Adaptation in the Presence of Exogeneous Information in an Artificial Financial Market

José L. Gordillo[1], Juan Pablo Pardo-Guerra[2], and Christopher R. Stephens[3]

[1] Dirección Gral. de Serv. de Cómputo Académico, UNAM.
jlgr@super.unam.mx
[2] Facultad de Ciencias, UNAM.
jpardog@prodigy.net.mx
[3] Instituto de Ciencias Nucleares, UNAM.
stephens@nuclecu.unam.mx

Abstract. In recent years, agent-based computational models have been used to study financial markets. One of the most interesting elements involved in these studies is the process of learning, in which market participants try to obtain information from the market in order to improve their strategies and hence increase their profits. While in other papers it has been shown how this learning process is determined by factors such as the adaptation period, the composition of the market and the intensity of the signals that an agent can perceive, in this paper we shall discuss the effect of external information in the learning process in an artificial financial market (AFM). In particular, we will analyze the case when external information is such that it forces all participants to randomly revise their expectations of the future. Even though AMFs usually use sophisticated artificial intelligence techniques, in this study we show how interesting results can be obtained using a quite elementary genetic algorithm.

1 Introduction

In recent years it has become ever more popular to consider financial markets (FMs) from an evolutionary, rather than the traditional *rational expectations*, point of view [1,2]. In particular, there has been a substantial increase in studies that use agent-based, evolutionary computer simulations, known as Artificial Financial Markets (AFM) [3]. In this paper we use a particular AFM — the NNCP [4] — whose design was motivated by the desire to study relatively neglected elements in other AFMs (for instance, the Santa Fe Virtual Market [5]), such as the effect of organizational structure on market dynamics and the role of *market makers* and information. All of these elements are crucial in the formation of market microstructure [6].

Among the most interesting aspects one can study in an AFM — and constituting the central topic of this work — is the process of learning, a set of mechanisms which allows agents to modify their buy/sell strategy with the aim of adapting to the conditions imposed by the market. In particular, in this paper

R. Monroy et al. (Eds.): MICAI 2004, LNAI 2972, pp. 342–351, 2004.

we study the effect of external information on the learning process. For this purpose we have analyzed the extreme case where the arrival of information forces all participants to change their perception about the state of the market and therefore their expectations about its future evolution.

Though the NNCP allows us to include many features that influence the behavior of the market, in this study we used only a small number of elements: informed and uninformed agents, adaptive agents, and information "shocks" in the market's development. Despite this relatively modest diversity of behaviors and simplicity of elements, it was possible to conduct experiments that produced significant results.

The structure of the next paper is as follows. In section 2 we describe the general form of the elements used in the NNCP, namely the market organization, the market participants, the information processes ("shocks"), and the learning mechanisms. In section 3 we explain the experiments conducted, along with a discussion of their main results. Finally, we give our conclusions, as well as some general ideas for future lines of research.

2 The Building Blocks of the NNCP

The workings of the NNCP are, in general terms, as follows: a simulation is carried out for a prescribed number of *ticks* on a single risky asset. An agent can divide his/her wealth between this risky asset and a riskless asset ("cash"). At each *tick* an agent —or a set of agents — takes a position (buy/sell/neutral). Shares are bought in fixed size lots of one share. Resources are finite and hence traders have portfolio limits associated with either zero cash or zero stock. Short selling is not permitted.

2.1 Market Organization

The market clearing mechanism we used for all our simulations in this particular study is a simple double auction, where at every *tick* each trader takes a position with an associated volume and at a given price, each trader being able to value the asset independently but with prices that are not too different. In this model price changes are induced only via the disequilibrium between supply and demand. Specifically:

1. At time (or "tick") t one lists all the positions taken by the agents and the associated volume and price. The agents' bids and offers are obtained via a Gaussian distribution with mean $\bar{p} = p(t-1)$.
2. A bid and an offer are matched only if they overlap, i.e. $p_b(t) > p_o(t)$. To realize a transaction we used: "*best bid/offer*", where the highest bid and the lowest offer are matched at their midpoint successively until there are no overlapping bids and offers.

After each *tick* price is updated exogeneously via a supply/demand type law as in Eq. (1).

$$p(t+1) = p(t)[1 + \eta(B(t) - O(t))] \tag{1}$$

In this equation, which is common to many AFMs, $p(t)$ is price at *tick* t and $B(t)$ and $O(t)$ are the demand and supply at t, while η is a tuning parameter. Note that $D(t) = (B(t) - O(t))$ depends not only on the positions taken by the agents but also on the mechanism used to match their trades, e.g. at what price two contrary trades will be matched. In this sense one may think of a "bare" $D(t)$, $D_B(t)$, that represents the imbalance in supply and demand associated purely with the desired trades of the agents while $D(t)$ represents the residual imbalance after matching those orders that can be matched under a given clearing mechanism.

2.2 Market Participants

We will divide traders into various classes. Two of the principal classes are informed and uninformed, or liquidity, traders. The latter make random decisions, buying or selling with equal probability irrespective of the market price. Informed agents on the other hand have a higher probability to buy than sell. One can try to rationalize this behavior in different ways, each rationalization being equally legitimate in the absence of further information. One can, for instance, imagine that informed agents have a better understanding of the market dynamics in that they "know" that in the presence of uninformed traders the excess demand the informed trader's bias generates will translate itself, via Eq. (1), into a price increase which will augment their portfolio values at the expense of the uninformed. Alternatively, one may simply imagine that the informed traders believe the market will rise. We will, in fact, consider a one-parameter family of informed traders described by a "bias", d, where the position probabilities are:

$$P(c) = \frac{2d}{3}, \qquad P(n) = \frac{1}{3}, \qquad P(v) = \frac{2(1-d)}{3} \tag{2}$$

where c represents Buy, v Sell and n Hold. For example, when $d = 1/2$ then the corresponding probabilities are $1/3$, $1/3$, $1/3$; which actually corresponds to an uninformed trader, i.e. a trader having no statistical bias in favour of one position versus another. In contrast, a trader with $d = 1$ has probabilities $2/3$, $1/3$, 0 and corresponds to a trader with a strong belief that the market will rise or, alternatively, to a trader who believes that there are many uninformed traders in the market that can be exploited by selling while the informed trader drives the price up. We will denote a trading strategy from this one-parameter family by the pair $(100d, 100(1-d))$. Thus, an uninformed, or liquidity, trader is denoted by $(50, 50)$ and a maximally biased one by $(100, 0)$. Essentially, the different traders have different belief systems about the market. As mentioned, due to the simplicity of the model we may not ask how it is that the different traders arrive at these different expectations. It may be that they have different information sets, or it may be due to the fact that they process the same information set in different ways, or, more realistically, a combination of these.

Given that we wish to compare the relative profitability of the different trading strategies we need to define profits. Here, the profits of agents are given in terms of a "moving target" where excess profit during timestep t is related to the increase in the market value of an active trading portfolio in the timestep t relative to the increase in the market value of a buy and hold portfolio in the same timestep. In this way an excess profit for a given trader over the timestep t can only arise when there has been a net change in the trader's portfolio holdings in the asset *and* a net change in the asset's price. This choice of benchmark always refers the market dynamics to a "zero sum" game, while with other benchmarks this is not the case. More concretely, we define the "excess" profit of a trader i in the time interval $t-1$ to t to be

$$e_i(t, t-1) = \delta n_i(t) \delta p(t), \tag{3}$$

where $\delta n_i(t)$ is the change in portfolio holdings over the timestep $\delta t = t-(t-1)$ of the trader i and $\delta p(t)$ is the change in asset price over this timestep. The excess profit earned between times t' and t is

$$E_i(t, t') = \sum_{n=t'}^{n=t} e_i(n, n-1) \tag{4}$$

2.3 Adaptation and Learning in the Presence of Endogeneous and Exogeneous Information

In the context of only informed and uninformed whose strategies remain static there can be no adaptation in the market, nor any learning. In order to caricature these elements we introduce adaptive agent strategies wherein an adapting agent may copy the strategy of the most successful agent currently in the market (the *copycat* strategy). In this sense the copycat agents have to both learn or infer what is the best strategy to copy, and then adapt their own strategy in the light of this new knowledge. The manner in which they do this is via standard "roulette wheel selection" as commonly used to represent the selection operator in Genetic Algorithms [7], to update their strategies using accumulated excess profits as the "fitness" function. In other words, a copycat copies the strategy of agent i with probability

$$P_i(t) = E_i(t, t') / \sum_i E_i(t, t') \tag{5}$$

where $E_i(t, t')$ is defined in Eq. (4). They observe the market, updating their information at a fixed frequency, for example, every 100 ticks, and copy the agent's strategy that wins the roulette wheel selection process. Given that the roulette wheel selection is stochastic it may be that a copycat does not copy the strategy with the most excess profit. It is important to clarify that we are not interested in pinpointing the agent which is copied in so far as we are interested in identifying the strategy that is picked during the process. This way we can conceive the selection mechanism as a roulette divided not in small regions representing individual agents, but rather in bigger slices that account for each of the strategies present in the market; Eq. (6) depicts this situation

$$P_i(t) = S_i(t, t') / \sum_{j=1}^{m} S_j(t, t') \tag{6}$$

where $S_j(t,t')$ is the sum of accumulated profits (Eq. (4)) of all the agents with strategy j, and where m is the total number of strategies present in the market.

The more successful a strategy is relative to others, the more likely it is that this is the strategy copied. The stochastic nature of the copying process reflects the inefficiencies inherent in the learning process.

In the above we described how copycats learn and adapt in the presence of endogeneous information, i.e. information intrinsic to the market itself, this information being the relative profits of the different trading strategies in the market. This information is dynamic and so the copycats change their expectations and beliefs about what will happen in the future. In real markets, however, new exogeneous information frequently arrives. In this context one must ask how the different traders will react to this information. Such exogeneous information is usually taken to be random. We will follow this paradigm here, imagining the exogeneous information to be in the form of information "shocks". In this case we assume that the participants are forced to change their perception of the state of the market — and therefore their buy/sell strategy. At this point, two interesting questions can be raised: First, how is the learning process affected by these information shocks? And second, is the information prior to a shock useful in the learning process?

3 Experimental Results

We will answer the above questions in the context of various simulations carried out using the NNCP artificial market. However, before presenting the experiments with their results, we will discuss some aspects of the copycat's adaptation process. As it has been mentioned previously, in order to adapt their strategies to the market's conditions, copycats must "play" a roulette formed from the profits of each strategy in the market (Eq. 6). It is convenient to recognize the stochastic effects of this game, in particular those produced by the composition of the market. We can illustrate this by thinking of a market where the copycats copy via roulette wheel selection the most popular strategy. Suppose the roulette is formed by the number s of agents that possess a strategy given by

$$P_i(t) = s_i(t)/\sum_j^m s_j(t)\,; \tag{7}$$

One question we can answer is how many copycats will adopt strategy i at time t. Let C be the number of copycat agents, I the initial number of agents with strategy i and T the total number of agents (i.e. $\sum_j^m s_j(t)$). When $t = 1$ (the first adaptation) it follows that

$$X_i(1) = C * P_i(1) = I * C/T\,; \tag{8}$$

where $X_i(t)$ is the average number of copycats that adopt strategy i at time t. Now, when $t = 2$, the number of agents with strategy i is $I + X_i(1)$, and in general, after K adaptations, one has

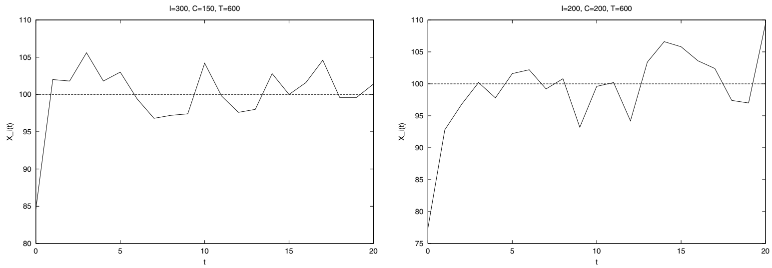

Fig. 1. Number of copycats that learn strategy i in simulations with different values of $\{I,C,T\}$. The values are {300,150,600} (left) and {200, 200, 600} (right).

$$X_i(K) = C*P_i(t-1) = I*C/T + I*C^2/T^2 + \ldots + I*C^n/T^n = I*\sum_j^k C^j/T^j \quad (9)$$

Obviously, $C < T$, and therefore

$$X_i(t)_{t\to\infty} = I * C/(T - C)\,. \quad (10)$$

which is the expected maximum number of copycats that will copy the most popular strategy. In Figure 1 we show the result of the first 20 adaptations in markets with different values of I, C and T, where copycats adapt every tick and the most popular strategy is associated with informed $(51, 49)$ agents. The graphs are a result of averaging over 10 different runs. In Figure 1 we see how the number of correct copycats asymptotes to a value close to that given by Eq. (10). There is a slight difference in that the graphs are for copycats that copy the most profitable strategy. However, for weak bias we see that Eq. (10) gives a good approximation. More generally, it gives a lower bound for the number of correct copycats.

Returning to the problem of learning: The objective of a copycat is to acquire the optimal strategy (i.e. the strategy that maximizes profits constrained to existing market conditions); conversely, the objective of the biased traders is to create an excess demand. This excess demand thus drives the price via the price evolution equation (Eq. 1) along with the profits of informed agents, as has been noted in previous work [4,8]. Additionally, both the excess demand and the profits of the informed traders depend on the composition of the entire population as well as on the distribution of biases. In this scenario, copycats try to copy informed traders to find the optimal strategy. This activates the learning process. However, complete learning is by no means guaranteed in the sense that they do not necessarily identify the best strategy. The quality of the learning depends on the signal to noise ratio (i.e. the size of the different regions in the roulette), which in its turn depends on the agent biases and the market composition. Note that the learning might be incomplete even in the case where there is only one other strategy to learn. As an illustration of the latter consider Figure 2, where we show the number of copycats that learn the correct strategy

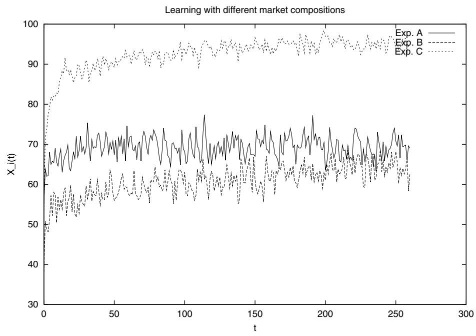

Fig. 2. Incompleteness of learning: number of copycats that learn strategy i in experiments with different market compositions

in three different experiments. In the first case (Experiment A), the market is composed of 20 agents of each of the following strategies: $(50, 50)$, $(60, 40)$, $(70, 30)$, $(80, 20)$, $(90, 10)$, 100 $(99, 1)$ agents, and 100 copycats. In Experiment B the market is formed by 100 uninformed agents (i.e. with a $(50, 50)$ strategy), 100 $(60, 40)$ agents and 100 copycats. Finally, Experiment C was composed of 100 $(50, 50)$ agents, 100 $(99, 1)$ agents and 100 copycats. The roulette at time t was built using $E_i(t, 0)$, that is, the profits calculated since the beginning of the experiment.

In Experiment A, the optimal strategy is $(99, 1)$. However, the presence of other strategies with lesser yields confuses the copycats in such a way that only about 70% of them present successful learning, i.e. that identify the optimal strategy, (the average due purely to the composition of the market is 50% in all cases; this can be derived through simple probabilistic arguments with the use of the roulette). In Experiment B we can observe that the interaction of the $(50, 50)$ and $(60, 40)$ strategies generates only a relatively small signal, hence explaining why the number of copycats that learn the best strategy is only slightly bigger than the average of the market's composition. Experiment C shows the imperfection of the learning process, even in a market with a very large difference in biases, i.e. that due to the stochastic nature of the roulette wheel complete learning cannot take place. We see in general then that the efficiency of learning depends on the market biases and the diversity of strategies in the market as well as the stochastic nature of the roulette wheel selection.

3.1 Simulating Markets with Exogeneous Shocks

As mentioned earlier, we model the effect of an information shock by changing the perception of all the market participants. This is done by means of a random re-selection of strategies among the agents at the moment of the shock. Specifically, by labeling the agents as either informed (I) or uninformed (U), we can visualize an information shock as a moment t_s in which all the market's participants select a strategy again, either U or I. Thus, an agent that prior to a shock is informed may become uninformed with a certain probability s or may remain in his original state with a probability $1 - s$[1]. With this scheme, over a period that contains several shocks an agent can no longer be tagged as being either informed or uninformed. However, we can label them with a chain that represents the different states they have occupied over a certain interval of time[2]. For instance, an agent may be defined by the sequence $IUUI$, which means that, starting out as an informed, it changed into a uninformed at the first shock, remained in this state after the following shock and returned to an informed strategy after the third shock. This simulation of information processes clearly creates a new level of complexity in the system since now we face many more options than the ones present in a static market.

Taking this into account, after a shock an evolutionary agent must re-learn what is the optimal strategy under the new market conditions. An interesting problem is determining how much endogeneous information an evolutionary agent needs to learn the best strategy, given the dynamic conditions of the market. In the examples presented in the first part of this section, each evolutionary agent used all the history of the trader's profits to make a decision, i.e. each agent had "long-term" memory. However, with the introduction of shocks into the system, the evolutionary agents now face new difficulties in processing the available information. In this sense, the amount of data used for inference becomes a crucial matter: it is not the same considering a multi-shock time window when the strategies can shift between informed and uninformed than when they remain static. And so, we may pose the following question: Will the same information be as useful in a system with changing perceptions? In Figure 3 we show the results of two experiments that shed some light on this question. In Experiment D the copycats try to copy a strategy using only the information generated by the market after each shock. Thus, they have only "short-term" memory as they do not keep in their memory any information prior to the shock. In contrast, Experiment E depicts the case in which copycats have long-term memory, preserving the entire information of the market's history without distinguishing data obtained before and after shocks. In these experiments the traders change their perception of the market to one of two possible states: either an informed or an uninformed strategy. In other words, despite the shocks, during all the periods of the experiment, the optimal strategy is invariant, i.e. to be constanstly informed; what changes are the perceptions of the participants which affects the strategy that they choose and therefore the profits they have accumulated. At the same time, after each shock an arbitrary agent has the same chance of ending in the set of the informed as in the set of the uninformed, in such a way that the composition of the market in each period is, on average, always the same.

---

[1] In the experiments presented in this paper that involve shocks, s takes a value of $\frac{1}{2}$, unless otherwise indicated.

[2] Since the definitions of informed and uninformed are mutually exclusive, we can consider them as states that describe the actual strategy of an agent. This way, if an agent is informed, we can say that he is in state I or, in other words, that he is occupying state I.

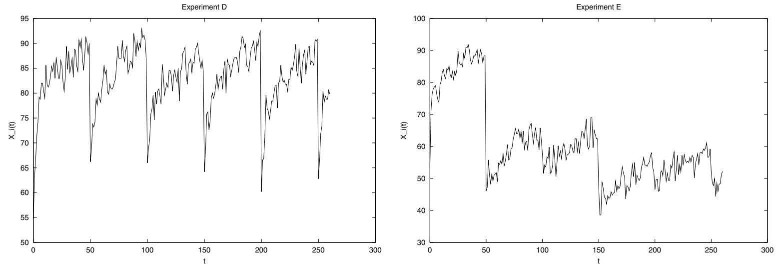

Fig. 3. Learning in markets with exogeneous shocks: copycats with long-term (right) and short-term (left) memory

We can see in this figure that the learning process where traders only use as their learning information set market information from the last shock until the present moment better adapt to the new market conditions. This is especially true in the present case where the first shock arrives when the learning process is almost finished; the shock produces a reset of the process but the learning process itself stays the same. After each shock the copycats realize that they must adjust their perceptions in the new market conditions by relearning everything. Meanwhile, in the case of learning with long-term memory, it is much more difficult for the copycats to identify the correct strategy after each shock as they are using past information that is no longer relevant. For example, they end up copying a strategy that was useful and therefore accumulated significant profit before the shock but is suboptimal after the shock. By the time the copycat has realized that the strategy is no longer optimal it has made significant losses.

4 Conclusions

We have shown here that it is possible to produce very interesting results utilizing only a fairly simple computational model. Although the model itself has little complexity and is relatively small — a significant benefit that expresses itself in small run times — our work on the NNCP has given us some interesting ideas on how financial markets might deal with external information. Among the most important findings, we can identify the role of memory during the learning process. As shown in section 3.1, the relevance of information considered during the selection of a strategy is of vital significance when it comes to optimizing profits: using irrelevant information from before the shock to determine the optimal post-shock strategy results in poor learning and a more efficient market.

And though the identification of shock-like structures in real financial markets is still somewhat controversial, results like these might translate into practical trading techniques in the future. We have also shown that learning is very much a statistical inference process in the context of a financial market and have exhibited some of the factors on which the efficiency of the learning depends.

There is, however, much more work ahead. The configuration of NNCP, as used in this paper, is quite simple, so it is not implausible to have left a large collection of behaviors out of the simulations. In this sense, we can consider several improvements to the model aimed at better describing the nature of financial markets, or at least at refining our approximations of their description. For instance, future models might consider more complex mechanisms of learning as well as sophisticated information shocks that do not affect the entire market. Equally important, the participation methods of the artificial agents could be enhanced into more than stochastic rules defined by a probabilistic bias. Nonetheless, all these approaches demand a better understanding of the system (from the mechanisms of learning to the role of market structure) as well as increased computational requirements.

Acknowledgments. This work was supported by DGAPA grants IN100201 and IN114302 and CONACyT grant 41221-F. We are grateful to the Dept. of Supercomputing, DGSCA, UNAM for access to computational resources. We are also grateful for useful conversations to Prof. Harald Benink.

References

1. Bedau M., Joshi S.: An Explanation of Generic Behavior in an Evolving Financial Market. Preprint 98-12-114E, Santa Fe Institute (1998).
2. Farmer J. D.: Market Force, Ecology and Evolution. Working Paper 98-12-117E, Santa Fe Institute (1998).
3. Chan N., Lebaron B., Lo A., Poggio T.: Agent Based Models of Financial Markets: a Comparison with Experimental Markets. Working Paper, Brandeis University (1999).
4. Gordillo J.L.: Análisis de Mercados Financieros mediante el Mercado Financiero Artificial NNCP. Tesis de Maestría. IIMAS-UNAM (2000).
5. Palmer R.G., Arthur W.B., Holland J.H., Lebaron B., Tayler P.: Artificial Economic Life: a Simple Model of a Stock Market. Physica D 73 (1994).
6. O'Hara M.: Market Microstructure Theory. Blackwell Publishers Inc (1997).
7. Goldberg D.G.: Genetic Algorithms in Search, Optimization and Machine Learning. Addison Wesley (1989).
8. Gordillo J.L, Stephens C.R.: Strategy Adaptation and the Role of Information in an Artificial Financial Market. Late Breaking Papers, GECCO (2001).
9. Gordillo J.L, Stephens C.R.: Analysis of Financial Markets with the Artificial Agent-based Model-NNCP. Encuentro Nacional de Ciencias de la Computación. Sociedad Mexicana de Ciencias de la Computación. Mexico (2001).