
The “mind or no-mind” dilemma in agents behaving in a market

Pietro Terna

*Dipartimento di Scienze economiche e
finanziarie G.Prato - Università di Torino
corso Unione Sovietica 218bis
10134, Torino, Italia
terna@econ.unito.it*

ABSTRACT: In computer simulation models based upon agents, what is the degree of sophistication that we have to put into the agents? Should we provide them or not with a “mind”? The answer ranges from Axelrod’s simplicity principle to the use of full BDI (Beliefs, Intentions, Desires) cognitive agents. To discuss the subject we introduce here three models: one with “no-mind” agents that operate in an unstructured market, the second with “minded” agents assuring some stability to an emerging unstructured market and, finally, the third with no mind agents, that show a sophisticated outcome in a structured market. No generalised results come from this presentation, but many useful doubts.

RÉSUMÉ: Quel degré de complexité faut-il introduire dans un modèle de simulation fondé sur des agents? Doivent-ils disposer d'un "esprit"? Nous présentons ici trois modèles qui utilisent des agents, avec et sans "esprit", et ou les résultats obtenus sont très différents suivant la structure du marché dans lequel opèrent les agents.

KEY WORDS: economic simulation, agent based models, Swarm

MOTS-CLÉS: simulation économique, modèles fondés sur agents, Swarm

1. Introduction

According to [GIL 00]:

Ostrom [OST 88] proposed that there are three different “symbol systems” available to social scientists: the familiar verbal argumentation and mathematics, but also a third way, computer simulation. Computer simulation, or computational modelling, involves representing a model as a computer program. Computer programs can be used to model either quantitative theories or qualitative ones. They are particularly good at modelling processes and although non-linear relationships can generate some methodological problems, there is no difficulty in representing them within a computer program.

The logic of developing models using computer simulation is not very different from the logic used for the more familiar statistical models. In either case, there is some phenomenon that we as researchers want to understand better. This is the “target”. We build a model of the target through a theoretically motivated process of abstraction (this model may be a set of mathematical equations, a statistical equation, such as a regression equation, or a computer program). We then examine the behaviour of the model and compare it with observations of the social world. If the output from the model and the data collected from the social world are sufficiently similar, we use this as evidence in favour of the validity of the model (or use a lack of similarity as evidence for disconfirmation).

The question now is: if our computer simulation model is based upon agents (e.g. built with Swarm¹, as the models presented here), to what extent must our agents be sophisticated? Should we provide them with a “mind”? The answer ranges from Axelrod’s simplicity principle [AXE 97] to the use of full BDI (Beliefs, Intentions, Desires) cognitive agents. To discuss the subject we introduce here, in Section 2, a model with rigidly “no-minded” agents that behave in an unstructured market, with unrealistic cycles and chaos; in Section 3 a model with learning “minded” agents, that assure some stability to the emerging unstructured market; finally, in Section 4, we present no mind agents operating in a structured market, with a sophisticated outcome. No generalised results come from this presentation, but many useful doubts. See Section 5 for a comment about the necessity of integrating the different cases treated here in a unique framework, to improve the possibility of comparing the different situations.

¹ See <http://www.swarm.org>

2. Chaos from agents: “no-mind” agents in an unstructured market

This first simulation comes from [TER 98], and is based upon agents that apply mechanically their rules, without learning or adaptive capabilities; the market in which they behave is completely unstructured. The experiment shows the emergence of chaotic price sequences in a simple model of interacting consumers and vendors, both equipped with minimal rules². (We are not seeking to produce chaos: it emerges as a side effect of the agents' behaviour.)

2.1. *Technical details*

We have consumers and vendors. Each agent is capable of reacting to messages, for example deciding whether it should buy at a specific offer price. We use a shuffler mechanism to change the order in which the agents operate and to establish random meetings of the members of the two populations. At every simulation step, artificial consumers look for a vendor; all the consumers and vendors are randomly matched at each step. An exchange occurs if the price asked by the vendor is lower than the level fixed by the consumer. If a consumer has not been buying for one (if the sensitivity parameter is set to 1) or more than one (if the sensitivity parameter is set to 0) step, it raises its price level by a fixed amount according to the counter rule and the sensitivity parameter introduced below. It acts in the opposite way if it has been buying and its inventory is greater than one unit. A simulated vendor behaves in a symmetric way (but without a sensitivity parameter): it chooses the offer price randomly within a fixed range. If the number of steps for which it has not been selling is greater than one, it decreases the minimum and maximum boundaries of this range, and vice versa if it has been selling.

In all the experiments, the result is that the mean price behaviour emerges as cyclical, with chaotic transitions from one cyclical phase to another. From a methodological point of view there are two kinds of emergence.

- Unforeseen emergence: while building the simulation experiment, we were only looking for the simulated time required to obtain an equilibrium state of the model with all the agents exchanging nearly at each time. The appearance of a sort of cyclical behaviour was unexpected (the inventory cycle having been undervalued).

- Unpredictable emergence: chaos is obviously observable in true social science phenomena, but it is not easy to make a reverse engineering process leading to it as a result of an agent based simulation.

² The Swarm code of this experiment can be obtained directly from the author, it is named ABCDE, from “Agent Based Chaotic Dynamic Emergence”.

2.2. Results

We have now to summarise the parameters used in the experiments reported in the Figures of this paragraph. Parameters: “theLevel” is the initial price below which consumers buy (it changes independently for each consumer while the simulation evolves); “agentNumber” is the number of consumers and vendors; “minStartPrice” and “maxStartPrice” are the initial limits within which vendors choose their selling price (they change independently for each vendor while the simulation evolves); the “reactivityFactor” is a multiplying factor that enhances the fixed values used by consumers and vendors to modify their prices or range of prices; “sensitivity” is zero or one, with the meanings introduced above. The series reported in the Figures are the mean of global prices (all prices offered in each day or cycle) or the min. or max. within global prices. (The “alternative way” of the title is related to an internal calculus problem).

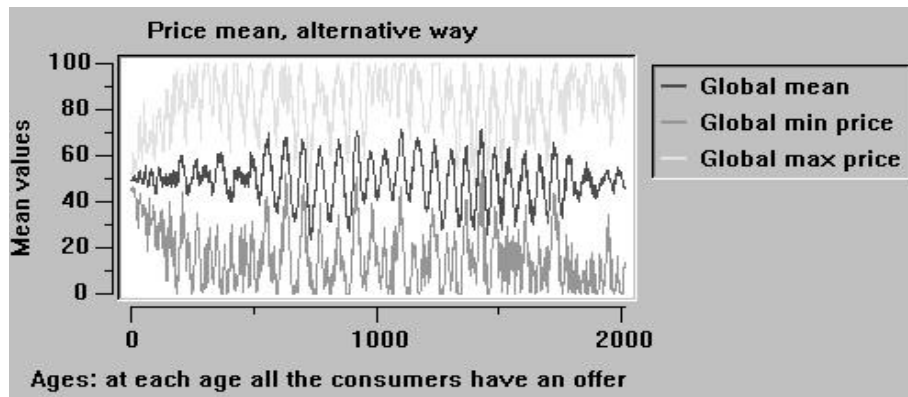


Fig. 1. Starting with a balanced situation, low reactivity and low sensitivity³.

The experiment reported in Fig.1 starts from a balanced situation of the buying and selling prices⁴. With a low reactivity factor and a low sensitivity parameter in the behavioural choices of the consumers, we have a relevant amplitude of the fluctuations, with chaotic appearance of the global mean of the prices offered in each cycle.

³ theLevel=50; minStartPrice=45; maxStartPrice=55; reactivityFactor=1; sensitivity=0.

⁴ Starting from an unbalanced situation the emergence of the cyclical behaviour and the chaos appearance are improved, with the presence of strange attractors in the singular value decomposition space

Technically: the FFT⁵ of the series of data shows only one peak, related to the constant value; Lyapunov exponents are in the range 0.6-0.7 and both capacity and correlation dimensions are less than 5. Also the experiment reported in Fig.2 starts from a balanced situation: the high reactivity factor improves the cyclical effect and the presence of the sensitivity parameter reduces the amplitude of the fluctuations; anyway, chaos appears.

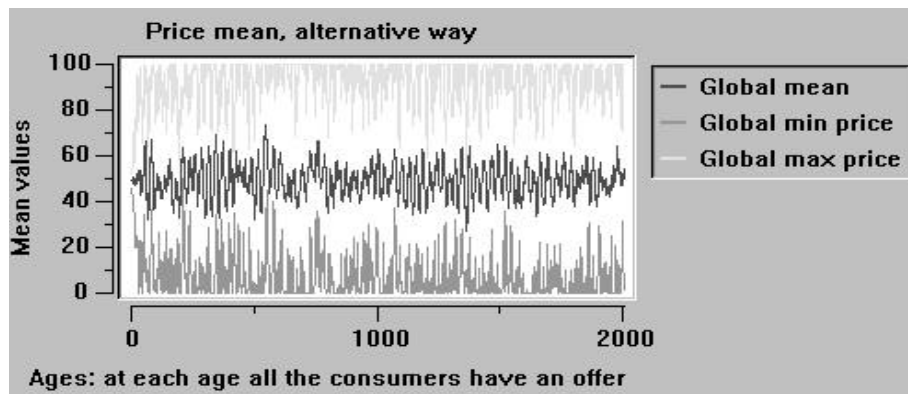


Fig. 2. Starting with a balanced situation, high reactivity and high sensitivity⁶.

The FFT shows the same results as above; Lyapunov exponents are in the range 0.5-0.8 and both capacity and correlation dimensions are less than 5.

2.3. Comments to the experiment

A general comment about chaos in these experiments: we are here in presence of endogenous chaos, emerging from agent interaction. We know, by construction, that here chaos is not related to a well specified structure, but to collective behaviour, which we are not able to describe ex ante in an analytical form. The price memory is diffused in the agents, having a collective effect of synchronisation or heterogeneous behaviour in the agents' reaction to the exchange prices.

The more interesting thing here, from the point of view of the “mind or no-mind” dilemma, is the rigidity of the cycles, without any form of stability in the market: our agents are performing their tasks without the aid of any form of adaptation; moreover, in a market that is not provided with any operating mechanisms.

⁵ Fast Fourier Transformation.

⁶ theLevel=50; minStartPrice=45; maxStartPrice=55; reactivityFactor=3; sensitivity=1.

3. Equilibrium from agents: minded agents in an unstructured market

The second model is a more sophisticated one, involving learning agents provided with goals, and able to develop an internal consistency between the actions they perform and the guesses on their effects [TER 00a]⁷.

In the Hayek sense of individualism and economic order [HAY 49] we are here dealing with agents behaving on a strictly individualistic basis, without the over-simplification of an artificial auctioneer market. A market with stable prices anyway emerges, with an empirical quasi-equilibrium. Note that we explicitly exclude from our model any form of prior description of this kind of equilibrium and that the agents have no knowledge about the mean price, which is only used by the observer to externally describe the experiment.

3.1. Technical details

To develop our agent based experiments, we introduce the following general hypothesis (GH): an agent, acting in an economic environment, must develop and adapt her capability of evaluating, in a coherent way, (1) what she has to do in order to obtain a specific result and (2) how to foresee the consequences of her actions. The same is true if the agent is interacting with other agents. Beyond this kind of internal consistency (IC), agents can develop other characteristics, for example the capability of adopting actions (following external proposals, EPs) or evaluations of effects (following external objectives, EOs) suggested from the environment (for example, following rules) or from other agents (for examples, imitating them). Those additional characteristics are useful for a better tuning of the agents in making experiments.

To apply the GH we are employing here artificial neural networks; we observe, anyway, that the GH can be applied using other algorithms and tools, reproducing the experience-learning-consistency-behaviour cycle with or without neural networks.

An introductory general remark: in all the cases to which we have applied our GH, the preliminary choice of classifying agents' output in actions and effects has been useful (i) to clarify the role of the agents, (ii) to develop model plausibility and results, (iii) to avoid the necessity of prior statements about economic rational optimising behaviour [BEL 96]. Economic behaviour, simple or complex, can appear directly as a by-product of IC, EPs and EOs. To an external observer, our Artificial Adaptive Agents (AAAs) are apparently operating with goals and plans. Obviously, they have no such symbolic entities, which are inventions of the observer. The

⁷ The Swarm codes - named bp-ct and ct-hayek - can be downloaded from the Swarm site (see <<http://www.santafe.edu/projects/swarm/users/user-contrib.html>>), or can be requested directly to the author.

similarity that we recall here is that the observations and analyses about real world agents' behaviour can suffer from the same bias. Moreover, always to an external observer, AAAs can appear to apply the rationality paradigm, with maximising behaviour.

We are considering learning agents, from the point of view of the bounded rationality research program [ART 90].

In designing a learning system to represent human behaviour in a particular context, we would be interested not only in reproducing human rates of learning, but also in reproducing the style in which humans learn, possibly even the ways in which they might depart from perfect rationality. The ideal, then, would not simply be learning curves that reproduce human learning curves to high goodness-of-fit, but more ambitiously, learning behaviour that could pass the Turing test of being indistinguishable from human behaviour with its foibles, departures and errors, to an observer who was not informed whether the behaviour was algorithm--generated or human-generated.

In one sense our agents are even simpler than those considered in neoclassical models, as their targets and instruments are not as powerful as those assumed in those models. From another point of view, however, our agents are much more complex, due to their continuous effort to learn the main features of the environment with the available instruments.

3.2. *The algorithm*

The name cross-targets (CTs) comes from the technique used to figure out the targets necessary to train the ANNs representing the artificial adaptive agents (AAAs) that populate our experiments. Following the GH, the main characteristic of these AAAs is that of developing internal consistency between what to do and the related consequences. Always according to the GH, in many (economic) situations, the behaviour of agents produces evaluations that can be split in two parts: data quantifying actions (what to do) and forecasts of the outcomes of the actions. So we specify two types of outputs of the ANN and, identically, of the AAA: (i) actions to be performed and (ii) guesses about the effects of those actions.

Both the targets necessary to train the network from the point of view of the actions and those connected with the effects are built in a crossed way, originating the name Cross Targets. The former are built in a consistent way with the outputs of the network concerning the guesses of the effects, in order to develop the capability to decide actions close to the expected results. The latter, similarly, are built in a constant way with the outputs of the network concerning the guesses of the actions, in order to improve the agent's capability of estimating the effects emerging from the actions that the agent herself is deciding.

We choose the neural networks approach to develop CTs, mostly as a consequence of the intrinsic adaptive capabilities of neural functions. The AAA has to produce guesses about its own actions and related effects, on the basis of an information set. Remembering the requirement of IC, targets in learning process are: (i) on one side, the actual effects - measured through accounting rules - of the actions made by the simulated subject; (ii) on the other side, the actions needed to match guessed effects. In the last case we have to use inverse rules, even though some problems arise when the inverse is indeterminate. For further technical explanations on the CT method see [BEL 96] or [TER 00a].

Some definitions: we have (i) short term learning, in the acting phase, when agents continuously modify their weights (mainly from the hidden layer to the output one), to adapt to the targets self-generated via CT; (ii) long term learning, ex post, when we effectively map inputs to targets (the same generated in the acting phase) with a large number of learning cycles, producing ANNs able to definitively apply the rules implicitly developed in the acting and learning phase.

A remark, about both external objectives (EOs) and external proposals (EPs): if used, these values substitute the cross targets in the acting and adapting phase and are consistently included in the data set for ex post learning. Despite the target coming from actions, the guess of an effect can be trained to approximate a value suggested by a simple rule, for example increasing wealth. This is an EO in CT terminology. Its indirect effect, via CT, will modify actions, making them more consistent with the (modified) guesses of effects. Vice versa, the guess about an action to be accomplished can be modified via an EP, affecting indirectly also the corresponding guesses of effects. If EO, EP and IC conflict in determining behaviour, complexity may emerge also within agents, but in a bounded rationality perspective, always without the optimisation and full rationality apparatus.

3.3. The hayekian market arising from the experiment

The experiment structure is briefly described below (a detailed definition of the structure is reported in the code ct-hayek, which is the adaptation of bp-ct to the specific experiment).

The consumer agent has the following information set: the expense of the previous period; the requirement of the previous period; the price of the unique good, as proposed by the agent in the previous period; the agent buying proposal in the previous period. The producer agent has in input: the revenue of the previous period; the production stream requirement of the previous period; the price of the unique good, as proposed by the agent in the previous period; the agent selling proposal in the previous period.

The production process is undefined; the producers are simply supposed to offer the quantity that they are producing in each time unit; consumers and producers

meet randomly each “day”. The exchange is possible if the producer proposed price is less or equal to the price proposed by the consumer; the price used in the exchange is that proposed by the producer and the quantity is the minimum between the consumer proposal and the producer proposal; if the exchange is not possible the reference price (used in the determination of the mean price of the day) is that of the consumer and the quantity is 0.

We introduce also external objectives (EOs) which are explained in detail in the distributed ct-hayek code. In the consumer case, the effect Expenditure is trained with a lowering target; the effect Requirement is trained with a constant target.

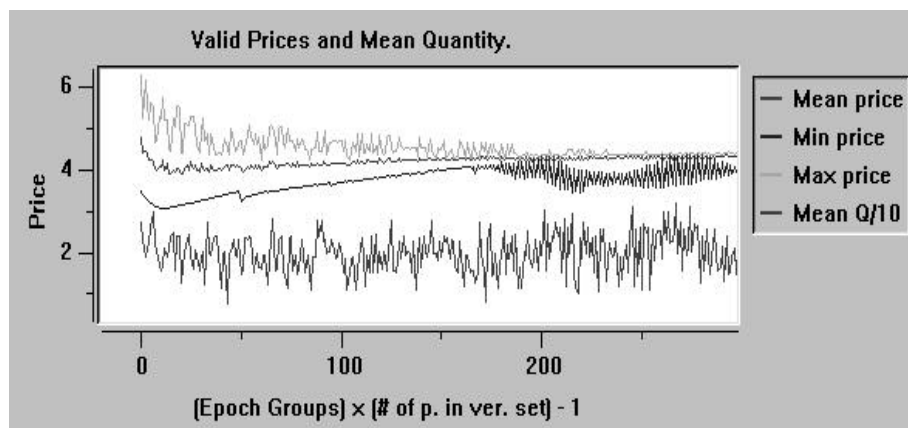


Fig. 3. Long term learning, without EOs.

The CT consequences are modification of the price guess, remembering that the EOs modifies the outputs of the effect side of the underlying ANN, but via CT also the action side. In the producer case we act symmetrically, increasing the Revenue target and keeping constant the other one; ProductionStream can be interpreted as the adequate production related to an existing fixed capital structure.

We have to read the lines of the graphics of the Figures as (bottom up): (i) mean quantity of the exchanges (divided by 10 for a scale necessity); (ii) Min price; (iii) mean price; (iv) Max price. Prices are obtained by the p variable described above. A market spontaneous equilibrium emerges immediately in our runs of the model. In the case with heavy relearning, reported in Fig. 3, we can see a robust equilibrium, with some individuals differentiating their behaviour. Operating simultaneously to diminish expenditures, raise revenues, stabilise requirements and production streams, we stabilise the market, with low oscillations in quantities, mainly when the agents are able to operate globally, after long term learning, as in Fig. 4.

3.4. Comments to the experiment

The more interesting thing here, from the point of view of the “mind or no-mind” dilemma, is the flexibility and adaptation ability of the agents that, while developing the internal consistency between guesses about their actions and guesses about the effects (of those actions), behave in a way that stabilise a market not provided with any form of operating mechanisms, as the previous one.

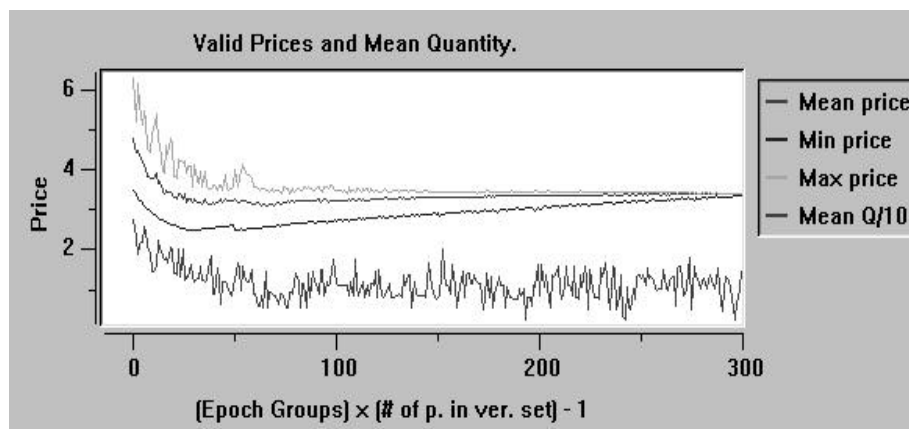


Fig. 4. Long term leaning, with EOs.

4. Bubbles and crashes from agents: “no-mind” agents in structured market

We are using here “no-mind” agents that operate in a rigid market mechanism. The Surprising (Un)realistic Market (SUM) model [TER 00b], deals with the micro-foundations of a stock market⁸. We avoid any artificially simplified solution about price formation, such as to employ an auctioneer to clear the market; on the contrary, our model produces time series of prices continuously evolving, transaction by transaction.

⁸ The Swarm codes - temporary named SUM - can be requested directly to the author.

4.1. Technical details

The core of the model is represented by a computational structure that reproduces closely the behaviour of the computerised book of a real stock market. The agents send to the book their buy and sell orders, with the related limit prices. The book executes immediately the orders if a counterpart is found in its log; otherwise, it records separately the buy and sell orders, to match them with future orders. The book is cleared at the beginning of each day. The book is cleared at the beginning of each day.

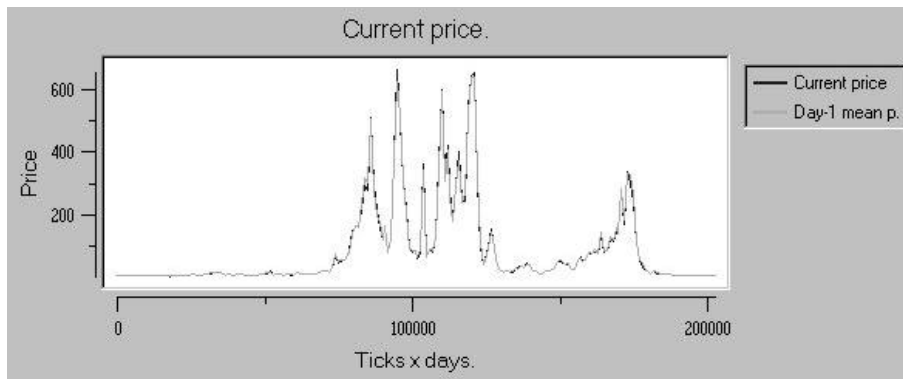


Fig. 5. Bubbles and crashes produced by no-mind agents.

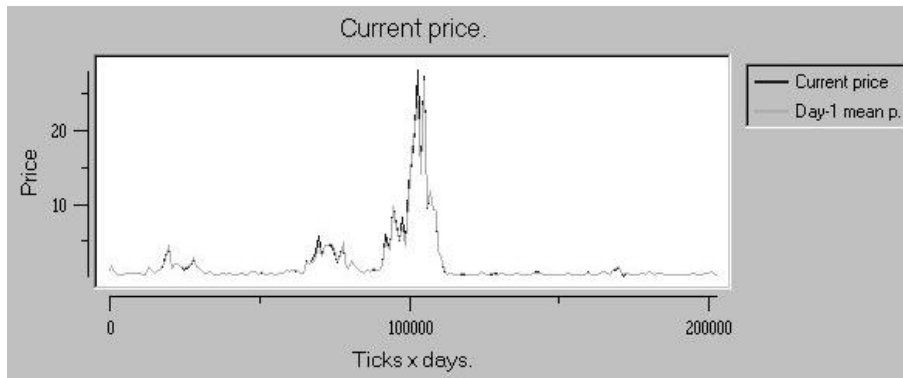


Fig. 6. Bubbles and crashes produced by no-mind agents, with a floor price.

Our (un)realistic market emerges from the behaviour of myopic agents that: (i) know only the last executed price, (ii) choose randomly the buy or sell side and (iii) fix

their limit price by multiplying the previously executed price times a random coefficient, as shown in Fig. 5 and 6.

In Fig. 6 we introduce also the rule of buying with a fixed probability (here $p = 0.2$) if the price falls below a specific floor. A bubble emerges again, but the market is more stable than in the first case.

4.2. Comments to the experiment

This structure generates increasing and decreasing price sequences with relevant volatility. Also bubbles and crashes appear in this market, generated within the market structure, without the need of exogenous explanations, as a consequences of the rules of the market.

From the “no-mind” in agents perspective, we show here that it is possible to generate complex patterns (as bubbles) without using BDI agents, if the structure of the market is highly sophisticated, and consequently able to generate internally sequences of prices linked to the agents’ actions in a non linear way.

5. Conclusions and further developments

The following statements may constitute a tentative conclusion:

- first of all we can fully accept the Axelrod simplicity principle, since it is easily shown that also very simple agents can generate complex emerging structures;
- more complicated agents facilitate the emergence of theoretical results, as the hayekian market;
- finally, external constraints are as strong as a mind proxy (the CT scheme) in determining the emergence of hard to explain patterns, like bubbles in a stock market.

Are human agents so far from the complexity of the economic system, as ants are from their anthill?

Further developments may arise, within the framework introduced in Section 4, relaxing hypothesis (i) for a small quota of the agents, in order to investigate the consequences of the presence either of subjects using technical trading rules to forecast the future market prices and of cognitive agents capable to learn from their experience. In some way, the last ones can correspond to the artificially intelligent agents behaving as econometricians proposed by Sargent [SAR 93], with the interaction of minded agents and structured environments or markets.

The framework of Section 4.1 is the natural candidate to develop the unified environment suggested in the Introduction, with the goal of comparing directly - in a unique structure - the four extreme situations of (i) no-mind agents behaving in an

unstructured environment, (ii) minded agents behaving in an unstructured environment, (iii) no-mind agents behaving in a structured environment and (iv) finally minded agents behaving in a structured environment (this last case is not treated in this paper).

6. References

- [ART 90] ARTHUR W.B., A Learning Algorithm that Mimics Human Learning., *Santa Fe Institute Working Paper* 90-026, 1990.
- [AXE 97] AXELROD R., Advancing the Art of Simulation in the Social Sciences, in R.Conte, R.Hegselmann and P.Terna (eds), *Simulating Social Phenomena*, Lecture Notes in Economics and Mathematical Systems 456, pp.21-40, Springer-Verlag, 1997.
- [BEL 96] BELTRATTI A., MARGARITA S., TERNA P., *Neural Networks for Economic and Financial Modelling*, ITCP, London, 1996.
- [GIL 00] GILBERT N., TERNA P., How to build and use agent-based models in social science, *Mind & Society*, no. 1, 2000.
- [HAY 49] HAYEK F.A., *Individualism and Economic Order*, London, Routledge, 1949.
- [OST 98] OSTROM T., (1988) Computer simulation: the third symbol system. *Journal of Experimental Social Psychology*, vol. 24, 1998, pp.381-392.
- [SAR 93] SARGENT T.J., *Bounded Rationality in Macroeconomics*, Oxford, Clarendon Press, 1993.
- TER 98] TERNA P., ABCDE: Agent Based Chaotic Dynamic Emergence, in J.Sichman, R.Conte e N.Gilbert (eds), *Multi-Agent Systems and Agent-Based Simulation*, LNAI series, vol. 1534, Berlin, Springer-Verlag, 1998.
- [TER 00a] TERNA P., Economic Experiments with Swarm: a Neural Network Approach to the Self-Development of Consistency in Agents' Behavior, in F. Luna and B. Stefansson (eds.), *Economic Simulations in Swarm: Agent-Based Modelling and Object Oriented Programming*. Dordrecht and London, Kluwer Academic, 2000.
- [TER 00b] TERNA P., *SUM: a Surprising (Un)realistic Market - Building a Simple Stock Market Structure with Swarm*, forthcoming, 2000.

Appendix

Agents avec et sans "esprit" dans un marché artificiel.