Macroeconomic Dynamics, **16**, 2012, 556–575. Printed in the United States of America.
doi:10.1017/S136510051000074X

A DYNAMIC ANALYSIS OF THE MICROSTRUCTURE OF MOVING AVERAGE RULES IN A DOUBLE AUCTION MARKET

Carl Chiarella and Xue-Zhong He
University of Technology, Sydney

Paolo Pellizzari
Ca' Foscari University

Inspired by the theoretically oriented dynamic analysis of moving average rules in the model of Chiarella, He, and Hommes (CHH) [*Journal of Economic Dynamics and Control* 30 (2006), 1729–1753], this paper conducts a dynamic analysis of a more realistic microstructure model of continuous double auctions in which the probability of heterogeneous agents trading is determined by the rules of either fundamentalists mean-reverting to the fundamental or chartists choosing moving average rules based on their relative performance. With such a realistic market microstructure, the model is able not only to obtain the results of the CHH model but also to characterize most of the stylized facts including volatility clustering, insignificant autocorrelations (ACs) of returns, and significant slowly decaying ACs of the absolute returns. The results seem to suggest that a comprehensive explanation of several statistical properties of returns is possible in a framework where both *behavioral* traits and *realistic microstructure* have a role.

Keywords: Microstructure Model, Continuous Double Auctions, Heterogeneous Agents, Stylized Facts

1. INTRODUCTION

In traditional economic and finance theory based on the assumptions of investor homogeneity and rational expectations, agents rationally incorporate all relevant information into their trading decisions. Hence the movement of prices is assumed to be perfectly random and to exhibit random walk behavior. The theory does provide a benchmark framework for our understanding of the dynamics of financial

This paper was written while Paolo Pellizzari was visiting the School of Finance and Economics at the University of Technology, Sydney (UTS). The financial support of the Paul Woolley Centre for Market Dysfunctionality and the Quantitative Finance Research Centre at UTS are gratefully acknowledged together with the Italian PRIN 2007 grant 2007TKLTSR, "Computational markets design and agent-based models of trading behavior." Chiarella and He acknowledge the ARC Discovery Grant DP0773776 in its support of this project. We thank the editor and two anonymous referees for their comments and suggestions. All remaining errors are ours. Address correspondence to: Paolo Pellizzari, Dipartimento di Matematica Applicata, S. Giobbe—Cannaregio 873, 30121 Venice, Italy; e-mail: paolop@unive.it.

© 2011 Cambridge University Press 1365-1005/11

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

asset prices, but it ignores some of the most important features of real-world economic agents, such as their heterogeneity and bounded rationality, and the impact of their interaction on the particular structure of financial markets. Also, empirical investigations of high-frequency (such as daily) financial time series in financial markets show some common stylized facts, including excess volatility, some skewness and excess kurtosis, fat tails, and volatility clustering; see Pagan (1996) for a comprehensive discussion of stylized facts characterizing financial time series and Lux (2006) for a recent survey of empirical evidence on various power laws. These facts are not entirely contradictory to traditional economic and finance theory, but standard results do not provide persuasive explanations for a large subset of these facts.

As a result, the literature has witnessed an increasing number of attempts at modeling financial markets by incorporating heterogeneous agents and bounded rationality, on which there is a nice overview in the recent surveys by Lux (2006, 2009), Hommes (2006), LeBaron (2006), and Chiarella et al. (2009a). This class of models characterizes the dynamics of financial asset prices as resulting from the interaction of heterogeneous agents having different attitudes to risk and different expectations about the future evolution of prices. One of the key aspects of these models is that they exhibit expectations feedback—agents' decisions are based upon predictions of future values of endogenous variables whose actual values are determined by equilibrium equations. In particular, Brock and Hommes (1997, 1998) proposed an *adaptive belief system* model of financial markets. The agents adapt their beliefs over time by choosing from different predictors or expectation functions, based upon their past performance. The resulting nonlinear dynamical system is, as Brock and Hommes (1998 and Hommes 2002) show, capable of generating a wide range of complex price behavior from local stability to high-order cycles and chaos. It is very interesting to find that adaptation, evolution, heterogeneity, and even learning can be incorporated into the Brock and Hommes type of framework; for details of such extensions, the reader should consult Gaunersdorfer (2000), Chiarella and He (2001, 2002, 2003), Hommes (2001, 2002), and Chiarella et al. (2002) for asset markets and De Grauwe and Grimaldi (2006) and Westerhoff (2003) for foreign exchange markets. Moreover, the recent articles by Westerhoff (2004) and Chiarella et al. (2005, 2006a) show that complex price dynamics may also result within a multiasset market framework. This broader framework of boundedly rational heterogeneous agents can also give rise to quite rich and complicated dynamics and so give a deeper understanding of market behavior. In particular, it is capable of explaining various types of market behavior, such as the deviation of the market price from the fundamental price, market booms and crashes, and quite a number of the stylized facts referred to earlier. More recently, He and Li (2007, 2008) used a simple market fraction model of fundamentalists and trend followers to analyze the mechanism generating the power-law distributed fluctuations. Their results provide a promising perspective on the use of these models to produce the observed characteristics of financial market time series.

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

In contrast to the theoretically oriented models discussed above, there is also a rapidly expanding literature on heterogeneous-agent models that is computationally oriented and for which we refer the reader to the recent survey by LeBaron (2006). These models are becoming increasingly important and have proved very powerful at generating the stylized facts. There are at least two important advantages of this approach compared to the theoretically oriented one. The first is that many behavioral aspects at the micro level, including the interaction of agents, can be aggregated at the macro level through computer simulations. The second is that more realistic market features, including budget or wealth constraints, no short selling, and irregular intraday trading of nonfractional shares, can be readily incorporated into the market microstructure of continuous double auctions and dealer and hybrid markets. It should be stated that earlier work on the computational class of models faced the problem of many degrees of freedom and many parameters, which made it difficult to understand and assess the main causes of the observed stylized facts. Here we shall follow the path of adding some realistic features to a model that is well understood from a theoretical perspective but involves a minimal, in some sense, set of parameters.

The advantages of both theoretically and computationally oriented heterogeneous-agent models have naturally led to some recent computationally oriented models that are based on theoretically oriented models but with more realistic market microstructure; for example, Chiarella and Iori (2002), Pellizzari and Westerhoff (2009), and Chiarella, Iori, and Perellò (2009b). The current paper falls into this category and conducts a dynamic analysis of a microstructure model of continuous double auctions based on the theoretically oriented work of Chiarella et al. (2006b) (CHH model hereafter) that gave a dynamic analysis of moving average rules under the market-maker scenario. The CHH model proposes a stochastic dynamic financial-market model in which demand for traded assets has both a fundamentalist and a chartist component. The chartist demand is governed by the difference between the current price and a long-run moving average. It shows that the moving average can be a source of market instability, and the interaction of the moving average and market noises can lead to a tendency for the market price to take long excursions away from the fundamental. The model, with the addition of noise, as outlined in in Chiarella et al. (2006c), reveals various types of market price phenomena, the coexistence of apparent market efficiency and a large chartist component, price resistance levels, and skewness and kurtosis of returns. In order to be able to conduct a theoretical analysis, CHH make some less realistic assumptions, including a fixed length for the moving average window, unlimited short selling and borrowing, trading of fractional amounts of shares, and no intraday trading. In particular, they assume homogeneity within the two groups of the model, fundamentalists and chartists. The current paper drops these unrealistic assumptions and also allows the market price to be determined by a market microstructure model of continuous double auctions (CDA) instead of by the stylized market-maker scenario used in the CHH model. In the CHH model, it is the probability of heterogeneous agents trading, rather than their demands,

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

that is determined by the rules of either fundamentalists mean-reverting to the fundamental or chartists choosing moving average rules based on their relative performance. We find that, with the inclusion of this realistic market microstructure, the model is able not only to obtain the essential features of the CHH model but also to characterize most of the stylized facts, including volatility clustering, insignificant autocorrelations (ACs) of returns, and significant slowly decaying ACs of the absolute returns. The results seem to suggest that a comprehensive explanation of several statistical properties of returns is possible in a framework where both *behavioral* traits and *realistic microstructure* have a role.

Chiarella et al. (2009b) is a computational model of an order-driven market that has some similarities to our model. Their work, however, is more focused on the description of the properties of the book and of the order flow, and they assume that every agent uses a strategy that blends three components (fundamentalist, chartist, noisy). Instead, our traders are either fully fundamentalist or fully chartist at any given time and have the chance to switch to another strategy when desired. A novel feature is that agents switch *individually* depending on their personal success. This is more general than what was presented in Chiarella et al. (2006b) and in many other models where *groups* as a whole can switch with some probability.

The plan of the paper is as follows. Section 2 introduces the model. Section 3 contains the main results and a sensitivity analysis. Section 4 concludes.

2. A CONTINUOUS DOUBLE AUCTION WITH HETEROGENEOUS TRADERS

Our model is inspired by that of Chiarella et al. (2006b), much of the structure of which is retained in term of types of agents and behavioral characteristics behind their demand functions. However, we extend the CHH model considerably in order to introduce a more realistic market microstructure through a CDA, which has become a widely used clearing device in many stock exchanges around the world.

It is convenient to briefly outline the salient features of the CHH model. Two types of agents, fundamentalists and chartists, populate the market, and trade depending on the value of the fundamental price and on the trading signal generated by a fixed-length moving average of past prices. The fractions of agents of types $h \in \{f, c\}$, standing for fundamentalist and chartist, respectively, evolve according to smoothed realized profits, as pioneered in Brock and Hommes (1998). The price reacts to the imbalance in demand and supply via a market maker, who clears the market and announces the price for the next period. The CHH model shows that trading based on moving averages can be destabilizing, especially if the window length is increased. The addition of noise, either in the fundamental process or market noise, has the potential to cause the model to generate bubbles, crashes, and environments where the chartist component is persistent. The price time series of the model with fundamental noise alone, however, does not exhibit

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

a number of the standard statistical features found in most financial-market time series, in particular the insignificant ACs of returns and significant slowly decaying ACs of the absolute returns. Chiarella et al. (2006c) further show that with careful selection of the sizes of the fundamental and market noise, the addition of additive (rather than multiplicative) noise to the price equation of the market maker brings these statistical features closer to what is observed in actual financial data.

We stress the fact that the CHH model abstracts from many details relating to the actual trading mechanism. Agents, for example, without considering endowments, can in principle buy or sell unlimited quantities, acting as price takers. Demands are filled by a market maker or an impersonal clearing device and all the chartists use the same moving average signal for trading purposes. Although it retains the main ideas of CHH, the main aim of this paper is to examine whether the introduction of more realistic microstructure features into the CHH model can maintain the main results of the CHH model and produce more realistic statistical features of the returns generated by simulation, as recently advocated by Lux (2009). There are two main differences of our model from the CHH model. First, we consider a CDA where agents who can neither borrow nor short-sell submit limit orders with no certainty that they will be filled. Exchange takes place only if two parties agree on quantity and price, and there is a layer of intraday activity that is missing from the original model. A book-based microstructure framework also requires that agents be endowed with (simple) ways to deal with the fundamental tradeoff in a continuous market: the more aggressive the limit order, the bigger is the probability of trading but, conversely, the smaller is the final profit if the transaction is executed. Second, our agent-based market allows the chartists to use individual lengths for their moving average. This was indeed one of the ideas left for future research in the original work. Based on the results obtained in this paper, we claim that a realistic microstructure model can perturb the original model in the "right" way, meaning that a book-based augmented CHH model is able to produce a host of common stylized facts.

2.1. The Heterogeneous Agents and Trading

We assume that agent i in a population of N agents is initially endowed with some units S_{i0} of a non–dividend paying risky stock and cash C_{i0}, with $i = 1, \ldots, N$. The endowments S_{it} and C_{it} at time t are updated in the obvious way whenever the agents trade. Agents cannot short-sell stocks nor borrow money. In other words, agents are banned from submitting bids with a limit price greater than their cash endowment and, symmetrically, cannot submit an ask if they have no stocks in their endowment. As they are restricted before submitting an order, bankruptcy is not possible. We assume that the interest rate $r = 0$ or, equivalently, that the interest rate payments are spent elsewhere.

As in the CHH model, agents are heterogeneous in that they can trade based on a fundamentalist or chartist strategy. In the first case, they seek to buy (sell) stock

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

when it is under(over)-valued with respect to an exogenously given, stochastically fluctuating fundamental price that evolves according to

$$p^*_{t+1} = p^*_t \exp(\sigma_f v_t),$$

where $\sigma_f \geq 0$ is the constant volatility of the fundamental return and $v_t \sim N(0, 1)$ follows a standard normal distribution. The chartists use a moving average price (computed over time windows of heterogeneous length) to obtain buy/sell signals.

In the following, the subscript $t = 1, \ldots, T$ will refer to calendar trading days when variables are constant over the tth "trading day," whereas $\tau \in \mathbf{R}^+$ is an intraday time subscript that will be used with variables that can assume different values in the same day, such as the price of stock traded in any continuous auction. Hence, p_τ is the last cleared price and it is not updated until a new transaction takes place.

The agents can submit one limit order per day valid for one unit of the stock.[1] A limit order is a quantity–limit price couple (q, l) that is submitted in a randomly selected instant $t < \tau < t + 1$ of a given day t. At the beginning of each day a random permutation $\mathcal{P}_t$ of $\{1, \ldots, N\}$ is drawn and agents take action in the order dictated by $\mathcal{P}_t$. In other words, orders are issued sequentially, in a random order that is independently sampled every day, so that agents have only one chance to trade each day, when it is their turn to "speak."

Each fundamentalist trader posts an order with some probability[2] ${}_t s_\tau = \min(1, \alpha | p^*_t - p_\tau |)$, with $\alpha > 0$ denoting the sensitivity to the deviations of the most recent price from the fundamental value, and refrains from posting with the residual probability. The chance of submitting orders increases when this deviation is large, and agents submit for sure whenever the mispricing is greater than $1/\alpha$, in either direction. Formally, the order $({}_t q_{i\tau}, {}_t l_{i\tau})$ of agent i, posted with probability ${}_t s_\tau$, is such that

$${}_t q_{i\tau} = \mathbf{sgn}(p^*_t - p_\tau),$$

where **sgn** denotes the sign function, so that that positive (negative) arguments lead to buy (sell) orders. We assume that the submitted limit price is uniformly drawn between the fundamental and the last available closing price p^{close}_{t-1}; thus

$${}_t l_{i\tau} = \begin{cases} U\left(p^{\text{close}}_{t-1}, p^*_t\right) & \text{if } p^{\text{close}}_{t-1} < p^*_t; \\ U\left(p^*_t, p^{\text{close}}_{t-1}\right) & \text{if } p^*_t < p^{\text{close}}_{t-1}. \end{cases}$$

The specification of the limit price is consistent with the idea that a fundamental agent would trade at any price below (above) the fundamental value, if he/she were a buyer (seller) within the trading day. Hence, he/she provides (limited) liquidity at price levels that are, at the same time, more favorable than the previous closing price and secure some random profit with respect to the fundamental value, if executed. We can also interpret this bidding strategy, which has similarities to that

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

of the zero intelligence constrained (ZIC) traders in Gode and Sunder (1993), as the choice of a random mark-up (mark-down) with respect to the reservation price of the fundamentalists that is given by the fundamental value p_t^*.

Each chartist submits an order with probability ${}_tr_\tau = |\tanh(a\ {}_t\psi_\tau^{L_i})|$, where $a > 0$ measures the strength of the extrapolation activity of the chartists and ${}_t\psi_\tau^{L_i} = p_\tau - ma_t^{L_i}$ defines a trading signal, namely the difference between the current price and a moving average of length L_i of the particular chartist, $ma_t^{L_i} = 1/L_i \sum_{j=1}^{L_i} p_{t-j}^{\text{close}}$. The limit order, submitted by the chartists with probability ${}_tr_\tau$, is $({}_tq_{i\tau}, {}_tl_{i\tau})$, where

$$
\begin{aligned}
{}_tq_{i\tau} &= \mathbf{sgn}\left({}_t\psi_\tau^{L_i}\right), \\
{}_tl_{i\tau} &= p_\tau(1 + \Delta\ {}_tz_\tau),
\end{aligned}
$$

${}_tz_\tau \sim N(0, 1)$, and the standard deviation $\Delta > 0$ is related to the aggressiveness of the agents in increasing the bid or in reducing the ask with respect to the current intraday price p_τ. A large Δ would produce bids exceeding the price by a large amount, and the same holds for aggressively low asks. Smaller values for Δ, conversely, would produce a lot of limit prices that are very close to the last price. Even if the proportion of improving bids or asks is constantly 50%, the latter case is likely to produce smooth price movements, whereas the former has the potential to induce large jumps in prices due to adjacent trades at rapidly increasing/decreasing levels.[3] This formulation captures the fact that the fundamental agents behave differently from the chartists. The former are anchoring their limit price to the difference of the intraday price from the last closing fundamental price available, whereas the latter are more sensitive to intraday dynamics and focus on the intraday price difference from the daily price moving averages with mixed short and long time scales, depending on their random draws on the moving average window lengths, fostering at times spectacular increases of the traded volume within the day. This behavior appears to describe rather convincingly some aspects of both fundamental and chartist trading, the latter being more hazardous and speculative.

From time to time, all the agents may end up on the same side of the market and there will be no counterparty for any outstanding order. In real markets, the absence of bids, say, would stimulate more aggressive offers that would in turn lead some agents to issue advantageous bids. In our model, however, agents do not look at the state of the book or at the order imbalance (as they focus on p_τ) and in order to ensure the correct functioning of the CDA in every situation, a small amount of random trading is introduced to facilitate trading near the latest market price. More precisely, with probability p_ϵ, agents will issue a random order (in place of what was described above) for a quantity ± 1 with equal probability and limit price given by $p_\tau + \sigma_\epsilon\ {}_tz_\tau$, where ${}_tz_\tau \sim N(0, 1)$ is newly sampled whenever needed. The limit price is obtained by offsetting the last observed price p_τ by a random amount, whose constant standard deviation is σ_ϵ.

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

2.2. The Switching of Trading Strategies

The traders can switch strategy at the end of each day, when the closing price becomes available and they can evaluate their realized profit, if any. Let $I_t \subseteq \{1, \ldots, N\}$ be the set of agents who traded on day t at some price p_{τ_i} and let $X_{it} \in \{f, c\}, i \in I_t$ be their state. They compute their realized profit according to

$$\pi_{it} = \begin{cases} p_t^{\text{close}} - p_{\tau_i} & i \in I_t \text{ is a buyer,} \\ p_{\tau_i} - p_t^{\text{close}} & i \in I_t \text{ is a seller.} \end{cases}$$

Then agents adjust an individual smoothed profit measure in their state as

$$U_{X_{it}}(i, t) = \pi_{it} + \eta U_{X_{it}}(i, t-1), \quad i \in I_t,$$

where $\eta (\in [0, 1])$ is a memory parameter. Observe that only agents who traded adapt the performance of the current state X_{it}, whereas the other agents $j \notin I_t$ do not alter their profit measure $U_{X_{jt}}(j, t) = U_{X_{jt}}(j, t-1)$.

Finally, all fundamentalists switch to chartism with probability

$$n_{c,i,t} = \frac{\exp[\beta U_c(i, t)]}{\exp[\beta U_f(i, t)] + \exp[\beta U_c(i, t)]}, \quad i = 1, \ldots, N.$$

Equivalently, the probability of any chartist switching to the fundamentalist strategy is given by $n_{f,i,t} = 1 - n_{c,i,t}, i = 1, \ldots, N$. Here $\beta > 0$ is a parameter related to the intensity of switching: small values of β make the agents insensitive to profits and prone to use the two strategies with equal probability. In contrast, large values of β make them more likely to switch at once to the most profitable strategy at time $t + 1$. Our switching mechanism differs in two important ways from the standard proposal pioneered in Brock and Hommes (1998). A first difference lies in the use of individual switching probabilities, rather than the global one that is used in CHH. In other words, agents of the same type can have rather different realized profit measures, due to different transaction prices (even on the same day) and different L_i's. As a consequence, their switching can be driven by different probabilities. In the second place, the smoothed realized profit measure is updated by an agent only if he/she succeeds in trading. In the standard treatment, trading always occurs, but in our setup, this does not necessarily happen, and we feel that there is no reason to update the accumulated profits if an agents fails to trade on a specific day. This means that, in such a case, we have $U_{X_{it}}(i, t+1) = U_{X_{it}}(i, t)$. A similar individual-based switching mechanism is used in Pellizzari and Westerhoff (2009).

2.3. Timing

A typical trading day t develops as follows and is illustrated in Figure 1:

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

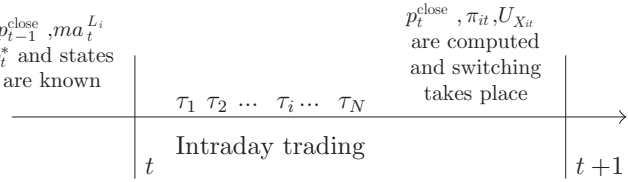


FIGURE 1. Schematic representation of the unfolding of a trading day in the model.

(1) At t^-, the end of day $(t-1)$, the closing price p_{t-1}^{close} of the previous session or trading day $t-1$ and all the moving averages $ma_t^{L_i}$, $i = 1, \ldots, N$, are available.
(2) At time t^+, the beginning of day t, the agents start trading at random times $t < \tau_i < t+1$, submitting their orders to the market (with no certainty that they will be executed).
(3) In $(t+1)^-$ the closing price p_t^{close} for day t is known and traders hence can compute profits π_{it} and adjust their performance measure $U_{X_{it}}(i,t)$. Notice that if an agent is unable to trade, his/her $U_{X_{it}}(i,t)$ remain unchanged.
(4) The complementary probabilities $n_{f,i,t}$ and $n_{c,i,t}$ are computed for all agents who possibly switch to the other strategy to be used starting at time $(t+1)^+$.
(5) New moving averages and new fundamental price can be computed to be used starting at time $(t+1)^+$.

2.4. The Market Protocol: CDA

We describe in this section the set of rules that govern the trading process or, in brief, the microstructure of our model. A CDA allows agents to submit orders at any time. We consider only limit orders, that is, quantity–price couples (q, l), to be interpreted as the binding promise to buy (sell) q units at a price no larger (smaller) than l. We stress that, once submitted, orders cannot be canceled or changed. Every positive real number l is a legal limit price, and we restrict the quantity to be $q = \pm 1$ to denote orders on the sell (-1) or buy $(+1)$ side. All the orders are sorted, kept in a book, and, if the book is nonempty, let $a_m \geq a_{m-1} \geq \ldots a_2 \geq a_1 > b_1 \geq b_2 \geq \ldots b_n$ be the sequence of asks and bids, denoted by a_i and b_j, respectively. The biggest bid b_1 is called the *best bid* and the smallest ask a_1 is called the *best ask*.[4] Whenever a new limit order (q, l) is submitted, it is matched against the opposite side to find the best compatible price. If possible, the order is executed at the price of the matched order (either a_1 or b_1). If no match is found, the order is inserted into the appropriate book. In more detail, if $q = 1$ and $l \geq a_1$ then one unit is exchanged at the price a_1 between the agents whose orders crossed and the best ask is removed from the book; otherwise the submitted order is inserted among the bids so that the buying queue becomes $b_1 \geq b_2 \geq \ldots \geq b_j \geq l \geq b_{j+1} \geq \ldots \geq b_n$. In the same way, if $q = -1$ and $l \leq b_1$, then one unit is exchanged at the price b_1 between the agents with crossing orders and the best bid is removed from the book; otherwise

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

the submitted order is inserted among the asks so that the selling queue becomes $a_m \geq a_{m-1} \geq \ldots \geq a_j \geq l \geq a_{j-1} \geq \ldots \geq a_1$.

We assume that at the end of the trading day, the book is completely erased, so that the next trading day begins with no queued orders. The Tokyo Stock Exchange, for example, adopts this procedure, whereas other markets let some orders survive across "days" according to different rules.

3. SIMULATION RESULTS

3.1. Parameter Selection

We discuss in this section the results obtained by simulation of the model previously described. For each parameter set we run the market 100 times for $T = 2{,}500$ trading days, constituting a period of about 10 years. The first 500 observations are discarded to avoid transient effects due mainly to the need to initialize the moving averages randomly.[5] We are then left with 100 series of 1,999 logarithmic returns that can be analyzed. The choice of the parameters is guided by the values that were used in CHH but still required some trial and error in order to get realistic time series, as in most time series analysis of heterogeneous agent-based models. In order to provide a comparison, we have resimulated 100 time series according to Chiarella et al. (2006c), focusing on the case of constant-amplitude noise, which is the most interesting as far as stylized facts are concerned. It is often the case in agent-based models that the number of parameters is large, as heterogeneity is allowed at the level of the single agent; for example, the length L_i of the moving average and initial endowments C_{i0}, S_{i0} differ across agents. However, we resort to the standard technique of sampling the individual parameters from a single distribution, thus enormously cutting down the number of effective parameters. All the length $L_i, i = 1, \ldots, N$, are sampled from the set of integers $\{20, 21, \ldots, 100\}$, which contains commonly used values for moving average lengths. Likewise, the distribution of the initial endowments is the same for each agent and it is chosen so that on the average agents have a number of stock whose value is equal to the cash they hold. Given that no short selling or borrowing is allowed, this means that there is no systematic bias in favor of buyers and sellers. The remaining parameters are listed in Table 1. Observe that σ_ϵ is a constant fraction of the initial fundamental price p_0^* and that α (the reaction coefficient of the fundamentalists) and β (the intensity of switching), which are not individual parameters indexed by i, can have different values in some of the 100 simulations, as they are uniformly drawn from a "small" set of nearby values. This is done to test for the local robustness of our results with respect to slight changes in α and β.

3.2. Time Series and Distribution Results

Some representative time series are shown in Figure 2. The graphs depict the results of two typical simulations in the two columns. The first row of the picture

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

TABLE 1. Parameters used in the simulation

Parameter	Value	Description
N	1000	Number of agents
S_{i0}	$\{1, 2, \ldots, 9\}$	Initial stock endowment
C_{i0}	$1000 S_{i0}$	Initial cash endowment
p_0^{close}	1000	Initial price
p_0^*	990	Initial fundamental value
α	$\{\frac{17}{1000}, \frac{18}{1000}, \frac{19}{1000}\}$	Reaction coefficient for fundamentalists
a	0.1	Reaction coefficient for chartists
L_i	$\{20, 21, \ldots, 100\}$	Length of MA windows
Δ	0.005	Aggressiveness parameter
β	$\{\frac{1}{16}, \frac{1}{17}\}$	Intensity of switching
η	0.2	Profit smoothing parameter
p_ϵ	0.05	Probability to issue a random order
σ_f	0.005	Volatility of fundamental value (daily)
σ_ϵ	$\sigma_f p_0^*$	Volatility of random offset by noise traders

shows the time series of the market price and the fundamental value. They share a common feature found in the CHH model, namely that the price is tracking the fundamental value but the deviations are persistent and sizable. The strong bimodality of the distribution of the density of the difference $p_t - p_t^*$ in the last row confirms that the price stays for long times either below or above the fundamental price, as can be verified by a careful study of the panels in the first row. It is worth noting that this result was already present in Chiarella et al. (2006c), hinting that some properties of the original and more theoretical model are robustly retained in our model where realism is added with bilateral trading, budget constraints, and no short selling.

The logarithmic returns $r_t = \log(p_t/p_{t-1})$ are represented for the same time span in the second row. There are episodes of large change in price (in absolute value) and visual evidence of some degree of volatility clustering when the price surges or crashes. The third and fourth rows illustrate the autocorrelation of r_t and $|r_t|$. The linear predictability of the returns is very weak, pointing to some form of efficiency of the market. At the same time, the absolute returns are significantly correlated and slowly decaying for up to 20–30 lags, thus confirming heteroskedasticity in returns. However, we do not find evidence of long memory in volatility, which sometimes extends for many lags in real time series.

Obtaining uncorrelated raw returns and long memory in the absolute returns are indeed the most difficult things to calibrate in most agent-based models [see Lux (2006) and He and Li (2007) for related discussion], and a subtle balance between α, β, and Δ in our model is needed. Loosely speaking, there must be occasional large returns and occurrences of "spikes," but these must not be too frequent, in order to avoid positive serial correlation at low lags. Price spikes can be induced by increasing β, which triggers more switches to the current best performing

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

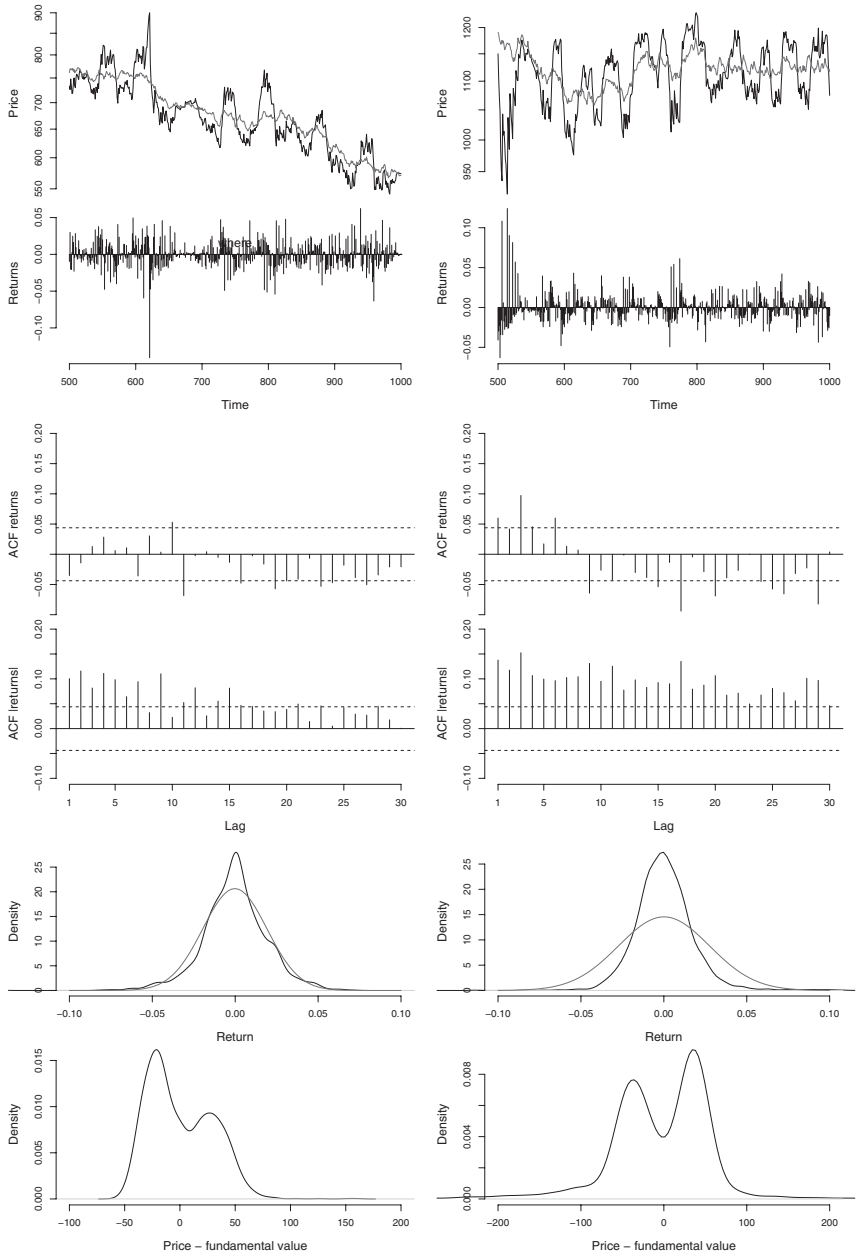


FIGURE 2. Some representative time series. From the top, the panels show the price and fundamental value; returns; autocorrelation of raw and absolute returns; density of returns (with a normal distribution with same mean and variance); and density of the difference between the price and the fundamental value.

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

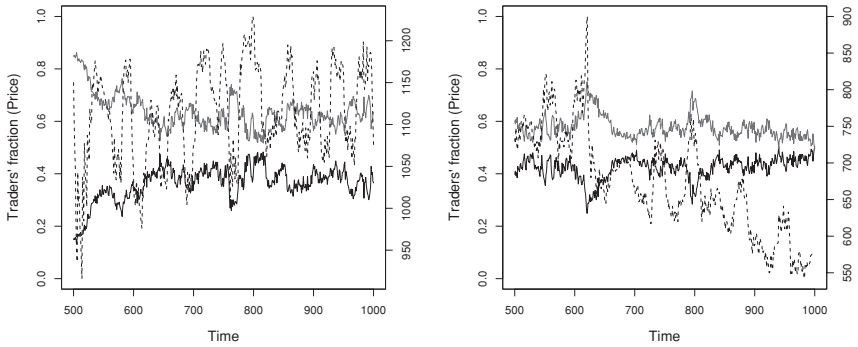


FIGURE 3. Evolution of the fraction of fundamentalist (black) and chartist (gray) traders for the same time series depicted in Figure 2. For convenience, the price is shown as a dashed line (see the right right-hand vertical axis).

strategy, and Δ, which can cause avalanches of rapidly surging or dropping prices. At the same time, α controls the activity of the fundamental traders, who step in and reverse the trend. The fundamentalists are also exploiting the fact that the finite budgets of the chartists force them to exhaust their resources after prolonged trends, as they cannot buy when their cash is over or (short-)sell when their stock holdings are null.

The fifth row of Figure 2 shows that the distribution of returns is leptokurtic and fat-tailed: the kurtosis for the depicted returns' time series is 5.34 (left) and 29.30 (right).

3.3. Market Fractions

Figure 3 plots the fractions of fundamentalists (black) and chartists (gray) and the market prices as functions of time. The technical component of the market is generically preponderant but, unlike the case in the CHH model, the correlation between the fraction of the chartists, and price changes is harder to spot. We feel that this is due to the heterogeneity of the chartists, who are no longer clearly responsible for bubbles and crashes. Remarkably, the market fractions of the fundamentalists are always below 50%, whereas the market fractions of the chartists are always above 50%. This implies that the market is dominated by the chartists most of the time. We do not observe the dramatic switching of all the agents to either one of the trading strategies, which is the striking feature of switching models such as the Brock and Hommes and CHH ones.

3.4. Statistical Results

The upper portion of Table 2 reports some descriptive statistics of the simulated returns averaged across the set of all our simulations. The maximum daily return, for example, exceeds 9.4% for half of our simulations, whereas the smallest return

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

TABLE 2. Descriptive statistics of the returns of the simulations of the model (upper panel) and of Chiarella et al. (2006c) (lower panel)

	Max	Min	Sd	Skew	Kurt	Volume
Median	0.094	−0.115	0.016	−0.032	9.631	162.2
Mean	0.218	−0.247	0.021	−0.537	55.343	160.4
Median	0.050	−0.051	0.014	0.001	3.380	—
Mean	0.056	−0.056	0.014	−0.002	3.484	—

Note: We report the median and mean value (across 100 independent simulations) of the biggest and smallest returns, volatility, skewness and kurtosis of returns.

is smaller than −11.5% in 50% of the cases. Some mild negative asymmetry is confirmed by the negative skewness in the time series of returns.

Table 2 confirms again that there is a fair amount of excess volatility and excess kurtosis. In all the simulations, the volatility of the fundamental process was fixed at 0.5% daily and hence trading is responsible for the inflated standard deviation of price returns (about 1.6%). The number of exchanged units (volume) is typically close to 160 per day. For comparison, the same statistics for the Chiarella et al. (2006c) data in the lower panel reveal that returns are slightly less volatile, have much smaller extreme values, are nearly symmetric, and have little excess kurtosis.

3.5. Behavior of Autocorrelation

In order to show effectively that many simulated time series exhibit very weak linear predictability but strong volatility clustering, we summarize the statistical properties of end-of-day autocorrelations by taking averages across the simulations. We denote by $\hat{\rho}_{jk}$ and $\hat{\gamma}_{jk}$ the estimated j-lag autocorrelations of raw and absolute returns of the kth simulation, respectively. The distribution of the sampled values $\hat{\rho}_{jk}$ and $\hat{\gamma}_{jk}$, $k = 1, \ldots, 100$, of the autocorrelation at j lags can then be graphically condensed for all $j = 1, \ldots, 100$, as shown in the top part of Figure 4. The graphs depict, for each lag j, a simplified version of the box-and-whisker plot of $\hat{\rho}_{jk}$ and $\hat{\gamma}_{jk}$; see Becker et al. (1998) for details. In particular, all the medians indicated by thick horizontal lines, and most "boxes," that lie between the first and third quartile are well within a 2σ confidence band (dashed lines) for the $\hat{\rho}_{jk}$. Our data display a slight negative medium-range correlation of returns (within the confidence bands at the 5% level) and, despite some outliers at specific lags for some time series, the left part of the figure shows that about 50% of the simulations satisfy the strict requirement of having *all* autocorrelation coefficients within the band, thus pointing to the fact that the model is robustly generating extremely low linear predictability in returns.

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

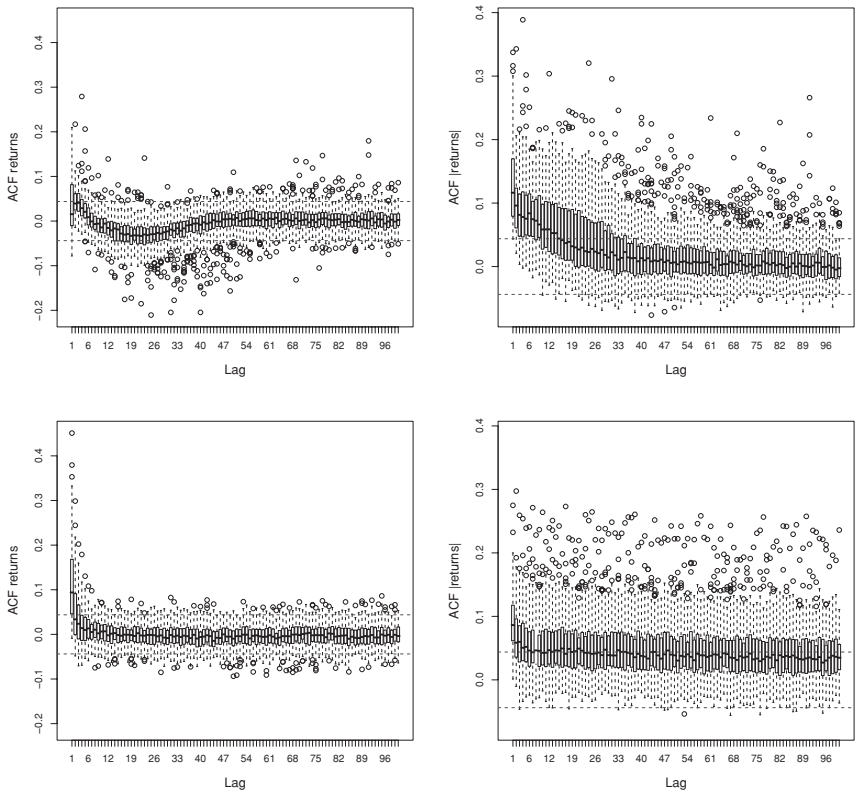


FIGURE 4. End-of-day autocorrelation of raw (left) and absolute returns (right) for the CDA (top) and the CHH (bottom) models. For each lag j on the horizontal axis we depict the box-and-whisker plot of the 100 coefficients $\hat{\rho}_{jk}$ and $\hat{\gamma}_{jk}, k = 1, \ldots, 100$. The median is shown by the thick line and the "boxes" depicts the interquartile range of the data. Outlying points extending beyond the "whiskers" are visible as single points (circles).

The top right graph of Figure 4 illustrates the distribution of the autocorrelations $\hat{\gamma}_{jk}$ of absolute returns. Evidence of volatility clustering is present when correlations exceed the upper dashed line. More than 50% of time series have significant heteroskedasticity for more than 15 lags and one quarter of the simulations strongly extend this features for 30 lags. Both the realizations depicted in Figure 2 are in this set. The reader should realize that we are not claiming that all our time series are equally realistic, and there are manifestly outlying "bad" cases in terms of autocorrelation in Figure 4. Instead, this aggregate examination is meant to validate a broad set of simulations showing that very often the results of the model are quite satisfactory.

The bottom part of Figure 4 shows the boxplots of autocorrelations of returns and absolute returns for 100 simulations of the Chiarella et al. (2006c) model.

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

The lower left panel of Figure 4 shows that the first few raw autocorrelations are significant and positive for a large number of series (more than 75% of times for the first lag and more than 25% of times for lags up to six). Hence, there is often residual predictability that is not found in empirical series. The lower right panel of Figure 4 shows many significant autocorrelations of absolute returns (which are, however, not uncorrelated, as just pointed out). The slow decay of autocorrelations lasts on average for many lags.[6]

To wrap things up, our model can deliver asymmetric and markedly leptokurtic returns, with virtually no linear predictability and slow decay of absolute autocorrelations for 20–30 lags; for comparison, the original CHH model, albeit able to display long memory in some cases, has symmetric returns with low excess kurtosis and several significant autocorrelations at first lags.

3.6. Noise and Sensitivity Analysis

In this section, we discuss the sensitivity to some key parameters, focusing in particular on the effects of the various sources of noise that operate in our setup.[7] We recall that randomness has a role in several components of the model. The first is the selection of the order (permutation) in which agents bid during a trading session. Second, the orders of the fundamentalists are randomly selected in the interval described by the fundamental value and the actual price. Third, chartists' bids or asks offset the price by a random amount with standard deviation Δ. Finally, the fundamental process changes are governed by σ_f. The model appears to be much more sensitive to the two last sources of noise and hence we limit our discussion to Δ and σ_f. We also try to give some impression of the impact of the fraction p_ϵ of noise traders.

Decreasing Δ reduces the intensity of the stylized facts. Setting $\Delta = 0.0025$, that is, halving the benchmark value, still produces uncorrelated, nonnormal, and leptokurtic returns, but only one-fourth of the simulations display significant autocorrelations of absolute returns for up to 10 lags. The median (mean) volatility of returns decreases to 0.01125 (0.01155) and the kurtosis drops to median (mean) values of 4.15 (4.42). Figure 5 depicts a typical time series when $\Delta = 0.0025$. Doubling Δ causes large and explosive swings in prices that break down the similarity to realistic time series.

Setting $\sigma_f = 0.25\%$, that is, reducing the changes in the fundamental price, results in smoother and somewhat pseudocyclic time series with pronounced autocorrelation and very mild traces of heteroskedasticity. The median (mean) volatility and kurtosis across 100 simulations with $\sigma_f = 0.25\%$ are 0.008955 (0.009061) and 3.87 (5.48). A typical time series in Figure 6 depicts only slightly nonnormal features. Similar to what happens for large Δ, increasing the size of σ_f to 1% daily produces excessive volatility and kurtosis, in the range 4–5% and 50–100, respectively.

To conclude this analysis, we investigate the importance of the proportion p_ϵ of noise traders in the model. As we argued before, they are needed to "move" the price when all other agents are on the same side of the market. One could ask

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

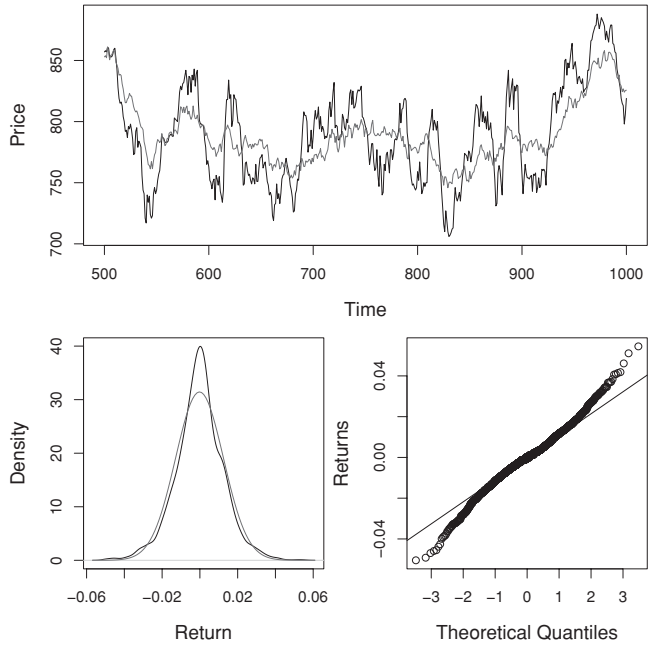


FIGURE 5. Price and fundamental value time series for one representative simulation obtained when $\Delta = 0.0025$. The density of returns (with a normal distribution superimposed) and a QQ-plot are also shown. The volatility and kurtosis of the depicted returns are 1.27% and 4.17, respectively.

about the extent of the statistical properties that can be generated by noise traders alone. Indeed, Maslov (2000) and LiCalzi and Pellizzari (2003) argue that the CDA can generate some stylized facts even with very simple behavior on the part of the agents. We run 100 simulations where only the fraction p_ϵ of noise traders submit orders (while the fundamentalists and chartists are inactive) to evaluate the incremental effect of the intertwined action of chartists and fundamentalists. The time series produced in these purely noisy markets are random walk–like and there is no connection between the fundamental value and the price of the stock. This is obvious, as no trader is using p_t^* to submit orders and, not surprisingly, the returns are white and there is no autocorrelation of absolute returns at any lag. The median (mean) volatility and kurtosis across all the simulations are 0.01214 (0.01211) and 4.480 (4.508). Hence, purely noisy markets exhibit insignificant correlations and some degree of leptokurtosis in returns, confirming the findings of the previously cited works, but there is no volatility clustering and smaller standard deviation. Moreover, no herding due to switching is detected, as active agents just bid or ask in an entirely random manner and the performance measure is erratic. All of this evidence indicates that the noise traders alone are not responsible for volatility clustering and the slowly decaying ACs of the absolute returns observed when

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

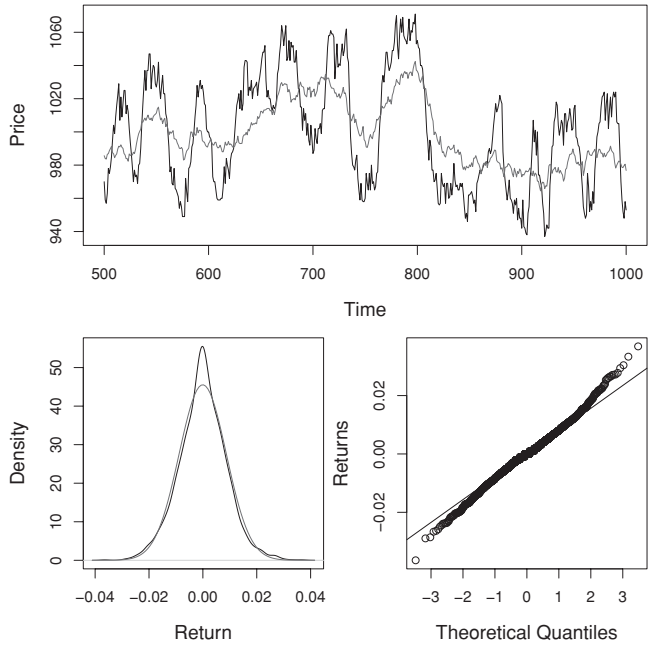


FIGURE 6. Price and fundamental value time series for one representative simulation obtained when $\sigma_f = 0.25\%$. The density of returns (with a normal distribution superimposed in red) and a QQ-plot are also shown. The volatility and kurtosis of the depicted returns for this realization are 0.88% and 3.73, respectively.

both the fundamentalists and chartists are active in the market. To summarize, our results are definitely richer than the ones that can be obtained with purely random traders, even if some stylized facts can still be obtained in a weak form in this case.

4. CONCLUSION

Inspired by the theoretically oriented CHH model under the market-maker clearing mechanism, this paper conducts a dynamic analysis of a microstructure model of continuous double auctions. The model removes some less realistic assumptions in the CHH model, including the fixed length of the moving average window, unlimited short selling and borrowing, trading of fractional shares, and no intraday trading, and in particular, the homogeneity among each type of fundamentalist and chartist. With a realistic market microstructure, the model is able not only to obtain the results of the CHH model but also to characterize many of the stylized facts, including slow decay of correlation of absolute returns. Our results seem to suggest that a comprehensive explanation of several statistical properties of the returns is possible in a framework where both *behavioral* traits and *realistic microstructure* have a role.

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

In this paper, we have not paid attention to the dynamics of the intraday price. Also, it is not clear whether and to what extent other clearing mechanisms than the CDA, such as automated dealerships or hybrid markets, would affect our results. We leave these issues for future research.

NOTES

1. Technically, they may submit no order in some circumstances, as they cannot sell short nor borrow. Allowing for cancellations of orders, resubmissions, multiple orders, or other features within a trading day would hugely complicate the model and would require much richer modeling of the agents' behavior. Moreover, we do not think that single-unit orders constitute a severe limitation for an agent willing to buy/sell multiple units, as he/she can keep trading for several days to reach the desired quantities. Chiarella, Iori, and Perellò (2009) provide a modeling framework of CDA markets where agents do form multiple unit-demand functions based on utility maximization.

2. Note that the form of the probability will be the same for each fundamentalist, but it will yield a different value for each of them, as the p_τ could possibly change at each instant, determined by the sequencing of the agents.

3. This mechanism is the reason that we do not model a ticked book; that is, we do not force limit prices to be on a discrete grid of finite step size δ. Sequences of improving orders can drive the price considerably up or down even if the δ is relatively small, whereas picking $\delta \to 0$ can basically lead to too little variation in prices and excessively peaked returns' distribution. Allowing a continuous limit price offers considerable flexibility, as it turns out that Δ can be tuned more easily than the "corresponding" discrete tick δ.

4. Slightly different orderings are possible with orders having the same limit price, based on the quantity or the time of submission. In this paper, as all orders have unit quantity, we resort to strict time preference to break ties.

5. Given that the length L_i of the moving average is at most $L = 100$ and we discard 500 observations, the initial randomization should have little or no effect.

6. One way to formally assess the presence of long memory is based on the computation of the Hurst exponent H for the returns. In the benchmark case of independent increments, the exponent takes the value $H = 1/2$, with values higher than $1/2$ pointing to some form of persistence. Estimated values are never higher than 0.4 for our simulated time series, whereas they exceed $1/2$ in 25 cases out of 100 for the Chiarella et al. (2006c) simulations. In close agreement with the visual evidence provided by Figure 4, long memory can be ruled out for our model, whereas it is present in one-fourth of the cases for the CHH model with added noise.

7. We have discussed previously the outcomes for other parameters, such as α and β. The initial amount of risky units and cash is also relevant. We feel that this is rather intuitive: for example, increasing the cash relative to the stock can fuel longer and more pronounced rises of prices, as agents can mechanically buy more stocks. In more realistic models, risk-related considerations are possibly limiting the positions, but our simple agents are "dumb" in this respect and prone to take extreme positions if they can.

REFERENCES

Becker, R.A., J.M. Chambers, and A.R. Wilks (1988) *The New S Language: A Programming Environment for Data Analysis and Graphics*. Pacific Grove, CA: Wadsworth and Brooks/Cole.

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

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

Chiarella, C., R. Dieci, and L. Gardini (2002) Speculative behaviour and complex asset price dynamics. *Journal of Economic Behavior and Organization* 49, 173–197.

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press

Chiarella, C., R. Dieci, and L. Gardini (2005) The dynamic interaction of speculation and diversification. *Applied Mathematical Finance* 12, 17–52.
Chiarella, C., R. Dieci, and X. He (2006a) Heterogeneous expectations and speculative behaviour in a dynamic multi-asset framework. *Journal of Economic Behavior and Organization* 62, 402–427.
Chiarella, C., R. Dieci, and X. He (2009a) Heterogeneity, market mechanisms and asset price dynamics. In T. Hens and K. Scheink-Hoppe (eds.), *Handbook of Financial Markets: Dynamics and Evolution*, pp. 277–344. Amsterdam: North-Holland.
Chiarella, C. and X. He (2001) Asset price and wealth dynamics under heterogeneous expectations. *Quantitative Finance* 1, 509–526.
Chiarella, C. and X. He (2002) Heterogeneous beliefs, risk and learning in a simple asset pricing model. *Computational Economics* 19, 95–132.
Chiarella, C. and X. He (2003) Heterogeneous beliefs, risk and learning in a simple asset pricing model with a market maker. *Macroeconomic Dynamics* 7, 503–536.
Chiarella, C., X. He, and C. Hommes (2006b) A dynamic analysis of moving average rules. *Journal of Economic Dynamics and Control* 30, 1729–1753.
Chiarella, C., X. He, and C. Hommes (2006c) Moving average rules as a source of market instability. *Physica A* 370, 12–17.
Chiarella, C. and G. Iori (2002) A simulation analysis of the microstructure of double auction markets. *Quantitative Finance* 2, 346–353.
Chiarella, C., G. Iori, and J. Perellò (2009b) The impact of heterogeneous trading rules on the limit order book and order flows. *Journal of Economic Dynamics and Control* 33, 525–537.
DeGrauwe, P. and M. Grimaldi (2006) Exchange rate puzzles. A tale of switching attractors. *European Economic Review* 50, 1–33.
Gaunersdorfer, A. (2000) Endogenous fluctuations in a simple asset pricing model with heterogeneous agents. *Journal of Economic Dynamics and Control* 24, 799–831.
Gode, D. and S. Sunder (1993) Allocative efficiency of markets with zero intelligence traders. *Journal of Political Economy* 101, 119–137.
He, X. and Y. Li (2007) Power law behaviour, heterogeneity, and trend chasing. *Journal of Economic Dynamics and Control* 31, 3396–3426.
He, X. and Y. Li (2008) Heterogeneity, convergence and autocorrelations. *Quantitative Finance* 8, 58–79.
Hommes, C. (2001) Financial markets as nonlinear adaptive evolutionary systems. *Quantitative Finance* 1, 149–167.
Hommes, C. (2002) Modeling the stylized facts in finance through simple nonlinear adaptive systems. *Proceedings of the National Academy of Science of the United States of America* 99, 7221–7228.
Hommes, C. (2006) Heterogeneous agent models in economics and finance. In L. Tesfatsion and K. Judd (eds.), *Handbook of Computational Economics*, pp. 1109–1186. Amsterdam: North-Holland.
LeBaron, B. (2006) Agent-based computational finance. In L. Tesfatsion and K. Judd (eds.), *Handbook of Computational Economics*, pp. 1187–1233. Amsterdam: North-Holland.
LiCalzi, M. and P. Pellizzari (2003) Fundamentalists clashing over the book: A study of order driven stock markets. *Quantitative Finance* 3, 470–480.
Lux, T. (2006) Financial Power Laws: Empirical Evidence, Models, and Mechanism. Economics Working Papers n. 2006.12, Department of Economics, Christian-Albrechts-University of Kiel.
Lux, T. (2009) Stochastic behavioural asset pricing and stylized facts. In T. Hens and K. Scheink-Hoppe (eds.), *Handbook of Financial Markets: Dynamics and Evolution*, pp. 161–215. Amsterdam: North-Holland.
Maslov, S. (2000) A simple model of an order-driven market. *Physica A* 278, 571–578.
Pagan, A. (1996) The econometrics of financial markets. *Journal of Empirical Finance* 3, 15–102.
Pellizzari, P. and F. Westerhoff (2009) Some effects of transaction taxes under different microstructures. *Journal of Economic Behavior and Organization* 72, 850–863.
Westerhoff, F. (2003) Speculative markets and the effectiveness of price limits. *Journal of Economic Dynamics and Control* 28, 439–508.
Westerhoff, F. (2004) Multiasset market dynamics. *Macroeconomic Dynamics* 8, 591–616.

https://doi.org/10.1017/S136510051000074X Published online by Cambridge University Press