

Alternative Distributed Models for the Comparative Study of Stock Market Phenomena

Alexandros Benos*

Dept. of Finance and Economics
HEC School of Management
1, rue de la Libération
F-78351 Jouy-en-Josas Cedex
FRANCE

Elpida Tzafestas

LAFORIA, Institut Blaise Pascal
Université Paris VI
4, Place Jussieu
75252 Paris Cedex 05
FRANCE

Abstract

In this paper we present a methodology of study of complex phenomena emerging in stock markets. This methodology is based on the use of distributed, multi-agent models with minimal knowledge representation and reasoning capabilities that has proven a powerful modeling tool for complex biological systems. Unlike neural and “neo-connectionist” models, ours’ allow a comparative and incremental evaluation of their validity and relevance to the observed phenomena. The possibility of their application to the modeling and study of stock market phenomena is demonstrated on a simple example of a central agency that regulates the behavior of the investors : we show how a “blind” or myopic behavioral model reproduces results found in the literature and how the mutation of the model according to the parameters’ values or the adaptation structures gives rise to a series of complex phenomena comparable to those observed in reality.

1 Introduction

Why do prices fluctuate and under what circumstances is it possible to control them ? What happens if some “irrational” investors enter the market and do they survive ? Is it possible to predict economic crises or other phenomena from observation of a market for a limited period ? The first step to answering any one of these questions is to find suitable economic agent models supporting experimental evidence (excess volatility, survival of technical analysts, etc.) while offering the possibility to control for behavioural investment factors. This paper attempts to build a structure which will allow the researcher to closely simulate real financial markets and highlight their drawbacks and their resilience to shocks.

Our methodology is based on the use of distributed, multi-agent models with minimal knowledge representation and reasoning capabilities that has proven a powerful modeling tool for complex biological systems. Unlike neural and “neo-connectionist” approaches, ours’ allows a comparative and incremental evaluation of their validity and relevance to the observed phenomena. Algorithmic models rely on the modeler’s knowledge about the modeled system, while evolutionary ones rely on the potential to discover structures fit to the systems at hand. While the latter appear more robust to system tuning, the former demonstrate both minimality in respect to the relevant aspects of the problem and incrementality in respect to modeling.

Standard financial economics literature is centered around the paradigm of homogeneity. Investors, consumers or traders are supposed to be identical in beliefs. They can vary in wealth endowments (rich and poor), utility functions (risk averse or risk neutral) and information sets (insiders versus the uninformed public) but they all pertain to the same knowledge about the economic environment. Common knowledge of the structure of the market is implicitly assumed and no diverse priors due to biological differences are allowed. In consequence, transactions occur only because of different approaches to risk (hedging) ; common knowledge of differential information leads to common posteriors and no trade results.

* Contacting author. Benos gratefully acknowledges financial support from the Groupe HEC. We would like to thank Patrick Chappuy of SupAéro in Toulouse for his excellent programming assistance.

Such homogeneity deprives (financial) markets of their essential feature : people trade because of different interpretations of the surrounding world. Traditional economic theory, coupled with the rationality paradigm, exempts agents from different beliefs since it argues that only the (hyper?)-rational agent can survive in a competitive market. Errors, misperceptions, lack of adaptation, are all heavily penalised, leading to groups of rational fully optimising economic agents present in the economy. Hence, a vast number of exciting questions are swept under the carpet and paradoxes such as the information acquisition paradox [23], the no trade theorem [28] and the fully revealing rational expectations equilibrium [21] appear. What is even more bothersome is that such claims are made with recourse to “common sense” without any modelling background. Thus, there exists no explicit model as of today which shows that misperceptions about the market environment lead eventually to bankruptcy. On the contrary, [6] show that, under certain circumstances, a rational strategy may not be the fittest and that the one may not be rational.

Most research in this area was abruptly blocked by the lack of the machinery to study complex dynamic processes in time such as learning, adaptation, reaction to other people's actions, etc. Game Theory has only addressed repeated games in their simplest form ; only lately has it attacked learning and experimentation in games [17]. Financial microstructure theory has dealt with asymmetric information models in imperfectly competitive markets (see Section 3) but has only allowed for perfect private signals. Unfortunately, its results rely heavily on that and change dramatically once competition between them is allowed or imperfect signalling is introduced. Analytic study of dynamic models is, however, intractable since stochastic processes with imperfect information are not Markovian and, therefore, the entire path of prices, trade volume, etc. must be used by rational optimising agents.

Our simulated economy uses analytic results from financial theory whenever possible and approximates optimal actions when not. We allow for a number of possible beliefs in the market, check their fitness and control their effects on macro-variables. We find that heterogeneity of agents produces excess price volatility and enriches the set of observable phenomena. The fitness of a specific agent or a behaviour class is only a relative measure ; one's survival depends on the nature of one's opponent, so generalisations become far from credible. Finally, too much information destroys the flavour of markets and too little makes them disorderly. Privately informed investors add to the stability of the system while uninformed speculators give it potential richness.

2 Multi-agent Modelling

2.1 From multi-agent systems...

Distributed artificial intelligence [7, 18, 25] is the subfield of artificial intelligence that is concerned with distributed rather than mono-agent systems and has been traditionally subdivided into distributed problem-solving (a good review may be found in [13]) and multi-agent systems [10, 11, 15]. The first domain emphasises in solving a particular problem by distributing it among a number of modules (or nodes) that cooperate by sharing knowledge/solutions etc., while the second deals with coordinating intelligent behaviour among a collection of (possibly preexisting) *autonomous intelligent agents*. The themes of distributed artificial intelligence and multi-agent systems research include task description, decomposition, distribution and allocation, interaction and organisation, coherence and coordination, modelling of other agents, interagent disparities, such as uncertainty and conflict etc.

An agent is a real or abstract autonomous entity [15] capable of acting on itself and on its environment, that holds a partial representation of this environment, that is allowed to coexist and communicate with other agents while in pursuit of its individual goals and whose behavior is therefore the consequence of its observations, its knowledge and its interactions.

Multi-agent systems are typically used for engineering and design purposes, to achieve highly performant systems whose behavioral complexity and efficiency is a result of interactions between the individuals, the agents. The gain lies in that the potentially complex goal of the system is achieved as a side-effect of the interactions between otherwise simple individuals. The state of the overall system is “distributed” over the agents and the overall state space may grow exponentially on the number of them. This implies as well that scaling of agent models multiplies the complexity of the achieved goals ; *complexity is therefore hidden during analysis and multiplied during synthesis*.

For more experimental purposes, such multi-agent modelling may be used as a *complexity management tool* for processes that involve multiple identifiable entities loosely-coupled through a shared interaction medium [16]. Such observations have been made for instance in the field of behavioral ecology where population models used in the past have been homogeneous and insensitive to data local in space and time, simply because the currently

available mathematical analysis tools are insufficient. A new research tool has been consequently invented and applied, that of multi-agent or individual-based modeling and simulation [26].

2.2 ...to multi-agent modelling

Multi-agent models assume the existence of many independent agents, that “live” and act in a simulated world. The power of approach lies in a web of possibilities that may be exploited but whose analytic formulation quickly becomes intractable. Namely, such agents are allowed to be initially *heterogeneous* in all respects (hence preserve an individuality) and to further specialize by adaptation and learning. Secondly, multi-agent modeling presents a strong methodological advantage, that of *incrementality* : agent -as well as world- models may be complexified step-by-step, what allows for the effects of successive constraints to be isolated for study. Last but not least, the complexity of a multi-agent system is generally *emergent*, i.e. phenomena that arise are generally *not* preprogrammed within the agents, so that minor modifications to the agent level show as major behavioral differences to the overall system level. This property also implies that *control, too, is emergent* : the critical parameters whose tuning/regulation is necessary to adjust the system to different initial environmental conditions are not directly related to those aspects of reality they “control”. This way, the (self)-control of the system is achieved through coupling with the world rather than through explicit representation of it : while this kind of modeling appears to demand a degree of intuitive understanding of the system by the modeller, the rise of observable complexity with only small-scale changes of the model is tremendous, and so is the gain in modelling time. An accompanying feature is that of integration of *heterogeneous time scales* within the same multi-agent system, when for example agents act at their own pace. For all these reasons, the analytic description of a multi-agent system becomes infeasible, so that multi-agent models have to be assessed and validated by simulation rather by analysis. The results reported in section 5 have been obtained by simulating the behavior of our market model. The simulations have been performed on a personal computer and the programming tool used is a multi-agent simulator written in Smalltalk-80, version 4.1.

3 Financial Market Modelling

Exchange mechanisms in financial markets are numerous, making a theoretical modelling of transactions for all capital markets an impossible task. We can distinguish mechanisms - and corresponding market structures - by their price quotation system, the transmission and execution procedures, the role of intermediaries or, lastly, the quantity and quality of information owned by traders during the exchange.

Costs and risks sustained by different market participants vary with the mechanism followed. It is hence straightforward to conclude that market structure exerts a paramount influence on order submission strategies. Since under liberal exchange, prices are determined by the meeting of supply and demand, asset prices depend on the structure of the market on which the assets are quoted. Market microstructure theory tries to explain the links between a particular mechanism of exchange, order placement strategies by traders and asset price characteristics.

After the early revolution of rational expectations equilibrium and the foundational problems associated with it, research turned to explicit modelling of trading mechanisms.¹ A number of alternative models have been developed, first focusing on the role and objectives of intermediaries [2, 8], and later on information issues.² It is now widely agreed that we can distinguish market structures, first, according to the moment of the exchange, and second, according to the counterpart of the exchange. The first criterion differentiates between call markets, where execution of orders takes place periodically, and continuous markets, where execution is continuous as long as a counterpart is found. The second criterion distinguishes between a quote-driven and an order-driven market. In the former, market makers guarantee liquidity by standing ready to service any order which arrives to them at announced bid and asked prices. The NASDAQ, the London Stock Exchange and the OTC derivatives market are the main examples of this type of organisation. The latter are markets where buy and sell orders are directly matched between customers with no intermediaries but the customer brokers. Liquidity is guaranteed by limit orders submitted by clients. Most centralised stock and derivative exchanges in Europe are organised in this way. Of course, there is a number of mixed structures like the NYSE where limit orders exist in tandem with a specialist or the pits of the CBOT and the CME where brokers may choose between going into bilateral negotiation among themselves and trading with a market maker.

¹ Most of the pioneering work in Rational Expectations Equilibrium has been done by Sanford Grossman in the late seventies and early eighties. All of his influential papers on the subject are collected in [22].

² For an excellent review of this literature, see [1] and the multiple references cited therein.

It is, nevertheless, essential, before we continue onto modelling the structure of a financial market, to emphasize the role of information in such markets. In most examples of economic markets, consumers and producers engage in transactions because of different preferences and/or wealth as expressed by diverse utility functions and initial endowments. Financial markets stand apart in the sense that transactions are often caused by different information and/or beliefs. Following the original approach [14], when a large number of rational profit-maximisers compete in a securities market where information is widely available, their competition will cause the effects of new information on intrinsic values of the assets traded to be “fully reflected” in actual prices. In its extreme but empirically invalidated form, even private information is instantaneously revealed by prices and privately informed investors earn only a competitive expected return from their trading.³

We obviously see that the two notions important in the concept of market efficiency are information and rationality. Market efficiency is inextricably linked to a particular information set : past prices, public information (i.e. available at zero cost) and private, proprietary information. In this paper, we study all three kinds of information by defining trading strategies with respect to each one and then running them against each other. Empirical research on real markets show that public information is more or less incorporated into current prices quickly enough leaving no abnormal returns for its holders.

On the other hand, rationality is casually assumed in the most basic definition of market efficiency but is far from innocuous. It goes beyond simple profit-maximisation into knowledge of the market (structure of trading, joint distribution of returns and signals, etc.), mathematical ability and a full hierarchy of beliefs concerning both natural and strategic uncertainty. Clearly, a mechanistic model of objective information driving prices does not capture heterogeneous beliefs, private information production / processing and behavioural hypotheses such as the “animal spirits” of Keynes. In this essay, we introduce a very light deviation from full rationality, allowing overconfidence of investors : once they learn something, they think they know it perfectly. A truly efficient market would notice such “imperfect” behaviour and by acting on it, remove it. Thus, overconfident investors would not survive in an efficient market.

The rest of this section takes a closer look at market structure theory. Most papers have focused around the simplest of both trading worlds : the quote-driven one with specialists quoting prices, and the order-driven one where traders exchange assets among them with no intermediaries. [27] develops in a seminal paper a version of the simple call auction market. In the latter, each agent submits to a central agency which announced a price, an order to buy or sell a number of securities based on information he might have privately acquired, the analysis he has previously done on fundamentals or technical details of the security, etc. These orders are then cleared by an auctioneer who uses a tatonnement process (raise the price if there are excess buyers, lower it if there are excess sellers) to compute an equilibrium price. [27] replaces the auctioneer with a market maker who after receiving the net order flow, services it from her own inventory. All orders are executed at a price which reflects the best estimation of the asset value by the market maker.

Suppose a unique investor receives a perfect signal on the intrinsic value of the security. He knows his demand will have an impact on the price at which he will trade. A part of his information will be revealed through his action. He hence submits a market order x , i.e., an order to buy or sell a certain number of shares, which is assumed linear in his private information :

$$x = \beta (s - p_0)$$

where s denotes his private signal and p_0 the initial price of the stock (its unconditional expected value).

The “efficient” price posted by the market maker is the expected value of the stock given the total order flow she observes. The total order flow is made up of the insider's demand x and of the uninformed public's demand u adding noise to the system. This noise is necessary in order for the insider to be able to exploit his informational advantage ; otherwise, a rational market maker would have been able to inverse the order flow and extract all of the private information. The price fixed by the market maker takes the form

$$p_1 = E(F|w)$$

³ There exists a huge amount of literature discussing the different theories of Efficient Markets and the multitude of empirical research on them. A very good survey is [4].

where F is the intrinsic value of the stock and w the net order flow. In a Gaussian world, the price is linear and can be written as $p_1 = p_0 + \lambda w$ with a positive constant λ expressing the inverse of liquidity.⁴ The lower is λ , the deeper is the market, that is the larger an order must be to influence price by a given amount.

If we increase the number of informed agents in the market and add personalised signals, we add a competition effect to the results : traders take into account the reactions of other agents. They act strategically. Nobody has perfect knowledge of the intrinsic value of the firm but instead they each receive a noisy signal s_i , and submit orders x_i , linear on this signal. The layout is the same with the market maker setting the price equal to a linear function of the net order flow.

The one-period model has been subsequently extended to many trading rounds. [27] extends his one-insider static model to many periods but his results are sometimes peculiar ; for example, the depth (or inertia) of the market stays constant throughout time and the insider exploits his signal in a cautious manner in order not to reveal an important part of it. [24] add a number of perfectly informed agents to it. In their model (HS), each insider receives the same perfect signal on the intrinsic value of the stock but has to conjecture his rivals' demands in order to estimate the price set by the market maker. It is now a question of dynamic optimisation over a finite horizon with perfect information.

There are other ways of modelling a financial market. [20] present a simplification of a continuous quote-driven market with a single market maker posting buy ("bid") and sell ("ask") prices at which investors can trade. They arrive one by one, following a predetermined stochastic process, and they can either have a precise information on the intrinsic value of the asset or decide to trade for liquidity purposes. The market maker cannot of course distinguish between the two motives for trade, so she once more revises her set prices according to the trades she executes. Apart from the arrival process and the execution mechanism, the revision of prices is the same as in the extensions of [27]. It is conditional on the order flow, possibly with a lag of one round. Both structures integrate asymmetric information among market participants into price dynamics. An extension of the call auction paradigm, outlined above, was nevertheless used as the core of the present simulation for simplicity reasons.

4 The Simulated Market

In the simulated economy, there is one risky asset and money. The liquidation or intrinsic value of the risky asset is a normal random variable, denoted F , with mean p_0 and variance ϕ . F is publicly revealed after the last trading round.

In Section 3, we have presented a number of approaches to model the exchange in a financial market. We have chosen to use [27] as the foundation of our approach because (i) it describes quite closely a number of real securities markets, namely those that are driven by orders without the explicit intervention of intermediaries ; (ii) it is implementable, that is we can use it to produce numerical values of parameters. This is not the case with Rational Expectations Equilibrium models like [23] for which no trading game exists. And (iii), it yields relatively simple analytical results in its basic form which are useful both to test the convergence of the simulation algorithm and to program. Finally, we can always complicate it by adding more agents, other sources of noise or extending it to many trading rounds. The quote-driven model of [20] is very hard to analytically solve. Therefore, in the market, traders will only submit orders "at best" (or market orders). They will not specify the price at which they are willing to exchange securities.

Finally, our model differs from other papers offering simulation results of financial markets. [5] uses neural networks to model the behaviour of artificially intelligent agents. Trade is bilateral in the sense that each agent is randomly matched with each other and a transaction occurs at the simple average of both agents expectation about next period's price. Formation of this expectation divides agents into three groups. "Smart" agents use the last average transaction price and its change, as well as the price they themselves last transacted and its change to form a prediction of next period's price. The procedure is modelled as a 3-layer, fully interconnected, feed forward neural network with back propagation learning. "Dumb" agents forecast only extreme values for next price and "naive" agents use only the last market price. Our model instead offers a complete specification of the environment, defines precise strategies for each type of investor (see next Section) which are founded on the widely accepted profit-maximising paradigm and allows for control of the updating of expectations. Trades are not the result of random matching but of calculated portfolio decisions of investors in a centralised market structure.

⁴ The exact analytic form of all coefficients and parameters can be found in the original articles.

[9] uses genetic algorithms to model imperfectly rational traders in a one-period futures market. They calculate a “fair” value for the contract according to a set of parameters (their “strategy”) and turn over this value to a central agency which calculates their median. This is the “equilibrium” price : anyone with a “fair” value above it is automatically a buyer and everyone with a “fair” value below it is automatically a seller. The spot price is calculated by an exogenous strategy both stationary and unknown to participants. Strategies are modified by a GA with traditional mutation and crossover operators. This model is clearly too far away from reality to serve as a simulation of a futures market, as its authors admit in their introduction. It only serves as a starting point to explore the effects of differential information and constraints on profitability of trading. The present paper adds a lot more structure to the market and to the algorithm, information and model sets of participants ; we regard it as a more complete vehicle for attacking the questions elegantly posed by [9].

4.1 Informed and Uninformed Agents

There are six types of agents in the model which are divided into three groups : the **uninformed public** (noise traders), the informed investors and the uninformed investors . The former buy and sell for liquidity reasons, submitting orders in a random way.

The informed investors receive private information about the intrinsic value of the asset. We assume such information is costless to them in order to abstain from information acquisition problems. One can think of it as fundamental information they learn because of their position in the firm issuing the risky security (insiders), or because of their employment as analysts in a financial analysis firm. Their labours lead them acquiring a noisy signal on the liquidation value F as follows :

$$s_i = F + e_i, \quad \text{with } e_i \approx \text{IIDN}(0, \varepsilon)$$

The personalised error e_i is independent of F . Informed agents are divided into two subcategories : the **rational** and the **overconfident** ones. The former know that their signal is noisy and take it into account in their trading strategy. The latter treat their signal as infinitely precise ; they think it equal to the realisation of F . Since signals are personalised, the transaction price reveals some information about the signals of other agents. In a multi period setting, a trader can therefore use past prices to infer his trading rivals’ signals. All informed agents behave strategically, that is they are aware of their impact on the equilibrium price.

Uninformed investors do not have any knowledge of the fundamental value of the security. They, however, trade for speculative reasons and may actually be a clear majority in the markets. We model this group as investors who use technical analysis as a decision tool. The only unacceptable outcome for them is a huge loss ; they are willing to take on a moderate loss since, in their opinion, it is on average phased out by moderate gains. If they achieve to eliminate big losses, their portfolio can only grow by the favourable effects of big gains. According to its proponents, fundamentals help to decide which security to buy ; technical analysis provides the correct timing by postulating that prices more or less follow trends. It is then a question of how sophisticated one wants to be and how much does one believe the trend is right.

In this paper, we model three distinct behaviours of uninformed investors. The first one follows **momentum** strategies : they follow market trends as they think “the market is right” and buy when the net order flow has been positive and sell otherwise. The second group follows **contrarian** strategies. They believe the trend will soon die off and go against the market, buying when the net order flow was negative. The last group is the **chartists** who use one of the graphical tools of traditional technical analysis (see last Subsection) All groups of participants submit market orders and do not condition on settlement prices.

4.2 The Rational Informed investor

There are n rational informed investors in the market, each receiving a private signal as outlined above. In a one-period static world, it is not hard to solve analytically the pricing schedule of the market maker and the optimal strategies of all rational agents. Since it is an extension of the Kyle (1985) model, we use the structure outlined in Section 2, with a linear price being set equal to the expectation of the intrinsic value of the asset conditional on the order flow the market maker observes :

$$\begin{aligned}
p_1 &= p_0 + \lambda \sum_{i=1}^n x_i \\
&= p_0 + \frac{\varphi \sqrt{n(\varphi + \varepsilon)}}{\sigma_u^2 [\varphi(n+1) + 2\varepsilon]} \sum_{i=1}^n x_i
\end{aligned}$$

We see market depth increases with the variance of the random orders submitted by the uninformed public, σ_u^2 , and decreases with φ . There are no monotone results for n and ε : a small number of poorly informed investors may decrease liquidity.

This is still not a multiperiod model. We can try and extend it to T trading rounds but it very quickly becomes intractable since the process of prices is no longer Markovian.⁵ For two trading rounds, we can get some (numerical) results though. When the game starts ($t=0$), each informed agent i receives his noisy signal s^i . He then submits a market order conditional on his signal, x_{1i} . In the second period ($t=2$) he can observe last period's price, p_1 , and condition on it his demand this period, x_{2i} . The market maker sets p_1 by conditioning on the order flow of the first period and $p_2 = p_1 + E(F/w_2)$. Conditional expectations are now much more complicated. Details are given in the Appendix. We can write

$$\begin{aligned}
x_{1j} &= \beta_1 (s_j - p_0) \\
x_{2j} &= \beta_2 (s_j - p_0) + \gamma_2 (p_1 - p_0) \equiv \beta_2 (s_j - p_0) + \alpha_2 \\
p_1 &= p_0 + \lambda_1 \sum_{j=1}^n x_{1j} \\
p_2 &= p_1 + \lambda_2 \sum_{j=1}^n x_{2j}
\end{aligned}$$

for constants $\beta_1, \beta_2, \gamma_2, \lambda_1$ and λ_2 . We can still not get analytic results but only numerically find solutions for the constants by providing explicit values for parameters like φ, ε and the like. We thus, faced two major problems. First, it is unpractical to ask the simulator to solve numerically for the constants every time we needed to change a parameter. Secondly, we only had solutions of any kind for two periods ; what if we let the simulation run for more ?

Both problems were solved by trying to express solutions to our model to the known analytical solution to the model of Holden and Subrahmanyam (1992) who deal with perfect information ($\varepsilon=0$).⁶ We therefore, searched for a relation of the form

$$\text{Results}(\varepsilon) = f[\text{Results}(0), \varepsilon]$$

By standard statistical analysis, we found

$$\frac{1}{\beta_t^2} = \frac{1}{(\beta_t^2)_{\varepsilon=0}} + k_t \varepsilon \quad \text{and} \quad \alpha_2 = g \varepsilon \beta_2$$

for $t=1, 2$ with a correlation coefficient of 0.9999 and slope coefficients of $k_1 = n/\sigma_u^2 = k_2$ and $g = 1/n$ (significant at the 1% level). We then assumed that, at trading round t , the insider considers only prices up to two periods in the past and neglects older ones. Under this assumption, we can write the demand strategies for the rational insiders as follows :

$$x_{ij} = \beta_t \left[(s_j - p_{t-1}) - \frac{\varepsilon}{n} (p_{t-1} - p_{t-2}) \right]$$

⁵ See articles by Genotte (1986) and Detemple (1986) and the book by Åström (1970).

⁶ The analytic details of their solution are given in the next Subsection which deals with the modelling of overconfident informed investors.

Given the coefficients calculated for ($\varepsilon=0$), we could update for the case of noisy information for any trading period t . We chose, however, to model the trading decisions of rational insiders in a recursive fashion over a 2-round horizon. This was done, firstly, because calculations of the perfect information coefficients was computationally heavy after two rounds, and, secondly, because the theoretical model of HS produced aberrant results, as will be shown below. Instead of the strategy above, we chose to program rational informed traders as submitting orders :

$$x_{2t+1,j} = \beta_1 (s_j - p_{2t})$$

$$x_{2t+2,j} = \beta_2 \left[(s_j - p_{2t+1}) - \frac{\varepsilon}{n} (p_{2t+1} - p_{2t}) \right]$$

Note, finally, that rational insiders take into account they are trading against other informed traders, both rational and overconfident, but not that there are uninformed investors as well in the market. That is why they are using n , the number of all informed, when calculating their orders.

4.3 The Overconfident informed investor

This type of insider treats his noisy signal as the exact intrinsic value of the security and not as an approximation of it. He overestimates the precision of his signal in an extreme way. We can therefore use the model of HS to describe him. HS solve analytically a multiperiod trading model with insiders receiving a perfect signal on the liquidation value. Our overconfident type behaves as if he had received a perfect signal, so their results are perfectly applicable here. They show there exists a unique linear equilibrium characterised by the following system of difference equations :

$$\beta_t = \frac{1 - 2\alpha_t \lambda_t}{\lambda_t [m(1 - 2\alpha_t \lambda_t) + 1]}$$

$$\lambda_t = \frac{m\beta_t \Sigma_t}{\sigma_u^2}$$

$$\alpha_t = \frac{1 - \alpha_{t-1} \lambda_{t+1}}{\lambda_{t+1} [m(1 - \alpha_{t-1} \lambda_{t+1}) + 1]^2}$$

$$\Sigma_{t+1} = (1 - m\beta_{t+1} \lambda_{t+1}) \Sigma_t$$

for trading periods 1 to T-1, where m is the number of perfectly informed (overconfident) and with boundary conditions :

$$\alpha_T = 0$$

$$\beta_T = \frac{1}{\lambda_T (m + 1)}$$

$$\lambda_T = \frac{1}{\sigma_u} \sqrt{\frac{m \Sigma_T}{m + 1}}$$

$$\Sigma_1 = (1 - m\beta_1 \lambda_1) \Sigma_0,$$

for constant Σ_0 and m the number of insiders. One can solve the system above by recurring to backward iteration of a cubic first-degree difference equation. In the numerical illustrations they include in their paper, the liquidity parameter, λ_t , decreases extremely rapidly with t . They do not show, however, that the trade intensity parameter, β_t , rises equally fast resulting in a change of one order of magnitude after three (!) trading rounds. It is, of course, mathematically correct but economically impossible, to attain an (almost) infinitely deep market with an infinite volume. This is why we decided to truncate the trading horizon of all informed investors to two rounds. Finally, note that the overconfident insiders take m to equal the total number of participants in the market. It is natural for an aggressive trader to consider he is competing with other aggressive traders as well.

4.4 Momentum and Contrarian strategies

Their decision rule is very simple. Following a momentum strategy is based on the premise that price continuations predominate, it implies positive correlation in returns and considers the market initially underreacts to news. Hence it proposes buying a stock that has risen in value and selling one that has fallen in value. On the contrary, a contrarian theory of the market thinks the market overreacts, that there is stock mispricing relative to fundamentals which will cause price reversals as errors are corrected. It proposes being long on “loser” stocks and short on “winners”.

We can model both strategies by looking at the absolute value of the last price change : if it is greater than a certain threshold, then the strategies activate (use of a threshold implies that the agents demonstrate some inertia in their decisions). Thresholds (Y in the following table) differ among agents and they are normally distributed around 3%.

	$\frac{ p_t - p_{t-1} }{p_t} \geq Y\%$		$\frac{ p_t - p_{t-1} }{p_t} < Y\%$
	$p_t \geq p_{t-1}$	$p_t \leq p_{t-1}$	
Momentum	Buy	Sell	Do nothing
Contrarian	Sell	Buy	Do nothing

We have normalised the quantities traded by these investors to avoid excessive price. Each one submits a buy (sell) order of

$$\frac{\text{Previous Net Order Flow}}{\sum \text{Buys} - \sum \text{Sales}}$$

Recall that we use negative values to represent sales.

4.5 Chartists

Chartists use graphical tools provided by technical analysis which, despite its condemnation by the academic world, is still widely used, at least in the short term. The simplest chartist rule is the “moving average” one. One first calculates two moving averages of prices, a shorter (5 days) and a longer (20 days) one. The former identifies short term trends and the longer, long-term ones. The trading signal is given by the crossing of the two. For example, when the 5-day (weekly) MA crosses the 20-day (monthly) one, it means that price is on a stronger short trend than the previous long period. One must buy since prices will continue rising. The end of the trend is signalled by a new crossing of the MAs, which gives a “sell” signal. Once again, quantities are normalised accordingly to avoid manipulations.

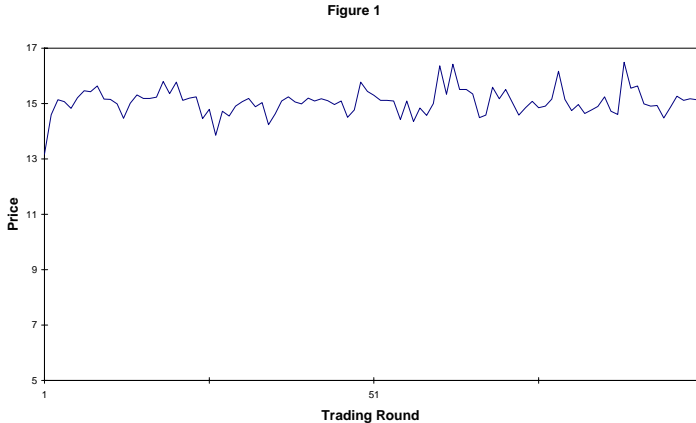
4.6 The Market Maker

The market maker has a very difficult signal extracting problem in our setting. She observes a net order flow coming from a wide range of trading strategies. Even had she have the entire market structure in mind (with the exact size of each group, their beliefs, their thresholds, etc.) the conditioning would not have been computable. Given this complexity problem, we postulate she is not able to correctly condition on it in order to set an efficient price. Instead, she considers dealing against the most aggressive of groups, the overconfident informed investors. She is hence using the parameters calculated in the HS model for a two-round horizon to set her pricing schedule. This is a shortcoming of the present set-up whose effects are clearly visible on the market maker’s wealth and stock inventory.


5 Main results

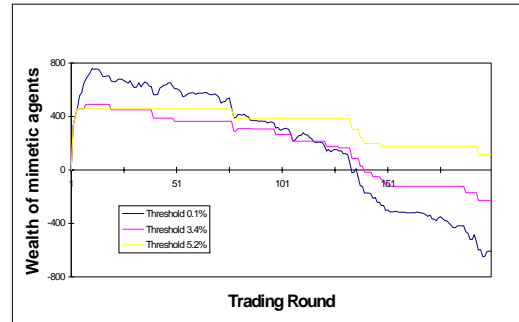

We start by introducing only informed agents in the market and we observe the equilibrium price converges quickly to the mean of private signals. The market is strongly efficient (see Figure 1), a result predicted by most of the theoretical literature [23]. Introducing a reasonable number of non informed investors perturbs price

efficiency and increases price volatility substantially (e.g. from 0.3 to more than 1). There emerge clear trends in the price process (see Figure 2) which does not converge to the mean of signals anymore.



The performance of the rational strategy depends on the quality of the specific signal and the mean of signals of its kin. The wealth of overconfident agents depends on the same parameters but exhibits a larger variability which may occasionally lead to higher returns. As far as the non informed group is concerned, gains depend on the relative size of each group and on the threshold used as shown in Figure 3.





Finally, chartists lose money consistently . This is due to their lagged response to price changes. In Figure 4, we see the chartist is selling one period after seeing the sell signal (point B) and the price has fallen ; he then buys one period after having seen the buy signal (point A) and the price has already gone up ! This is due to the inertia of the MA.

Figure 3




Figure 4

6 Conclusions and perspectives

We have sketched the multi-agent modelling methodology and justified its suitability for financial modelling research. We have demonstrated our claims on a simulated market where traders are modelled individually and have discussed a number of observations : diversity of agent types induces excess volatility and absence of information perturbs the market efficiency. Furthermore, the fitness of individual agents has been shown to be a relative measure that depends on the exact market setting : the same agent may lose or gain money in different settings according to the proportions of agent types in the market and the relations between agents' parameters. All these results are substantial in that they have real-life analogues, although the models are still primitive and our study far from complete.

The next stage of the modelling will therefore aim at evaluating the correspondence between agent models and obtained phenomena : under what circumstances do irrational agents manage to survive and how can we characterize the evolution at the market level ? These get us to the question of adaptivity : may any behavior calibration or adaptation mechanism allow to irrational agents to survive in the market and does it make sense to change completely one's behavioral profile (and from mimetic become, say, contrarian) ? Are previous models still operational when the market has multiple assets and which is the reaction of the market to abrupt perturbations, such as a macroeconomic shock ? We believe that the presented methodology has the right potential for this kind of research and our investigations seem to support our belief.

References

- Admati, A.R. (1991). The informational role of prices, *Journal of Monetary Economics*, 28: 347-361.
- Amihud, Y. and H. Mendelsohn (1980). Dealership market: Market making with inventory, *Journal of Financial Economics*, 8: 31-53.
- Åstrom, K.J. (1970). *Introduction to Stochastic Control Theory*, Academic Press, NY.
- Ball, R. (1994). The development, accomplishments and limitations of the Theory of Stock Market Efficiency, *Managerial Finance*, 20 (2/3): 3-48.
- Beltratti, A. and S. Margarita (1993). Evolution of trading strategies among heterogeneous artificial economic agents, in J.-A. Meyer, S.W. Wilson & H.L. Roitblat (Eds.), *From Animals to Animats (2), Proceedings of the 2nd International Conference on Simulation of Adaptive Behaviour (SAB '92)*, MIT Press.
- Blume, L. and D.Easley (1992). Evolution in market behaviour, *Journal of Economic Theory*, 58: 9-40.
- Bond A. H. and L. Gasser, Eds. (1988). *Readings in Distributed Artificial Intelligence*, Morgan Kaufmann, San Mateo, CA.
- Cohen, K., S. Maier, R. Schwartz and D. Whitcomb (1979). Market makers and the market spread : A review of recent literature, *Journal of Financial and Quantitative Analysis*, 14: 813-833.
- de la Maza, M. And D. Yuret (1994). A futures market simulation with non-rational participants, in R.A. Brooks & P. Maes (Eds.), *Artificial Life IV, Proceedings of the 4th International Workshop on the Synthesis and Simulation of Living Systems*, Bradford/MIT Press.
- Demazeau Y. and J.-P. Müller, Eds. (1990). *Decentralized A.I., Proceedings of the First European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW '89)*, Elsevier/North-Holland.
- Demazeau Y. and J.-P. Müller, Eds. (1991). *Decentralized A.I. 2, Proceedings of the Second European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW '90)*, North-Holland.
- Detemple, J. (1986). Asset pricing in a production economy with incomplete information, *Journal of Finance*, 41 (2): 383-392.
- Durfee E. H., V. R. Lesser and D. C. Corkill (1989). Trends in cooperative distributed problem solving, *IEEE Transactions on Knowledge and Data Engineering*, 1(1):63-83, March 1989.
- Fama, E. F. (1965), Random walks in stock market prices, *Financial Analysts Journal*, 21 (5): 55-59.
- Ferber J.(1989). *Des objets aux agents : Une étude des structures de représentation et des communications en Intelligence Artificielle*, Thèse d'État, Université Pierre et Marie Curie, Paris, June 1989.
- Ferber, J. and Drogoul, A. (1992). Using reactive multi-agent systems in simulation and problem-solving, in N. M. Avouris & L. Gasser (Eds.), *Distributed Artificial Intelligence : Theory and Praxis*.
- Fudenberg, D. and D.M.Kreps (1988). A theory of learning, experimentation, and equilibrium in games, Working Paper, MIT and Stanford.
- Gasser L. and M. N. Huhns, Eds. (1989). *Distributed Artificial Intelligence*, Vol. 2, Pitman Publishing/Morgan Kaufmann, London/San Mateo, CA.
- Genotte, (1986). Optimal portfolio choice under incomplete information, *Journal of Finance*, 41 (3): 733-746.
- Glosten, L. and P. Milgrom (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics*, 14: 71-100.
- Grossman, S.J. (1981). An introduction to the theory of rational expectations under asymmetric information, *Review of Economic Studies*, 48: 541-559.
- Grossman, S.J. (1989). *The Informational Role of Prices*, MIT Press.

- Grossman, S.J. and J.E. Stiglitz (1980). On the impossibility of informationally efficient markets, *American Economic Review*, 70: 393-408.
- Holden, C.W. and A. Subrahmanyam (1992). Long-lived private information and imperfect competition, *Journal of Finance*, 47 (1): 247-270.
- Huhns M. N., Ed. (1987). *Distributed Artificial Intelligence*, Pitman Publishing/Morgan Kaufmann, London/San Mateo, CA.
- Judson, O.P. (1994). The rise of the individual-based model in ecology, *Trends in Ecology and Evolution*, 9(1 January):9-14.
- Kyle, A. (1985). Continuous auctions and insider trading, *Econometrica*, 53: 1315-1335.
- Milgrom, P. and N. Stokey (1982). Information, trade, and common knowledge, *Journal of Economic Theory*, 26: 17-27.

Legends of Figures

Figure 1. The price process in a simulation over 100 trading periods with 5 rational and 5 overconfident informed investors who solve a two-round problem recursively. Parameters are set as follows: $\phi = 4$, $\epsilon = 1$, $p_0 = 10$, $F = 15$, Noise variance = 1. The standard deviation of the price is 0.49 and the process resembles white noise around a mean of 15.04.

Figure 2. The price process in a simulation over 100 trading periods with 2 rational and 2 overconfident informed investors as well as 3 chartists, 2 contrarian and 1 momentum uninformed investors. Parameters are set as in Figure 1 and the standard deviation of the price is 0.80. There are clear trends in the process.

Figure 3. Wealth of uninformed investors following momentum and contrarian strategies. The results are taken from a 200-period simulation with 5 rational and 5 overconfident investors facing 3 momentum (or mimetic) agents and 2 contrarian ones. Parameters are set as in Figure 1. We observe that the finer the threshold, the better the performance.

Figure 4. The price process and moving averages from one simulation. It is clear buy (pt. A) and sell (pt. B) signals can be executed one round later. This inertia penalises the chartist.