and Yamamoto, R., 2006. Long-memory in an order driven market. Physica A, 383, 85 –89.

Levy, M., Levy, H. and Solomon, S., 2000. Microscopic Simulation of Financial Markets. Academic Press, London.

LiCalzi, M. and Pellizzari, P., 2003. Fundamentalists clashing over the book: a study of order-driven stock markets. Quantitative Finance, 3, 1–11.

LiCalzi, M. and Pellizzari, P., 2007. Simple market protocols for efficient risk sharing. Journal of Economic Dynamics & Control, 31, 3568–3590.

Lucas, R.E., 1978. Asset prices in an exchange economy. Econometrica, 46, 1429–1445.

Lux, T., 1995. Herd behaviour, bubbles and crashes. The Economic Journal, 105, 881–896.

Maslov, S., 2000. Simple model of a limit order-driven market. Physica A, 278, 571–578.

Mike, S. and Farmer, J.D., 2008. An empirical behavioral model of liquidity and volatility. Journal of Economic Dynamics and Control, 32, 200–234.

Pagan, A.R., 1996. The econometrics of financial markets. Journal of Empirical Finance, 3, 15–102.

Sargent, T.J., 1993. Bounded Rationality in Macroeconomics. Oxford University Press.

Shiller, R., 1981. Do stock prices move too much to be justified by subsequent changes in dividends? American Economic Review, 71, 421–436.

Weisbuch, G., Kirman, A. and Herreiner, D., 2000. Market organisation and trading relationships. Economic Journal, 110, 411–436.