Universidad de
San Andres

DEPARTAMENTO DE ECONOMIA

Learning and Contagion Effects in Transitions Between Regimes: A Schematic Model of Bank Runs

D. Heymann (CEPAL, UBA, ITDT), R.P.J. Perazzo (Center for advanced studies), A. Schuschny (Center for advanced studies)

CICLO DE SEMINARIOS 1997

Día: Martes 21 de Octubre

9:00 hs.

Learning and Contagion Effects in Transitions Between Regimes: A Schematic Model of Bank Runs

Heymann, D. †, Perazzo, R.P.J. ‡ and Schuschny, A. ‡

† CEPAL (Comisión Económica para América Latina, Un. Nations) and
University of Buenos Aires, School of Economics
Paraguay 1178, (1057) Buenos Aires, Argentina
‡ Center for Advanced Studies, University of Buenos Aires
Uriburu 950, (1114) Buenos Aires, Argentina

Abstract

We present a schematic model involving adaptive agents who deposit resources in a "bank". Their behavior is determined through a learning process based on two types of information. One (global) is shared by all agents; the other (local) is derived from data that individuals obtain by observing their close neighbors. The system is able to self-organize: the learning mechanisms allow agents to coordinate their deposit strategies. We then analyze the effects of an exogenous disturbance which increases the likelihood that the bank is "illiquid". Given this "fundamental" shift, the system can show contagion effects which can be interpreted as a "panic". Although the model is highly stylized, it is possible to tune its parameters to mimic approximately the evolution of deposits during a recent (1995) banking crisis in Argentina.

1 Introduction

The actions of groups of economic agents sometimes appear to undergo sudden changes, as if there was a collective shift to a new mode of behavior. Such changes can emerge in several contexts; the case of bank runs seems a particularly noticeable one.

Banking panics have been interpreted in two ways [1, 2, 3, 4]:

1) As events based on the diffusion of information about "economic fundamentals", which induces the revision of beliefs about the prospects of the banking system ("informationally based" events),

2) As pure coordination phenomena in a system with multiple potential equilibria, in which an "extraneous" event causes agents to change the state on which they coordinate their expectations, leading to a shift in the observed configuration.

Both explanations need not be mutually exclusive. It may happen, for example, that some macroeconomic news, which "fundamentally" modify the perceived return on deposits, induce a response which in turn triggers a self-reinforcing avalanche of withdrawals, thus amplifying the effect of the impulse. In this instance, the panic is not simply generated by a "sunspot" (since in "normal times" agents would disregard the possibility of multiple equilibria) but it does have a component whereby the observation that people are drawing their assets from the banks leads to a cascade of more withdrawals.

We want to study in a very schematic setting the interplay between two influences on behavior: (i) the beliefs of the agents about the "fundamental state" of the system, and (ii) the effect of what other individuals are observed to be doing. The first mechanism is modeled through a learning procedure (a "genetic algorithm") which is based on the available economy-wide information that all agents share. Even though this information is the same for all the agents, each of them adapts his decisions starting from different initial "interpretations" of the flow of data. Therefore, some diversity in decisions is to be expected. The second influence which we take into account corresponds to the information that the individual has about the (good or bad) outcomes corresponding to the group of agents that happen to be his neighbors. We may interpret a response to this type of local information as a contagion effect.

The model considers a large number N_{ag} of agents, who receive one unit of resources at the beginning of each "day". The unit can be deposited in a "bank" or kept by the agent. The choice must be made each market day. The depositors bear a cost if the aggregate level of deposits is lower than a certain threshold value S_o (e.g. the bank cannot meet its contractual obligations) or higher than the value S_1 (e.g. due to congestion costs). If the level of deposits stays between those two limits the depositors receive the benefit of a positive return. Individuals decide in advance a deposit strategy for the subsequent "week" of N_d days, determining which days they will make a deposit and which days they will keep the unit. These features are sufficient to generate a self-organization of the system in such a way that in the aggregate the daily deposits of agents are near the upper bound S_1, except in some particular situations where the deposit pattern gets locked into a suboptimal state.

We also incorporate a "local" coordination mechanism by allowing agents to choose a behavior that imitates that of their neighbors instead of using the strategy derived from their own internal mechanism of learning. We find that both influences (learning from the global information and contagion from the close neighborhood) adjust to each other when the level of deposits is high; however, when deposits happen to fall "too much", the local effect may lead the system to a new behavior (similar to the transition to the magnetic phase in a ferromagnet) which has the characteristic of a bank run.

In order to study the interaction of both coordination mechanisms the system is disturbed. To do this we vary (exogenously) the lower threshold level of deposits, implying that a lower amount of potential withdrawal demands can be satisfied. In this case the "local" component takes over after a few periods, giving rise to a short transient which can be interpreted as the diffusion of an "awareness of a dangerous state". If the exogenous shock is reversed quickly, the system soon returns to the undisturbed state. However, if the new value of the lower threshold is maintained, a "panic wave" builds up, with a large fall in the supply of funds to the bank. If now the exogenous parameter returns to its initial level, thus putting an end to the "fundamental" shift, a reorganization of the system takes place. This can happen in two possible ways. If the "bad fundamental state" has lasted for a long time, the model generates the possibility that all individuals retain their funds, leading to an autharchic equilibrium in which the banks do not receive deposits. Alternatively, if the crisis is short, the system again converges to a stationary state with a high level of deposits.

In the last section we "calibrate" the parameters of the model to describe qualitative features of the recent (1995) banking crisis in Argentina. Although the model is extremely schematic and stylized, it is possible to identify the three periods mentioned above as well as some elements of collective behaviour that seem characteristic of crisis episodes and their aftermath.

2 Coordination through global information

The learning procedure that takes into account the globally available information acts by improving strategies which are based on the history of the returns obtained in successive visits to the bank. Each agent of the system has assigned a population of N_p "deposit plans" for the subsequent "week" of N_d days. To update the plans we use a "Genetic Algorithm" [6]. Each strategy is encoded in a "genome" that specifies each day whether the agent make a deposit or retain the unit of resources. The genomes are initially selected at random, and are updated by selection, crossover and mutation [6], so that less successful strategies are progressively eliminated.

The formal steps are the following: Each agent possesses N_p deposit strategies. Each strategy consist of a chain of N_d bits where N_d stands for the number of days of the "week". Each of the N_{ag} agents receives each "day" an income unit which can be deposited in the "bank" or retained, according to the current strategy. The "bit" ${}^p\#_k^t$, for the t-th "day" of the p-th strategy, of the k-th agent, takes the value $+1$ (-1) in case that such agent, acting according to that strategy, chooses to make (withhold) a deposit on that day.

$$p(1 \leq p \leq N_p) \left\{ \underbrace{\begin{array}{|c|} \hline +1-1-1+1-1\cdots-1-1+1-1+1 \\ \hline -1+1-1+1-1\cdots-1+1-1-1+1 \\ \hline \vdots \quad\quad \vdots \quad\quad \vdots \quad\quad \vdots \\ \hline +1-1+1-1\cdots{}^p\#_k^t\cdots+1-1+1-1 \\ \hline -1-1+1-1+1\cdots-1-1+1+1-1 \\ \hline \end{array}}_{N_t(1\leq t\leq N_d)} \right.$$

When a "week" finishes, although the agent has effectively used only the best strategy (chosen from the data of the preceding week), he compares the performance of all the strategies in his population. These are then ranked by the utility that they would have produced. The population of strategies is then updated using an "elitist" version of the Genetic Algorithm: those plans ranked within the upper 50% survive to the next generation, while those in the lower 50% are replaced by strategies obtained by crossover. All strategies are next subject to random mutations.

Once each agent has selected one strategy, he daily implements the action of either making a deposit or withholding the unit of resources. The bank records the daily level of deposits. Let c be the current strategy of each agent. The fraction of resources deposited, D_t, in the t-th day is:

$$D_t = \frac{1}{N_{ag}} \sum_{k=1}^{N_{ag}} \theta({}^c\#_k^t) \tag{1}$$

where $\theta(x) = 1$ if $x > 0$ and $\theta(x) = 0$ otherwise
Once the bank calculates the level of deposits, it next determines the rate of interest r_t that corresponds to that day. The possible outcomes are the following:

- If $D_t < S_o \Rightarrow r_t < 0$ (We take $r_t = -1$). The agents who have made a deposit on that day make a loss because the level of deposits is less than the lower threshold level S_o[1].

- If $S_o \leq D_t \leq S_1 \Rightarrow r_t = f(D_t)$, where $f(D_t)$ is a positive decreasing function: the interest rate decreases with the flow of deposits.

- If $D_t > S_1 \Rightarrow r_t = -r_{\text{cost}}$, with r_{cost} is some small positive value (we take $r_{\text{cost}} = 0.01$). When the level of deposits exceeds the value S_1, the agents bear a cost due to congestion at the bank.

The fitness of each strategy is recorded daily and during the period of N_d "days" with the following rules:

- If $r_t = r(D_t)$ and ${}^p\#_k^t = 1 \Rightarrow U_k^p(t+1) = U_k^p(t) + (1 + r_t)$, with earnings or losses according to sign of r_t determined by the thresholds.

- If $r_t = -1$ and ${}^p\#_k^t = 0 \Rightarrow U_k^p(t+1) = U_k^p(t) + 1$, since if the agent has not made a deposit he mantains his unit of revenue.

[1]Clearly, the assumption that $r(t) = -1$ in this case is an extreme one. However, the results of the model are robust as long as agents make a loss if the total amount of deposits is small (for example $r_t \sim -0.01$).

- If $r_t = r(D_t) > 0$ and ${}^p\#_k^t = 0 \Rightarrow U_k^p(t+1) = U_k^p(t) + (1 - r_t)$, since the bank paid a positive rate and, if the agent has not made a deposit, he will bear an opportunity cost.

At the end of the week all the available strategies can be ranked according to their utility and the updating procedure can be evoked.

3 Coordination through local contagion

In order to incorporate a local mechanism of "contagion", each "day" a choice is offered to the agents of adopting one of two options: (i) to continue with their individual deposit plans or (ii) to imitate the strategy used by their nearest neighbors. Each agent determines (*ex - post*) if the decisions of his nearest neighbors have been better or worse than the one implicit in his current strategy. The neighborhood is determined by locating agents in an orthogonal lattice with periodic boundary conditions. The label k of each agent is thus replaced by the row and column indices (i, j).

Each agent is allowed to abandon his current strategy whenever he finds that his performance is not satisfactory. It may be noticed that, owing to this fact, one has to distinguish between the *current action* ${}^c a_{(i,j)}^t$ taken by the agent on the t-th day and his current *plan* for the same day ${}^c\#_{(i,j)}^t$. To decide a shift away from his plan, the individual determines the daily outcome obtained with his strategy and compares it with the one which would have applied if he had imitated the actions of his nearest neighbors. For this purpose, the agent calculates a "local field":

$$h_{(i,j)} = \frac{1}{4}\left({}^c a_{(i,j+1)}^t + {}^c a_{(i,j-1)}^t + {}^c a_{(i+1,j)}^t + {}^c a_{(i-1,j)}^t \right) \quad (2)$$

Once the local field $h_{(i,j)}$ is known, the agent can determine the action that he would take by following his neighbors. The procedure is similar to the simulations of the ferromagnetic interactions with the *Ising Model*. In our case, the option of either making a deposit or retaining the unit of income corresponds to the two possible orientations of a "magnetic" domain (neighborhood). The choice is not entirely deterministic. We assume that there is some "noise" in the agent's decision, represented by a Glauber dynamics with a "thermal agitation" characterized by a temperature $T = 1/\beta$. The agent has some probability of either imitating his neighbors or acting in the opposite way as the local field. The probability is[2]:

$$P(\text{action} = \pm 1) = \frac{1}{1 + e^{\mp 2\beta h_{(i,j)}}} \quad (3)$$

During a market "day", each agent computes the utility that would have been obtained using the action suggested by the local average field. If this is greater than the one computed with its current strategy, the agent changes his deposit plan (with the above probability) and uses the action prescribed by the local field. Otherwise, the agent sticks to his individual plan, using the next "bit" in the genome of his current strategy.

[2]The system is assumed to be in a *heat-bath* that acts as a background noise characterized by a temperature T. Each site of the lattice updates its vale through a stochastic dynamics with a probability that depends on the local field given by 3. In the absence of noise ($T = 0$), the transition probabilities are:

- If the local field $h_{(i,j)} \neq 0$ then $s_{(i,j)}(t + \delta t) = sign(h_{(i,j)})$ with probability 1.
- If $h_{(i,j)} = 0$ then $s_{(i,j)}(t + \delta t) = \pm 1$ with probability $\frac{1}{2}$.

It is well known that the lowest-energy state of a system of spins with attractive interactions between nearest neighbors has all the spins aligned in the same direction. There are two (degenerate) ground states, one with $s_i = +1 \forall i$ and the other with $s_i = -1 \forall i$. These states display a *ferromagnetic ordering* that is characterized by a global positive magnetization. The fact that these states are also the attractors of the noiseless dynamics rerults from the interaction of the spins of the system. The ferromagnetic ordering is a *cooperative phenomenon* in the sense that all the spins of the magnetic system contribute in the same way to the global magnetization. The dynamics correspond to a gradient flow along the directions that reduce the total energy of the system (this is a *Lyapunov function* for the system that is monotonically reduced at every time step [7]).

4 Results

Global coordination

The outcome of the processing of the global information often gives rise to a asymptotic coordinated state in which the agents self-organize to make as many deposits as tolerated by the upper threshold S_1 (see figure 1). The result is not affected by the detailed choice of the function $f(D_t)$ which defines the interest rate as a function of the total volume of deposits. This is so because each agent finds that his individual utility grows as long as he increases the number of days in which he suceeds to make a deposit. The coordinated state is not unique. It may happen that the system gets trapped in a *sub-optimal* configuration in which, during certain days, there are no deposits (figure 1). This configuration can emerge if, in an early stage of the evolution, and simply because of random effects, the level of deposits in those days turns out to be lower than the lowest threshold S_o. During the following steps of the adaptation, any deposit made on such days causes a loss, and therefore the event of making a deposit is discarded in many investment strategies. In this type of states, the system is confined to a sub-optimal situation by "entropic barriers", which can only be overcomed in the improbable case that a large enough fraction of agents simultaneously changes the individually chosen action for that particular day.

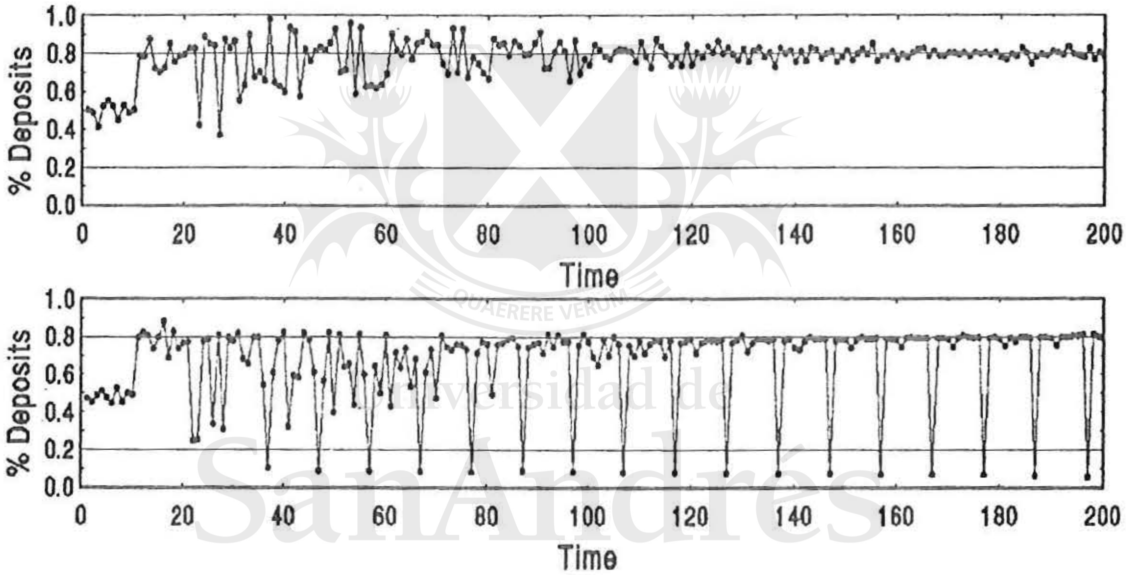

Figure 1: Simulations without local interactions with 256 agents, $r_{\text{cost}} = -0.01$, $p_{mut} = 0.005$, $S_o = 0.2$, $S_1 = 0.8$, $N_d = 10$ and $N_p = 20$. The first case shows an optimum convergence of the system while in other figure the self-organization results a in sub-optimal state

The asymptotic state of the system with deposits S_1 is a stationary state that is *dynamically* stable. All agents continue to explore new strategies that correspond to random mutations of their current plans. The system therefore "visits" configurations that are in the neighborhood of the maximum threshold S_1. The stability is *due to the fact that the system is diverse*: the agents tend to use complementary strategies because all of them should not deposit every day, since then total deposits would exceed de highest threshold S_1. The system evolves an "ecology" of strategies by which different agents deposit in different days. In addition, some agents make more deposits than others.

The coordinated state that is reached in this way "visits" (dynamically) the neighborhood of a Nash equilibrium. No agent can change its strategy unilaterally to improve his profit without disturbing some other agent of the system. When the state has been reached with the maximun admissible level of deposits, individuals who have not made a deposit that day cannot change their action and improve their profit. There is a large multiplicity of equilibrium configurations: any such configuration is equivalent to another one in which a depositor exchanges his position with an agent who has retained his funds. The problem that we address here has features in common with the "El Farol" model of B. Arthur, where restaurant patrons coordinate their visits to the establishment [5]. In both cases, there is no need for a "social planner" solving an NP-hard problem in order to reach the coordinated state. In our example, the actions chosen by the individuals evolve in such a way that, every day, the "appropriate" number of agents do make a deposit.

One can test the robustness of the coordinated state by studying the response to a sudden change in the global conditions. We can induce an exogenous "crisis" by increasing the lower threshold level S_o to a higher value, thus compelling the whole set of agents to revise their decisions. The crisis can be finished by returning S_o to its original value some time later. The evolution of the system is displayed in figures 2, 3 and 4.

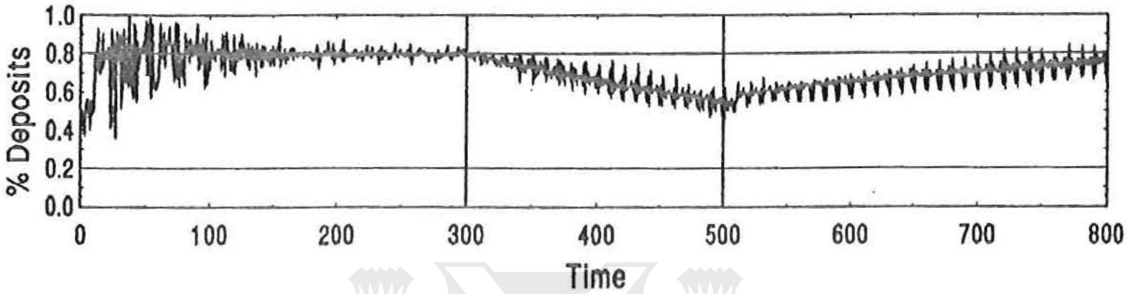

Figure 2: Evolution during a "exogenous crisis" when the local information processing is discarded. All the parameters are the same as those of the previous figures.

After the shock, agents are seen not to make a sudden switch in behavior; rather, they gradually adapt their strategies to the new situation as they learn that it is convenient to keep their funds. The adjustment is basically governed by the mutation probability of the genetic algorithm (figure 2). The simulated evolution approximates an exponential decay with a time constant which is long compared with the size of the business "week", and therefore looks gradual. If the shift in S_0 persists, agents eventually learn to keep their funds away from the bank.

Learning and local contagion

The first time steps of figure 3 show the evolution of the system when both coordination mechanisms are active. We have fixed the thermal fluctuations of the Glauber dynamics at $T = 0.1$, to avoid the singularity for $T = 0$. This parameter represents random fluctuations within the system. For this value of T, in stationary conditions, approximately 15% of the agents constantly switch strategies, basing their decisions on imitating their neighbors. A smaller value of T tends to hinder fluctuations; a very high value makes learning impossible, and the system does not self-organize.

As long as T stays below a critical value, both coordination mechanisms reinforce each other. This happens in two ways. In the first place, the agents "realize", through their own learning procedure that a certain level of noise is present and that, therefore, a fraction of the agents is mistaken. The genetic algorithm makes them act as if the actual value of the upper theshold S_1 was somewhat higher, thus exploring configurations that involve a higher level of deposits. In the second place, the presence of fluctuations tends to lower the "entropic barriers"; consequently, the system gets less often trapped in a suboptimal configuration.

The relevance of the the interplay of both coordination mechanisms becomes noticeable in the event of an external shock. The main effect of the local contagion is to favor cooperative effects, inducing abrupt changes in the collective behaviour of the agents. Once the "artificial crisis" begins, the performance of the individuals shows two features. First, agents who were imitating their neighbors no longer do so, and choose to use their own strategies. This process takes place during a time interval in the order of a couple of "weeks". Second, if the threshold S_0 is maintained at its new (higher) level, the agents increasingly learn to keep their units of resources; at a certain moment, a"critical" fraction of individuals acts in that way, thus inducing an avalanche in which the agents massively imitate each other by withholding the funds. The period of "awareness" (based on the agents's own learning) leads at some point to a "contagion cascade", with an abrupt decline in the volume of deposits (figures 3); this can be interpreted as a bank run. The cascade is triggered only if a large enough fraction of agents have decided "by themselves" to withhold their funds. If the "crisis" is short, so that this critical fraction is not reached, the coordinated state of high deposits is rapidly re-established, because the agents do not lose the memory of the previous successful strategies and start using them again "as if nothing had happened".

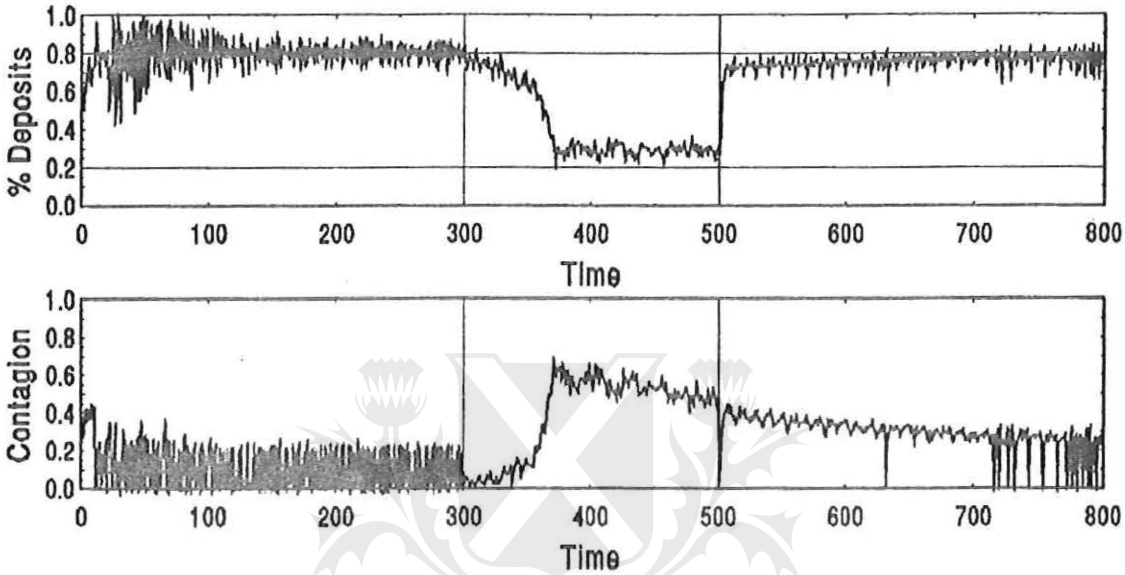

Figure 3: Self-organization process with both decision mechanisms when the temperature is $T = 0.1$. We also show the evolution during an "exogenous crisis". The parameters are the same as those of the previous figure.

When the bank run behavior emerges, the transition takes place in a very short time, and it is entirely similar to a ferromagnetic phase transition to a state with positive magnetization. In the limit of an infinite number of agents or when $T = 0$, the transition is instantaneous. If the "crisis" is long, and the the agents finally end up adopting a strategy that completely ignores the banking system. Both learning and contagion mechanisms act coherently to generate a similar behavior. In the model, once the banking system "collapses", it is difficult to reconstruct it. A highly improbable circumstance has to occur, namely that many agents decide *at the same time* to trust the bank again. The system is again trapped between high entropic barriers, that can only be lowered allowing high fluctuations (for instance setting T to a high value), so as to induce the system to visit configurations in which a sizeable fraction of the agents deposit their revenue.

Actual banking crises may show an asymmetry between declines and recoveries, with sudden runs and more gradual increases in deposits. In figure 4(a) we plot the total daily level of deposits (in current accounts, savings and time deposits, both in pesos and foreign currencies) during an interval (November 1st. 1994 - October 24th, 1995) which includes a severe banking crisis in Argentina. We have also shown in the picture the estimated values obtained with the present model. The free parameters have been callibrated as explained in the figure caption. The vertical scale has been normalized in such a way that the upper level corresponds to the beginning of the crisis and the lowest to the level of deposits on the day when the government announced that deposits were guaranteed. Each day of the time series is equivalent to a computer "day" of the "artificial banking system". The periods of "normal performance", awareness, panic and exit from the crisis can be identified clearly.

The model can mimic the sudden fall in deposits, but it predicts a much more rapid recovery than was actually observed. This shows that there is no memory built into the system to account for the asymmetry in the data. In order to extend the model, we have tried the following changes. The first is to draw new random initial strategies, letting the Genetic Algorithm start again with its adaptive work in a "fundamentally new" state of the system. This would correspond to the introduction of "new rules" of behavior. We have chosen to trigger this new procedure whenever a threshold of 50 % of the agents act according to contagion. The second element that helps to smooth out the exit from the crisis is to assume that in a certain number of cases (say, 40 %), agents are "cautious", i.e. they decline to follow contagion when the action suggested by the neighborhood is to trust the banking system and make a deposit. The evolution obtained by introducing these changes is presented in figure 4(b).

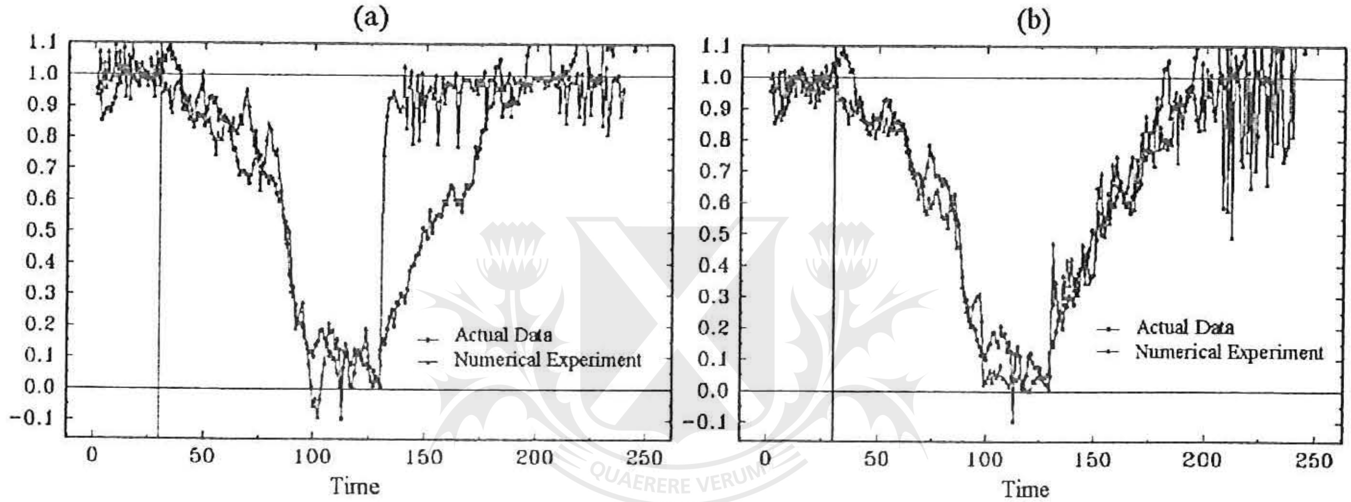

Figure 4: Normalized deposits demand. The numerical experiment has the following parameters: $N_{ag} = 256$, $N_p = 20$, $T = 0.001$, the mutation probability of the genetic algorithm is 0.005. The crisis begins in $t = 30$ (December 20th, 1994) and finishes in $t = 120$. We overlapped the series of the total daily level of deposits during the banking crisis in Argentina. In the case shown in (b) we have assumed that, if the fraction of agents acting by contagion is greater than 50 %, a new population of "predictors" is drawn again, and 60 % of the agents are inhibited to deposit under the inducement of contagion.

5 Conclusions

The system that we have considered possesses two stationary regimes. In one of them, the agents have learned to accomodate their deposit strategies to avoid exceeding an upper limit of deposits determined by the (parametrically given) "physical size" of the system. The model also contemplates the possibility of an "autarchic equilibrium" in which each agent retains all the funds and refuses to use the banking system. The adjustment and learning rules are capable of describing the self-organization process leading to an equilibrium, and they *also* describe the transitions that occur as a consequence of an exogenous shock: it is not necessary to make any particular assumption about the behavior of the agents before, during or after the crisis; neither is necessary to taylor the model for specific durations of the shock. The rules that govern the adaptation and the processing of information are the same before and during the "crisis" episode. Even though the model has two characteristic times (the "day" and the "week"), the adaptive processes give rise to intermediate situations in which these characteristic times are insufficient to fully describe the behavior of the system.

The study of the crisis also makes it possible to analyze in some detail (and with references to actual phenomena) the efficacy of the adaptive processes and the learning procedures of the agents. This point has a particular relevance because, in general, under stationary circumstances there is no way to consider the transients that may take place while the system self-organizes.

The crisis, which has its own characteristic times, places the system in a changing environment, where the evolution depends on the learning *rate* of the agents. Representing the recovery from the Argentine banking crisis requires the introduction of several learning processes in more than one time scale, suggesting that some type of "long memory" of previous events may operate in certain conditions as an influence on current behavior.

References

[1] Diamond, D. and Dybvig, P. (1983), *Bank Runs, Deposit Insurance and Liquidity*, Journal of Political Economy, **91**, 3.

[2] Jacklin, C. and Bhattacharya (1988), *Distinguishing Panics and Information Based Bank Runs*, Journal of Political Economy, **96**, 3.

[3] D'Amato, L. Grubisic, E. and Powell, A. (1997), *Contagion, Bank Fundamentals in Macroeconomic Shock? An Empirical Analysis of the Argentine 1995 Banking Problems*, mimeo.

[4] Temzelides, T. (1994), *Evolution, Coordination, and Banking Panics*, mimeo, june 1994.

[5] Arthur, W.Brian (1994). *Complexity in Economic Theory: Inductive Reasoning and Bounde Rationality*, AEA Papers and Proceedings, **84**,2 , May.1994 , 406-411.

[6] Goldberg, D.E. (1989). *Genetic Algorthms in Search, Optimization, and Machine Learning*

[7] Amit, D. J. (1989), *Modeling Brain Function: The World of Attractor Neural Networks*, Cambridge University Press, chapter 3.