



Laboratory of Economics and Management
Sant'Anna School of Advanced Studies

Piazza Martiri della Libertà, 33 - 56127 PISA (Italy)
Tel. +39-050-883-341 Fax +39-050-883-344
Email: lem@sssup.it Web Page: <http://www.sssup.it/~LEM/>

LEM

Working Paper Series

**Institutional Achitectures and Behavioural
Ecologies in the Dynamics of
Financial Markets:
a Preliminary Investigation**

*Giulio Bottazzi **

*Giovanni Dosi **

*Igor Rebesco **

** Sant'Anna School of Advanced Studies, Pisa, Italy*

2002/24

November 2002

Institutional Architectures and Behavioural Ecologies in the Dynamics of Financial Markets: a Preliminary Investigation*

Giulio Bottazzi Giovanni Dosi Igor Rebesco

S. Anna School of Advanced Studies, Pisa, Italy

November 5, 2002

Abstract

The paper compares the properties of market dynamics, under different trading protocols. At an empirical level, we present some evidence stemming from the comparison between different intra-daily trade regimes within the world largest Stock Exchanges. Such evidence also motivates the investigation of the properties of an agent-based model under three alternatives market mechanisms, namely a Walrasian auction, a batch auction and an ‘order-book’ double auction. The results highlight the importance of market mechanisms *per se*, even when holding constant the behavioural characteristics of the agents.

1 Introduction

In this work we explore the impact of different institutional structures governing financial market interactions upon market dynamics. Motivated also by some suggestive comparative evidence drawn from the world largest Stock Exchanges, we compare the statistical properties of three alternative market models, namely simplified versions of (i) Walrasian auctions; (ii) batch auction-type markets and (iii) ‘order-book’-type markets.

*Among the many insightful comments which helped in shaping the present draft, we would like to mention in particular those by Dooyne Farmer, Thorsten Hens, Blake LeBaron and Yi-Cheng Zhang. The usual disclaimers applies.

The roots of the investigation do ramify well beyond the confines of financial markets. They concern indeed one of the most controversial questions which the economic discipline has faced since its origins, that is: What determines the relatively orderly aggregate properties – if any – and the degrees of efficiency of market exchanges? Are they mainly due to what goes on in the agents minds, or, conversely, are they primarily the outcome of some organizing processes which market mechanisms themselves impose?

Ultimately, one may think of four basic interpretations and combinations thereof.

A *first* one emphasizes the purported equilibrating features of the fine understanding agents supposedly hold both of their environment (possibly including the strategies of other agents) and of the means to pursue their interests. Obviously, “rational expectations” are the extreme version, but – in much milder forms – the emphasis on the equilibrating (or disequilibrating) role of agents beliefs and behavioural rules dates back at least to Adam Smith’s *Theory of Moral Sentiments*.

A *second* almost symmetric opposite perspective points are the orderly properties of market selection – in the roughest version, of any market – rewarding successful behaviours and weeding out failing ones, irrespectively of the degrees of *ex ante* ‘rationality’ which individual agents display. Again, it is a conjecture with a respected pedigree, featuring in contemporary economics Milton Friedman’s famous “as if” argument and a few more sophisticated reassessments in evolutionary games frameworks¹.

A *third* major view focuses upon the properties of particular distributions of budget-constrained behaviours over heterogeneous, possibly ‘bounded rational’, populations. Prominent examples are – away from the financial arena – Werner Hildenbrand’s path-breaking investigations of the statistical conditions yielding ‘well behaved’ demand functions², and – nearer our concerns here – Gode and Sunder’s analysis of ‘zero-intelligence’ agents³.

Finally, a *fourth* perspective, which one could call the “grand evolutionary view”, conjectures that relatively orderly market dynamics are emergent properties stemming from far-from-equilibrium interactions amongst heterogeneous learning agents⁴.

¹The classic reference on the “as if” argument is Friedman [1953]. Many evolutionary analyses do involve explicit accounts of market selection processes, but reject claims that any selection in any type of market will necessarily yield aggregate ‘ordered’ properties and even less so optimal ones (more on these issues in Dosi and Winter [2002] and the references therein). More specifically on financial markets, see Blume and Easley [1992] and De Long *et al.* [1991] on the possibility of long-term survival of purportedly “irrational” behaviours.

²Hildenbrand [1994]; see also Aversi *et al.* [1999] for an attempt to link such an analysis with an underlying evolution of preferences.

³Gode and Sunder [1993]; see however also the critical remarks in Cliff and Bruten [1997].

⁴Such a ‘grand program’ is spelled out at greater detail in Dosi and Winter [2002].

Come as it may, institutions governing the *physics of exchanges* – including centralized *vs.* decentralized trading mechanisms, the frequency of trading, the rules for price formation and those prescribing who is trading with whom and when – are central to this latter interpretation, but are also likely to be important parts of the other three, except for their simplest versions. After all, market institutions shape also the information agent access, the processes of competition and selection, the mechanisms of aggregation and price formation, etc.

However, surprising as it sounds, not much work has gone into the study of the aggregate implications of different architectures of both financial and real markets⁵. As LeBaron ends his survey of agent-based models of financial markets, one of the major open issues ahead concerns the study of the properties of different trading set-ups (LeBaron [2000], p. 698). This is also the point of departure of this work which tries to identify some distinctive properties of diverse market mechanisms.

More precisely, one begin to address two challenging questions, namely,

- i*) what happen to market dynamics if one changes market interaction mechanisms, while holding constant individual characteristics (including the distribution of cognitive and behavioural patterns), and, conversely,
- ii*) holding constant institutional set-ups governing information diffusion and interaction patterns, what happens as one varies the “ecology” of behavioural types of agents?

We address this questions making use of an agent-based simulation environment, ‘The Financial Toy Room’ (FTR)⁶. On purpose we tackle the task by incremental steps. For the time being, we freeze all learning processes and we focus on the comparative properties of diverse institutional architectures, of trading mechanisms nesting different (budget-constrained) populations of ‘fundamentalist’ and ‘chartist’ agents.

In section 2 we present some novel evidence on the properties of daily time series under the two alternative market protocols which distinguish the opening and closing phases of major Stock Exchanges. While a few statistical ‘stylized facts’ hold across countries and across market regimes, finer properties are seemingly influenced by the latter. How can one

⁵Among the remarkable exceptions these are the studies by Alan Kirman and collaborators: see for example Kirman and Vignes [1991] on the fish market. Concerning financial markets, ‘microstructure’ studies (see Goodhart *et al.* [1996] and a few other contributions in Frankel *et al.* [1996]) certainly represent a major step in the right direction, although one still falls short of any explicit account of the dynamics of exchanges.

⁶The original version of the FTR – now available at the site <http://ftrsim.sssup.it> – was developed by Francesca Chiaromonte and collaborators at the International Institute of Applied System Analysis (IIASA), Laxenburg, Austria: cf. Chiaromonte and Dosi [1998].

model such dependencies of market dynamics upon institutional set-ups? As one discusses in section 3, multiple-agent “artificial market” models have not paid so far much attention to the issue. On the contrary, this is what we formalize there, based on the FTR simulation environment. Results (section 4) robustly support the importance of specific institutional arrangements, and, together, hint at a thread of often non-linear interactions between market institutions and behavioural ecologies which represent a puzzle on their own.

2 Generic *stylized facts* and institution-dependent phenomena

A good deal of current research on the statistical properties of financial time-series has gone – indeed for sound reasons of precedence – into the identification of robust, generic properties which appear to hold across markets and across different temporal windows of observation. As well known, such stylized facts include fat-tailed distributions of returns; ARCH effects; autocorrelation of volumes and cross-correlation volumes/volatility (for detailed discussions, cf. Brock [1997], Brock and de Lima [1995], Guillaume *et al.* [1997], Levy, Levy and Salomon [2000], Dacorogna *et al.* [2001]). However, quite a few studies, broadly in the ‘microstructure perspective’ have been also devoted to the degrees to which particular market organizations contribute to parameterize the foregoing generic properties and/or yield further institution-specific phenomena (cf. among others Amihud and Mendelson [1987], Stoll and Whaley [1990], Madhavan [1992], Biais *et al.* [1999] and the survey in Calamia [1999]).

One way of tackling the topic, which we share here, is by exploiting the fact that most Stock Exchanges daily undergo the transition between two diverse sets of market protocols. A first opening section is typically organized as a periodic batch auction in which orders are collected during a call period to form demand and supply schedules that are crossed to determine the unique equilibrium price (corresponding to the maximum executable volume) at which all transactions occurs. This is followed by a trading section characterized by continuous double auctions in which each agent can post bid and ask prices. In turn such continuous auctions may involve a special category of agents – the market makers – surrogating the auctioneer and making public either firm bid and ask prices (under the “quote-driven” system) or buy/sell intentions (under the “order-drive” system) and an order book in which limit orders are stored. The trading phase terminates with a closing section yielding the fixing, i.e. the closing price, often obtained by the weighted average of transaction prices over the last period of the trading section.

Given the significant difference in the architectures of exchanges between the opening and trading phases, their comparison might reveal precious albeit circumstantial information on their comparative properties.

So, for example, Amihud and Mendelson [1987] and Stoll and Whaley [1990] compare the open-to-open and close-to-close daily series on the NYSE highlighting higher volatility and negative autocorrelation of returns in the former.

Expanding on such a line of inquiry, here we compare the open-to-open and close-to-close daily series for “blue-chips” over the period 1/1/1997 - 14/4/2002 for the Stock Exchanges of Paris, Frankfurt, Milan, New York, Madrid, London, Tokyo and Toronto⁷.

In particular one is interested in comparatively assessing linear predictability; autocorrelation of price returns; skewness and kurtosis in the distribution of returns themselves;; cross-correlation volume/returns; ARCH effects and cross correlation returns/volatility⁸.

The evidence reveals that some properties do not discriminate between the two market regimes:

- Dickey-Fuller tests highlight the presence of a unit root in all price series and annual volatilities⁹ of returns are nearly the same across markets and trading phases;
- first-order autocorrelation in volumes are all significantly positive and quite high;
- positive cross-correlation volume/returns and volume/volatility are general phenomena (even if they appear somewhat more pronounced in the close-to-close series);
- kurtosis is much higher than that associated with Gaussian distribution, revealing fat tails in the return distribution of all series.

At the same time, two properties stand out as regime-related (see Fig. 1 and Fig. 2):

- opposite autocorrelation patterns (with the exception of Tokyo Stock Exchange) occur. Open-to-open price series display a negative autocorrelation while close-to-close series entail positive (although not statistically significant) autocorrelation¹⁰);

⁷Our series are simple averages over a subset of blue chips from CAC 40, DAX 30, MIB 30, DJIA 30, IBEX 35, UKX 100, TPXC 30, and TSE 35, after eliminating those stocks which have a life shorter than our period of observation.

⁸The full analysis can be obtained from authors upon request.

⁹This is measured as the standard deviation computed on daily date times $\sqrt{250}$.

¹⁰A priori, one should just expect a modest positive autocorrelation on all indexes due to the so-called “small cap effect” (cf. Brock [1997] and Campell *et al.* [1997]).




Figure 1: The autocorrelation of returns for open-to-open and close-to-close series.

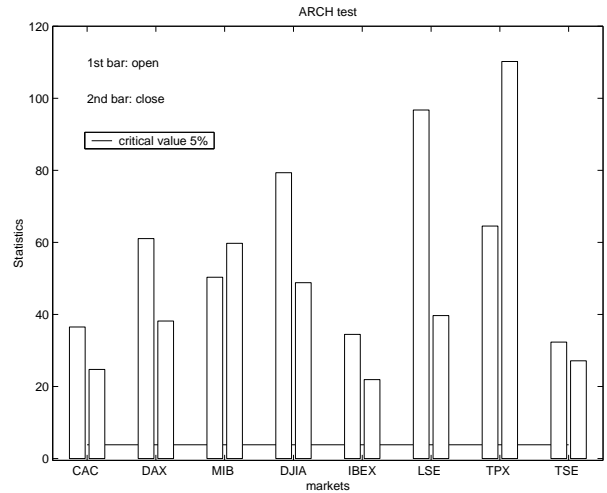


Figure 2: Testing for ARCH effects for open-to-open and close-to-close returns.

Table 1: Statistics from SETS of LSE.

Returns	Skewness	Kurtosis	Abs. deviation	Annualized volatility
Close-to-next-open	-2.51	37.7	0.0034	0.084
Open-to-close	0.46	8.4	0.006	0.13

- ARCH effects, while generally present appear to be more pronounced in opening prices (with the exception of the Italian and Japanese Stock Exchanges).

A complementary exercise concerns the impact of different markets phases within the same market. In order to investigate the phenomenon one has to consider close-to-next-open and open-to-close returns. Some problems arise with the latter statistics because the trading section is often a complex mix of ‘order book’ protocols and market maker intervention (sometimes interrupted by intra-day batch auctions) which is impossible to separate in their effects. In order to overcome this drawback we have singled out a segment of stocks traded in a pure ‘order book’ protocol¹¹: the Stock Exchange Electronic Trading Services (SETS) of the London Stock Exchange (LSE) of which we have analyzed daily data for 80 equities in the time window 1/1/1997 to 14/4/2002.

The statistics summarized in table 2 are sufficient to highlight the difference between the two phases.

Note in particular that the skewness of returns is significantly negative for the opening

¹¹That follows the usual batch auction opening the daily section.

phase, whereas it is not present in open-to-close returns.

We do not have any ready interpretation of these phenomena. For the time being, let us just consider them as adding to the circumstantial evidence on the impact of the forms of market organization upon market dynamics.

Clearly one of the paramount difficulties in disentangling such a relationship rest in the impossibility of making real world, historical, experiments changing market institutions and leaving the rest constant.

However, one way of partially overcoming such an obstacle is via laboratory experiments (for influential discussions of the role of market set-ups in experimental economics, cf. S. Sunder [1995], V. Smith [1982] and Plott and Sunder [1988] among others).

Another complementary way is through the “thought experiments” entailed in the comparison between computer-simulated “artificial markets”, diverse by construction in their market set-ups. This is what we shall do next.

3 Different market architectures and heterogeneous behaviours: an agent based model

The spirit of the model which follows is to a large extent akin that inspiring already existing computer-simulated “artificial financial markets”, such as those by Arthur *et al.* [1997], Beltratti and Margarita [1992], LeBaron [2001], Lux and Marchesi [1999] and Marengo and Tordjman [1996] (see the review in LeBaron [2000]).

Obvious common points of departure are (*i*) the acknowledgment of the limitations of models of market dynamics centered upon the behaviour of a mythical representative agent endowed with unbiased forward-looking expectations and, conversely, (*ii*) the challenge of founding the theory into an explicit account of heterogeneous interactive agents.

Within such a common perspective, however, differently families of models significantly differ in the ways they model both the agents behavioural repertoires and the mechanisms of interaction.

Concerning the former, one happens to find a whole range of modeling commitments going from relatively realistic stylizations of the trading rules of actual practitioners all the way to “zero intelligence” agents just subject to budget constraints¹².

¹²Palmer *et al.* (1994) is nearer to first extreme while Gode and Sunder (1993) are of course archetypes of the second one. Moreover, in a few models agents are allowed to endogenously differentiate through adaptive learning: cf. Arthur *et al.* (1997) and Lux and Marchesi (1999).

Regarding trading protocols one finds two basic modeling styles and variations thereof. A first one compresses collective interactions into “law” governing price responses to excess demands (cf. Farmer [2002]). Conversely a second family of models attempt to model explicit trading mechanisms. Remarkably, however, even those models taking this latter route have hardly addressed a systematic comparison of the properties of different mechanisms. This is the task of the present work.

Our model describe a population of N heterogeneous agents acting as speculative investors and participating in the discrete-time trading of two assets: a riskless bond B that pays a constant interest rate r at each time steps and a risky asset A paying a constant dividend d . The value of r and d are common knowledge and the price of the asset A in term of B is fixed by the market.

On the behavioural side we stick to a rather simple ecology of agents, just comprising “noisy fundamentalists” and trend following chartists. In that, our agents – unlike “zero intelligence ones” – do have some reaction algorithm to the dynamics of the environment, but – alike them – are prevented from adaptively refining their understanding of the market.

Concerning market architectures, we compare three institutional set-ups, namely,

- A **Walrasian auction** in which all agents transfer their whole demand curves to the auctioneer that matches them in order to clear the market.
- A **batch auction** in which each agents can simultaneously post a buy or a sell order. Demand and supply schedules are then derived and crossed in order to determine the equilibrium price at which all agents exchange.
- An **order book** in which agents can post both market and limit orders that are matched following a price priority.

3.1 The market model: trading protocols

In each round of the **Walrasian auction** each agent $i \in [1, N]$ is supposed to provide the auctioneer with its complete personal demand curve $\Delta A_{i,t}(p)$, i.e. the amount of the asset ΔA that it is willing to buy ($\Delta A > 0$) or sell ($\Delta A < 0$) for each possible price p . The auctioneer then compute the aggregate excess demand $\Delta A_t(p) = \sum_i \Delta A_{i,t}(p)$ and fixes the asset price p_t at the value that clears the market: $\Delta A_t(p_t) = 0$. Notice that, in general, the individual demand functions are time dependent as agents react to changing market conditions. However as long as personal demand curves are well-behaving decreasing functions, the existence and uniqueness of p_t is guaranteed.




Figure 3: An example of supply and demand schedules for a batch auction. Suppose that buy orders are all limit orders (the DD' line). Conversely sell orders are composed by both limit orders (the SS' , dotted line) and an amount M of market orders which shift upward total supply to the line SS'' . Correspondingly, market price shifts downward from P' to P .

This stylized exchange protocol provides a price fixing mechanism that forces the market to equilibrium at each iteration. This features makes it an excellent analytical benchmark. At the same time, since the amount of information that the agents must provide to the auctioneer is infinite, encompassing all the possible desired positions for all the possible prices, this can hardly represent a sound approximation to any real trading mechanism wherein the handling of (finite) information generally entails a non-zero cost, however small.

The others two trading protocols that we will consider are indeed based on the processing of a finite amount of information.

Consider a **batch auction**. In each round each agent provides a (finite) number of “orders” that can be considered as statements concerning the conditions for its participation to the market. The following two basic instances of an order market are analyzed, namely:

limit orders represented as ordered couples (p, q) of a price $p > 0$ and a quantity q . For “buy limit orders” ($q > 0$) the price p stands for the maximum price at which the order issuer is willing to buy the asset quantity q . For “sell limit orders” ($q < 0$) p is the minimum price at which the order issuer want to sell a quantity $-q$ of the asset.

market orders $(., q)$ express only a quantity that should be sold (“sell market orders”), if $q < 0$, or bought (“buy market orders”), if $q > 0$ at the best available price (the lower for buy orders and the higher for sell orders) on the market.

Once all the limit and market orders are collected, the auctioneer uses the former to build a demand and a supply schedules, computing the total amount of asset notionally demanded and offered at a given notional price. The market orders are “priced” at the price that is more likely to guarantee their fulfillment among the prices of the limit orders of the same side, that is the largest price for buy orders and the smallest for sell orders. Their quantities are simply added to the schedules obtained with limit orders, which are in this way shifted vertically: see Fig. 3 for an illustration.

The crossing point between the two curves so obtained identifies both the price p_t and the total traded quantity. The orders whose price p is consistent with p_t (i.e. buy orders with price $> p_t$ and sell order with $p < p_t$) are totally or partially fulfilled. To this aim the orders are considered in sequence, with higher (lower) prices and higher quantities taking precedence for buy (sell) orders. This rule guarantees that the market orders are the first to be fulfilled.

Notice that, when the largest buy price is less than the smaller sell price, a total traded quantity of 0 is obtained¹³.

The two protocols presented so far are both characterized by some price fixing mechanism that determines a unique price at which all the transactions take place. An alternative “physics of interaction” has to be devised in order to approach a market settings similar to the decentralized “continuous trading” phase of a stock exchanges, characterized by a continuous flow of arriving orders that can be immediately executed or stored for possible delayed execution. The name **book** comes actually from the list of unmatched orders that stands on the market and are “written” in a file in order to be possibly taken in consideration in future transactions.

In our implementation, during a **book** session, agents are sequentially selected at random and are requested to place a given amount of orders, both of limit and market types, to the market. The arriving orders are managed according to the following simple rules:

limit orders are immediately executed against orders of the opposite side of the book (i.e. buy side for sell orders and sell side for buy orders) if the price of the buy order is not lower than the price of the sell order. The execution price is the one associated with the order already in the book and the quantity is obviously the smaller between the two orders¹⁴.

¹³This event indicates that no trading can take place during this session so that the price is considered fixed at the value of the previous session and the next session begin.

¹⁴If the first matching does not completely fulfill the arriving order, this operation is continued until either:

i) no more suitable order are on the opposite side or

If there are not orders on the opposite side of the book or if the order is not completely fulfilled, it is stored on the “book” on the basis of a price/quantity/time priorities.

market orders are treated analogously, after assigning them a price equal to the best order on the opposite side. If there are no order on the opposite sides, a *reference price* is used instead.

The *reference price* is equal to the price of the last transaction or the fixing price of the last trading session if no transactions have yet take place. The fixing price of the session is the last reference price of the session.

3.2 The Agents

The implementation of the foregoing trading protocols basically induces just a single requirement concerning the way in which one models the behaviour of agents: they should be able to provide the auctioneer (or “the market”) with a well-behaved demand function or, alternatively, with limit/market orders. However, in order to compare market dynamics under different protocols, it is mandatory to model as much as possible agents behaviours as governed by the same rules and shaped by the same kind of information. Ultimately, this means that the way in which agents generate the orders to be posted in the batch auction and book protocols should be consistent with the demand function they transmit to the auctioneer in the Walrasian auction¹⁵.

3.2.1 The demand functions

Let us start with the description of a generic agent i ¹⁶.

ii) the arriving order is completely fulfilled

¹⁵Incidentally note that the requirement is far from trivial. For example, suppose that individual demand functions come from a utility maximization procedure. In such a framework each agent has an infinite number of market positions which lay on its demand curve. But then, how can one choose the finite number of orders which traders deliver in batch auctions or continuous trading? Through some heuristic procedure? But then how can these heuristics be extended to the generation of a complete demand function? Here one cannot tackle these intricate issues. Suffice to say that the heuristic rules agents are assumed here to employ in the batch auction and continuous trading can always be rationalized in terms of (myopic) wealth maximization conditional on the “models of the world” agents hold.

¹⁶For the time being the suffix will be dropped for convenience, but it should be clear that all agent-specific parameters can in principle be different within the population of agents. The specification of their values is provided below.

At the beginning of the session each agent construct its individual demand function, and determines the amount of wealth it wants to invest in the risky asset for any possible value of the hypothetical transaction price p . The basic idea is that the agent willingness to invest in the risky asset is based on its estimate of its own wealth in the next time step, i.e. on its forecast of the stock price at time $t + 1$. As customary such willingness increases with the expectations of returns in excess of the riskless rate and decreases with the risk involved in taking long positions on the asset. If W_A is the wealth the agent wants to have invested in the risky asset at the end of the trading session, one can then write

$$W_A(t) \sim \frac{\text{excess return}}{\text{risk}} \quad (1)$$

Suppose now that the price of the asset is p . If $E_{t-1}[h]$ is the agent expected return on asset price, tomorrow's wealth with a portfolio solely made of the asset is $p(1 + E_{t-1}[h]) + d$ where d are the constant dividends. Tomorrow's wealth of a portfolio of the same current value, but invested in the bond would be $p(1 + r)$. The agent evaluation of the excess return of asset portfolio can then be written $E_{t-1}[h] + d/p - r$. As an estimate of the risk involved in the speculative activity, the agent takes its expectation about variance of price returns $V_{t-1}[h(t)]$. Remembering that $W_A(t) = A(t)p$ where A is the amount of risky asset possessed by the agent, one obtains from (1) the demand curve

$$\Delta A(p, t) = -A(t-1) + \frac{E_{t-1}[h(t)] - r + d/p}{\beta V_{t-1}[h(t)] p} \quad (2)$$

with $\Delta A(t) = A(t) - A(t-1)$, where $A(t-1)$ stands for the quantity of risky asset at the end of the previous trading session, and β is a constant capturing the agent risk aversion. When $\beta \gg 1$ even relatively low uncertainty on the future asset price is sufficient to strongly decrease its attractiveness for the agent; for $\beta \ll 1$ the opposite holds ¹⁷.

We add to (2) the restriction that agents cannot hold short positions in terms of both assets and bonds. This means that $\Delta A(p, t) \geq -A(t-1)$ and if the offered quantity of asset resulting from (2) is greater then $-A(t-1)$, it is replaced with $-A(t-1)$. At the same time, the demanded quantity of asset at price p cannot be larger then B/p where B denotes the quantity of bonds possessed by the agent. The resulting demand function is in general not continuous but presents a monotonically decreasing behaviour. Its shape depend on sign of the difference $y - r$: illustrations are presented in Fig. 4 and Fig. 5.

¹⁷Notice that the demand function in (2) is basically grounded on a mean-variance evaluation of portfolio consistent with a Markowitz-type approach to asset pricing. It cannot however be obtained, in general terms, via an expected utility maximization procedure.




Figure 4: An example of personal demand function for $y < r$.




Figure 5: An example of personal demand function for $y > r$.

We assume a demand function, consistent with agent budget-constraints, that depends only on the parameter β , measuring the agent risk aversion, and on the agent forecast about future asset prices, captured by the two expectations $E_{t-1}[h(t)]$ and $V_{t-1}[h(t)]$. For sake of simplicity, we suppose that agents obtain their forecasts using exponentially weighted moving averages (EWMA) predictors. The recursive expression for the expected returns and variance then becomes:

$$\begin{aligned} E_t[h(t+1)] &= (1-\lambda) \sum_{\tau=0} \lambda^\tau h(t-\tau) \\ V_t[h(t+1)] &= (1-\lambda) \sum_{\tau=0} \lambda^\tau h(t-\tau)^2 - E_{t-1}[h]^2 \end{aligned} \quad (3)$$

where $E_{t-1}[\cdot]$ is the agent forecast of returns $h(t) \equiv p(t+1)/p(t) - 1$ for period $t+1$; $V_{t-1}[\cdot]$ is the estimation of the risk (referred to period $t+1$) and $\lambda \in [0, 1]$ is a weighting coefficient setting the “time scale” on which the averaging procedure is performed. Note that setting $\lambda = 1$ simply turns off the updating algorithm: initial values are retained forever¹⁸.

3.2.2 The generation of orders

Whatever the form of the demand function, to be able to participate to trading under batch auction and book protocols, agents must be able to express single orders of both market and limit types. To obtain that we follow a simple procedure: when an agent is required to provide an order, it picks at random a price that is not too far from the last price of the asset p_{t-1} and then decide the associated quantity on the ground of its demand function. The support of this distribution is decided by the distance from the last price of the asset and the agent’s p^* i.e. the price at which its present portfolio would face no rebalancing. The exact procedure depends on the side of the order, i.e. whether of the buy or sell type.

For a **sell order** the order price p is obtained from

¹⁸This feature will be exploited below to model a particular class of agents.

$$\log(p) = \log(p^*) + \epsilon |\log(p_{t-1}) - \log(p^*)| \quad (4)$$

where ϵ is random variable drawn from a uniform distribution in $[0, 1]$. The associated quantity is derived from the demand function $\delta A(p)$. Clearly, according to (4) as $p \geq p^*$ the quantity is negative, as it should be for a sell order. If $\epsilon < \eta$, with $\eta \in [0, 1]$ a parameter specific of the agent, the supplied order is a market order $(., \delta A(p))$; otherwise it is a limit order $(p, \delta A(p))$. The parameter η defines the agent propensity to submit market orders: if $\eta = 1$ all the orders are market orders while if $\eta = 0$ all the orders are of the limit type.

Analogously, for a **buy order** the price is obtained by

$$\log(p) = \log(p^*) - \epsilon |\log(p_{t-1}) - \log(p^*)| \quad (5)$$

where ϵ has the same meaning as above. Notice that the buy order can in principle be issued by agents having no share of asset, i.e. $A = 0$. In this case p^* does not exist and the following rule is used instead

$$\log(p) = \log(p_{t-1}) - \epsilon |\log(p_{t-2}) - \log(p_{t-1})| \quad (6)$$

The subsequent determination of the order quantity and the decision whether to submit a limit or a market order is done following the same rules as of the sell-type order.

Notice that, according to this whole procedure, the generated limit orders $(p, \Delta A(p))$ are points always laying on the agent demand function. Moreover, the prescription that agents take in consideration the last price of the asset to generate prices for their orders is consistent with the idea of maximizing the probability of the order to be fulfilled.

3.3 Behavioural heterogeneity

To sum up, the behaviour of any generic agent in our model can be completely specified with only three parameters: the risk aversion β , the time horizon λ on which it evaluates past asset performances in order to build its own forecasts about price movement, and its propensity η to submit market orders instead of limit orders (this last parameter being meaningful only in markets characterized by a batch auction or order book protocol).

In this framework different forms of inter-agent heterogeneity can be easily obtained, as we shall do below, by varying these parameters within the population.

However, in order to check the consistency of the posting rules described above let us first consider the particular case with homogeneous agents, i.e. agents with the same value for

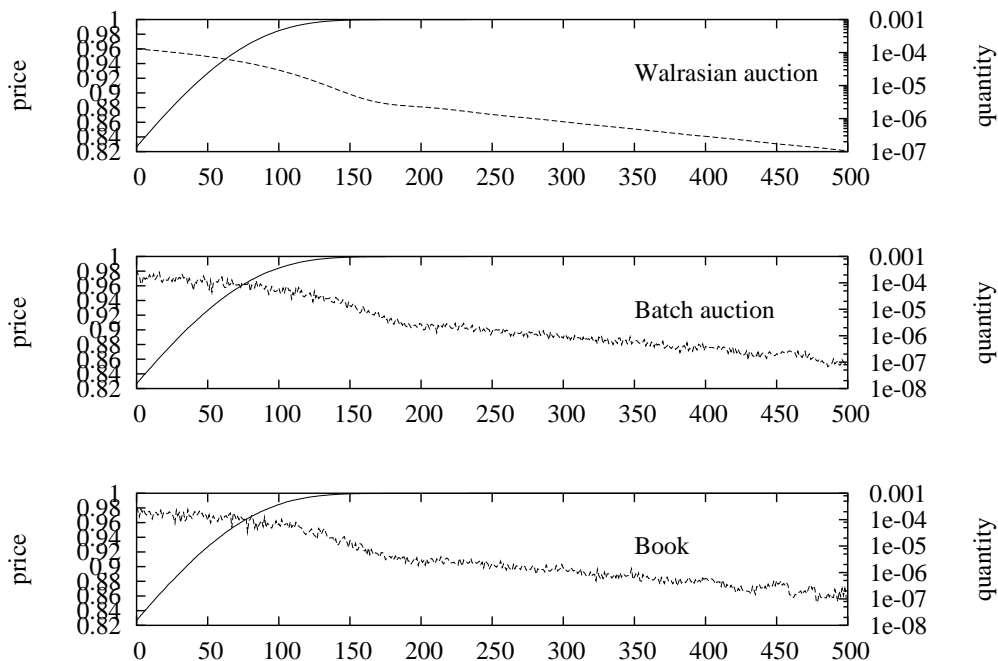


Figure 6: Price and quantity behaviour with homogeneous agents for the three market protocols. The model parameter are $r = .01$, $d = .01$, $\lambda = .991$, $\beta = 100$ and $\eta = .1$ and agents are initialized with heterogenous forecasts for $E_{t-1}[h(t)]$ and $V_{t-1}[h(t)]$.

the parameters β , λ and η , differing only in the initial conditions concerning their estimated forecast $E_{t-1}[h]$ and $V_{t-1}[h]$. For suitable choices of the parameters the system possesses a unique stable equilibrium with the asset price given by $p = d/r$. It is immediate to verify that this price clear the market when $E_{t-1}[h] = 0$, $V_{t-1}[h] = 0$ in (2): this represents a fixed point for the market dynamics (for a complete discussion about the stability of this fixed point, see Bottazzi (2002)). As can be seen in figure 6, all the trading protocols produce very similar results: the price tends asymptotically toward the equilibrium price while the traded quantity goes to zero (the batch and order book protocols generate, as expected, some fluctuations in traded quantities and prices whose amplitude tends however to zero when $t \rightarrow \infty$). The same kind of result for the homogeneous agents case is also obtained for choices of the parameters that lead the system to periodic or aperiodic cycles (c.f. Bottazzi (2002)). In all cases, the dynamics under the three protocols is almost identical (with some deviation essentially due to the possibility to have zero trading for some time steps in the batch auction and order book protocols).

The central question then becomes: what happens when one introduces heterogeneity?

There are two sources of heterogeneity in our model. A first one is “intrinsic” to the

agents in so far as they differ from the start in their parameters β , γ and η . A second one is endogenous to market dynamics and concerns idiosyncratic shocks on the estimates agents make.

In the following analysis we will use both forms of heterogeneity. In particular we introduce two distinct groups of agents. A first group is formed by **noisy fundamentalists**. They are characterized by $\lambda = 1$ and by a dynamic noise in their forecasted excess return: hence they do not update their expectations according to market dynamics, but extract them from a given distribution independently at each round. A second group, which can be referred to as **trend following chartists**, are parameterized with a $\lambda = 0.97$ ¹⁹. They are basically “naive econometricians” who obtain their forecasts from last market trends. Notice also that the simple fact that these agents also constantly update their evaluation of risk, destroys the possibility for them to generate a “rational bubble” dynamics of prices, even if they are the only agents in the market (for a discussion of this point see Bottazzi (2002)).

Concerning endogenous heterogeneity, we add idiosyncratic shocks to agents forecasted price returns extracted independently from a uniform distribution of support $[-.01, .01]$ for trend followers and $[-0.5, 0.5]$ for noisy fundamentalists²⁰.

Finally we experiment with different values of η (0, .1, .2, .3, .4, .5), homogeneous over the population²¹.

In our simulations we have experimented, within every market structure, with different ecologies of types of agents with the following proportions between ‘noisy fundamentalists’ and ‘chartists’: 0/100, 25/75, 50/50, 75/25, 100/0.

4 Simulation results

In line with the empirical evidence discussed in section 2 we study the statistics concerning skewness, kurtosis, autocorrelation of returns and autocorrelation of volumes.

We compute statistics using 2000 steps of simulation, after discarding the first 2000 in order to avoid possible transient effects²².

¹⁹This value is near to the one actually used by practitioners as discussed in RiskMetrics Technical Manual.

²⁰We set the value of $\beta = 10^6$ for the whole population. This seemingly very large number is due to the fact that for convenience we normalize the total amount of asset to 1 while the dynamic depends on the product of β and the total number of assets.

²¹Recall that the parameter captures the propensity to deliver market orders: for values of η higher than .5 an overwhelming number of market orders tends to reduce market liquidity, possibly leading to the absence of trade for very long periods.

²²We checked that the results obtained are independent with respect to the random number generator seed.




Figure 7: Skewness for the batch and the order book returns. The statistics are computed with different population fractions of noisy fundamentalists and with different values of η .

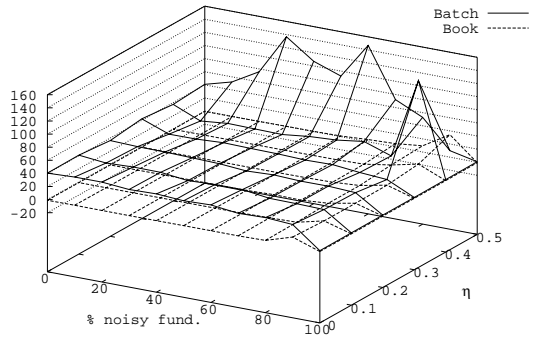


Figure 8: Excess kurtosis for the batch and the order book returns. The statistics are computed with different population fractions of noisy fundamentalists and with different values of η .

It is remarkable that the simulation results, notwithstanding the utter simplicity of the model, most often display comparative properties in tune with those of actual financial markets, both in terms of sign and even in the (rough) orders of magnitude.

Consider first the skewness in the distribution of returns, shown in Fig. 7. It is, rather surprisingly, very low in the ‘book’ protocol – irrespectively of both the proportions in the types of agents and the market order propensity – and it becomes just slightly negative when the market is populated almost only by noisy fundamentalists. On the contrary, in the batch auction the skewness is significantly negative and its value seems to depend on the population composition but not on the parameter η . In fact it increases together with the proportion of noisy fundamentalists and converge to zero when only noisy traders are present. In this case, as expected, the distribution of returns tends toward a Gaussian one.

Kurtosis in return distribution is higher in batch auction than in the order book trading mechanism, as shown in Fig. 8. In the latter case, it is almost independent from the different forms of heterogeneity (except when only noisy fundamentalists are present). Conversely in the batch auction it is very sensitive to η (showing an explosive behaviour when $\eta \sim 0.5$), but not to proportion of agents types.

Of course, there is no way to match our ecologies of behaviours with empirical, unobserved, ones. However, inter-institutional comparisons batch auctions vs. order books concerning skewness and kurtosis in our simulations are in line with the evidence from actual data (see table 2).

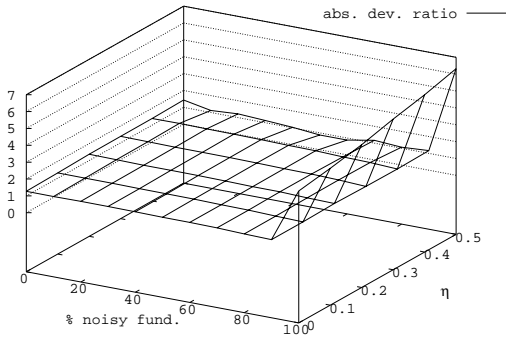


Figure 9: Ratio of returns absolute deviations in the batch and order book protocols. The statistics are computed for different proportions of agents types and for different values of η .




Figure 10: Excess kurtosis and skewness for Walrasian auction for different shares of agents types. They converge to 0 when the market is populated only by noisy fundamentalist.

In the case of Walrasian auction, statistics confirm the results obtained for the batch auction (see figure 10): the excess kurtosis is around 50 and the skewness is negative (about -8). As expected, when the market is comprised of only noisy traders these two statistics converge towards Gaussian values.

Returns autocorrelations are shown in table 2. They display a very short memory in both batch auctions and order book markets (5 and 3 lags respectively), qualitatively in line with empirical markets (although not in terms of orders of magnitude: our simulations yield a far too high one-step autocorrelation). Moreover it is striking that the one lag negative autocorrelation for the order book seems independent from both the ecology of agents and the types of order posted. The same does not apply to the batch auctions that are apparently very sensitive to both sources of heterogeneity.

Autocorrelations of volumes for the batch auction are reported in table 3. They show a rather long memory, significant until ~ 20 time lags. No data are reported for the order book where autocorrelations never move away significantly from 0. Again heterogeneity seems matter in the former trading mechanisms, but not in the latter.

Finally the ratio of absolute deviations between order books and batch auctions reported in Fig. 9 is around 1.4, meaning that the former produces a more volatile dynamics than the latter ²³.

²³The absolute deviation is used for its robustness in measuring the distribution width when skewness and kurtosis are large.

% noisy	η	Batch auction					Book				
		1	2	3	4	5	1	2	3	4	5
0	0.0	0.840	0.592	0.333	0.154	0.054	-0.252	-0.124	0.030	-0.009	-0.033
0	0.1	0.841	0.591	0.337	0.152	0.048	-0.321	-0.082	0.021	-0.001	0.021
0	0.2	0.839	0.587	0.326	0.146	0.050	-0.381	-0.074	0.077	-0.082	0.090
0	0.3	0.853	0.653	0.428	0.239	0.107	-0.341	-0.082	0.003	0.017	0.008
0	0.4	0.852	0.651	0.425	0.243	0.119	-0.321	-0.081	0.025	0.002	0.002
0	0.5	0.823	0.622	0.413	0.247	0.129	-0.260	-0.098	0.000	0.029	-0.033
50	0.0	-0.589	0.124	-0.033	-0.009	0.003	-0.763	0.423	-0.235	0.104	-0.043
50	0.1	0.816	0.615	0.380	0.189	0.062	-0.315	-0.027	0.037	-0.058	0.006
50	0.2	0.813	0.621	0.383	0.186	0.067	-0.336	-0.070	0.059	-0.040	-0.023
50	0.3	0.811	0.665	0.464	0.266	0.119	-0.322	-0.031	-0.014	-0.029	0.010
50	0.4	0.767	0.678	0.523	0.348	0.202	-0.306	-0.095	0.032	0.007	0.006
50	0.5	-0.076	-0.092	-0.094	-0.023	-0.054	-0.236	-0.082	-0.006	0.001	0.001
100	0.0	0.816	0.615	0.380	0.189	0.062	-0.315	-0.027	0.037	-0.058	0.006
100	0.1	-0.529	0.010	0.034	-0.003	-0.018	-0.778	0.466	-0.319	0.213	-0.136
100	0.2	-0.578	0.113	-0.052	0.007	0.026	-0.767	0.430	-0.257	0.149	-0.099
100	0.3	-0.539	0.044	0.024	-0.044	0.024	-0.745	0.383	-0.209	0.116	-0.070
100	0.4	-0.484	-0.021	0.001	0.014	-0.030	-0.694	0.275	-0.108	0.045	-0.027
100	0.5	-0.419	-0.071	-0.012	0.005	-0.003	-0.624	0.163	-0.068	0.055	-0.045

Table 2: Returns autocorrelation coefficients for the Batch auction and the book protocols. The first 5 time lags are shown. Higher lags never yield significant correlation coefficients. Bold figures are statistically significant over two standard deviations.

5 Conclusions

The institutional arrangements governing trading mechanisms, we have shown, bear significant influences upon price and quantity dynamics. Comparative empirical evidence hint at this property. And this is robustly confirmed by inter-institutional comparisons based on our multi-agent model of ‘artificial’ financial markets. Moreover, while each market mechanism exerts its distinct impact on market dynamics, the actual properties of the latter stem from the interactions between market set-ups and the ecologies of behaviours within the population of agents.

Given these non-linear interactions between institutional structures and behavioural ecologies a puzzling issue regards the normative properties of different set-ups. For example, can one identify robust criteria for “efficiency” in such heterogeneous worlds? This is indeed one major research task on the agenda ahead.

References

- [1] AMIHUD Y., MENDELSON H. (1987). “Trading mechanisms and stock returns: An em-

%	η	1	10	20	30	40	50	60	70	80	90	100
0	0.0	0.841	0.535	0.194	-0.069	-0.260	-0.362	-0.386	-0.355	-0.245	-0.081	0.133
0	0.1	0.841	0.577	0.250	-0.021	-0.223	-0.363	-0.428	-0.426	-0.350	-0.204	-0.022
0	0.2	0.851	0.641	0.329	0.053	-0.157	-0.307	-0.381	-0.402	-0.381	-0.288	-0.165
0	0.3	0.740	0.572	0.255	-0.021	-0.202	-0.304	-0.336	-0.299	-0.182	-0.009	0.157
0	0.4	0.739	0.531	0.236	-0.025	-0.216	-0.347	-0.391	-0.369	-0.271	-0.111	0.075
0	0.5	0.772	0.418	0.209	0.038	-0.100	-0.190	-0.233	-0.228	-0.196	-0.094	0.013
50	0.0	0.685	0.447	0.213	-0.005	-0.124	-0.222	-0.242	-0.214	-0.107	-0.017	0.063
50	0.1	0.691	0.497	0.217	-0.024	-0.192	-0.272	-0.286	-0.236	-0.126	0.006	0.072
50	0.2	0.689	0.530	0.209	-0.027	-0.207	-0.306	-0.328	-0.267	-0.195	-0.070	0.050
50	0.3	0.686	0.344	0.144	-0.024	-0.131	-0.201	-0.211	-0.192	-0.125	-0.044	0.049
50	0.4	0.759	0.107	0.023	-0.059	-0.096	-0.126	-0.130	-0.107	-0.079	0.009	0.028
50	0.5	0.604	0.467	0.258	0.130	0.033	-0.039	-0.069	-0.091	-0.099	-0.100	-0.103
100	0.0	0.214	-0.033	-0.027	0.001	0.009	0.024	-0.000	0.010	-0.009	-0.084	-0.014
100	0.1	0.136	0.003	0.017	0.065	0.000	-0.031	0.036	0.024	-0.026	0.028	-0.010
100	0.2	0.170	0.004	0.010	0.044	-0.008	-0.022	-0.030	0.025	-0.025	0.020	0.010
100	0.3	0.142	-0.002	-0.031	0.022	-0.028	-0.021	0.035	0.009	0.001	-0.020	0.000
100	0.4	0.125	0.000	-0.026	-0.008	-0.012	-0.028	-0.006	0.003	-0.018	-0.020	0.018
100	0.5	0.035	0.041	0.009	-0.003	-0.042	-0.054	-0.037	0.017	-0.043	-0.012	0.004

Table 3: Volume autocorrelation coefficients for the Batch auction for different fractions of noisy fundamentalist (in percentage) and different values of η (market order propensities). Bold figures are statistically significant over two standard deviations

pirical investigation”, *Journal of Finance*, 42, 533-553.

- [2] ARTHUR W. B., DURLAUF S. N., LANE D. A. (1997). *The Economy as an Evolving Complex System II*, Addison-Wesley.
- [3] ARTHUR W. B., HOLLAND J. H., LEBARON B., PALMER R., TAYLER P. (1997). Asset Pricing under Endogenous Expectations in an Artificial Stock Market, in ARTHUR W. B. *et al.* (1997).
- [4] AVERSI R., DOSI G., FAGIOLO G., MEACCI M., OLIVETTI C. (1999). Demand Dynamics with Socially Evolving Preferences, *Industrial and Corporate Change*, 8, 353-408.
- [5] BELTRATTI A., MARGARITA S. (1992). Simulating an Artificial Adaptive Stock Market, Mimeo, Turin University.
- [6] BIAIS B., HILLION P., SPATT C. (1999). Price discovery and Learning During Preopening Period in the Paris Bourse, *Journal of Political Economy*, 107, 6, 1218-1248.
- [7] BLUME L., EASLEY D. (1992). Evolution and Market Behaviour, *Journal of Economic Theory*, 58, 9-40.

- [8] BOTTAZZI G. (2002). A Simple Micro-Model of Market Dynamics, *LEM Working Paper*, Sant'Anna School of Advanced Studies, forthcoming in *WHEIA 2002 Proceedings*.
- [9] BROCK W. A. (1997). Asset Price Behavior in Complex Environment, in ARTHUR W. B. *et al.* (1997).
- [10] BROCK W. A., DE LIMA P. (1995). Nonlinear Time Series, Complexity Theory and Finance, in *Handbook of Statistics 12: Finance*, edited by G. S. MADDALA, RAO H. AND VINOD H., North Holland, Amsterdam.
- [11] BROCK W.A., HOMMES C.H., WAGENER F.O.O. (2001). Evolutionary Dynamics in Financial Markets with Many Traders Types, Working Paper, University of Wisconsin, Madison.
- [12] CALAMIA A. (1999). Market Microstructure: Theory and Empirics, *LEM Working Paper*, Sant'Anna School of Advanced Studies.
- [13] CAMPBELL J. Y., LO A. W., MACKINLEY A. C. (1997). *The Econometrics of Financial Markets*, Princeton University Press, Princeton.
- [14] CHIAROMONTE F., DOSI G. (1998). Modelling a Decentralized Asset Market: an Introduction to "Financial Toy-Room", Interim report #98015, IIASA, Laxemburg.
- [15] CHIAROMONTE F., DOSI G., JONARD N. (1999). A Speculative Centralized Asset Market with Adaptive Heterogenous Agents, *LEM Working Paper*, S. Anna School of Advanced Studies.
- [16] CLIFF D., BRUTEN J. (1997). Zero is not enough: On the Lower Limit of Agent Intelligence for Continuous Double Auction Markets, Hewlett-Packard Laboratories, Bristol.
- [17] DACOROGNA M. M., GENÇAY R, MÜLLER U., OLSEN R. B., PICTET O. P. (2001). *An Introduction to High-Frequency Finance*, Academic Press, New York.
- [18] DANIELS M. G., FARMER D. J., IORI G., SMITH E. (2001). "Dynamic storage, market liquidity and price volatility", Santa Fe Institute Working Paper.
- [19] DE LONG J. B., SHLEIFER A., SUMMERS L., WALDMANN R. (1991). The Survival of Noise Traders in Financial Markets, *Journal of Business*, 64, 1-20.

- [20] DOSI G., WINTER S. G. (2002). Interpreting economic change: evolution, structures and games, in *The economics of Choice, Change and Organization* edited by AUGIER M. and MARCH J. G., Edward Elgar, Northampton.
- [21] FARMER J. D. (2002). Market Force, Ecology and Evolution, *Industrial and Corporate Change*, forthcoming.
- [22] FRANKEL J., GALLI G., GIOVANNINI A. (1996). *The Microstructure of Foreign Exchange Markets*, University of Chicago Press, Chicago.
- [23] FRIEDMAN M. (1953). *Essays in positive economics*, University of Chicago Press, Chicago.
- [24] GODE D. K., SUNDER S. (1993). Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality, *The Journal of Political Economy*, 101, 119-137.
- [25] GOODHART C., FIGLIUOLI L. (1991). Every Minute Count in Financial Markets, *Journal of International Money and Finance*, 10, 23-52.
- [26] GOODHART C., O'HARA M. (1995). High Frequency Data in Financial Markets: Issues and Applications, Unpublished Manuscript, London School of Economics.
- [27] GOODHART C., ITO T., PAYNE R. (1996). One day in June 1993: a study of the workings of Reuters D2000-2 electronic foreign exchange trading system, in FRANKEL *et al.* (1996).
- [28] HENS T., SCHENK-HOPPÉ K. R. (2001). Evolution of Portfolio Rules in Incomplete Markets, Working Paper, University of Zurich.
- [29] HILDENBRAND W. (1994). *Market demand*, Princeton University Press, Princeton.
- [30] KIRMAN A., VIGNES A. (1991). Price dispersion. Theoretical considerations and empirical evidence from the Marseille fish market, in *Issues in Contemporary Economics*, edited by ARROW K. J., Macmillan, London.
- [31] LANE D. (1993). Artificial Wolds in Economics, Part I and II, *Journal of Evolutionary Economics*, 3, 89-107 and 177-197.
- [32] LEBARON B. (1999). Building Financial Markets with Artificial Agents: Desired Goals and Present Techniques, Brandeis University Working Paper.

- [33] LEBARON B. (2000). Agent based computational finance: suggested readings and early research, *Journal of Economic Dynamics and Control*, 24, 679-702.
- [34] LEBARON B. (2000). Empirical Regularities from Interacting Long and Short Memory Investors in an Agent Based Stock Market, Brandeis University Working Paper.
- [35] LEBARON B. (2001). A Builder's Guide to Agent Based Financial Markets, *Quantitative Finance*, 1(2), 254-261.
- [36] LEVY H., LEVY M., SOLOMON S. (1994). Microscopic Model of the Stock Market, *Economics Letters*, 45, 103-111.
- [37] LEVY H., LEVY M., SOLOMON S. (2000). *Microscopic Simulation of Financial Markets*, Academic Press, New York.
- [38] LUX T., MARCHESI M. (1999). Volatility Clustering in Financial Markets: A Micro-Simulation of Interacting Agents, *International Journal of Theoretical and Applied Finance*, 3(4), 675-702.
- [39] LUX T., MARCHESI M. (1999). Scaling and Criticality in a Stochastic Multi-Agent Model of a Financial Market, *Nature*, Vol. 397, 498-500.
- [40] MADHAVAN A. (1992). Trading Mechanisms in Securities Markets, *Journal of Finance*, 47, 607-642.
- [41] MARENGO L., TORDJMAN H. (1996). Speculation, Heterogeneity and Learning: A Simulated Model of Exchange Rate Dynamics, *Kyklos*, 49, 407-437.
- [42] PALMER R., ARTHUR W. B., HOLLAND J. H., LEBARON B., TAYLER P. (1994). Artificial Economic Life: a Simple Model of a Stock Market, *Physica D* 75, 264-274.
- [43] PLOTT C. R., SUNDER S. (1988). Rational expectations and the aggregation of diverse information in laboratory securities markets, *Econometrica*, 56, 1085-1118.
- [44] PRING M. J. (1995). *Technical Analysis Explained*, McGraw Hill, New York.
- [45] *RiskMetrics Technical Document* (1981), JPMorgan.
- [46] SUNDER S. (1995). Experimental Asset Market: A Survey, in *The handbook of experimental economics* edited by KAGEL J. H., ROTH A. E., 445-495.

- [47] SMITH V. L. (1982). Microeconomic Systems as an Experimental Science, *The American Economic Review*, 72(5), 923-955.
- [48] STOLL H., WHALEY R. (1990). Stock market structure and volatility, *Review of Financial Studies*, 3, 37-71.