



Network Topology and the Behaviour of Socially-Embedded Financial Markets

Olivier Brandouy^{1,2} and Philippe Mathieu^{1,2}(✉)

¹ GREThA UMRS 5113 – U. Bordeaux, Bordeaux, France

² CRISTAL UMR 9189 – U. Lille, Lille, France

philippe.mathieu@lifl.fr

Abstract. We study the impact of the network topology on various market parameters (volatility, liquidity and efficiency) when three populations or artificial trades interact (Noise, Informed and Social Traders). We show, using an agent-based set of simulations that choosing a Regular, a Erdős-Rényi or a scale free network and locating on each vertex one Noise, Informed or Social Trader, substantially modifies the dynamics of the market. The overall level of volatility, the liquidity and the resulting efficiency are impacted by this initial choice in various ways which also depends upon the proportion of Informed *vs.* Noise Traders.

1 Introduction

Financial markets have a central role in modern economies and receive from both the academic and the politic world an important amount of interest. Among others, some technical aspects in their behaviour are still discussed and scrutinized: for example, volatility is presented as a normal outcome of their behaviour, mainly (but not uniquely) driven by the “real world” stochastic they reflect. Another important aspect is their liquidity, which, at coarse grain, determines the ability of an investor to resell (to short) immediately or at very short notice one of his positions. Other technical parameters are observed, while they more or less all pertain to the same central hypothesis, the efficient market hypothesis (see [1, 7, 8]). If market are efficient, they reflect a certain degree of the available information regarding the economy. As such, they allow money to be allocated where it is needed and appropriately rewarded.

However, the social nature of Financial Markets, the fact they are made of human beings connected by networks (even if they seem to fade behind more-and-more automated processes) has been taken into account relatively recently for adding an extra layer of complexity in the analysis was not recognized as necessary.

In this paper we tackle this complexity (following in that other researchers like, for example, [11] or [4]) and propose a set of analysis geared at understanding their role, if any, in relation to a set of market parameters. As such,

our research question is the one of a possible effect of the underlying network topology of financial markets in the emergence of price motions and beyond, market regimes. We thus follow a line that recognizes the centrality of information contagion (and reputation) in the behaviour of stock markets.

The strength of our approach derives from the agent-based philosophy we adopt, which is more and more common in computational economics (see for example [10, 12, 13], or [2]). Studying such a complicated set of questions can be done efficiently using an agent-based artificial market (see, in an other context [5]), for this latter is designed so to mimic some essential features of the real world (experiments over several days but with access to an intra-day information, populations of agents but heterogeneity within this population, social spacialization etc. . .).

The paper is organized as follows: in a first section, we describe the overall architecture of our financial market. The second section is dedicated to the agent behaviour description. The third section presents our empirical strategy and our initial results.

2 The Model

We study a simplified Economy populated with a large number of atomistic traders where only two assets can be traded: a risk-free one and a risky one A , respectively yielding a rate of return r_f and paying a random dividend. Instead of interacting uniquely through the price system, as usually proposed in the literature, these are placed within a social network and as such are embedded in a social neighbourhood. The topology of the social network can either be a regular network (RN), such a ring for example (Fig. 1a, a Barabasi-Albert [3] (BA) scale-free network (Fig. 1b) or a Erdős-Rényi random network (ER, see 1c, see [6]), with a given level of connectedness. The characteristics of these networks blatantly differ: the RN is constructed such as all agents have exactly the same number of neighbours. In the BA case, some agents can be seen as “hubs” (with lots of neighbours), while others have a very limited social neighbourhood. Finally, for the ER case, agents do have different amounts of neighbours, but by opposition with the BA case, we cannot guaranty the connectedness of the network. Through these topologies, they can share some information regarding the dividend linked to the risky asset. During the trading process, each trader receives a piece of information allowing to estimate the level of the dividend with a certain level of accuracy. Part of the simulation will consist in observing how this information can be refined through social exchanges and how communication can affect the efficiency of the market. In addition, each trader is granted a level of trust by her/his direct neighbours which may evolve regarding the accuracy of their predictions regarding the value of this asset. These aspect are described further in Subsects. 2.1 and 2.2.

2.1 Assets, Portfolio of Assets

In our artificial world, time is discretized into “periods” and “ticks”. Periods can be analysed as “days”, and within each of these periods, ticks can be seen as a measure of intra-day time.

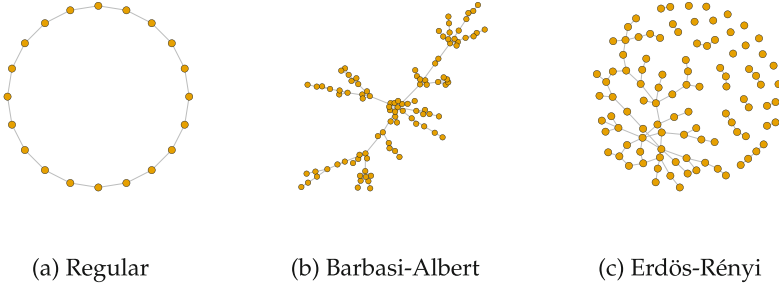


Fig. 1. Three social networks structuring our OTC market

The random dividend paid to the traders allows to compute a fundamental value for the asset. For doing so, we refer to a modified version of the Gordon-Shapiro model [9]: let p_t be the price of the risky asset A ; p_t can be modelled as the sum of further dividends that will be paid over time till infinity: $p_{t=0} = \sum_{i=1}^{+\infty} \frac{d_t}{(1+ke)^i}$, with ke the required rate of return for the risky asset at time t (see for example [1]).

We introduce some volatility in the dividend process: $\tilde{d}_t = d_{t-1} \times (1 + G_t)$, with $G_t \sim N(\mu, \gamma)$. In addition, we impose that $ke > G_t \forall t$.

In doing so we preserve the intuition behind the GS model: on average, dividends grow at a fixed rate, this rate being above the required rate of return ke . G_t is not observable by the agents and is never common knowledge to anyone in the economy. However these can access to an approximation of G_t , denominated $g_t \in [g_t, \bar{g}_t]$ allowing them to perform the computation of an estimate of p_t . The framing imposed by g_t around G_t creates some level of uncertainty around G_t , the latter being ruled by a parameter in further simulations.

$$p_t \in \left[\frac{d_0(1 + \underline{g}_t)}{ke - \underline{g}_t}; \frac{d_0(1 + \bar{g}_t)}{ke - \bar{g}_t} \right] = [p_t; \bar{p}_t] \quad (1)$$

So far, each investor has his own representation of the current price of the risky asset p_t and the possible next dividend through a set of information Ω_t . As we will see later, this representation may be more or less accurate, depending upon the category of agents and the information search made by each of these. In the population of artificial traders, each agent i is initially endowed with arbitrary quantities of both assets, $Q_{A,i} \geq 0$ and $Q_{r_f,i} \geq 0$.

Portfolio Construction: At the beginning of each period, agents want to maximize their expected utility in choosing θ , the proportion of risky asset A in their portfolio, so to have the highest possible value for their expected utility

$$E(U(w(1 + r_f) + \theta\tilde{\varepsilon})) \text{ with } w \text{ the financial wealth} \tag{2}$$

Let λ_i be the Arrow-Pratt absolute risk aversion coefficient, we fall in the classical case where the solution to the optimisation problem is

$$\theta^* = \frac{\mu}{\lambda_i \sigma_i^2} \tag{3}$$

In Eq. 3, μ and σ respectively denote the expected returns and the volatility of the risky asset. Knowing p_t and using g_t , investors can compute the possible boundaries for their return in one period of time conditional to the information they get. The return for the risky part of their investment:

$$r_{A,t} \in \left[\log\left(\frac{p_t}{p_{t-1}}\right); \log\left(\frac{\bar{p}_t}{p_{t-1}}\right) \right] = \left[\underline{r}_t; \bar{r}_t \right] \tag{4}$$

As such, they can equally compute the expected return and, henceforth, the volatility σ of the risky asset conditional to their estimate of G_t (see Eq. 5); if we replace the population of possible returns for the asset A in t $r_{A,t}$ by X :

$$\mu(X) = \int x f(x) dx \quad \text{Var}(X) = \sigma^2(X) = \int x^2 f(x) dx - \mu^2 \tag{5}$$

with $f(x)$ the probability density function for x . For the sake of simplicity, we use in this research a uniform PDF.¹ As such we have:

$$\mu(r_t) = \frac{1}{2}(\bar{r}_t + \underline{r}_t) \quad \text{Var}(r_t) = \frac{1}{12}(\bar{r}_t - \underline{r}_t)^2 \tag{6}$$

2.2 Agents Behaviour, Reputation and Trust

Three categories of traders co-evolve within this framework:

- “*Noise Traders*” (henceforth NT). NT only rely on a public signal ps characterized by the widest framing around G_t .
- “*Informed Traders*” (similarly “IT”) who differ from the preceding NT in receiving a narrower estimate of G_t . So to speak, they receive a signal $is_i = N(\mu, \rho_i)$ allowing them to observe the true value of the risky asset with a reduced level of variance. By definition $\rho_i < \sigma^{PS}$, with σ^{PS} the standard deviation of the public signal.

¹ Notice that if g_t , which stands for the approximation of G_t follows a Normal distribution, it only allows agents to determine a framing for the possible price. As such, these latter must choose a “target” between the upper and the lower boundaries. This is the reason why they rely on a Uniform PDF for doing so.

– “*Social Traders*” (similarly “ST”) who receive exactly the same signal as NT but can screen within their network environment, at a social distance equal to 1, if another trader, whatever his type, has a reputation large enough to be trusted in his own estimate of G_t . Said differently, ST can receive a (possibly) noisy information from their neighbours in addition to the public signal $ss_i = N(\mu, (\sigma^{PS} + \phi_t) \text{ with } \phi_t \in] - \sigma^{PS}, +\infty[$. The information they receive can either reduce the variance of the price estimate because it contains at least a piece truth $\phi_t \in] - \sigma^{PS}, 0]$, or it does not include anything instead of noise $\phi_t \in]0, +\infty[$. Of crucial importance here is the way ST select these information in their neighbourhood. For doing so, ST synthesize the latter following these steps:

1. Collect and archive the estimate of each neighbour in terms of possible deviation $cs_{i,t} = \left\{ \sigma^{PS}, \rho_{j,j \neq i}, \phi_{j,j \neq i} \right\}_t$; note that using the reverse relations coming from Eq. 6, ST can reconstruct the boundaries within which their neighbours believed the next dividend will be. In addition, each ST is endowed with a limited memory of k item: $cs_{i,t=1,\dots,k}$.
2. Make a choice in terms of information structure to use; in every case, the information chosen in the set derives from the strategy of the ST; we focus here on a single strategy where the decision is made with respect to the confidence or trust the agent has established with his/her N neighbours in time: $T_{i,j \neq i} = \left\{ T_{i,1}, T_{i,2}, \dots, T_{i,N} \right\}_t$. One could imagine that the information in the neighbourhood is averaged, or any other aggregation technique.
3. Once the choice of information is made, the agent computes the optimal quantity of risky asset to hold θ^* and the upper and lower bounds for the price of the risky asset using the reverse relations coming from Eq. 6. The agent will then send a limit order to buy or to sell r_A in choosing respectively the upper and the lower bound as reservation prices.
4. Trust is updated at each tick, according to the appropriateness of the information transmitted by the neighbour. For doing so, the agent compares the actual dividend d_t perceived at date t and the expectations of its neighbours concerning this dividend at time $t - 1$. Let $[a, b]_{i,t}$ be the interval proposed by neighbour i at date $t - 1$. $m_{t,i} = \frac{a+b}{2}$ is the mid-point in this interval. The variation of trust linked to this information is equal to:

$$\Delta T_i = \frac{1}{(d_t - m_{t,i})^2} \times \begin{cases} -1 & \text{if } d_t - a < 0 \vee b - d_t < 0, \\ 1 & \text{if } d_t - a > 0 \wedge b - d_t > 0. \end{cases} \quad (7)$$

3 Empirical Settings and Results

In this section we present the way experiments have been designed, the nature of the collected data and we present our results.

3.1 Experiments

Since we are interested in understanding the role of the topology of the social network structuring the relationships among traders, we explore several network settings. The other elements of the simulations are described below:

- The total population of agents is set at $N = 500$.
- Each experiment is developed over 30 artificial days, which means that 30 dividend signals are generated and used by the agents so to decide the optimal composition of their portfolio.
- Within each day, agents are allowed to express their orders 20 times, but once by round-table cycle. Said differently, the scheduling system is such that agents are allowed to speak one time per round, 20 rounds in a day being organized. This generates pseudo intra-day prices and quotes.
- The Preferential-Attachment algorithm generates scale-free networks; the number of edges in the network is equal to $2 \times N$, N being the number of vertices.
- Our Erdos-Renyi and Regular Networks are generated so to have the same average number of edges per vertex (*i.e.*): as such, the connectedness rate c_r of a ER network for a given population of N agents is equal to $r_c = \frac{2}{N}$. The RN is, in this particular case, a circle.
- Concerning the way information is disclosed and processed:
 - NT receive a signal around the next (unknown) dividend with boundaries set at $\pm 50\%$.
 - IT have a better accuracy with regard to the dividend: their boundaries are set to $\pm 10\%$
 - ST receive the same information as NT but they select an information in their social neighbourhood with respect to the highest level of the trader's reputation at a social distance of 1. For each network setting, we run 12 experiments and we collect prices, quotes, and the trading volumes. All these information are linked to a time stamp.

For each network setting, we run 12 experiments and we collect prices, quotes, and the trading volumes. All these information are linked to a time stamp.

A “typical” experiment produces market motions that are *qualitatively* realistic (see Fig. 2).

3.2 Results

We are not interested in the mean return observed in our various experimental settings: this parameter should not be affected by the set of initial parameters but by the random process driving the dividend. However, the first step in our data analysis process consists in transforming observed prices p_t to returns $r_t = \ln(p_t) - \ln(p_{t-1})$ at a tick-by-tick level (each time a new price is observed in the market). We do not consider the return produced by a new signal emission (end-of-the day effect) so to neutralize the volatility it may trigger: we are not interested in the volatility provoked by the dividend process but rather by

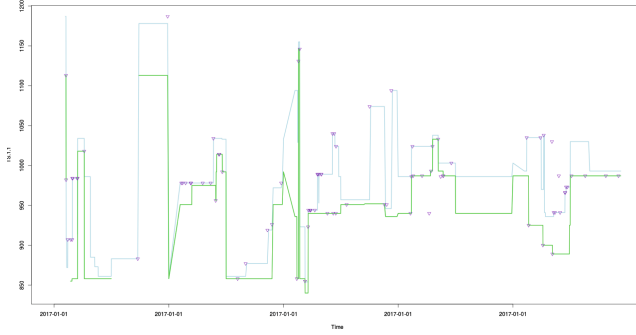


Fig. 2. Prices and quotes observed over a couple of seconds within the same day (i) Green line: best bid, (ii) Blue line: best ask, (iii) Purple triangles: prices (Color figure online)

the knowledge of each artificial agent and the one that is spread in its direct neighbourhood.

More interesting are the parameters that are used to gauge the efficiency of the market at the intra-day level:

Volatility – σ : we compute the average volatility of the market within each day using the standard deviation of the returns. We expect a higher volatility when the proportion of NT is important and when the impact of ST might be limited by the nature of the Network over which they behave (for example, in the RN setting where full connectedness is not granted).

Liquidity – λ : is simply in this case the average of stocks that are traded within a day. If a new estimate of the next dividend has little impact on the agents, whatever the reason, the exchanged volume should be relatively low compared to one observed if this information deeply modifies the composition of their optimal portfolio.

Efficiency – Ω : is the average of the absolute value of the deviation of prices with regards to the fundamental value. The lower this metric, the more efficient the market.

We first propose some graphical illustrations of our results.

In Figs. 3, 5 and 7, three matrix are nested, one for each network setting (the order is systematically the same (i): ER, (ii): PA and (iii): RN) these matrix represent using a colour (associated to a given level reported on the gauge at the right side of the matrix) one of the parameters we study (σ , λ , Ω).

For these matrix, rows are organized by mix of NT *vs.* IT. The total proportion of these sub-populations is always equal to 80% of the total population, but their proportions vary from 80% to 0%. ST always account for 20% of the population. All these figures have been chosen arbitrarily and may be discussed and challenged in an other study.

Columns simply report the values obtained during each run of the same experiment (*e.g.* 12 columns). We also provide an aggregate appreciation of the

results obtained over all the experimental rounds (in Figs. 3, 6 and 8, statistical tests not being displayed in this paper).

Volatility. We first observe how volatility seems linked to the nature of the network over which the information diffusion occurs and to the proportion of NT/IT. All results are presented in Fig. 3. Not surprisingly, the higher the proportion of IT the lower the volatility (remember that Ask and Bids are constrained by the boundaries around the estimate of the dividend: since IT have a narrower estimate, prices should vary less when they populate the market). Figure 4 summarized these values in calculating the average value observed within each network setting row by row (over 12 runs). Overall, the ER setting generates the highest average level of volatility 8.926% (all population mix and runs considered). By decreasing order of volatility come respectively the RN (8.463%) and the Preferential Attachment Network (8.328%). This may be linked to a better information diffusion in the PA network, provided at least some ST are located on vertice with a high degree of connectivity. However, as we can see in Fig. 4, this appears to be particularly true when the proportion of NT becomes relatively important (30% and more). One can imagine that in this case, IT become less determinant in the price emergence and that the nature of the Network tends to play a more important role in the diffusion of the best price estimate.

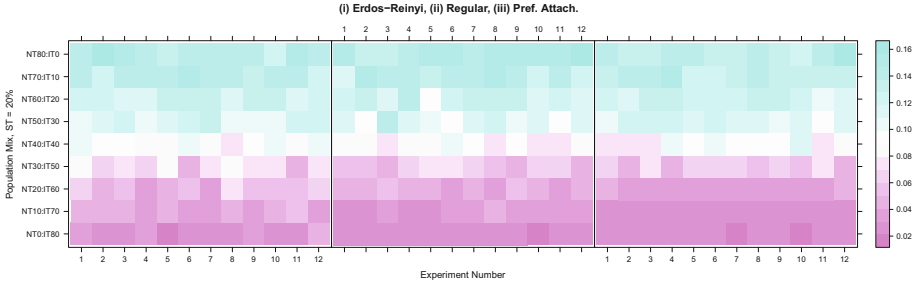


Fig. 3. Level of volatility over 20 days per experiment and mix of NT vs. IT (20% IT in any case) (i) By ranking orders, the overall highest level of volatility: (1) Erdős-Rényi, (2) Regular and (3) Preferential attachment networks (ii) Volatility is a monotonic function of the proportion of NT (iii) A similar threshold of 30% NT emerges whatever the network structure

Liquidity. The liquidity of the market is, on average, higher for the RN and the PA (resp. 471 and 464) than for the ER (446). Here again we observe a quasi linear effect: when the number of NT decreases, the liquidity of the market increases (see Fig. 5). The increase in the market liquidity is nonetheless different for each network setting: it requires 60% of IT for establishing a monotonic relationship for the ER, 50% for the RN, and 40% for the PA (see Fig. 6). Here

again, if the proportion of IT vs. NT appears to have an impact of this parameter, the Network structure also seems to play a role in establishing different levels of liquidity.

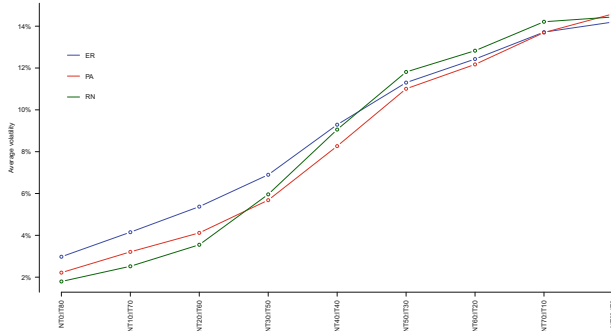


Fig. 4. Average level of volatility over all 12 experiments per mix of NT vs. IT (20% IT in any case) – The rate at which volatility increases the most can be observed for the RN case when the proportion of *NT* > 20%

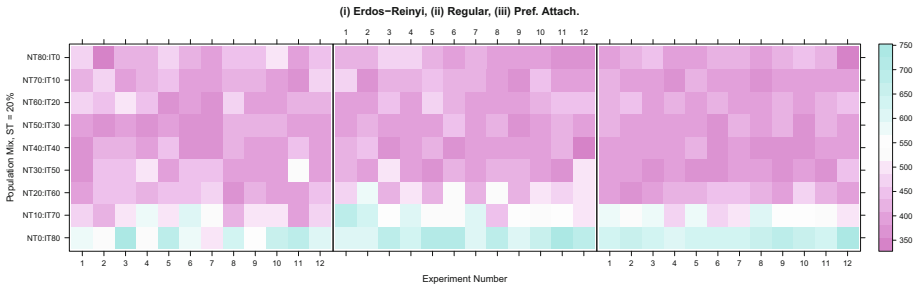


Fig. 5. Level of liquidity over 20 days per experiment and mix of NT vs. IT (20% IT in any case) (i) By ranking orders, the overall highest level of liquidity: (1) Regular and (2) Preferential Attachment and (3) Erdős-Rényi networks (ii) Liquidity decreases with higher proportions of NT (iii) The higher the availability of information, the higher the liquidity of the market

Efficiency. When it comes to efficiency, the best network structure is clearly the RN ($\Omega = 143$) while PA and RN obtain respectively values for Ω equal to 150 and 157. The overall impression in analysing Fig. 7 is that if some aggregate difference can be noticed, this seems not to be linked monotonically to the agent population mix as it was clearly the case for λ and σ . This result, although it must be confirmed and observed on larger experiments involving hundred of runs, would indicate that the single factor at play in the efficiency level (as we compute it her), would be the nature of the Network over which our artificial traders behave. This impression is also confirmed in Fig. 8.

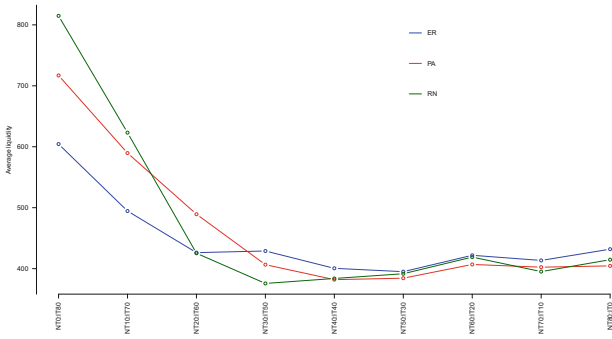


Fig. 6. Average level of liquidity over all 12 experiments per mix of NT vs. IT (20% IT in any case) – The rate at which liquidity increases monotonically varies among network settings

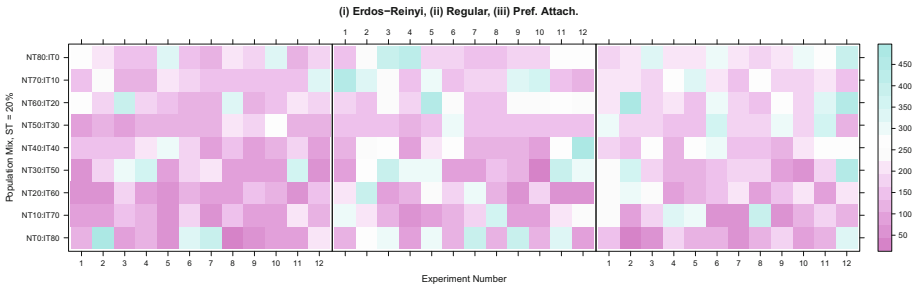


Fig. 7. Level of efficiency over 20 days per experiment and mix of NT vs. IT (20% IT in any case) (i) By ranking orders, the overall highest level of efficiency: (1) Erdős-Rényi, (2) Preferential attachment and (3) Regular networks (ii) Efficiency does not exhibit an evident relationship with the proportion of NT

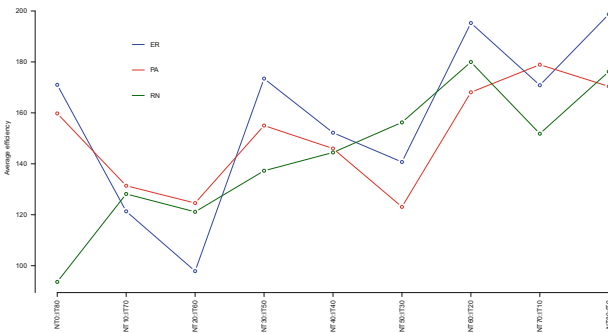


Fig. 8. Average level of efficiency over all 12 experiments per mix of NT vs. IT (20% IT in any case) – The behavior of Efficiency appears to be linked to the proportion of IT more strongly in the case of the RN setting (quasi monotonic relationship) w.r.t. the other settings where monotonicity cannot be identified

4 Conclusion

Considering that a huge proportion of financial transactions occur in OTC markets, we propose an investigation of the topology of the underlying social network over which they operate. Actually, we do not mimic a real OTC market but only focus on what we believe to be an essential feature of their nature: the underlying network. We thus address the question of the impact of this network over three essential parameters that are considered as essential for gauging market dynamics: liquidity, volatility and efficiency.

For doing so, we study how information is spread over market participants using a social-influence mechanism between three categories of agents (Noise, Informed and Social Traders). In this simplified world, traders reputation is a public signal allowing agents to estimate the reliability of the information located in their social neighbourhood, and eventually to prefer this latter to the one they own themselves. As such, the market network topology, in relation with the mix of traders populations and the signals they receive create an artificial framework geared at analysing information propagation and its subsequent, possible effects on prices, returns, and exchanged volumes.

Although simplified, this framework cannot be studied without a powerful multi-agent system. We do adopt this approach and show that, at least at some level, choosing a Regular, a Erdős-Rényi or a scale free network and locating on each node one Noise, Informed or Social Trader, substantially modifies the dynamics of the market. The overall level of volatility, the liquidity and the resulting efficiency are impacted by this initial choice in various ways which also depends upon the proportion of Informed vs. Noise Traders.

We believe these preliminary results could be explored and extended further, notably in addressing the way reputation, which is one main characteristics of the information propagation, actually evolves at fine grain and precisely modifies the dynamics within the network.

References

1. Agosto, A., Moretto, E.: Variance matters (in stochastic dividend discount models). *Ann. Financ.* **11**(2), 283–295 (2015)
2. Bajo, J., Mathieu, P., Escalona, M.J.: Multi-agent technologies in economics. *Intell. Syst. Acc. Financ. Manag.* **24**(2–3), 59–61 (2017)
3. Barabasi, A.-L., Albert, R.: Emergence of scaling in random networks. *Science* **286**(5439), 509–512 (1999)
4. Benhammada, S., Amblard, F., Chikhi, S.: An artificial stock market with interaction network and mimetic agents (short paper). In: International Conference on Agents and Artificial Intelligence (ICAART), Porto, Portugal, 24–26 February 2017, vol. 2, pp. 190–197. SciTePress (2017). <http://www.scitepress.org/>
5. Brandouy, O., Mathieu, P., Veryzhenko, I.: On the design of agent-based artificial stock markets. In: Filipe, J., Fred, A. (eds.) ICAART 2011. CCIS, vol. 271, pp. 350–364. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-29966-7_23

6. Erdős, P., Rényi, A.: On random graphs I. *Publ. Math. (Debr.)* **6**, 290–297 (1959)
7. Fama, E.F.: The behaviour of stock market prices. *J. Bus.* **38**, 34–105 (1965)
8. Fama, E.F., Fisher, L., Jensen, M.C., Roll, R.: The adjustment of stock prices to new information. *Int. Econ. Rev.* **10**(1), 1–21 (1969)
9. Gordon, M.J.: Dividends, earnings, and stock prices. *Rev. Econ. Stat.* **41**(2), 99–105 (1959)
10. Mathieu, P., Secq, Y.: Using LCS to exploit order book data in artificial markets. In: Nguyen, N.T., Kowalczyk, R., Corchado, J.M., Bajo, J. (eds.) *Transactions on Computational Collective Intelligence XV*. LNCS, vol. 8670, pp. 69–88. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44750-5_4
11. Nongaillard, A., Mathieu, P.: Agent-based reallocation problem on social networks. *Group Decis. Negot.* **23**(5), 1067–1083 (2014)
12. Phelps, S., Parsons, S., McBurney, P.: Automated trading agents versus virtual humans: an evolutionary game-theoretic comparison of two double-auction market designs. In: *Proceedings of the 6th Workshop on Agent-Mediated Electronic Commerce*, New York, NY (2004)
13. Veryzhenko, I., Arena, L., Harb, E., Oriol, N.: A reexamination of high frequency trading regulation effectiveness in an artificial market framework. *Trends in Practical Applications of Scalable Multi-Agent Systems, the PAAMS Collection*. AISC, vol. 473, pp. 15–25. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-40159-1_2