Traders' long-run wealth in an artificial financial market

Marco Raberto [a,1], Silvano Cincotti [a], Sergio M. Focardi [b], Michele Marchesi [c]

[a]Dipartimento di Ingegneria Biofisica ed Elettronica, Università di Genova, Via Opera Pia 11a, 16145 Genova, Italy

[b]The Intertek Group, 96 rue de Javel, 75015 Paris, France

[c]Dipartimento di Ingegneria Elettrica ed Elettronica, Università di Cagliari, Piazza D'Armi, 09123 Cagliari, Italy

[1]Corresponding author. E-mail: raberto@dibe.unige.it Tel: +39 010 3532797 Fax: +39 010 3532290

Abstract:: In this paper, we study the long-run wealth distribution regarding different trading strategies in the framework of the Genoa Artificial Stock Market. The Genoa market is an agent-based simulated market able to reproduce the main stylised facts observed in financial markets, i.e., the fat-tailed distribution of returns and the volatility clustering. Various populations of traders have been introduced in a "thermal bath" made by random traders, who makes random buy and sell decisions that are constrained by the available limited resources and depend on the past price volatility. We study both trend follower and trend contrarian behaviour; fundamentalist traders, i.e., traders believing that stocks have a fundamental price depending on factors external to the market are also investigated. Results show that the population which prevail and which lose cannot be decided on the basis of the strategy alone. Trading strategies yield different results in different market conditions. Generally, in a closed market, i.e., a market with no money creation process, we found that trend followers lose relevance and money to other populations of traders and eventually disappear; whereas in an open market, i.e., a market with money inflows, trend followers can prevail and make the behaviour of other traders' populations less and less profitable.

Keywords: artificial financial markets, market simulations, wealth distribution, trading strategies, trading behaviour, asset prices, econophysics.

1. Introduction

Understanding if there are winning and losing market strategies and determine their characteristics is an important question from the point of view of investors and regulators alike. On one side, it seems obvious that different investors exhibit different investing behaviour which is, at least partially, responsible for the time evolution of market prices. On the other side, it is difficult to reconcile the regular functioning of financial markets with the coexistence of different populations of investors. If there is a consistently winning market strategy than it is reasonable to assume that the losing populations disappear in the long run.

It was Friedman who first advanced in 1953 the hypothesis that in the long run irrational investors cannot survive as they tend to lose wealth and disappear. Offering an operational definition of rational investors, however, presents conceptual difficulties as all investors are boundedly rational. No agent can realistically claim to have the kind of supernatural knowledge needed to formulate rational expectations. The fact that different populations of agents with different strategies prone to forecast errors can coexist in the long run is a fact that still requires an explanation.

In this paper, we will first describe the Genoa Artificial Stock Market (GASM) and introduce the general setting of our computational experiments. We will then present experimental results on the price processes and the wealth distribution of agents as regarding their trading behaviour, so offering a possible interpretation of results.

2. The Genoa Artificial Stock Market

The Genoa Artificial Stock Market is a computational laboratory conceived to offer a simulated experimental facility with realistic trading features.

In the last decade, an increasing number of agent-based computer-simulated artificial financial markets have been proposed, for a review see LeBaron 2000 and Levy, Levy, Solomon 2000. A special mention is devoted to the pioneering work done at the Santa Fe Institute [Palmer et. al. 1994, LeBaron et al. 1999]. Generally speaking, the goal of building artificial financial markets is to explain the emergence of the characteristic statistical properties of asset prices on the basis of hypotheses on traders' behaviour, market microstructure and economic environment. These problems are usually too complex to be treated analytically, so the methods of microscopic simulations are employed.

We build an agent-based computer simulator of an artificial financial market and called it the Genoa Artificial Stock Market (GASM[1]). While, in the past the attention of researchers has been mainly focuses on modelling the agents' optimisation and learning capabilities, little effort has been devoted in our opinion in studying how the market microstructure and the macroeconomic environment impact market prices. The GASM has been conceived to address these problems. Indeed, in this paper we study the interplay of different trading strategies in a changing market environment. Results show that a trading strategy cannot be judged only on the basis of the strategy itself, but its success depends also on the market conditions.

An early release of GASM has been already presented in Raberto et al. 2001. Indeed, the GASM simulator has been conceived as an evolving system, able to be continuously modified and updated. It has been implemented using object-oriented technology and extreme programming [Beck 1999, Succi & Marchesi 2001] as development process. These technologies make possible to develop complex systems and to make substantial modifications very quickly without jeopardizing quality.

We present here the current release of GASM. A number of features deserve special mention:

- The trading mechanism of the GASM is based on a realistic auction-type order matching mechanism that allows to define a demand-supply schedule;
- Agents have only limited financial resources;
- The number of agents engaged in trading at each moment is a small fraction of the total number of agents.

The trading mechanism that will be described in detail in the next section permits to explicitly construct a demand-supply schedule. This is an essential feature since price fluctuations are due to an imbalance between demand and supply. As market must clear, this means that somehow the

[1] The name is devoted to the beautiful city where most of this work was performed. Moreover, in the Middle Age, Genoa was a major financial centre, where the I.o.u. and the derivatives were inventes, see Briys & De Varenne 2000.

"intentions" to buy or sell must be modified to allow orders to match. This is the essence of the demand-supply schedule. As demonstrated in many studies on market microstructure, e.g. see O'Hara 1995, the details of the order matching process have a bearing on both price setting and price-volume relationships.

The finiteness of agents' financial resources is another essential feature of the GASM. It is a realistic feature of real markets. It poses significant constraints on possible trading strategies if different populations of agents are to coexist indefinitely. Actually, the study of the interplay between different trading populations is one of the major scientific objectives of the GASM.

One of the features of real markets is that prices are set by transactions that involve only a small fraction of the market population. The notion that the entire population of investors is continuously engaged in trading is simply unrealistic as trading cost would skyrocket. Even professional fund managers tend to limit their trading activity to only a few trades per day and this only to fine tune portfolios that remain substantially stable for periods of the order of weeks.

The above implies that the consideration of the "thermodynamic limit" of markets, i.e. an infinite number of traders, is simply unrealistic and possibly misleading. Finite size effects in real markets are not an artifact but a real feature. The interaction of a small set of agents sets the "wealth" of the entire market.

Let's now consider the trading strategies implemented in the GASM. Let's first observe that the notion that financial markets are purely speculative markets is somewhat misleading. A significant fraction of market trading is done for real needs and not only for speculation. We claim that, in the long run, is the interplay between the flow of cash in and out financial markets and the creation/destruction of stocks that determine the price trends of the market.

Throughout this work we assume that there is only one security in the market plus cash. At every instant, agents hold a fraction of their wealth in cash and the remaining part in stock. The assumption of our first experimental setting is that there is no net inflow or outflow of cash in the market. At every trading period we assume also that only a small fraction of traders (2% of the total market), randomly selected, is engaged in trading. This is a simple representation of a background trading.

This background trading has only two structural elements: a volatility feedback and the wealth distribution of agents. Though trading is random and trading strategies are random too (see next section for details) we allow the market to be in different states of volatility. Volatility is uncertainty. The state of uncertainty of the market is modelled by the order limit prices issued by traders. We impose that in periods of high volatility traders are more nervous and therefore allow for wider price limits in order to get their trades done quickly. The assumption is that trades have to be done for exogenous reasons and therefore traders want to execute trades at the best possible price. Fearing large market movements, they allow more freedom in the setting of limit prices.

The other important structural element of the GASM is the agents' wealth distribution. It is well known empirically that wealth distribution tends to follow a Pareto inverse power law [Pareto 1897, Levy & Solomon 1997]. It has been demonstrated theoretically that auto catalytic processes naturally lead to inverse power law distribution of wealth [Montroll & Shlesinger 1982, Bouchaud & Mezrd 2000, Huang & Solomon 2000 (a), Huang & Solomon 2000 (b)]. The wealth of agents in the GASM is indeed governed by an auto catalytic process which explains the emergence and conservation of inverse power law distributions.

In the presented model, each trader is endowed with an initial wealth which follows a Zipf law and spanning four orders of magnitude. We observe that this distribution is conserved throughout the simulation. Moreover, the differences in wealth among traders may yield big imbalances in trading orders, thus giving rise to big price movements and consequently to a return distribution with fat tails, as empirically observed in real markets.

Using this initial experimental setting of the GASM, we then add populations of agents characterized by specific trading strategies. In the first instance we add two populations of trend followers and of fundamentalists, and we then proceed to add contrarian strategies.

3 Traders' populations

At every discrete trading moment, h, a generic trader, i, holds an amount $c_i(h)$ of cash and an amount $a_i(h)$ of the stock. Traders are segmented into four population-types depending on their respective trading behaviour: random, momentum, contrarian and fundamentalist traders.

Random traders are characterized by very simple trading strategies: no intelligence and random trading constrained by limited resources and past volatility. Though very simple, these ingredients are sufficient to build an artificial financial market, see Raberto et al. 2001, which is able to reproduce the main stylised facts of real markets: volatility clustering and fat tails in the distribution of price returns.

In this paper, we introduce, in the framework of the GASM, three population-types which have been already described in the literature and which represent more realistic trading behaviours. The aim is twofold: first, we want to study the behaviour of these stylised populations in a realistic environment characterized by limited resources and a market clearing mechanism; second, we want to address the important issue about the existence or not of winning strategies.

3.1 Random traders

At each simulation step, each random trader issues an order with probability equal to 2%. The order can be a buy or a sell order with probability 50%.

Suppose the i-th trader issues a sell order at time h+1. We assume that the quantity of stock offered for sale by the i-th trader at time h+1, a_i^s, is a random fraction of the quantity of stock he or she owned at time h. A limit price s_i is associated to the sell order, and we stipulate that the order cannot be executed at prices below s_i. The limit price s_i is computed as follows: $s_i = p(h)/N_i(m, s_i)$ where $N_i(m, s_i)$ is a random draw from a Gaussian distribution with average $m = 1.01$ and standard deviation s_i.

The value of s_i is proportional to the historical volatility $s(T_i)$ of the asset price through the equation $s_i = k * [s(T_i)]^q$, where k is a constant and $s(T_i)$ is the standard deviation of log-price returns [Raberto et al. 1999], calculated in a time window T_i time steps long. Tuning the system by performing numerous simulations with different parameters, we found that the values for k and q, which give the best fit with real world market data, are respectively 2.2 and 1.1. T_i is randomly drawn from a uniform distribution of integers in the range from 10 to 100 for each trader at the beginning of the simulation. As highlighted in section 2, linking limit orders to volatility takes into account a realistic aspect of trading psychology: when volatility is high, uncertainty as to the "true" price of a stock grows and traders place orders with a broader distribution of limit prices.

Buy orders are generated in a fairly symmetrical way with respect to sell orders. If the i-th trader issues a buy order at time $h+1$, the amount of cash employed in the buy order is a fraction of his or her available cash at time h.

A limit price b_i is associated to each buy order and we stipulate that buy orders cannot be executed at prices higher than the limit price.

In the case of buy orders, limit prices are computed as follows: $b_i = p(h) * N(m, s_i)$, where $N(m, s_i)$ is a random draw from a Gaussian distribution with average m and standard deviation s_i. As for sell orders. $m = 1.01$ and $s_i = k * [s(T_i)]^q$, where $k = 2.2$ and $q = 1.1$. T_i varies randomly from 10 to 100 time steps for each trader. The quantity of assets ordered, a_i^b, is therefore given by the integer part of the ratio between the amount of cash employed in the buy order and the limit price b_i.

It is worth noting that, as $m = 1.01$, the mean value of all b_i is likely to be greater than $p(h)$, while the mean value of all s_i is likely to be smaller than $p(h)$. In other words, we introduce a spread between the average value of buy/sell orders to represent the fact that a trader placing an order wants to increase the chance of the order being executed. Hence, for a buy order the trader is likely to be willing to pay slightly more than $p(h)$; conversely for a sell order, the trader is likely to offer the stock at a lower price than $p(h)$.

In conclusion, each random trader exhibits heterogeneous random behaviour subject to two constraints: the historical price volatility (which is included in s_i) and the finiteness of the resources available to him.

Random traders represent the bulk of traders who perform trades for reasons not linked to the market status, but to their needs. For instance, they represent investors who need to sell stocks to pay personal expenses, or people who got a compensation and decided to invest this money in the stock market. For this reason, their trades are completely random.

3.2 Momentum traders

The momentum trader is a trend follower who makes decision depending on the trend of past prices. The momentum trader speculates that if prices are raising, they will keep raising, and if prices are decreasing, they will keep decreasing. They represent, in a simplified way, traders following technical analysis rules, and in general traders following a herd behavior. In the literature, they are often referred as noise or chartist trader, see Black 1987, De Long et al. 1990, De Long et al. 1991, Lux 1997, Lux & Marchesi 1999.

At each simulation step, each momentum trader places an order with probability 2%. The order is a buy order if the past prices trend is positive and it is a sell order if the trend is negative.

To be more precise, at time step $h+1$, the i-th momentum trader computes the trend of past prices in the following way: $trend(h+1,T_i)=[p(h)-p(h-T_i)]/T_i$. T_i is the time window used in the calculation of trends, i.e., the number of time steps backward from the last simulated price. If $trend(h+1,T_i)>0$, the trader issue a buy order whose limit price b_i is given by $b_i = p(h)+trend(h+1,T_i)*T_i$. Conversely, if $trend(h+1,T_i)<0$, the order is a sell order and its limit price s_i is: $s_i = p(h)+trend(h+1,T_i)*T_i$. Therefore, in the case of buy (sell) orders, we have: $b_i > p(h)$ ($s_i < p(h)$) and the difference $b_i - p(h)$ ($p(h)-s_i$) is greater when the past trend is sharper. This is to say that when prices are increasing (decreasing) in a fast way, the momentum trader forecasts an identical increase (decrease) in the future.

A time window T_i is assigned to each trader at the beginning of the simulation through a random draw from a uniform distribution of integers in the range from 10 to 50. If T_i would be the same for all traders, momentum traders would behave in the same way.

In the case of a buy order, the trader's order size is a random fraction of his or her available cash and the quantity of stock demanded, a_i^b, is given by the ratio between the cash employed in the order and the value of b_i.

If the trader issues a sell order, the quantity of stocks offered for sale, a_i^s, is a random fraction of the amount of stocks owned.

3.3 Contrarian traders

The contrarian traders are structured in the same way as the momentum traders. The difference consists in their trading behaviour. A contrarian trader speculate that, if price is raising, it will stop raising soon, and will decrease, so it is better to sell near the maximum, and vice versa. This trading behavior is present among traders, albeit probably less popular than trend-following strategies.

In a closed market, the rationale of a contrarian strategy is that, when prices steadily increases, the total value of stocks exceeds the total value of cash in the market. This leads to an imbalance which will eventually bring prices down again. The opposite happens when the price steadily decreases. It is worth noting that both momentum and contrarian traders compute the trend on a time interval randomly chosen between 10 and 50 time steps, so they don't work on immediate price variations, but on steadier trends.

In detail, the contrarian behaviour is the following. If $trend(h+1,T_i)>0$, the contrarian trader issues a sell order whose limit price s_i is determined by: $s_i = p(h)-trend(h+1,T_i)*T_i$. If $trend(h+1,T_i)<0$, the order is a buy order and its limit price b_i is the following: $b_i = p(h)-trend(h+1,T_i)*T_i$. Therefore, even in this case we have that $b_i > p(h)$ and $s_i < p(h)$, and the chance of executing the order is higher than for orders at the past price.

The distribution of parameter T_i and the mechanics of order placement is the same of momentum traders, reported in the previous section. At each simulation step, also each contrarian places an order with probability 2%.

3.4 Fundamentalist traders

The fundamentalist trader believes that stocks have a *fundamental price*, due to factors external to the market, like price-earning, profitability, ROI, ROE and economic characteristics of the firm business. They believe that, in the long run, the price of the stock will revert to its fundamental price. Consequently, they sell stocks if the price is higher than fundamental price, and buy stocks in the opposite case.

In real markets, fundamentalists represent traders who do not follow short-term speculative behaviors, but try to make steady gains filtering out speculative movements of the market. In this model, fundamentalists at each time step "decide" whether to trade with a probability of 2%.

If a fundamentalist decides to trade, he will place a buy or a sell order , if the present price is lower or higher than the fundamental price, p_f , respectively. The fundamental price is the same for all fundamentalists.

The order limit price is always p_f , while its amount, in stocks for sell orders and in cash for buy orders, is computed randomly as a fraction of the present amount of stocks or cash owned by the trader, in the same way as for momentum and contrarian traders.

4 Market clearing

The price process is determined at the intersection of the demand and the supply curves. We compute the two curves at the time step $h+1$ as follows. Suppose that at time $h+1$ traders have issued U buy orders and V sell orders. For each buy order, let the pair $(a_u^b, b_u), u = 1,2,...U$ indicate respectively the quantity of stocks to buy and the associated limit price. For each sell order in the same time step, let the pair $(a_v^s, s_v), v = 1,2,...V$ denote respectively the quantity of stocks to sell and the associated limit price. Let us define the functions:

$$f_{h+1}(p) = \sum_{u|b_u \geq p} a_u^b \;; \tag{1}$$

$$g_{h+1}(p) = \sum_{v|s_v \leq p} a_v^s \;; \tag{2}$$

$f_{h+1}(p)$ represents the total amount of stocks that would be bought at price p (demand curve). It is a decreasing step function of p , i.e., the bigger p , the fewer the buy orders that can be satisfied. If p is greater than the maximum value of $b_u, u = 1,2,...U$, then $f_{h+1}(p) = 0$. If p is lower than the minimum value of $b_u, u = 1,2,...U$, then $f_{h+1}(p)$ is the sum of all stocks to buy. Conversely, $g_{h+1}(p)$ represents the total amount of stocks that would be sold at price p (supply curve) and is an increasing step function of p . Its properties are symmetric with respect to those of $f_{h+1}(p)$.

The clearing price computed by the system is the price p^* at which the two functions cross. We define the new market price at time step $h+1$, $p(h+1)$ as $p(h+1) = p^*$. Buy and sell orders with limit prices compatible with p^* are executed. Following transactions, traders' cash and portfolio are updated. Orders that do not match the clearing price are discarded.

5 Market initialisation

At the beginning of the simulation ($h = 0$), the price $p(0)$ is set at $ 100.00 and each trader i is endowed with a certain amount of cash $c_i(0)$ and a certain amount of stocks $a_i(0)$. The distribution of $c_i(0)$ and $a_i(0)$ follow a Zipf law. A random variable X is said to follow a Zipf law if: $P(X \geq x) = 1/x$, for $x \geq a$ where a is an appropriate positive constant. To create this type of distribution we used the ranking property of the Zipf law. This property states that in any sample randomly drawn from a Zipf law, the size of each item is approximately inversely proportional to its rank. This means that if we order the samples in decreasing order, the second sample is half the size of the first, the third one third, the n^{th} sample is $1/n^{th}$ of the first.

We therefore artificially created populations of agents whose wealth at time step 0 follows the ranking property of the Zipf law.

The population of random traders is made of 10,000 individuals and their cash and stock portfolio are initialised as follows: $c_i(0) = 10{,}000{,}000/i$ and $a_i(0) = 100{,}000/i$ for $i = 1,...,10{,}000$. As there are 10.000 agents in the market, the range of wealth spans four orders of magnitude.

The populations of momentum, contrarian and fundamentalist traders are made of 500 traders each. In this case, we have: $c_i(0) = 500{,}000/i$ and $a_i(0) = 5{,}000/i$ $i = 1,...,500$.

For the all four populations, the poorest trader has 10 stocks and $ 1,000 of cash, while the richest random trader is far wealthier than the richest momentum, contrarian or fundamentalist trader. Moreover, the number of random traders is twenty times the number of traders belonging to the other three populations. The reasons of these choices is that we want random traders to act as a "thermal bath" where particular trading strategies can be studied with a minimum impact on the stability of the system and on the statistical properties of price, which are very similar to those of real markets.

Traders's portfolios initialisation follows approximately[2] the identity: $c_i(0) \cong p(0) * a_i(0) \ \ \forall i$. According to that, the aggregate amount of cash at beginning of the simulation $C(0) = \sum_i c_i(0)$ equates approximately the aggregate value of stocks: $p(0) * A(0)$ where $A(0) = \sum_i a_i(0)$.

The equality between the aggregate value of cash and the aggregate value of stocks, i.e., $C(h) \cong p(h) * A(h)$ holds on average during the entire simulation. This is due to the order matching process that is responsible for setting the stock price.

[2] We say "approximately" due to rounding

Observe that, in the closed market, the total value of cash and the total number of stocks are constant during the simulation, i.e. $C(h) \equiv C(0)$ and $A(h) \equiv A(0)$. Let $\overline{p}$ be the price at which the aggregate value of stocks equals the total value of cash: $\overline{p} = C(0)/A(0)$. Then, the total value of stocks at time h depends only on the current price, $p(h)$.

Moreover, considering only random traders, who are the majority and are responsible for background trading, the trading mechanism used implies that, at every time step, the average number of stocks available for selling orders and the average amount of cash available for buying orders are constant.

Suppose that at time h the price is higher than $\overline{p}$: $p(h) > \overline{p}$. In this case, the aggregate value of stocks to sell is higher than the value of cash available to buy them. As a consequence, average sell orders for stocks will be larger than average buy orders. The order matching mechanism will therefore tend to reduce the stock price to match orders (on average). The reverse holds if the price is higher than $\overline{p}$.

The price will therefore oscillate around $\overline{p}$. Simulations confirm this mean-reverting behavior. We have initialized traders' wealth so that $p(0) = 100{,}00$ is the equilibrium price for the closed market.

In the case of external inflows or outflows of cash or stocks (open market), the equilibrium price varies and at every h can be determined as follows: $\overline{p}(h) \approx C(h)/A(h)$.

For the reasons said above, we set the fundamental price used by fundamentalists, p_f, to $\overline{p}$. This is the most plausible and unbiased value that a trader adverse to speculative behavior can give to the fundamental price in our closed market.

6 Simulation results and interpretation

In this section we describe the results of the numerical experiments performed on the GASM. First we explored the market behaviour when only random traders are present and there is no inflow or outflow of cash We then added different populations of traders. Finally we simulated a progressive increase in the amount of cash available to traders. For all these different market conditions, we explored the evolution of the distribution of wealth of the populations involved.

6.1 Closed market with random traders only

We let the market run with a population of random traders for 10,000 time steps. To give an idea of the timing involved, assuming that a time step is of the order of an hour, 10.000 time steps roughly correspond to 5 years of market activity.

In agreement with our previous work [Raberto et al, 2001], the price process exhibits the main features of real markets: fat tails of return distribution, zero autocorrelation of prices, slow decay of the autocorrelation function of the absolute values of logarithmic returns.

Figure 1 shows the cumulative distribution of standardized logarithmic returns $R(h)$, i.e., logarithmic returns $r(h) = \log p(h) - \log p(h-1)$ detrended by their mean and rescaled by their standard deviation. For comparison, the solid line represents the cumulative distribution of the standard normal distribution $N(0,1)$. It is possible to observe a clear deviation from the Gaussian behaviour with approximate power law scaling in the tail. A log-log regression of points which satisfy the condition $|R| > 2$ gives the slope: -3.40 ± 0.01 in agreement with values found in various financial daily time series [Mantegna & Stanley 2000, Bouchaud & Potters 2000]. This result is lower than the value: -3.69 ± 0.02 found in an early release of GASM [Raberto et al, 2001]. In that work, the leptokurtic shape of the log-returns distribution was originated both by an agent aggregation mechanism, similar to the herding model by Cont and Bouchaud 2000, and a functional dependence of limit prices on past prices volatility. In this new release of GASM, we have not included the herding phenomena. Fat tails are due to both the link between market volatility and limit prices and the distribution of wealth between agents which follows a Pareto law with exponent equal to 1. The dependence of limit prices on market volatility is a sort of microscopic implementation of the GARCH model, which notoriously can yield fat-tailed distributions [Roman et al, 2001]. The Zipf law type initial distribution of wealth is a realistic assumption [Levy and Solomon 1997], and contributes to the leptokurtosis of the returns distribution. As there are 10.000 agents in the market, the range of wealth spans four orders of magnitude. Because only 2% of traders act in the market at the same time, when a rich trader "decides" to issue an order, there is a low probability that another rich trader is chosen by the system to make an opposite order of the same size. This fact can cause big price variations.

In Figure 2, we present the autocorrelation function $C(t)$ of the absolute returns $|r|$ and of the raw returns r at different time lags t. While the autocorrelation of raw returns decays immediately, the autocorrelation of the absolute value of returns shows the presence of long-range correlations. Both an exponential and a power law decay have been tested for the descend curve of absolute returns. The exponential fit give an exponent equal to: $(0.20 \pm 0.03)10^{-2}$ with 0.088 as prediction error[3]. The value found in the power law fit is: 0.10 ± 0.01 with a prediction error equal to 0.043. While in the early release of GASM [Raberto et al. 2001] the exponential fit was more satisfactory, here the power law fit is found to perform better, in agreement with real markets [Liu et al 1999]. The reason lies in the memory of random traders that varies in a range between 10 and 100, so allowing the interplay between many time scales.

In Figure 3, we show the distribution of wealth among traders. The wealth distribution is presented at $h = 0$ and after 10,000 steps. The wealth $w_i(h)$ of trader i at time step h is defined as follows: $w_i(h) = c_i(h) + p(h) * a_i(h)$. It is worth noting that, though the global amount of cash and the total number of stocks are constant in a close market, the global wealth $W(h) = \sum_i w_i(h) = C(0) + p(h) * A(0)$ is a time-varying quantity which is obviously a linear function of the stock price. In Figure 3, the solid line represents the values of initialisation of wealth for $h = 0$, i.e., $w_i(0)$ versus the agent index i. Let us note the characteristic straight line in a log-log plot representing the Zipf law. Dots presents the values $w_i = c_i(10{,}000) + p(0) * a_i(10{,}000)$. The final value of wealth would be: $w_i(10{,}000) = c_i(10{,}000) + p(10{,}000) * a_i(10{,}000)$. We used $p(0)$ instead of $p(10{,}000)$ in order to have a direct comparison between the initial and the final distribution of wealth between agents, without the bias of the difference between $p(0)$ and $p(10{,}000)$.

[3] The prediction error of a linear fit is the standard deviation of the difference between the observed value and the theoretical value given by the result of the fit, see Wild & Seber 2000.

The main result of Figure 3 is that the Zipf law wealth distribution is conserved after 10,000 steps. This conservation is an important computational finding which is in agreement with the theoretical analysis of autocatalytic processes, see Montroll & Schlesinger 1983, Huang & Solomon 2001.

6.2 Closed market with four populations of traders

Using a population of random traders as a "thermal bath" we then added fundamentalist, momentum and contrarian traders, and studied the behavior of the wealth distribution of these populations under the assumption of a closed market without any inflow or outflow of cash. We added 500 fundamentalist, 500 momentum and 500 contrarian traders in a "thermal bath" made by 10,000 random traders. The ratio between the number of traders and their initial wealth permits the maintenance of the "thermal bath". Momentum, fundamentalist and contrarian trader are not able to influence the price process and the market produces prices characterised by the same statistical properties already discussed in the previous section.

As regarding the average dynamics of wealth, the main result is that on average momentum traders lose wealth, while fundamentalists and contrarian gain wealth. This result is independent from the separate or contemporary presence of single populations, as expected under the approximation of "thermal bath". It is worth noting that the wealth of random traders remain approximately the same during the simulation.

The distribution of wealth of single populations continues to follow approximately a Zipf law which is shifted downward or upward depending on the gaining or loosing of wealth.

In Figure 4, we present this result in the case of a typical simulation 10,000 steps long where all the four populations are present. The figure shows the average wealth $\sum_i w_i(h)/\sum_i$ of all traders i belonging to a single population. The wealth is computed at the current price $p(h)$ and is plotted versus the simulation time h.

A qualitative, heuristic explanation of these results is the following. The population of fundamentalist wins and gain wealth because their strategy is exactly in line with the mean reverting behavior of the market. We recall that the fundamental price, p_f, is set to the mean reverting price, $\overline{p}$. Fundamentalists always sell stocks at prices higher than $\overline{p}$ and buy at prices lower than $\overline{p}$. Given the trading mechanism and given that they issue rare orders randomly situated in the price path, the mean reverting behavior of stocks make them earn a profit on all trades.

The success of the strategy to sell if $p(h) > \overline{p}$ and to buy if $p(h) < \overline{p}$ is empirically confirmed by Figure 4, and by the other simulations we run in a closed market with random traders and with any choice of the other trader populations.

The reason why momentum traders progressively lose wealth while contrarian traders gain wealth is subtler. Let us recall that both kinds of traders compute a price trend subtracting from the current price, *p(h)*, the price at time $h - T_i$. Due to the mean reverting behaviour of the market around the price $\overline{p}$, generally, the probability that the price $p(h)$, randomly drawn from the experimental sequence, be higher or lower than $\overline{p}$ is the same and is equal to 50%. On the contrary, if the trend at time *h* is positive, i.e. if $p(h) - p(h - T_i) > 0$, the probability that $p(h) > \overline{p}$ is higher than 50%. This is due to the fact that there has been an upward price trend in the previous T_i time steps. Of course, if $p(h) - p(h - T_i) < 0$, the probability that $p(h) > \overline{p}$ is lower than 50%. Figure 5 shows the price distribution for a simulation of 38,000[4] time steps, together with the distributions of prices under the conditions that the preceding price trend is positive and negative. We computed the trend for a time lag of 30 time steps, i.e., $T_i = 30$. Let us note first that the average of prices $\langle p \rangle$ is not equal to the theoretical mean reverting price $\overline{p}$. This is due both to rounding effects in the initial wealth distribution and to the limitedness of the sample. This bias has no effects on the overall market behaviour. Clearly, all the three distributions shown in Figure 5 are approximately log-normal, and the average prices of the conditional distributions are well over or under $\langle p \rangle$ (the average of the cumulative distribution), depending on the trend preceding their prices.

Consequently, if the price trend is ascending, it is more likely that the price is higher than $\overline{p}$, and the correct strategy is to sell, as said above, but momentum traders will buy, thus making, on average, an error. Contrarians, on the other hand, will correctly sell a part of their stocks. If the trend is downward, the opposite is true.

6.3 Open market

We opened the market allowing cash to increase progressively by a fixed amount every period. New cash is randomly distributed to agents in proportion to their total wealth. We noted first that in this case prices exhibit a strong trend. This confirms that the injection of cash in a market with a fixed amount of stocks produces asset inflation. We assumed also that the fundamental price p_f does not vary with time.

We performed two different simulations, depending on the amount of cash injection. In the first simulation, the cash inflow was small. At every time step, each trader has the probability of 1% to receive a cash inflow equal to the 0.1% of his current wealth. In this case, see Figure 6, fundamentalists increase wealth just because of the injection of cash. In fact, due to the trend, fundamentalists soon sell all stocks they own and remain with cash only. All other populations increase their average wealth, but with different rates. Contrarians increase their wealth more than random traders, which in turn increase their wealth more than momentum traders.

In the second simulation, the cash injection is on average twenty times greater. Each trader has the probability of 2% to receive an amount of cash equal to the 1% of his current wealth. In this case, the situation of contrarian and momentum traders is temporally inverted. Figure 7 shows the first 1,000[5] of a typical simulation 10,000 steps long. Now momentum traders win the wealth race and gain more than random and contrarian ones.

---

[4] The simulation is longer than usual in order to give better estimates of the probability density functions.
[5] The time window of the simulation is shorter because the high cash inflow produces remarkable effects even at few time steps.

The explanation of this finding is that in a growing market, if the cash injection is greater than a certain threshold, which is difficult to be determined analytically, the trend caused by the money inflow overcomes the noise due to the background trading of random traders. Therefore, momentum traders, i.e., trend followers, make on average the right choice. On the other hand, contrarians which speculate against the trend, do wrong because, with the time constants involved (traders memories and rate of cash injection), the trend does not revert sufficiently to allow them making a profit.

7 Conclusions

The computational experiments performed in this work show a number of important results. First, they demonstrate that the average price level and the trends are set by the amount of cash present and eventually injected in the market. In a market with a fixed amount of stocks, a cash injection creates an inflation pressure on prices.

The other important finding of this work is that different populations of traders characterized by simple but fixed trading strategies cannot coexist in the long run. One population prevails and the other progressively lose weight and disappear. This result is in agreement with the observations of Friedman (1953), but appears in contrast with the papers by De Long et al. (1990,1991). Indeed, De Long at al, followed a different approach to the problem, presenting a analytical model in which irrational or noise traders are characterized by erroneous stochastic beliefs.

Which population will prevail and which will lose cannot be decided on the basis of the strategies alone. Trading strategies yield different results in different market conditions.

In real life, different populations of traders with different trading strategies do coexist. These strategies are boundedly rational and thus one cannot really invoke rational expectations in any operational sense. Though market price processes in the absence of arbitrage can always be described as the rational activity of utility maximizing agents, the behaviour of these agents cannot be operationally defined.

This work shows that the coexistence of different trading strategies is not a trivial fact but requires explanation. One could randomize strategies imposing that traders statistically shift from one strategy to another. It is however difficult to explain why a trader embracing a winning strategy should switch to a losing strategy. Perhaps market change continuously and make trading strategies randomly more or less successful. More experimental work is necessary to gain an understanding of the conditions that allow the coexistence of different trading populations.

References

K. Beck. Extreme Programming Explained. Embrace Change. Addison-Wesley, Reading, Massachusetts, 1999.

F. Black, Noise, J. Finance 41 (3) 1986, pp. 529-543.

J.-P. Bouchaud & M. Potters. Theory of Financial Risk. From Statistical Physics to Risk Management. Cambridge University Press, Cambridge, 2000.

J.-P. Bouchaud & M. Mezard, Wealth condensation in a simple model of economy, Physica A 282 (3-4) 2000, pp. 536-545.

E. Briys & F. De Varenne, 1000 years of risk management, Risk, January 2000, pp. 47-48.

R. Cont & J.-P. Bouchaud, Herd behaviour and aggregate fluctuations in financial markets, Macroecon Dyn 4 (2) 2000, pp. 170-196.

J. B. De Long, A. Shleifer, L. H. Summers, R. J. Waldmann, Noise tradr risk in financial markets, J. Pol. Econ. 98 1990 pp. 703-738.

J. B. De Long, A. Shleifer, L. H. Summers, R. J. Waldmann, The survival of noise traders in financial markets, J. Business 64 (1) 1991 pp. 1-19.

M. Friedman. The case for flexible exchange rates, in: Essays in Positive Economics, M. Friedman (Ed,), University of Chicago Press, Chicago, 1953.

Z. F. Huang & S. Solomon (a), Finite market size as a source of extreme wealth inequality and market instability, Physica A 294 (3-4) 2001, pp. 503-513.

Z. F. Huang & S. Solomon (b), Power, Levy, exponential and Gaussian-like regimes in auto catalytic financial systems, Eur Phys J B 20 (4) 2001, pp. 601-607.

B. LeBaron, Agent-based computational finance: Suggested readings and early research, J. Econ. Dyn. Control 24 (5-7) 2000, pp. 679-702.

M. Levy, S. Solomon, New evidence for the power law distribution of wealth, Physica A 242 (1-2) 1997, pp. 90-94.

M. Levy, H. Levy, S. Solomon. Microscopic Simulation of Financial Markets. From Investor Behaviour to Market Phenomena. Academic Press, San Diego, 2000.

Y. Liu, P. Gopikrishnan, P. Cizeau, M. Meyer, C.-K. Peng, H. E. Stanley, Statistical properties of the volatility of price fluctuations, Phys. Rev E 60 (2) 1999, pp. 1390-1400.

T. Lux, Time variations of seconds moment from a noise trader/infection model, J. Econ. Dyn, Control, 22 (1) 1997 pp. 1-38.

T. Lux & M. Marchesi, Scaling and criticality in a stochastic multi-agent model of a financial market, Nature 397 1999 pp. 498-500.

R. N. Mantegna & H. E. Stanley, An Introduction to Econophysics. Correlations and Complexity in Finance, Cambridge University Press, Cambridge, 2000.

E. W. Montroll & M. F. Schlesinger, On 1/f noise and other distributions with long tails, Proc. Nat. Acad. Sci. 79 1982 pp. 3380-3383.

M. O'Hara, Market Microstructure Theory, Blackwell, Cambridge MA, 1995.

R. G. Palmer, W. B. Arthur, J. H. Holland, B. LeBaron, P. Tayler, Artificial economic life – A simple model of a stock market, Physica D 75 (1-3) 1994, pp. 264-274.

B. LeBaron, B. W. Arthur, R, G. Palmer, Time series properties of artificial stock market, J. Econ. Dyn. Control. 23 (9-10), 1999, pp. 1487-1516.

V. Pareto, Cours d'Economie Politique, Lausanne 1897.

M. Raberto, E. Scalas, G. Cunibeti, M. Riani, Volatility in the Italian stock market: an empirical study, Physica A 269 (1) 1999, pp. 148-155.

M. Raberto, S. Cincotti, S. M. Focardi, M. Marchesi, Agent-based simulation of a financial market, Physica A 299 (1-2) 2001, pp. 320-328.

H. E. Roman, M. Porto, N. Giovanardi, Anomalous scaling of stock price dynamics within ARCH-models, Eur. Phys. J B 21 (2) 2001, pp. 155-158.

G. Succi & M. Marchesi. Extreme Programming Examined. Addison-Wesley, Reading Massachusetts, 2001.

C. J. Wild & G. A. Seber. Chance Encounters. A First Course in Data Analysis and Inference. John Wiley & Sons, New York, 2000.

G.K. Zipf, Human Behaviour and the Principle of Least Action, Addsion-Wesley, Cambridge MA, 1949.

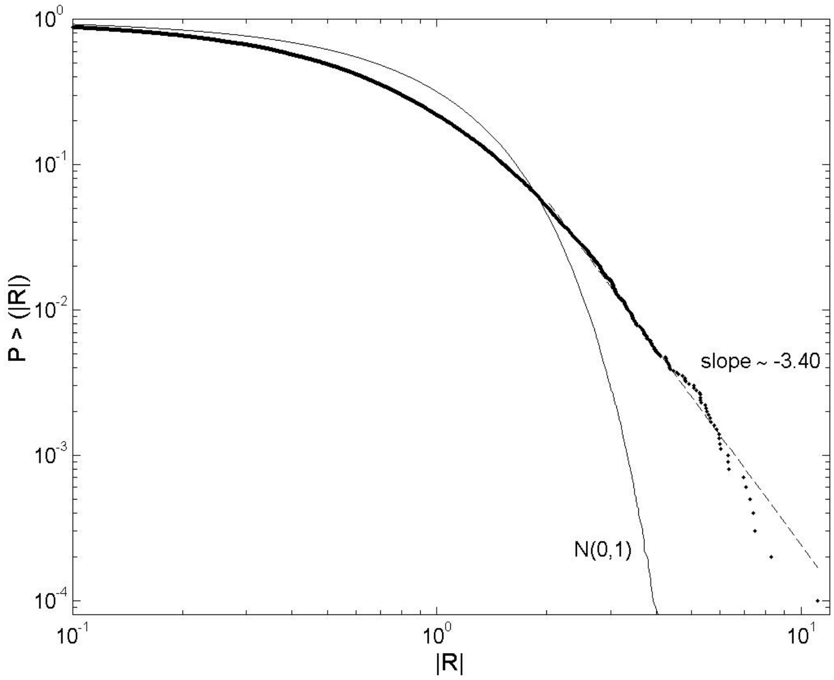

Figure 1. Cumulative distribution of standardised logarithmic returns $R(h)$. Considering logarithmic returns $r(h) = \log p(h) - \log p(h-1)$, $R(h)$ are computed as follows: $R(h) = (r(h) - m_r)/s_r$ where m_r and s_r are respectively the mean and the standard deviation of

$r(h)$ **over** h**. Dots represent an estimate of the cumulative distribution of** $R(h)$ **related to a simulation of GASM. The solid line represents the cumulative distribution of a random variable drawn from a normal distribution. The positive and the negative tails were merged by employing absolute returns. The dashed line is the power law fit** $P = K\,|R|^{-t}$ **with** $t = 3.40$ **of the tail of the empirical cumulative distribution for** $|R| > 2$ **.**

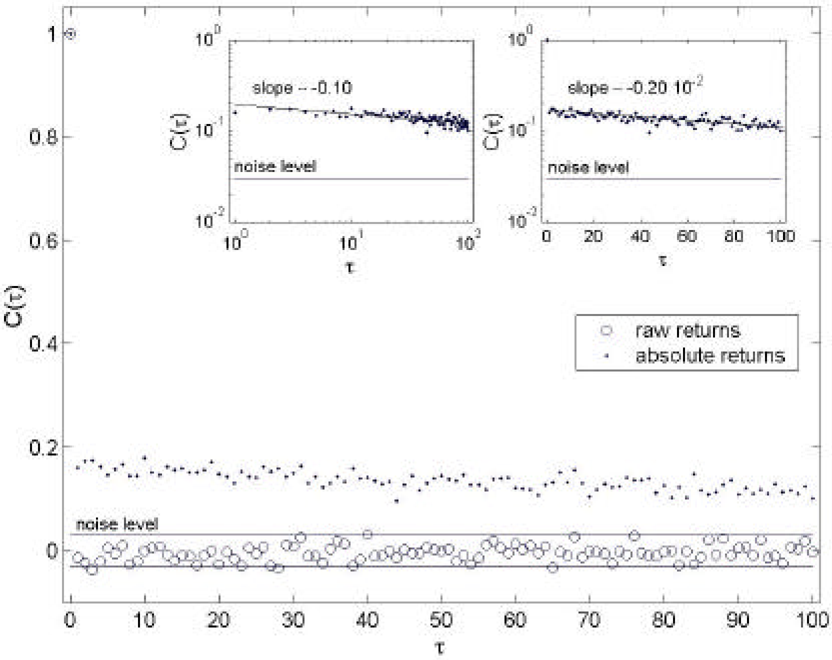

Figure 2. Estimate of the autocorrelation function of logarithmic returns. Dots represents the autocorrelation of absolute returns $|r(h)|$**, circles are related to the autocorrelation function of raw returns** $r(h)$**. Noise levels are computed as** $\pm 3/\sqrt{M}$ **where** M **is the length of the time series (** $M = 10{,}000$ **). Absolute returns are fitted with an exponential decay (inset on the right) and a power law decay (inset on the left).**

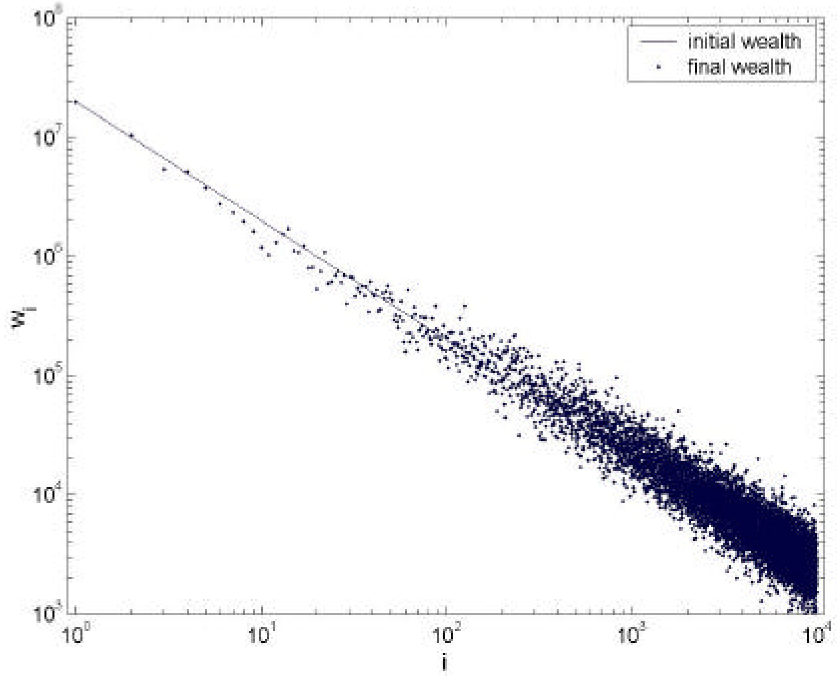

Figura 3. Distribution of wealth w_i between random traders i at the beginning and at the end of the simulation (after 10,000 steps).

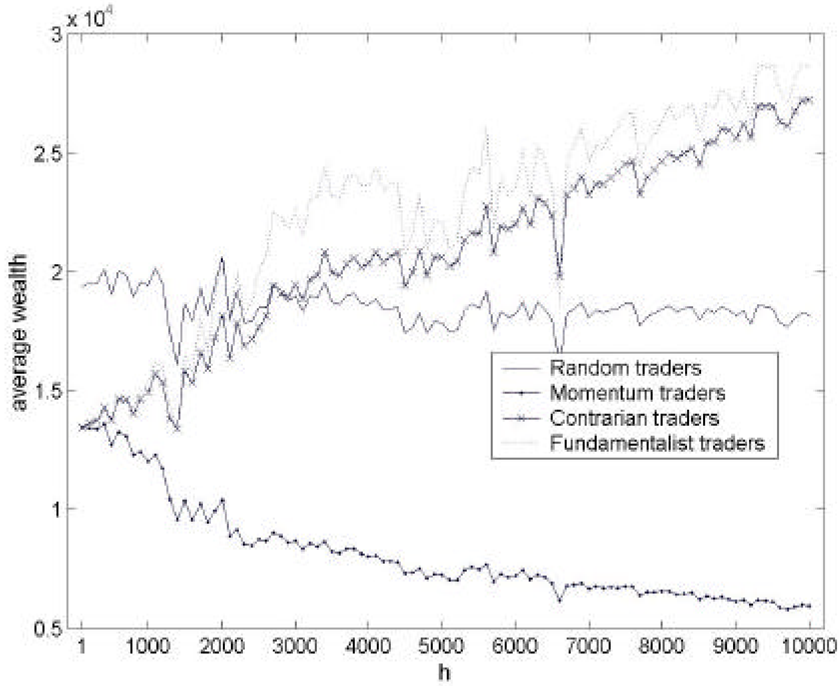

Figure 4. Average wealth dynamics of the four population of traders for a typical simulation 10,000 steps long with no cash inflow or outflow.

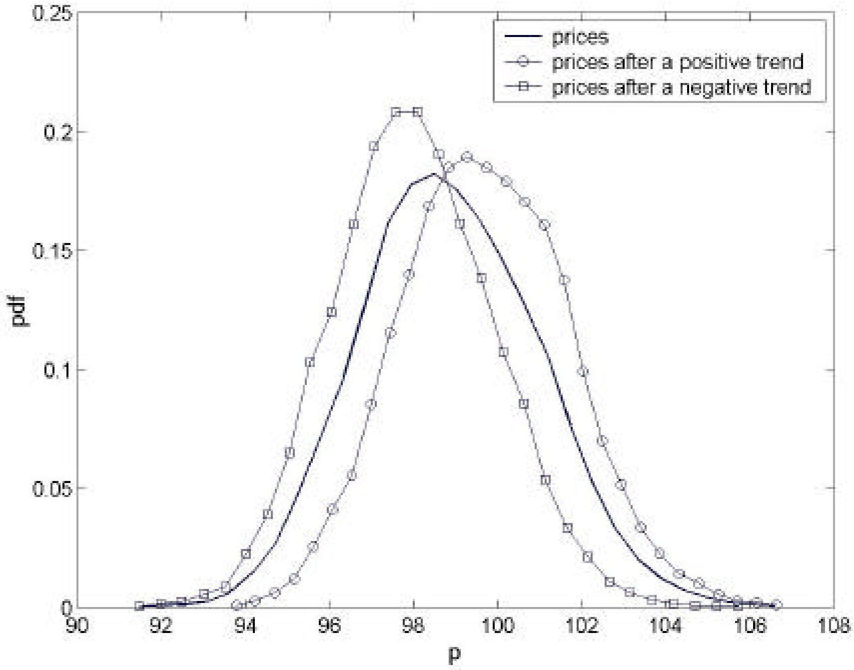

Figura 5. Probability density function (pdf) estimate of prices for a typical simulation 38,000 steps long. The distribution of all the 38,000 prices $p(h)$ is represented by the thicker solid line, the mean of the distribution is equal to $\langle p \rangle$ = 98.84. The solid line with circles is related the 18,915 prices $p(h)$ that satisfy the condition: $p(h) - p(h-30) > 0$ (positive trend), now the mean is 99.70. The solid line with squares represents the 18,985 prices satisfying the condition $p(h) - p(h-30) < 0$ (negative trend) whose mean value is: 97.99.

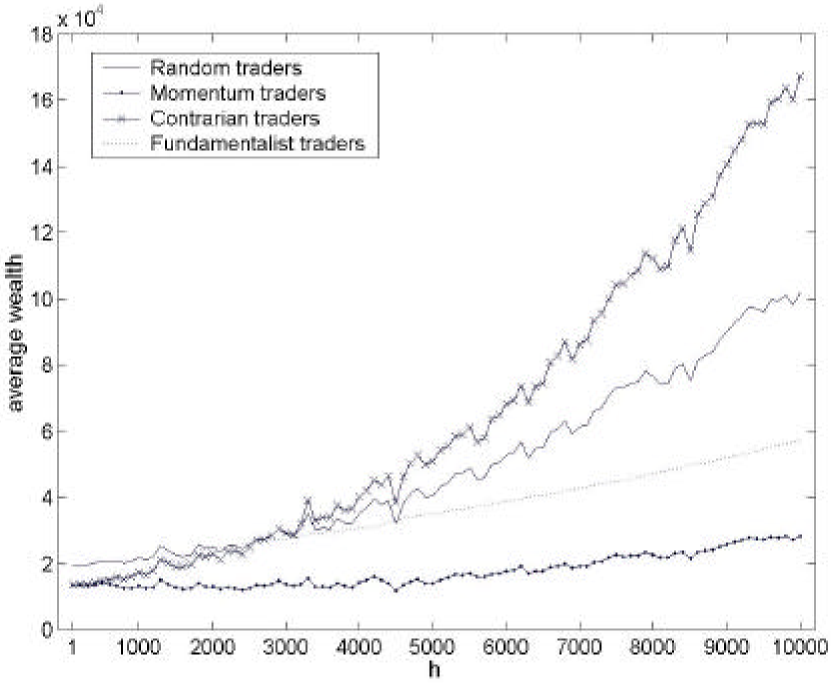

Figure 6. Average wealth dynamics of the four population of traders for a typical simulation 10,000 steps long with moderate cash inflow.

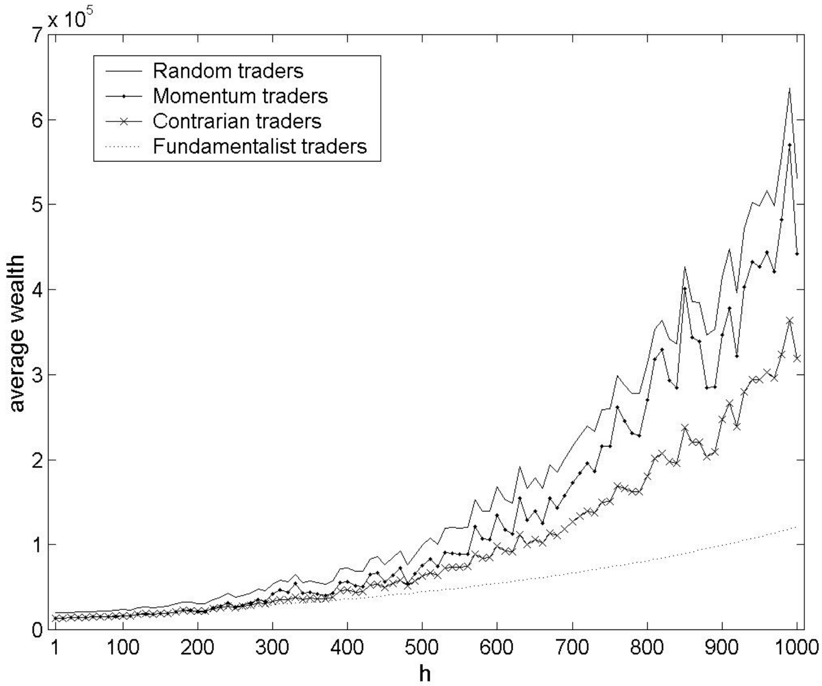

Figura 7. Average wealth dynamics of the four population of traders for a typical simulation 1,000 steps long with high cash inflow.