Herding effects in order driven markets

Giulia Iori, Department of Economics, City University London

Joint work with Gabriele Tedeschi and Mauro Gallegati, Department of Economics, Universitá Politecnica delle Marche, Ancona.

March 1, 2011

Objectives

Introduce a mechanism of imitation among traders that

- It relies on the limit order book to determine the price and not on ad-hoc price formation rules.
- The communication network among agents emerges endogenously
- reproduces and could potentially explain some of the stylized facts of real markets

- Agents trade based on their private signal and the signal of one other agent whom they chose to follow according to a stochastic rule.
- An agent is more likely to follow a rich agent than a poor agent. This is a sensible rule of thumb if we imagine that agents believe that other agents become rich because of their skill at investing and not just because they were lucky.
- Agents become more confident in the quality of their signal as they gain followers. Again this makes intuitive sense.
- First result is that some investors are followed by many other investors - they become "gurus". This is observed in actual markets.
- The main result is that gurus and agents who follow gurus earn higher returns on average than agents who follow non gurus. Thus the rule of thumb survives if agents consider whether to stick to it based on average returns.

In a earlier paper[1] the communication structure was taken as hexogenus and we assessed how imitations among otherways noise traders, can give rise to well known stylized facts such as fat tails and volatility clustering.

In the current paper we endogenize the network formation mechanism, through a fitness mechanism based on agents wealth. In this setting we study under which conditions a guru may rise and fall over time and the effect of imitation on prices and the distribution of agents' wealth.

[1] G. Tedeschi, G. Iori and M. Gallegati The role of communication and imitation in limit order markets, Eur. Phys. J. B, 1, 489-497 (2009)

The market

The model set up is similar to Chiarella et al. (2009)[2] but the synchronization of agents' strategies here is driven by herding and not by chartist/fundamentalist common trading rules. The market is order-driven. A population of N traders can either place market orders, which are immediately executed at the current best listed price, or they can place limit orders.

Limit orders are stored in the exchange's book and executed using time priority at a given price and price priority across prices. A transaction occurs when a market order hits a quote on the opposite side of the market.

[2] C. Chiarella, G. Iori, J. Perello, The impact of heterogeneous trading rules on the limit order book and order flows, Journal of Economic Dynamics and Control, Volume 33, Issue 3, March 2009, Pages 525-537.

Schematic representation of the order book

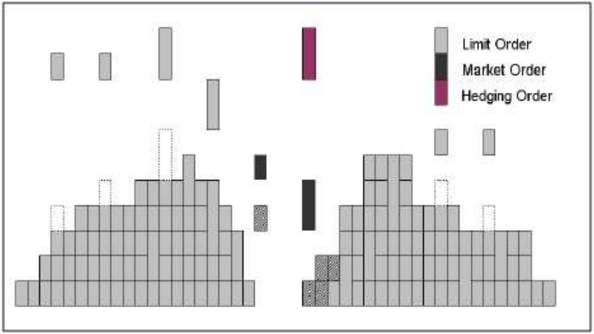

The interaction network

Agents start with the same amount of cash $C(0)$ and stocks $S(0)$, so that all agents have the same initial wealth $W(0) = C(0) + p(0)S(0)$.

Agents also starts with one outgoing link with random agent j, and possibly some incoming links with other agents. As time goes by, some traders may become richer than others.

As a measure of agents' success we define their fitness at time t as their wealth relative to the wealth W_t^{max} of the richest agent $i_{\max}$ at that time:

$$f_t^i = \frac{W_t^i}{W_t^{\max}}. \quad (1)$$

We introduce an endogenous mechanism of preferential attachment based on agent's fitness.

Links are rewinded at the beginning of each period, in the following way: each agent i cuts its outgoing link, with agent k, and forms a new link, with a randomly chosen agent j, with a probability

$$p_r = \frac{1}{1 + e^{-\beta(f_t^j - f_t^k)}}$$

This allows for only partial information to be available to agents at any point in time.

CITY UNIVERSITY LONDON

The rewind algorithm is designed so that successful traders gain a higher number of incoming links and thus have a higher probability of being imitated.

Nonetheless the algorithm introduces a certain amount of randomness, and links with more successful agent have a finite probability to be cut in favour of links with less successful agents. In this way we model imperfect information and bounded rationality of agents.

The expectation formation mechanism

At the beginning of each trading period t_k, agents make idiosyncratic expectations about the spot return, $\hat{r}^i_{t_k,t_k+\tau}$ in the interval $(t_k, t_k+\tau)$. We assume that agents are not informed and have random expectation of future returns. We also assume that agents are heterogeneous in that they have different forecasts of the returns' volatility, σ^i_t. Expected returns are thus given by

$$\hat{r}^i_{t_k,t_k+\tau} = \sigma^i_{t_k}\epsilon_{t_k} \qquad (2)$$

where $\sigma^i_{t_k}$ is the positive, agent specific, time dependent volatility and $\epsilon_t \sim N(0,1)$ is a normal noise.

After individual expectation are generated, a consultation round starts during which agents sequentially, and in a random order, revise their expectation. The revised expected return is obtained by weighing agent *i*'s own expectation with that of agent *j* to which *i* is linked to such as

$$r^i_{t_k,t_k+\tau} = w\hat{r}^i_{t_k,t_k+\tau} + (1-w)\hat{r}^j_{t_k,t_k+\tau}. \tag{3}$$

At the beginning of each period t_k agents expectations are reset to random values.

We stress that in the model imitation is purely expectation based, and agents do not imitate the actions of others. This choice is motivated by the fact that in a real market normally the order book is not fully visible to traders, and the order submission is anonymous.

CITY UNIVERSITY LONDON

While agents are noise traders in our model, we assume that they correctly anticipate the impact of herding on asset prices. In particular if an agent has several incoming links, and w is large (in which case the agent expects to be able to influence the decisions of others), the agent forecasts a larger price volatility. This effect is captured by setting

$$\sigma^i_{t_k} = \sigma^i_0(A + I^{\%}_{i,t_k}(1 - w)) \tag{4}$$

where $I^{\%}_{i,t_k}$ in the percentage of existing links that point to agent i at time t_k and A is a constant parameter. The values of σ^i_0 are chosen, with uniform probability, in the interval $(0, \sigma_0)$.

Trading happens over a number of periods t_k, with $k = 1, \cdots T$. At the beginning of each period, traders make expectations about the price at the end of a given time horizon τ (that we take to be the same for all traders). The future price expected at time $t_k + \tau$ by agent i is given by

$$\hat{p}^i_{t_k, t_k+\tau} = p_{t_k} e^{\hat{r}^i_{t_k, t_k+\tau} \sqrt{\tau}} =$$

$$p_{t_k} e^{w \sigma^i_{t_k} \epsilon_{t_k} \sqrt{\tau} + (1-w) \sigma^j_{t_k} \epsilon_{t_k} \sqrt{\tau}}$$

where $\hat{r}^i_{t_k, t_k+\tau}$ is the agent's expectation on the spot return and p_{t_k} is the reference price observed by all agents at the beginning of each period.

Risk Averse Agents with Exponential Utility Functions

- The optimal composition of the agent's portfolio is determined by trading-off expected return against expected risk.
- We assume an exponential utility of wealth $U(W_t^i, \alpha) = -e^{-\alpha W_t^i}$, where the coefficient α measures the risk aversion of traders.
- the optimal portfolio composition at a given price level p is independent on the agent's wealth and given by

$$\pi^i(p) = \frac{\ln(\hat{p}^i_{t_k+\tau}/p)}{\alpha V^i_{t_k} p}, \tag{5}$$

where $V^i_{t_k}$ is the risk perceived by agent i, normally taken as the unconditional variance of returns.

- We assume that those agents who are highly imitated, think that they must be excellent forecasters if many people follow their word. Thus, success leads to confidence.
- To these popular agents the risky asset appears as less risky, because they are good at forecast it. Thus their assessment of risk is far below the unconditional variance of the risky asset returns.

This effect is captured by setting

$$V^i_{t_k} = V_{t_k}(A - (1 - w)I^{\%}_{i,t_k}) \tag{6}$$

The unconditional variance of returns, V_{t_k} is assumed to be calculated in the same way by all the agents as

$$V_{t_k} = \frac{1}{N}\sum_{j=1}^{N}[r_{t-j} - \bar{r}_N]^2$$

where

$$\bar{r}_N = \frac{1}{N}\sum_{j=1}^{N}\log\frac{p_{t-j}}{p_{t-j-1}}.$$

CITY UNIVERSITY LONDON

- If the amount $\pi^i(p)$ is larger than the number of stocks already in the portfolio then agents decide to buy, if smaller they decide to sell.
- Following Bottazzi et al (2003) we first estimate the price level p^* at which agents are satisfied with the composition of their current portfolio

$$\pi(p^*) = \frac{\log(\hat{p}^i_{t+\tau_i}) - \log(p^*)}{\alpha^i \sigma_t p^*} = S^i_t.$$

- This equation admits a unique solution with $0 < p^* < \hat{p}^i_{t+\tau_i}$.
- Agents buy at any price $p < p^*$ and sell at any price $p > p^*$.
- Note that agents may wish to sell even if they expect a future price increase.

- Budget constraint:

$$p \leq \hat{p}^i_{t+\tau_i} = p_M$$

to ensure $\pi(p) > 0$ and rule out short selling.

- To ensure that an agent has sufficient cash to purchase the desired stocks, the smallest value of p we can allow is

$$\pi(p_m) - S^i_t = \frac{C^i_t}{p_m}.$$

- This equation admit a unique solution.

- The possible value at which an agent can satisfactorily trade are in the interval (p_m, p_M).
- Agents randomly pick a price p in the interval (p_m, p_M)
    - if $p < p^*$ submit a limit order to buy an amount

$$s^i = \pi^i(p) - S^i_t$$

    - if $p > p^*$ submit a limit order to sell an amount

$$s^i = S^i - \pi^i(p).$$

- If $p > p^*$ and $p \leq b_t^q$ the selling order can be executed immediately at the bid. Agents then submit market orders to sell an amount

$$s^i = S_t^i - \pi^i(b_t^q)$$

- if $p < p^*$ and $p \geq a_t^q$ the buying order can be executed immediately at the ask. Agents then submit market orders to buy an amount

$$s^i = \pi^i(a_t^q) - S_t^i.$$

- If the depth at the bid (ask) is not enough to fully satisfy the order, the remaining volume is executed against limit orders in the book at quotes above (below) p.
- If there are not enough orders in the book to execute the incoming market order fully the remaining volume is converted into a limit order at price p.

The price is recalculated as trading goes on as follows: p_t is given by the price at which a transaction occurs, if any;

if no new transaction occurs, a proxy for p_t is given by the average of the quoted ask a_t^q (the lowest ask listed in the book) and the quoted bid b_t^q (the highest bid listed in the book), $p_t = (a_t^q + b_t^q)/2$;

if no bids or asks are listed in the book a proxy for p_t is given by the previous traded or quoted price.

Bids, asks and prices are positive and investors can submit limit orders at any price on a prespecified grid, defined by the tick size Δ[3].

[3] The tick is the smallest possible change of p_t.

Simulations and results

The model is studied numerically for different values of the parameter w. In the simulations the number of traders is set at $N = 150$. Each agent is initially given the same amount of stock $S_0 = 100$ and cash $C_0 = 100$. The initial stock price is chosen at $p_0 = 1000$. We fix $\tau = 200, \alpha = 0.01$, and β^i uniformly distributed in the interval $[5, 45]$ The results reported here are the outcome of simulations of $T = 1000$ periods and $N_t = 300$ trades per period. Simulations are repeated $M = 100$ times with a different random seed[4]

[4] We have tested the stability of our results and verified that the model shows a qualitatively similar behaviour for a range of values of the parameters.

Network

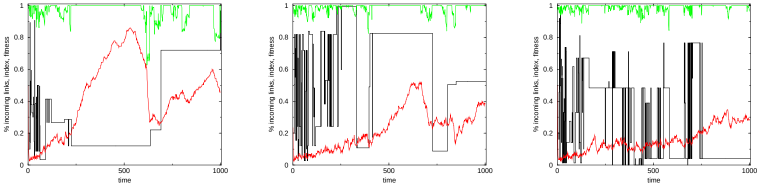

Figure: The index of current guru(black), the percentage of incoming link to current guru (red) and fitness of current guru(green) for $w = 0.1$ (left side), $w = 0.5$ (centre) and $w = .9$ (right side). The figure shows that agents alternate as the guru during the simulation (black line). In fact, as the guru acquires an increasing number of links (red line), one or more of his followers may become richer than the guru himself, as signalled by the fact that the fitness (green line) of the guru becomes, at times, smaller than 1. As other agents become rich they start to be imitated more and more and eventually one of them becomes the new guru.

CITY UNIVERSITY LONDON

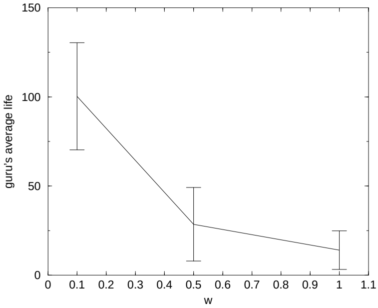

Figure: Average Guru's live as a function of w. The stability (or average life) of the guru becomes longer as imitation increases

Prices

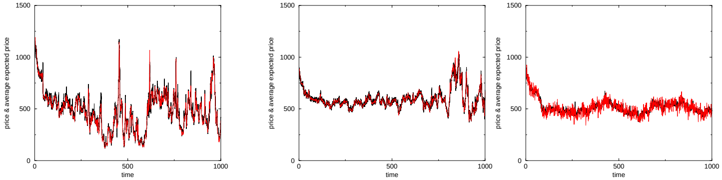

Figure: Prices (black line) and average expected prices (red line) with *w*=0.1 (left), *w*=0.5 (center) and *w*=.9 (right). We can notice that coordination causes wider excursions of price movement. Further, we observe that prices and expected prices follow each other closely when *w* is small.

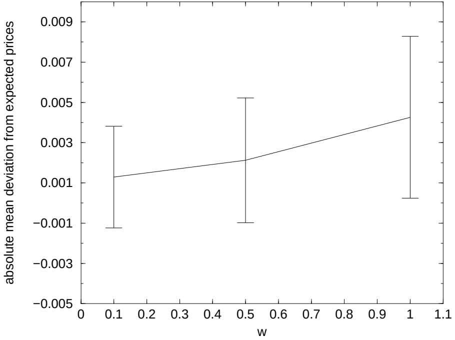

Figure: Deviation between realized prices and expected prices as a function of w. The deviation coefficient is defined as $|\frac{p_{t_k} - \hat{p}^i_{t_k}}{p_{t_k}}|$, and we average it over time and the number of simulations.

CITY UNIVERSITY LONDON

Positive feedback has several implications on the correlation between asset price behaviour and traders' action. In particular, experimental investigations and theoretical models (see De Long et al. (1989), Hommes et al. (2003), Heemeijer et al. (2007)) show that, when noise traders follow positive feedback strategies, they buy when prices increase and sell when prices fall. A natural way to assess the co-movement between the increase (decrease) in prices and increase (decrease) in purchase orders is to study their correlation.

CITY UNIVERSITY LONDON

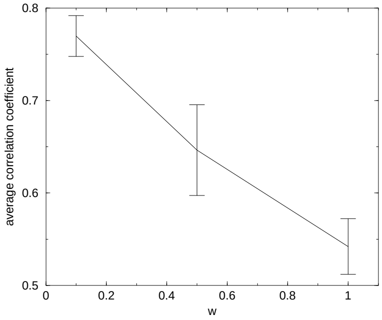

Figure: Average correlation coefficient between price changes and number of buyers as a function of w. The correlation coefficient (the Person correlation coefficient significant at 1% level) is an decreasing function of w and reaches a value above 0.7 when imitation is significant.

Wealth distribution

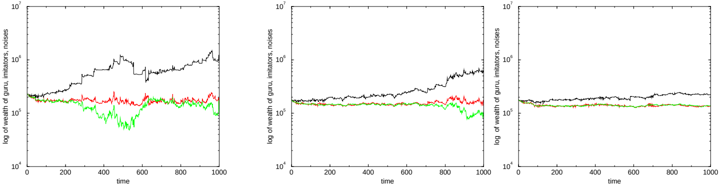

Figure: Wealth time series of guru (black line), followers (red line) and rest of the system (green line) for $w = 0.1$ (left side), $w = 0.5$ (center) and $w = 1.0$ (right side).

CITY UNIVERSITY LONDON

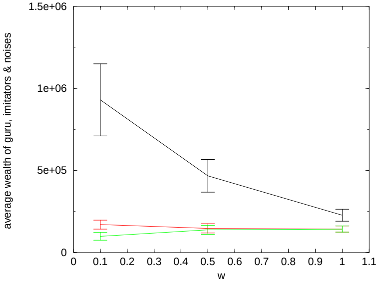

Figure: Average wealth, over all times and all simulations, of the guru (black line), followers (red line) and rest of the system (green line). The wealth of the guru increases with *w* and the gap between the wealth of the guru and the wealth of the rest of the system (both followers and non followers) widens with the level of imitation.

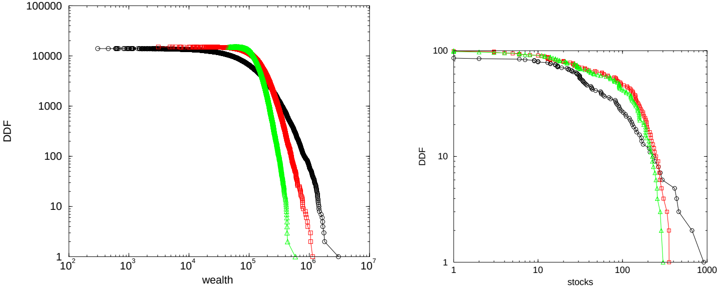

Figure: The decumulative distribution function (DDF) of the wealth (left side) and the decumulative distribution function (DDF) of the stocks (right side) for $w = 0.1$ (black line), $w = 0.5$ (red line) and $w = 1.0$ (green line). Raising imitation, the model generates heterogeneity, as indicated by the fatter tails in the distribution of agents' wealth and stock.

Discussion

To explain the above results we first need to show that the imitation of expectations translates into imitation of actions.

- An expected price increase (decrease) in our model does not necessarily lead to a decision to buy and, even if so, buy order could be submitted as limit orders. Market orders are more likely to be submitted when agents are very optimistic or pessimistic. In fact in this case the interval $[p_m, p_M]$ over which orders can be placed is wider and it becomes more likely that a price level is chosen such that the order can be immediately executed.
- In our model it is the agents with many incoming links who forecast a high volatility σ_t^i and are more likely to submit market orders. If a popular agent has enough connections it can influence several others to overestimate price changes and submit market orders in turns. Thus the coordination of expectation does leads to a coordination of actions (submission of market orders).

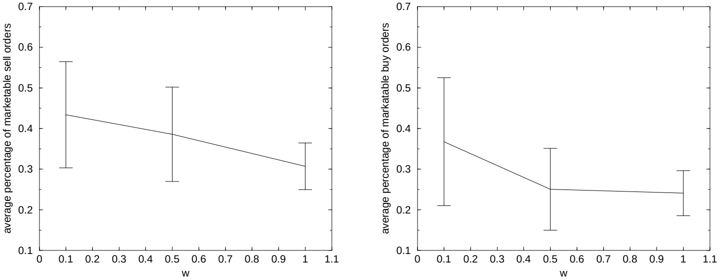

Figure: Average fraction of the volume of market orders to sell(left) buy (right) over the total volume of orders in the same direction for different value of w. The average is taken over each trading period and plotted for different values of w. The result shows, as anticipated, that, when increasing imitation, a higher fraction of market orders is submitted. Thus the coordination of expectation leads to a coordination of actions and the model generates an expectations feedback system.

- Furthermore, while agents are risk adverse, highly connected agents underestimate risk. Consequently these traders, when w is small and their percentage of incoming links, $I^{\%}$, is high, invest more (on average) in the risky asset than others. Followers in turn invest on average more than non followers because they, like the guru, overestimate returns.
- A series of market orders in the same direction can generate considerable price changes. Thus, the forecasts from highly connected agents of an overall high volatility are self-fulfilling, providing an ex-post justification for equation (13).

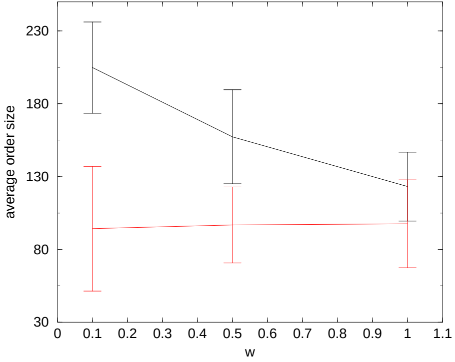

Figure: Mean size of orders of the guru (black) and the rest of the system (red) as a function of w.

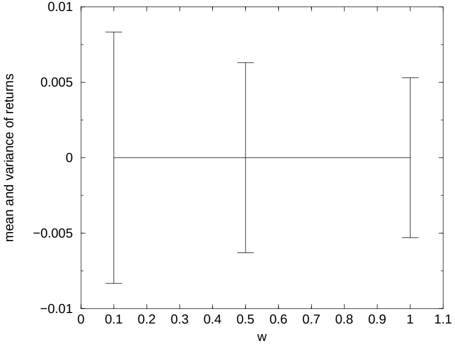

Figure: Mean and variance of returns as a function of w. Volatility increases with imitation.

- As long as the guru is not the last to trade (we assume a random entrance to the market for all agents including the guru) he will consistently gain on the trades that follow, in the same direction, his trade.
- By investing more, gurus and followers earn, on average, higher profits than no followers.

These results are in line with other studies on noise traders risk with positive feedback in financial markets. Particularly, De Long et al. (1990a) show that noise traders can earn higher returns solely by bearing more of the risk that they themselves create. Our results are in line with previous studies(see Haltiwanger and Waldman (1985), Heemeijer et al. (2007)) and confirm that traders have an incentive to imitate and be imitated, since predicting a price close to the predictions of other players turns out to be most profitable.

Conclusions

- Our results result allows us to conclude that profit is a good mechanism of links formation.
- The endogenous attachment mechanism introduced in our model allows a guru to emerge spontaneously in the system, rise and fall in popularity over time, and possibly be replaced by a new guru. A few gurus could also co-exist and compete among themselves for popularity.
- For the endogenous attachment mechanism to be capable of creating, sustaining and destroying a guru, agents need to benefit from imitating and being imitated. In fact, if an agent profits from being imitated, he becomes richer, which induces an even larger fraction of agents points to him.
- On the other side, if followers also profit from imitating the guru, their could eventually over-perform the guru, and become guru in turn.

- The fact that our unsophisticated investors, trivially driven by imitative behaviour, can earn very high profits implies that Friedman's hypothesis is inadequate.
- The assumption that noise traders quickly go bankrupt and are eliminated from the market is unrealistic in presence of herding and positive feedback. In fact we have shown that noise traders can earn very high profits from herding and cause large price fluctuations.
- These results should not be underestimated, particularly in those situations when market prices exhibit large fluctuation. In these cases in fact is unlikely that prices incorporate true information and the idea of full rationality is implausible.