

LEARNING AND IMITATION: TRANSITIONAL DYNAMICS IN VARIANTS OF THE BAM

D. HEYMANN

*CEPAL (Comisión Económica para América Latina y el Caribe, United Nations) and
Facultad de Ciencias Económicas, Universidad de Buenos Aires
Paraguay 1178, (1057) Buenos Aires, Argentina
dheyman@cepal.org.ar*

R. P. J. PERAZZO*

*Centro de Estudios Avanzados and Departamento de Física
Facultad de Ciencias Exactas y Naturales, Universidad de Buenos Aires
Ciudad Universitaria, Pabellon 1, (1428) Buenos Aires, Argentina
rperazzo@itba.edu.ar*

A. R. SCHUSCHNY

*CEPAL (Comisión Económica para América Latina y el Caribe, United Nations)
Av. Dag. Hammarskjöld 3477, Vitacura, Santiago de Chile, Casilla 179-D, Chile
schuschny@yahoo.com*

We study the dynamics of self-organized systems when disturbed by shocks. For this purpose, we consider extensions of the “Bar Attendance Model” [1] (BAM), which provides a stylized setting for the analysis of the emergence of coordination in the behavior of a large collection of agents. We represent the learning process of the agents through genetic algorithms, which respond to global (publicly available) information. In addition, we allow the actions of agents to be influenced by local information, as expressed in the behavior and performance of neighboring individuals. In the context of the BAM, we show that, in the event of a shock, the imitation behavior may become widespread and generate a contagion cascade which mimics a collective panic. We use this framework to represent features of the dynamics of an actual bank run.

Keywords: Coordination problems; self-organization; transitional dynamics; contagion effects; genetic algorithms; multi-agent models.

*Present address: Instituto Tecnológico de Buenos Aires, Departamento de Investigación y Desarrollo, Av. Madero 399, Buenos Aires, Argentina.

1. Introduction

The self-organization processes in multi-agent systems have been receiving considerable attention in a variety of contexts, from the routing of messages in information networks to the behavior of ensembles of economic agents. The models that study such processes try to understand how and when individual decisions induce collective effects which show up in the performance of the system as a whole and give rise to some sort of macroscopic order. A canonical instance of this line of research is the “Bar Attendance Model” (BAM) (also referred to as “El Farol Model”) proposed by Ref. 1, which provides a simple illustration of a system where coordination is achieved through the inductive learning of agents. In that model, a (large) number of individuals have to decide whether or not to go to a bar at a certain date. The physical capacity of the shop is taken as given. The prospective customers share the perception that the bar becomes unacceptably crowded when the number of people who are present exceeds a critical value, which is smaller than the total population. Otherwise, if attendance is below that threshold, they derive utility from going to the bar.

The model is such that agents choose an action individually, based on information that can be considered global (i.e. it refers to aggregate values of the system) and public (i.e. it is available to all the agents of the system). This information is assumed to be the total attendance in the previous days. In the model agents are assumed to perform inductive reasoning to predict the number of customers at the shop in the next day. The system self-organizes in such a way that the attendance converges to the maximum acceptable level (plus or minus random fluctuations) [14].

In real life, the actions of groups of agents sometimes appear to undergo sudden changes, as if there was a collective shift to a new mode of behavior [13]. This paper explores the dynamics in such transitions, when the internal order of a coordinated, locally stable system breaks down, and a new configuration emerges. We carry out the analysis with variants of the BAM. In all cases, the learning process of agents is represented by a genetic algorithms (GA). In the context of the BAM, the basic assumption is that the actions of an agent are generated by a population of attendance strategies, each of which determines a pattern of presence and absence for a period covering a certain number of days. Those alternative attendance programs are revised through the GA in view of the utility that the strategies would have generated to the agent during the period. The interplay of the strategy revisions of the ensemble of agents produces a co-adaptive behavior which results in a coordinated state. We concentrate on instances where that state is disturbed by a shock.

We add several ingredients to the standard BAM. One is the introduction of a lower bound for utility-generating attendance (so that agents, say, dislike a bar with too few customers). This opens the possibility of obtaining different macroscopic equilibria. In this regard, we show that for some initial conditions, the system may get trapped in states where the shop remains empty in some days: this corresponds to situations where agents develop the self-fulfilling anticipation that total

attendance will not be sufficient to reach the lower threshold. A zero-attendance equilibrium in every period generated in this way seems to bear an analogy to an expectations-driven bank run.

The second variant concerns the learning procedure of agents, by incorporating information arising from the local environment of a particular individual. This is done in order to contemplate possible contagion effects.^a We assume that agents have access to information regarding the behavior of others who are close to them, and can condition their decision on such information when they find that their own strategies do not generate reliable or satisfactory outcomes.

We first analyze the robustness of the self-organization process in this modified setting. Then, we consider the effect of a fundamental shock that makes visiting the shop no longer convenient for agents. In the absence of contagion, no abrupt transitional dynamics emerges. However, when the imitation channel for modifying behavior is introduced, the evolution of the system shows definite features. The simulations that we perform suggest that, in such a situation, the decline in the quantity of customers is gradual at first, as individuals, based on their own specific learning procedures, start deserting the shop. This can be interpreted as the diffusion of an awareness of a change in conditions. After a few time periods, the number of clients starts to fall faster, while contagion becomes widespread. A situation of panic builds up. If the exogenous shock is reversed soon, the system rapidly returns to the original, undisturbed state. However, in the case where the new value of the lower threshold is maintained, a large-scale movement develops, with a drastic fall in attendance. If now the exogenous parameter returns to its initial level, reversing the fundamental shift, a gradual reorganization of the system takes place. This can happen in two ways. If the bad fundamental state has lasted for a long time, the model generates the possibility that all individuals stay away from the shop, leading to an autarchic equilibrium. Alternatively, if the crisis is short-lived, the system again converges to a stationary state with a high level of attendance.

While a “run” in the BAM has only loose analogies with a banking crisis, we show that the model including the contagion effect can account for some qualitative characteristics of an actual run, that which took place in Argentina at the beginning of 1995. This may suggest that the volume of bank deposits in such an episode was driven both by learning about the fundamental conditions of the system, and by imitative responses to what other agents were observed to be doing.

2. Coordination through Public Information

In this section we present a modification to the BAM to consider a situation in which it is common belief that the attendance should not exceed a given value (denoted by S_1) for customers to be comfortable and, at the same time, it should not fall below a lower bound (S_0) to make attendance convenient. The adaptation procedure is based upon the past attendance recorded at the bar, an information

^aThis type of argument has been used in models of herd behavior; cf. Refs. 2 and 3.

that is assumed to be available to all the agents of the system, the individual decisions make use only of public, global (aggregate) information. We model the relaxation to equilibrium using genetic algorithms [10].

The learning procedure works by improving strategies taking into account the performance during a set of consecutive visits to the bar. Each agent has assigned a population of N_p attendance plans for the subsequent period of N_d days. Each strategy is encoded in a genome that specifies each day whether the agent goes to the bar or not. All the initial populations of genomes are selected at random, and are updated by selection, crossover and mutation [10], so that less successful strategies are progressively eliminated.

The formal steps of the algorithm are the following:

1. Each agent possesses N_p attendance strategies. Each strategy consist of a chain of N_d bits where N_d stands for the number of days of the attendance 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 she chooses to go (not to go) to the bar in that day.^b
2. Once each agent has selected one strategy (chosen on the basis of the data of the preceding period of N_d days), she implements daily the action of either going or not going to the bar. The bar records the daily level of attendance. Let c be the current strategy of each agent. The fraction D_t of the public who has attended the bar in the t -th day is:

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

where $\theta(x) = 1$ if $x > 0$ and $\theta(x) = 0$ otherwise.

3. Each agent calculates the utility obtained that day by comparing the attendance D_t with the two bounds S_0 and S_1 . The agents who have made a correct guess of the attendance level add a positive contribution to their utility, $U_k^p(t+1) = U_k^p(t) + (1+C)$ regardless whether they have attended the bar or not. The agents who have made a wrong guess of the attendance level bear a cost that lowers their utility, $U_k^p(t+1) = U_k^p(t) + (1-C)$. The cost is the same whether they have attended the bar or not. The utility is recorded daily during the period of N_d days.
4. When the period finishes, the agent compares the performance of all her strategies.^c The fitness of each strategy is determined by the utility that could have been achieved with the corresponding attendance policy. At the end of the period of N_d days all the available strategies are ranked according to their utility (fitness).

^bThe model only involves pure strategies, i.e. it does not consider the case in which agents attend the bar with some, possibly time varying probability.

^cIt is assumed that each individual change of attendance strategy causes a negligible change in the total attendance of each day, otherwise the coadaptation process requires to test each possible combination of individual strategies i.e. a total of $N_{ag}^{N_p}$ combined situations.

5. The population of strategies associated with each agent is updated according to the usual procedure of the genetic algorithm. We have used an *elitist* version of the genetic algorithm: those genomes ranked within the upper 50% survive to the next generation, while those in the lower 50% are replaced by new ones obtained by crossover of the surviving strategies. All the genomes (strategies) are next subject to random mutations, i.e. a bit is changed from 1 to 0 or vice versa with probability p_{mut} . According to the common practice p_{mut} is chosen to be small (typically $p_{mut} \sim 0.001$) since otherwise the evolution process does not converge to any stable state because successful strategies do not survive.
6. The fittest strategy is selected as the current attendance schedule for the next period of N_d days and the whole procedure is iterated.

In the usual settings of the BAM (with only an upper bound of tolerable attendance) the outcome is an asymptotic coordinated state in which the agents self-organize to go to the bar as frequently as tolerated by the upper crowding threshold S_1 (see Fig. 1). Its stability is *due to the internal diversity of the system*: the agents tend to use complementary strategies that avoid crowding. By contrast, if all the agents chose the same schedules the resulting attendance level would fall outside the acceptable range. As a consequence the system evolves an ensemble of strategies by which different agents go to the bar in different days. Such configurations are *dynamically* stable. All agents continue to explore new strategies that correspond to random mutations of their current plans and the system only visits configurations that are in the close neighborhood of the state that has been reached.

The case that we consider in the present paper involving an upper *and* a lower bound of acceptable attendance, does not give rise to a unique coordinated state. In addition to the full attendance case it may also happen that the ensemble of agents gets trapped in a configuration in which, during certain days, the bar receives no customers (Fig. 1). After few time steps of evolution and simply due to the random fluctuations of the adaptation process, the attendance at one particular day turns out to be below the lowest threshold S_0 . When this happens all the clients that choose to go to the bar on those days get a negative contribution to their utility. The choice of attending the bar is thus progressively discarded in the following adaptation steps and therefore it becomes a common belief that the bar will be empty those days. The occurrence of a configuration of this sort is an example of path dependence in the self-organization process, because it becomes stable in an early stage of adaptation, conditioning all the future evolution of the system [9].

When the system is confined into one of these configurations, the adaptation dynamics allows a transition to the macroscopic ordered state of full attendance. Such transition is however highly improbable because it requires that a large enough fraction of agents *simultaneously* coincide in changing their individually chosen actions for a particular day. In statistical mechanics, this situation is usually referred to as confinement by *entropic barriers* [17].

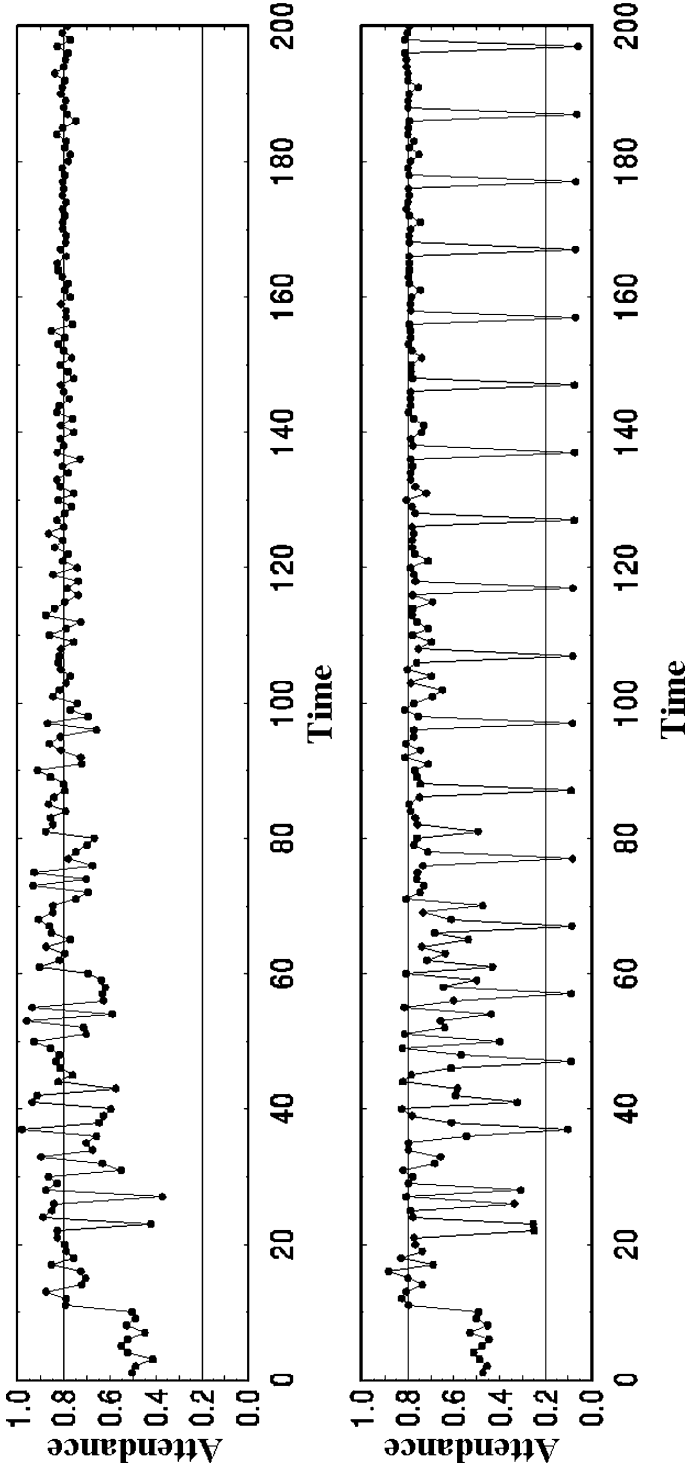


Fig. 1. Simulations without local interactions with 256 agents, $C = -0.01$, $p_{mut} = 0.005$, $S_0 = 0.2$, $S_1 = 0.8$, $N_d = 10$ and $N_p = 20$. The upper panel shows an optimum convergence while the lower one shows how the self-organization results a in sub-optimal state. Each day 10% of the agents update their strategies through the genetic algorithm.

In either of the cases presented above (full attendance or empty days) the system visits the neighborhood of a Nash equilibrium. No agent can improve her utility changing unilaterally her attendance strategy. There is a large multiplicity of equilibria: any two configurations in which a client of the bar exchanges her position with one who has not attended, are equivalent. One can check that the cumulative attendance of the agents for the whole period of N_d days follows a binomial distribution. In the stationary state the (time) average of the fraction of the clients that go to the bar S_1 , is the same as the probability that a given agent can attend the bar. Thus the probability to reach a weekly cumulative attendance of A_w days is

$$P(A_w) = \binom{N_d}{A_w} S_1^{A_w} (1 - S_1)^{N_d - A_w}. \quad (2)$$

3. Coordination in the Presence of Contagion

3.1. *The model for local contagion*

In the previous section we have prepared the stage to include a new information channel in the self-organization process. We can now incorporate into the model the effect of local information, in the sense that it reflects the situation of a small set of customers belonging to the neighborhood of any given agent. This can be considered a way to model a contagion mechanism. Here every agent can choose between two alternatives: (i) to continue with her individual attendance plan or (ii) to imitate the strategy followed on the average by her nearest neighbors. The neighborhood is determined by placing the agents in a square grid 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 her current strategy whenever she finds that her 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 located in the site (i, j) , on the t -th day, and her *current plan* for the same day $^c \#_{(i,j)}^t$. To decide a shift away from her plan, the agent determines the daily outcome obtained with her strategy and compares it with the one obtained if she had imitated her nearest neighbors. For this purpose each agent calculates the local “mean action” (field):

$$h_{(i,j)} = \frac{1}{4} \left(^c a_{(i,j+1)}^t + ^c a_{(i,j-1)}^t + ^c a_{(i+1,j)}^t + ^c a_{(i-1,j)}^t \right). \quad (3)$$

Once this is known, the agent can determine if she will follow her neighbors.^d The alignment of each agent with the average local attendance policy is however not entirely deterministic. We assume that there is some noise in the agent’s decision that is characterized by a control parameter $T = 1/\beta$. This is introduced through a

^dThe arrangement of agents in a grid is borrowed from the *Ising Model* of ferromagnetism.

Glauber dynamics [11] in which the probability of either imitating or ignoring the local policy of her neighborhood is

$$P(\text{action} = \pm 1) = \frac{1}{1 + e^{\mp 2\beta h_{(i,j)}}}. \quad (4)$$

Thus, each agent computes every day the utility that would have been obtained following the local mean action. If this is greater than the one computed with her current strategy, the agent changes her attendance plan *for the next day* and uses (with the above probability) the action prescribed by the local field. Otherwise, the agent sticks to her individual plan, using the next bit of information of her current strategy.

3.2. *Self-organization with local effects*

The first question that we have to answer is whether the self-organization process presented in the previous section does not break down by the addition of local information into the learning process of the agents. The results that we obtain are shown in Fig. 3. The self-organization is indeed preserved in the presence of contagion effects.^e However, the new organization process involves larger level of fluctuations which are to be attributed not only to the random search performed in the individual optimization, but also to the occasional contagion that is induced by the Glauber dynamics.

In the first 300 time steps of Fig. 3 we show the evolution of the system when both information channels are active. We have fixed the random fluctuations of the Glauber dynamics at $T = 0.1$. As follows from the figure, for this value of T and under stationary conditions, approximately 15% of the agents constantly switch strategies, imitating their neighbors. From the genetic algorithm point of view, that effect is equivalent to an upward revision of the value of p_{mut} . With such moderate values of T , the presence of local contagion does not prevent the system from self-organizing. The learning procedure partially compensates the new instabilities induced by fluctuations, by promoting attendance strategies with a somewhat higher attendance than the highest tolerable level. The presence of these larger fluctuations has, on the other hand, an important consequence, namely the observed disappearance of metastable states with zero attendance in some specific days: the buildup of the entropic barriers mentioned in the preceding section is prevented because in the initial adaptation steps the dynamics forces the system to explore a wider space of possible strategies.

^eIt should be noted that the system is initialized with random attendance of agents. Thus the initial conditions allow for a considerable number of individuals going to the shop at a given day: it is as if there has been already a trial period in which agents have known the shop, so that the lower threshold of attendance can be exceeded. The problem of restoring attendance once has been reduced below S_0 is studied in the context of the exercise where the system is disturbed by a shock.

With the above choice of parameters learning is still possible because the fluctuations are not large enough to destroy the information content of the selected strategies of each agent. For even larger values of T all agents follow an increasingly random attendance policy that amounts to deciding to go to the bar by tossing a coin. In this situation learning becomes impossible and the system fails to self-organize.

4. The Transitional Dynamics

The relevance of the interplay of both coordination mechanisms gives rise to noticeable effects in the event of an external shock representing a sudden change in the global conditions. An exogenous disturbance can be simulated by using the lower threshold S_0 as a control parameter and changing it to a value greater or equal to S_1 . This compels the whole set of agents to revise their decisions. The crisis can be finished by returning S_0 to its original value some time later.

It is important to compare the behavior of the system with and without contagion. In Fig. 2 we show the evolution during the disturbed period in the absence of the local information channel. After the shock, the population does not switch behavior abruptly; rather, agents gradually adapt their strategies to the new situation as they learn that it is convenient to avoid going to the bar. The rate of adjustment of the average attendance is governed by the mutation probability of the genetic algorithm. The evolution has an exponential time constant given by $1/p_{mut}$ which is long compared with the length of the period of N_d days. If the shift in S_o persists, agents eventually learn to desert the bar.

The main effect of the local contagion is to open the possibility of coherent effects in which agents imitate each other, giving rise to rapid changes in aggregate behavior. Once the crisis begins, the collective performance of the clients of the bar displays two features. First, there is a period of awareness in which the agents who were imitating their neighbors no longer do so, and choose to use their own strategies (notice the dip in the contagion effect right after $t = 300$ in Fig. 3). The intuitive reason for this is that on average agents make more wrong decisions and get poor utilities; therefore, contagion offers no advantage. For the settings chosen for the simulations, this process may last for a period of the order of a couple of business periods of N_d days.

Second, as the threshold S_0 is maintained at its new high level, the agents increasingly learn to avoid going to the bar. A moment is then reached in which a critical fraction of individuals acting in that way induce an avalanche in which agents massively imitate the strategy of not going to the bar, because acting in this way produces a good utility to their neighbors. This has all the features of a panic wave. Thus, the period of awareness (in which the agents act according to their own learning) leads at some point to a contagion cascade or panic wave, with an abrupt decline in the attendance level (see Fig. 3). This collective behavior is triggered only if a large enough fraction of agents have decided by themselves to

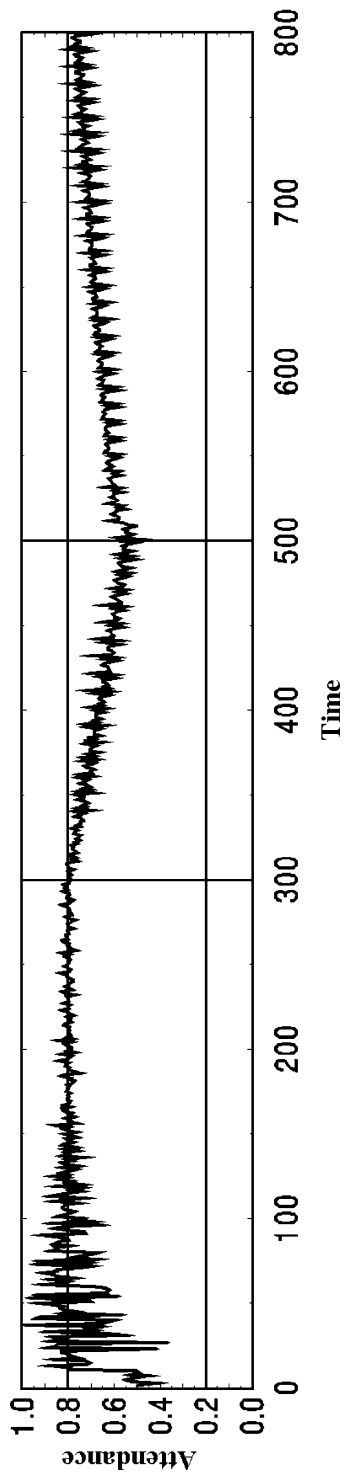


Fig. 2. Evolution during a exogenous crisis without contagion. All the parameters are the same as before.

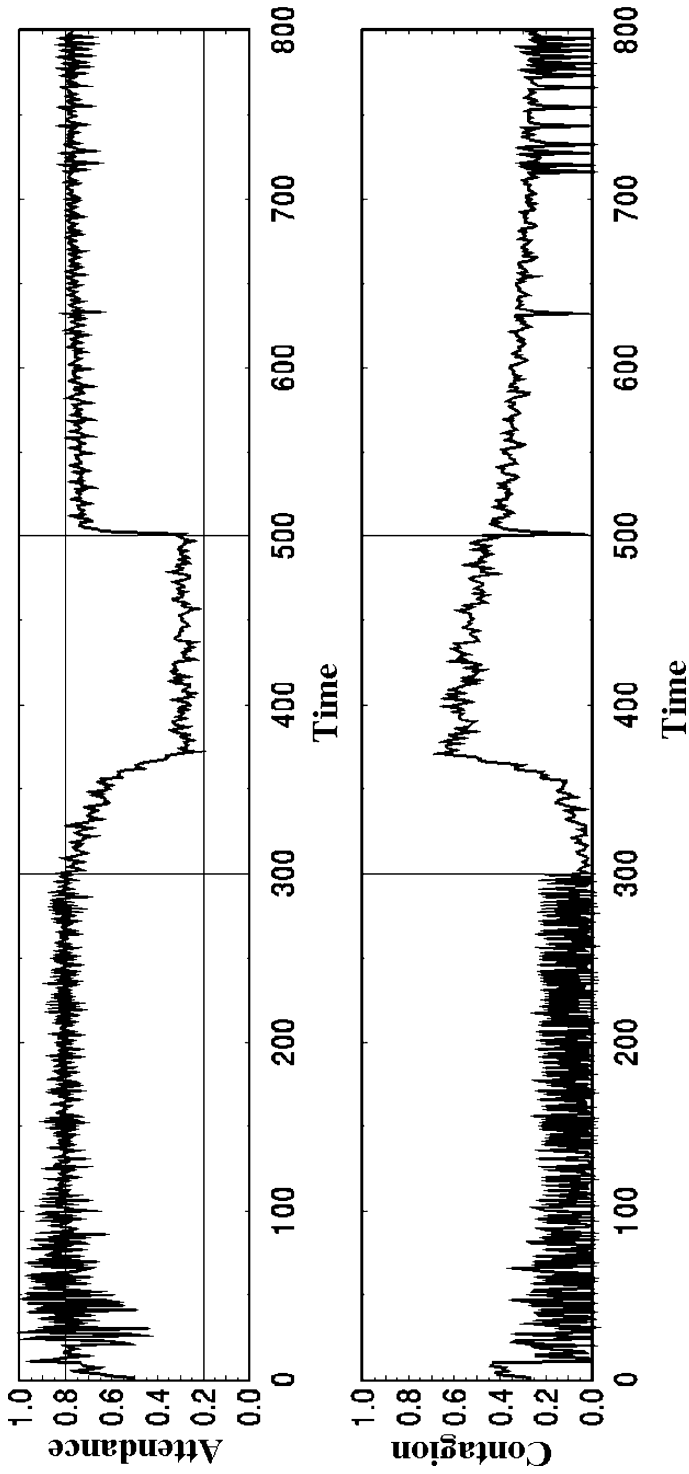


Fig. 3. The self-organization process and the effects of a crisis are shown when both information channels are active and with $T = 0.1$. The lower panel gives the effect of contagion i.e. the fraction of the agents that act according to the local field. The parameters are the same as those of the previous figure. Notice the awareness, panic and recovery periods.

avoid going to the bar. When the panic emerges, the transition takes place in a very short time.

If the disturbance is short-lived, the possibility of reaching a coordinated state of high attendance remains essentially unaffected because the agents do not stop considering previously successful attendance strategies. Recovery in this case is also quite abrupt requiring little re-training of the agents. The reason is that after few days in which the lower threshold has regained its previous value, most of the agents have not yet lost memory of old strategies and find that they are again suitable, so that many of them go back to using those strategies “as if nothing had happened”. On the other hand, if the crisis is long enough to make agents forget the once successful strategies, the new equilibrium corresponds to the bar with no clients.^f

After a longer time the only successful overall strategy that survives is never going to the bar. In this stage both (genetic) learning *and* contagion mechanisms act coherently to generate the same behavior. In the model, once attendance collapses, it is highly difficult to reconstruct it. The system finds itself trapped again between high entropic barriers. A very improbable circumstance has to occur to escape this situation, namely that many agents decide *by chance and at the same time* to go to the bar again. If this happens that strategy would produce a high utility and stabilize again a regular attendance. As already observed, such barriers can only be overcome allowing a high level of fluctuations (setting for instance T to a high value), so as to induce the system to visit configurations which are far from the current state.

5. An Application of the Model: Stylized Features of a Bank Run

In this section we show how some of the above discussed qualitative features can be found in the case of a severe banking crisis that took place in Argentina in the aftermath of the Mexican devaluation by the end of 1994. This is a particularly interesting case because a quantitative record of the transitional dynamics is available.

Banking panics have been interpreted in two ways (c.f. Refs. 7, 8, 12 and 19): as events based on the diffusion of information about economic fundamentals, which induces a revision of beliefs about the prospects of the banks, or as pure coordination phenomena in a system with multiple potential equilibria, when some extraneous shock causes agents to change the state on which they focus their expectations. Both explanations need not be mutually exclusive. It may happen, for example, that some macroeconomic news, which modify the perceived return on deposits,

^f Once the shock has taken place and before zero attendance is reached, its value stays for a long time at a value that is close to $1 - S_1$ (0.2 in the case shown). This is so because all successful strategies required that all clients, at least 20% of the time, should not go to the bar. Thus at least 20% of the time the agents would follow their strategy (and mistakenly go to the bar at some future date) instead of imitating the neighbors. This effect is reinforced by the fact that the neighborhood strategy is followed only if it is strictly better than one's own.

induce a response which in turn triggers a self reinforcing avalanche of withdrawals, thus amplifying or accelerating the effect of the impulse. In this case, the panic is not simply generated by a sunspot (since without the news, agents would disregard the multiplicity of 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.

Although the environment that we depict is highly schematic, it is tempting to use it to consider instances for which an actual quantitative record is available such as is the case with bank deposits. As we see below, the model can capture stylized features of the collective behavior that is found in a crisis episode and its aftermath.

At a certain level, one can find intuitive analogies between the action of going to the bar in the BAM, and that of making a deposit in a bank. It seems clear that people space their visits to banks in such a way that, in a stationary state, the aggregate attendance level remains more or less constant while the identity of the individuals who decide to go changes from day to day. This behavior could be rationalized, by analogy with the BAM, by defining a maximum attendance level beyond which, given the physical capacity of a bank, customers would incur some congestion costs. However, that somewhat strained analogy would provide a weak basis for a representation of the behavior of financial activities. In any case, our interest in the use of this self-organization model in a financial context has little to do with the way in which customers routinely choose the days in which they will be physically present at banks. Rather, what we want to consider is how the system reacts to a fundamental shock that creates doubts about the solidity of banks. In this regard, the existence of strategic complementarity in the decisions to place resources in banks is a standard feature of the relevant literature.

One may well assume, as a simplifying hypothesis, that there is a lower bound in the volume of bank deposits such that, if the aggregate value goes below that bound, prospective depositors suffer losses due to the inability to fully recover their funds; that lower limit would be a function of parameters such as the profitability of banks. Thus, bad news would correspond to an increase in that lower bound of deposits such that an individual deposit is no more profitable. In the limit, the value of the bound would be so high that well informed agents would desert the bank. This is the scenario that we want to analyze. As stressed before, the self-organization in the timing of attendance (and thus, the nature of the upper bound on attendance) is of very secondary importance for our purposes: nothing would be changed if we just predetermined the days in which a given agent would go to the bank in case he wished to make a deposit; the issue is whether he would want to trust the bank at all.

Our experiment consists in shocking the system so that making deposits in the bank is fundamentally inconvenient, and investigating the dynamics of the process through which agents learn to abstain from making deposits, and how (and if) they again return to the bank if the fundamental parameter goes back to its initial value. Within this picture the attendance schedule of each agent, reflects the planned

deposit schedule for the next N_d market days, i.e. it would reflect the planned decision of making a deposit in the bank or not doing it. It is implicitly assumed that at the end of each day the agent gets back her deposit, collects the corresponding interests and continues with her investment schedule for the following day. The fact that the level fluctuates reflects the random events that occasionally induce one agent to retain the money or make a deposit. When an agent changes a 0 into a 1 in her attendance schedule, she switches programmed action so as to make a deposit that day, and the opposite when the change is from 1 to 0. In either case there is a (small) change in the aggregate value of the deposits for that day.

The origin of the crisis may be attributed to the arrival of an exogenous information that casts doubt on the fundamental reliability of the financial system; the end of the disturbance can be associated to news (e.g. of government actions) which in some way determine an enhancement of the ability of the banks to honor deposit contracts.

A severe banking crisis took place in Argentina in the aftermath of the Mexican devaluation by the end of 1994 [15]. In Fig. 4(a) we plot the total daily level of deposits (in current accounts, savings and time deposits, aggregating accounts in Argentine pesos and foreign currencies) during the period between November 1994 and October 1995. The picture also shows values generated with the present model. The free parameters have been calibrated as explained in the figure caption. The vertical scale has been normalized in such a way that the upper level corresponds to the volume of deposits at the beginning of the crisis and the lowest to the amount of funds at the banks 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 attendance model.

The phases of normal performance, gradual withdrawal of deposits (awareness), panic and recovery can be clearly identified in the figure. The simulated series describes quite well the first stages of the crisis, but not the recovery. This happens because in the mean, agents have not forgot completely strategies that were successful prior to the crisis and quickly start using them again. This can therefore be attributed to the fact that no memory is built into the artificial system to account for the asymmetry in the data which displays a sudden run and a slower return of deposits to the banks.

In order to include this kind of long term memory, we have elaborated a slightly modified algorithm. We first choose a situation in which new initial strategies are drawn, letting the genetic algorithm start again with its adaptive work in a fundamentally new state of the system. This modification would correspond to the introduction of new rules of behavior. For the purpose of the simulation, this procedure was triggered whenever more than half of the agents act according to contagion. The second variant that we introduced 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

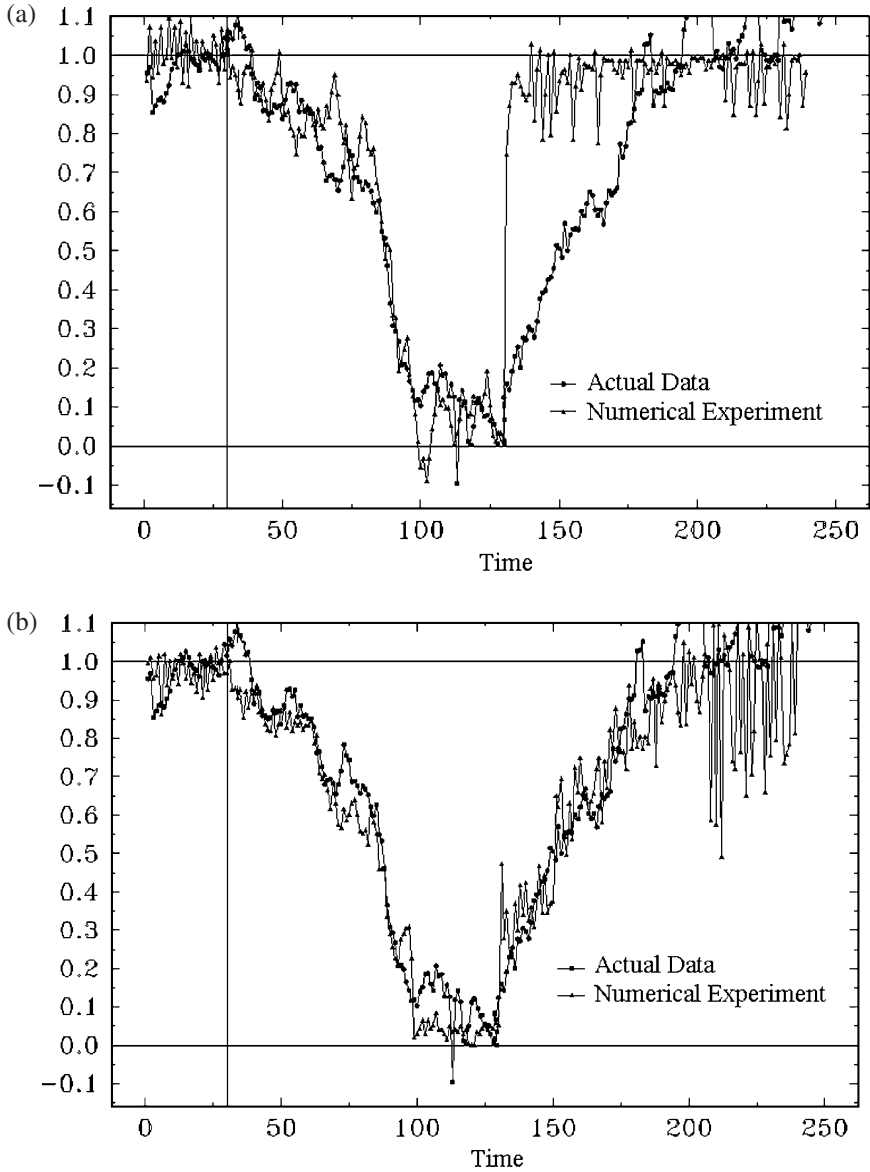


Fig. 4. Normalized deposits during the Argentine banking crisis. The numerical experiment has the following parameters: $N_{ag} = 256$, $N_p = 20$, $T = 0.001$, the mutation probability of the genetic algorithm is 0.005. The crisis begins in $t = 30$ (December 20th, 1994) and finishes in $t = 120$. We overlapped the numerical experiment and the series of the total daily level of deposits. In the panel (b) we have included in the experiment two threshold effects: (i) if the fraction of agents acting by contagion is greater than 50%, a new population of strategies is drawn again, and (ii) 60% of the agents are inhibited to deposit when induced by contagion.

Fig. 4(b): it can be observed that there is a reasonably good fit, both in the period where deposits declined and in the upswing.

6. Conclusions

We have considered several extensions of the BAM in which individual agents mutually adapt their individual actions to reach a coordinated, macroscopically ordered configuration. In these models agents learn to adjust their attendance schedules to avoid exceeding and staying below certain limits of attendance. The multiagent system displays several possible collective orderings. One is an equilibrium in which all agents coordinate their individual attendance schedules to saturate the maximum accepted capacity of the bar. Another is a kind of autarchic equilibrium in which agents refuse to go to the shop. There is also a family of suboptimal, metastable configurations in which the bar is deserted in some specific isolated days. All these states can be connected to the full attendance configuration, but moving from one to the other would require the highly improbable situation of a simultaneous change of behavior of a large collection of agents. This is a qualitative feature that can be recognized in a wide number of everyday situations in which overcoming the barriers amounts in practice to a global reorganization of the system [18].

We have also considered the use of information that concerns the immediate neighborhood of each agent. This provides a local perspective in the decisions of the agents. We found that the self-organization mechanism based only on public information (i.e. the aggregate attendance level) can be considered robust in the sense that addition of contagion effects does not modify the end result of the adaptation process. However the presence of the two alternative information channels was found to play an important role in shaping the transitions between regimes. The adaptation in the original version of the BAM can only give rise to gradual changes. The interplay between the two information channels can instead generate abrupt changes, which are commonly observed in real circumstances. The composite effect of the two alternative mechanisms is to turn the ensemble of multiple agents into what is known as an *excitable system*. These are complex systems found in a wide variety of fields that encompass social, biological and physical situations displaying a special kind of bounded homeostasis (or “corridor effects” as named by Ref. 16). They are stable against small perturbations, but if the shock is larger than a given threshold the system only turns back to its original equilibrium after performing a long excursion away from that situation. We found that such excursion (or transitional dynamics) in the present framework clearly displays differentiated moments that may be taken to correspond to awareness, panic and recovery.

It is important to remark that the *same* learning (adaptation) rules are capable of describing the widely different situations that correspond to the self-organization leading to an equilibrium on the one hand, *and* the out-of-equilibrium transition that follows exogenous shocks of various intensities on the other. It is not necessary to make any other particular assumption about the behavior of the agents with

reference to any particular global situations that may arise either before, during or after the crisis; neither is necessary to tailor the model for a specific duration or intensity of the shock. The rules that govern the adaptation and the processing of information are the same before, during and after the transition episode and constitute an essential part of the multiagent system.

We have also used the model to understand on purely qualitative grounds, the transitional dynamics generated during the bank run that took place in Argentina as a consequence of the devaluation of the Mexican peso. It is abundantly clear that the metaphor of the BAM should not be stretched: the model can certainly *not* be taken to describe the workings of a financial system. However the ability of the model in reproducing stylized features of the empirical quantitative data of the transitional regime, and the enormous complexity of the interactions that are involved in the real world dynamics, suggests that the transitions between regimes may be roughly described by simple schemes of self-organization in spite of the fact that the real life, detailed internal mechanisms may be very intricate.

In that application the emergent collective pattern of behavior could be described with a simple structure that involves the interplay of only few degrees of freedom concerning public (global) information and the contagion of individual decisions. If considered as arising from a learning or adaptive mechanism, abrupt changes can emerge in a natural way by including this second information channel. On the other hand the interpretation of the gradual recovery after the crisis may have different interpretations; as either a long term memory of events that happened before the crisis or as a co-adaptative, recovery of confidence in the banking system.

Acknowledgments

The authors wish to acknowledge an anonymous referee for helpful comments. R. Perazzo acknowledges partial support from the research contract EC 9000XX/27 and for the warm hospitality at the ZIF of the University of Bielefeld, where part of this work was carried out.

References

- [1] Arthur, W. B., Complexity in economic theory: Inductive reasoning and bounded rationality, *AEA Papers and Proceedings* **84**, 406–411 (1994).
- [2] Banerjee, A. V., A simple model of herd behavior, *Quarterly Journal of Economics* **110**, 797–817, (1992), Banerjee, A. V., The economics of rumours, *Review of Economic Studies* **60**, 309–337 (1993).
- [3] Bikchandani, S., Hirshleifer, D. and Welch, I., A Theory of Fads, Fashion, Custom and Cultural Change as Information Cascades, *Journal of Political Economy* **100**, 992–1026 (1991).
- [4] Burgos, E., Ceva, H. and Perazzo R. P. J., Quenching and annealing in the minority game, cond-mat/0007010 (2000) and Dynamical quenching and annealing in self-organization multiagent models, cond-mat/0012412 (2000).
- [5] Cavagna, A., Comment on “Adaptive competition, market efficiency, and phase transitions”, *Phys. Rev. Lett.* **84**(5), 1058 (2000); and R. Savit, Reply to Comment on

- “Adaptive competition, market efficiency, and phase transitions”, *Phys. Rev. Lett.* **84**, 1059 (2000).
- [6] Challet, D. and Zhang, Y.-C., *Emergence of cooperation and organization in an evolutionary game*, *Physica A* **246**, 407–417 (1997).
- [7] D’Amato, L., Grubisic, E. and Powell, A., *Contagion, Bank Fundamentals in Macroeconomic Shock? An Empirical Analysis of the Argentine 1995 Banking Problems*, mimeo (1997).
- [8] Diamond and Dybvig, Bank Runs, Deposit Insurance and Liquidity, *Journal of Political Economy* **91**, 3–15 (1983).
- [9] Giardina, I. and Bouchaud, J.-Ph., *Bubbles, Crashes and Intermittency in Agent Based Market Models*, mimeo (2002).
- [10] Goldberg, D.E., *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley (1989).
- [11] Glauber, R. J., Time dependent statistics of the Ising Model, *Journal of Mathematical Physics* **4**, 294–310 (1963).
- [12] Jacklin, C. and Bhattacharya, Distinguishing panics and information based bank runs, *Journal of Political Economy* **96**, 3–14 (1988).
- [13] Jefferies, P., Lamper, D. and Johnson, N. F., Anatomy of extreme events in a complex adaptive system, preprint (2002).
- [14] Johnson, N. F., Jarvis, S., Jonson, R., Cheung, P., Kwong, Y. R. and Hui, P. M., Volatility and agent adaptability in a self organizing market, *Physica A* **256**, 230 (1998).
- [15] Kaminsky, G. and Reinhart, C., On crises, contagion and confusion, *Journal of International Economics* **51**, 145–168 (2000).
- [16] Leijonhufvud, A., Effective demand failures, *Swedish Journal of Economics* **75**, 27–58 (1973).
- [17] Palmer, R., Broken ergodicity, in *Lectures in the Sciences of Complexity, Vol I*, D. Stein (ed.); Proceedings of the 1988 Complex Systems Summer School; June–July 1988 Santa Fe New Mexico, USA. Addison-Wesley, 1989.
- [18] Savit, R., Manuca, R. and Riolo, R., Adaptive Competition, Market Efficiency, Phase Transitions, *Phys. Rev. Lett.* **84**, 4, 5 (2000).
- [19] Temzelides, T., *Evolution, Coordination, and Banking Panics*, mimeo, June 1994.