

Market Organisation

GERARD WEISBUCH ⁽¹⁾

ALAN KIRMAN ⁽²⁾

and DOROTHEA HERREINER ^(2,3)

(1) Laboratoire de Physique Statistique
de l'Ecole Normale Supérieure,
24 rue Lhomond, F 75231 Paris Cedex 5, France.

(2) Department of Economics
European University Institute

Florence , Italy.

(3) The College of Wooster
Department of Economics

Wooster, OH 44691, USA

email:weisbuch@physique.ens.fr

November 9, 1995

1 Introduction

Markets can have very different structures. They may be based on an auction mechanism, on a system of sellers who post prices or on bilateral trading and bargaining. Within each of these structures a considerable amount of organisation may develop. In particular certain trading relationships will be established. Some individuals in a market will systematically deal with certain others and this gives rise to a graph of trading or communication links. In most economic models this feature is ignored and, for example, in the standard Walrasian model, an equilibrium is considered to be attained if aggregate excess supply is zero at certain prices. How the exchanges necessary to equilibrate the market actually take place, that is who trades with whom, is not analysed. In other models such as the standard "search models" (see, for example Diamond (1989) agents visit sellers and buy from the cheapest according to some rule. While this is a step towards realism the idea that buyers will become attached to certain stores is, in general, not taken into consideration. All sellers are anonymous and are searched with equal

probability. While such models may be plausible for transactions which take place infrequently they do not seem adequate for transactions that are made often and on a repeated basis.

A networks of preferential links between sellers and buyers is especially appropriate to markets of perishable goods as opposed to durable goods. When goods are perishable, sellers have to calculate with care the amount they wish to bring to the market since no surplus can be kept as inventory. For this they try to obtain an accurate estimate of how many customers they expect to attract. Buyers are aware also that sellers are trying to minimize excess supply and will therefore prefer stores at which their demand will be satisfied. From their experience they may learn that these will not necessarily be those which offer the lowest prices.

Thus what interests us is how agents may learn to develop certain links with each other rather than simply match at random.

There are several models in which there is a network of links between agents but where the trading structure represented by these links is fixed and given. A typical example is a model in which economic agents "live" on a lattice and only interact with their neighbours. This can be thought of as representing some sort of spatial organisation (see Follmer [1974], Durlauf [1990], Benabou [1992], Blume [1993] and Ellison [1993]). In this case one is interested to know whether pockets or clusters with certain behaviour or characteristics may form. The spatial connotation is by no means necessary however and alternative structures of links can be considered (see Kirman, Oddou and Weber [1986], Ioannides [1990], and Gilles et al.(1994).) Indeed in many problems in economics, links are created as a result of community of interest or relationship between characteristics rather than mere geographical proximity. It may also be the case that although agents react only to their neighbours the consequence of their choices may have more widespread effects (see Weisbuch et al. (1994)). Yet, in all of these contributions the communication or trading links are taken as given exogenously. Our aim, in this paper, is to try to bridge the gap between the two extremes, on the one hand random matching and on the other, fixed exogenous links between individuals.

We wish to examine how links between agents develop and are reinforced. A market which has strong links which persist over time can be thought of as one which has organised itself whereas one in which agents simply search at random might be thought of one which is lacking in any organisation. A small number of authors have examined the problem of this sort of self organisation in markets although from a rather different point of view. Durlauf [1990] introduces something of this sort when he considers not the network itself as changing but rather that agents may choose when to place themselves in the network, and this recalls an older model of neighbourhood preferences due to Schelling. Another example is that of Stanley et al. (1994). They develop an evolutionary model of the repeated prisoners dilemma in which, as agents learn from experience, they may refuse to play with certain others and one can examine the distribution and local concentration of communication links and of strategies that develop. Lesourne (1992) discusses at length the self-organisation of markets and looks at the emergence of privileged relations between certain buyers and sellers. A lot of his emphasis is, however, on self organisation through prices something which will not feature directly in our model.

More directly related to the present paper is Vriend's [1994] contribution which presents a first step to simulating a model in which either the links themselves or the probability that they will be used over time evolve. He constructs a model of a market in which buyers learn when to shop and firms learn, from experience, when to buy. The term 'learn' has a very particular meaning here. What agents do is to select between rules on the basis of the profitability, in the past, of having

used those rules. Thus learning in Vriend's context does not involve updating within a parametric model as it will in this paper. In his model firms sell indivisible units of a homogeneous good, the price of this good is fixed and agents demand at most one unit. Nevertheless it is particularly interesting to note the development and persistence of a non degenerate size distribution of firms even though all firms are identical to start with. Furthermore some buyers always return to the same store whilst others continue to search. Thus both loyalty and searching behaviour coexist. The approach adopted by Vriend and pursued in Kirman and Vriend (1995) is to look at a market in which agents have no knowledge of the market structure and they can only learn by observing the payoffs from the actions they take. In such a model which is in the spirit of the "artificial life" approach agents know and calculate nothing a priori and one observes to see what sort of organisation emerges. The advantage of this approach is that very little structure is imposed, a priori. The difficulty is that one has to rely essentially on simulations since such a model is not analytically tractable and furthermore it is difficult to attribute emergent features to any one particular feature of the model.

Here we adopt a rather different approach. We analyse several simple models of a market for a perishable good, in which buyers, (retailers), meet sellers, (wholesalers) once or several times in a day and buy quantities of the homogeneous good to resell on their own local market. However, unlike the models of Vriend and Kirman and Vriend, we build quite a lot of calculating ability for the agents into the model in that prices are chosen as a result of specific calculations about the amount buyers, behaving in an intelligent way will be prepared to purchase. Learning here involves updating by agents within a parametric model and we are able to examine analytically the consequences of changes in the values of the parameters of the model. We concentrate our analysis on the problem of whether links will be established between particular buyers and sellers and what is the nature of these links are they permanent, "marriage", do buyers occasionally have "affairs" with other sellers while remaining basically "faithful" to their regular partners, do buyers "divorce" their sellers and then remarry or do they continually "shop around"? The important parameter in this is the sensitivity of buyers to past profits from sellers when making their choice as to whom to visit in the next period.

We find that the change from order, settled marriage patterns to disorder is surprisingly abrupt as a function of the learning parameter for the simplest one session model. We show analytically why this is so using the "mean field" approach (for other applications to economics see, for example, Aoki(1995)), and then do simulations of more complex models. These show that the same patterns of behavior persist. We finally compare our theoretical predictions with empirical data from the wholesale fishmarket in Marseille.

Our model is a springboard for future research in which many other features can be the subject of learning but by starting in this way we are able to obtain simple and clear characterisations of the basic relationships between aggregate structure and the parameters of the individual learning process.

2 The simplest Model

Let us consider a set of n buyers i and a set of m sellers j .

2.1 Prices and quantities

The buyers will wish to buy a quantity from sellers given the price per unit p which is proposed and are able to resell in their own local market where they face a demand function $p(q)$, which determines the relationship between the price they obtain and the quantity q that they bring to the local market. Let us suppose in order to simplify matters that $p(q)$ is known by the buyers, is the same for all buyers and that it is a simple function of q such as:

$$p(q) = \frac{b}{q + c}, \quad (1)$$

where b/c is the maximum price at which a non-negative quantity can be sold and c is a "characteristic quantity" which is sold at half the maximum price. The buyer's profit is then:

$$\pi_b = q \left(\frac{b}{q + c} - p \right), \quad (2)$$

We suppose then that the buyer knows the demand curve he faces and is thus able to compute the quantity that will maximise his profit for a given price p proposed by the seller. This quantity is:

$$q = \sqrt{\frac{bc}{p}} - c, \quad (3)$$

We make similar assumptions for the sellers, in particular that they know the behavior of buyers described by the three equations above, and they can therefore maximize their own profit per transaction:

$$\pi_s = q(p - p_a) = \left(\sqrt{\frac{bc}{p}} - c \right) (p - p_a), \quad (4)$$

with respect to the price p that they charge to the buyers, where p_a is the price at which the sellers themselves purchase the fish. (Since the price p which maximises equation 4 is the solution of a third degree equation, its expression is rather complicated and not given here).

In order to simplify assumptions as much as possible, let us suppose that:

- All sellers propose a single price p to all their customers.
- Customers choose one shop everyday. As long as the shop has supplies, they are offered goods at price p ; customers then purchase q given by equation 3. Whether they do get supplies when they visit the shop depends upon the time they visit the shop, which is randomly chosen in the simulations. It also depends on how much the sellers bring each day to the market.
- Each day sellers bring $n_b q$ to the market, where n_b is the expected number of customers for that day. In the simplest version of the model, this is simply how many customers visited yesterday.

2.2 Preferences and learning

Now that quantities and prices are determined, what happens in the market depends on the actual visits of buyers to sellers: sellers are chosen by buyers according to preferences that are built upon the history of each buyer's past profits with each seller. More precisely, the fidelity J_{ij} of buyer i to seller j determines the probability P that buyer i visits seller j as follows:

$$P = \frac{\exp(\beta J_{ij})}{\sum_j \exp(\beta J_{ij})}, \quad (5)$$

where β , a term taken from statistical physics, measures the non-linearity of the relationship between probability and fidelity. From the results of statistical physics, one might infer that large values of β result in "market organization" with strong and stable preferences of buyers, while small values of β result in a "disordered market" where buyers select sellers at random.

Preferences for sellers evolve according to a variant of Hebb's rule used in cognitive science. J_{ij} are increased according to profits π_b realized by buyer i following transactions with seller j :

$$J_{ij} = J_{ij} + \pi_b, \quad (6)$$

Furthermore at each time step all J_{ij} are decreased by a constant factor $1 - \gamma$:

$$J_{ij} = (1 - \gamma)J_{ij}. \quad (7)$$

This decrease, which can be interpreted as forgetting, results in limiting the maximum amplitude of J_{ij} to π/γ where π is a maximum profit per time step. (in the above expressions the sign = means "assign the right hand-side of the expression to variable J_{ij} ", as in programming).

3 Results

3.1 Indicators of order

Simulations generate a large number of data about individual transactions such as which shop was visited, purchased quantities, and agents' profits. The organization process itself is harder to monitor. We used two methods to do this.

Firstly, adapting a measure used in statistical physics (Derrida [1986]), we defined an order parameter y by

$$y_i = \frac{\sum_j J_{ij}^2}{(\sum_j J_{ij})^2}, \quad (8)$$

In the organized regime, when the customer is faithful to only one shop, y_i gets close to 1 (all J_{ij} except one being close to zero). On the other hand, when a buyer visits n shops with equal probability, y_i is of order $1/n$. More generally, y_i can be interpreted as the inverse number of shops visited. We usually monitor y , the average of y_i over all buyers.

Secondly, when the number of shops is small, 2 or 3, a simplex plot can be used to monitor on line the fidelity of every single buyer. Figures 1a and 2a, for instance, display simplex plots at different steps of a simulation. Each agent is represented by a small circle whose colour or shade is specific to the agent. The circle's position is the barycenter of the triangle for a choice

of weights proportional to the fidelity of the agent to the 3 shops each of which corresponds to one of the 3 apexes of the triangle. Proximity to one corner is an indication of fidelity to the shop corresponding to that corner. Agents represented by circles which are close to the center are undecided. This representation is commonly used in chemical physics for instance to describe phase diagrams of ternary compounds.

3.2 The simplest model

The simplest model, with fixed price, was run with 3 sellers and 30 buyers, for a large variety of parameter configurations and initial conditions. The following parameter configuration was chosen to give simple time charts for shop performance and to highlight differences in buyers behavior. Price parameters b , c , and pa where respectively 1, 1, and 0.3, which corresponds to a price of $p = 0.579$ for a purchased quantity per transaction $q = 0.314$ (according to equation 3), a profit per transaction for the buyer $\pi_b = 0.0572$ (equ. 2) and a profit per transaction for the seller $\pi_s = 0.0877$ (equ.4). Our choice of the memory constant of equation (7) is $\gamma = 0.1$. Initial J_{ij} were all 0.

Depending on the value of the non-linear parameter β , two distinct behaviors are observed.

3.2.1 Disorganized behavior

For low values of the non-linear parameter β , e.g. $\beta < 5.2$, buyers never build-up any fidelity. This is observed in figure 1, which describes the dynamics obtained with $\beta = 0.8$. The daily profit of buyers averaged over all buyers and over 100 days after a transition period of 100 days, is 0.0502. This result is lower than the average profit per transaction for the buyer $\pi_b = 0.0572$. This is due to all those occasions when a buyer visited an empty shop. The daily profit of sellers averaged over all sellers and over 100 days after a transition period of 100 days, is 0.6532. This result is lower than 10 times the average profit per transaction for the seller $\pi_s = 0.877$ (the factor 10 corresponds to the average number of buyers per shop). This difference was also generated indirectly by buyers who visited empty shops: meanwhile some shops with supplies were not visited, and this resulted in losses for their owner.

The order parameter, y fluctuates well below 0.50 and thus corresponds to randomly distributed J_{ij} . This feature is also clear from the simplex plots of the J_{ij} . Figure 1 shows that the performance of shop number 1 exhibits large fluctuations. The same is true for the two other shops.

3.2.2 Organized behavior

In sharp contrast, the same analysis performed with $\beta = 10$ shows a great deal of organisation (see Figure 2).

The order parameter, y , steadily increases to 1 in 200 time steps. As seen on the simplex plot at time 50, each customer has built-up fidelity to one shop. Performance of shop number one also stabilizes in time, and variations from stationarity are not observed after 20 time steps.

The daily profit of buyers averaged over all buyers and over 100 days after a transition period of 100 days, is 0.0572, exactly the average profit per transaction for the buyer. Because buyers have not changed shops during the last 100 days, sellers learned to purchase the exact

exact quantity needed to satisfy all their buyers, and they had no losses themselves: their daily average profit of sellers is 0.877.

By avoiding daily fluctuations in the number of customers visiting a shop, the ordered regime is beneficial to both customers and sellers, that is both obtain higher profits than in the disorganised situation.

4 Mean Field Approach

The simple model and the results discussed in the previous sections can be formally analysed within the framework of the Mean Field approach. This approach is commonly used in the statistical theory of magnetism to understand qualitatively and predict the statistical properties of a magnet as a function of temperature and external field. It consists in replacing randomly fluctuating quantities by their average, thus neglecting fluctuations. It is only an approximation, which is often convenient to obtain at least a qualitative understanding of the behavior of the system.

4.1 The order/disorder transition

A generalization of this approach to our problem is to write a differential equation describing the time evolution of one J_{ij} :

$$\frac{dJ_j}{dt} = -\gamma J_j + \langle \pi \rangle \quad (9)$$

where $\langle \pi \rangle$ is the average profit, related to π the profit obtained from one actual transaction as follows:

$$\langle \pi \rangle = \pi P(j) \frac{\exp(\beta J_j)}{\sum_j \exp(\beta J_j)}, \quad (10)$$

where the fraction represent the probability of that buyer i visits seller j and $P(j)$ is the probability that the shop still has goods to sell when he comes. We suppress here the i index corresponding to the buyer. In other words, the above set of equations couples the evolution of all the J_j . Equilibrium values are obtained by equating the derivatives to zero.

Let us consider the simplest case of two shops and to further simplify computation, let us suppose that $P(j) = 1$, which happens when sellers bring to the market a lot of extra goods to satisfy unexpected customers. The equilibrium relations are:

$$\gamma J_1 = \pi \frac{\exp(\beta J_1)}{\exp(\beta J_1) + \exp(\beta J_2)}, \quad (11)$$

$$\gamma J_2 = \pi \frac{\exp(\beta J_2)}{\exp(\beta J_1) + \exp(\beta J_2)}, \quad (12)$$

Subtracting equation 12 from equation 11, we see that the difference between the two fidelities, $\Delta = J_1 - J_2$, obeys the following implicit equation:

$$\gamma \Delta = \pi \frac{\exp(\beta \Delta) - 1}{\exp(\beta \Delta) + 1}. \quad (13)$$

The right hand side of the equation is in fact the hyperbolic tangent of $\beta\Delta/2$. The above equation has either one or three solutions according to the slope of the hyperbolic tangent at the origin (see Figure 3, where the graph of the two sides of equation 13 are shown).

By developing the hyperbolic tangent in series for small values of $\beta\Delta/2$, it is easily seen that for:

$$\beta < \beta_c = \frac{2\gamma}{\pi} \quad (14)$$

there is only one solution $\Delta = 0$ and $J_1 = J_2 = \frac{\pi}{2\gamma}$. Since in this case the average J_j are small and equal, the probabilities of visiting either shop simply fluctuate. No order is observed.

In the opposite situation, when β is above β_c , the zero solution is unstable and one obtains two symmetrical solutions where one fidelity is larger than the other one by a factor which is exponential in $\frac{\beta\pi}{\gamma}$. The transition between the two regimes is abrupt as can be observed in figure 4. A development in series of the hyperbolic tangent around 0 shows that the larger fidelity increases in β as the square root of the distance to the transition:

$$\Delta = \sqrt{\frac{24(\beta - \beta_c)}{\beta^3}} \quad (15)$$

Fidelities are then continuous across the transition, but they rise (or decrease) with an infinite slope at the transition.

Expression (14) can be generalized to any number n of shops:

$$\beta_c = \frac{n\gamma}{\pi} \quad (16)$$

The above expression is obtained by investigating the stability of equation 9 linearized in the neighborhood of the symmetric fixed point corresponding to all $J_j = \frac{\pi}{n\gamma}$. The eigenvalues of the characteristic equations are obtained by the diagonalisation of an $n \times n$ matrix which necessitate some cumbersome but standard linear algebra.

Interpreting β and equation 16 is now straightforward. $\frac{\pi}{\gamma}$ is a cumulated profit with an effective discount rate of γ . It is also the maximum fidelity J_j obtained when shop j is visited with probability 1 as seen from equation 9. Equations 14 and 16 show that β has the dimension of an inverse profit. β also acts as a discrimination rate as seen on figure 3: large values of β correspond to shifting from shop 2 to shop 1 for a small change in differences of fidelity around 0 since the hyperbolic tangent is the difference of the probabilities of visiting either shop. One can then consider $\frac{1}{\beta}$ as the width of a probability distribution of visiting either shop as a function of Δ : When this width is smaller than the integrated profit divided by the number of shops, order is achieved.

The above analysis shows that as long as the mean field approximation remains valid, the qualitative behavior of the dynamics, ordered or disordered, only depends on one parameter, namely the ratio between β and β_c . All other parameters simply change the scale of profits, prices, numbers of shops and customers. The time scale of learning depends on γ : order, when achieved, is reached faster for larger values of γ . But when profits are increased in proportion, (since β_c is the parameter controlling the transition, the influence of γ can only be checked independently by changing γ and π in proportion to keep the ratio β over β_c constant), new

phenomena to be described in the next section appear; they are based on fluctuations in the number of buyers and cannot be dealt with using the mean field approximation.

The three parameters which control the transition depend upon sellers (especially π , buyers profit which is determined by the price which sellers set) and buyers (especially β and γ). We might assume that all agents are not identical and that their characteristic parameters might vary. Prices might not vary widely because of competition among sellers. On the other hand, memory (characterised by γ) and discrimination rate (characterised by β) might differ between buyers. If these variations are large enough, we might expect to observe two distinct classes of buyers: faithful buyers, who most of the time visit the same shop, would be those whose parameters are such that $\beta > \beta_c$, while searchers with parameters such that $\beta < \beta_c$ would wander from shop to shop. Indeed precisely this sort of "division of labour" is observed on the Marseille fish market which was the empirical starting point for this paper.

4.2 Hysteresis

Another important qualitative result of the mean field approach is the existence of hysteresis effects: buyers might still have a strong preference for one shop that offered good deals in the past, even though the current deals they offer are less interesting than those now offered by other shops.

Let us come back once more to the case of two shops 1 and 2, now offering different profits π_1 and π_2 . Replacing profit π in equations 11 and 12 by respectively π_1 and π_2 , equation 13 becomes:

$$2\gamma\Delta - (\pi_1 - \pi_2) = (\pi_1 + \pi_2) \frac{\exp(\beta\Delta) - 1}{\exp(\beta\Delta) + 1}. \quad (17)$$

Coming back to the graphical representation of figure 3, we see that a difference in profits results in shifting the straight line representing the left hand side of equation 13 with respect to the hyperbolic tangent. The three intersections remain only as long as the difference in profits is not too large. Which of the two extreme intersections is actually reached by the learning dynamics depends on initial conditions.

Thus, as illustrated on figure 5, buyers can remain faithful to a shop asking for a higher price (which results in a lower profit for the buyer), provided that they became attached to this shop when it practiced a lower price. When the most often frequented shop changes its prices, the fidelity to that shop describes the upper branch of the fidelity versus profit curve (figure 5). The fidelity remains on the upper branch as long as it exists, i.e. until the point where the slope is vertical. When profits carries on decreasing a sudden and discontinuous transition to the lower branch occurs. This is the condition when customers change their policy and visit the other shop. But, if the first shop reverses its high price/low buyer profit policy when fidelity is on the lower branch, the transition to the higher branch only occurs when the slope of the lower branch becomes vertical, i.e at a higher profit than for the downward transition.

A consequence of this phenomenon, is that in order to gain customers who are faithful to another shop, a challenger has to offer a profit significantly greater than the profit offered by the well established shop: when connections have reached equilibrium in the ordered regime, customers switch only for differences in profits corresponding to the vertical slopes of the curves $J(\pi)$ in figure 5 (i.e. not when profits are equalised!!). In other words, economic rationality (i.e. choosing the shop offering the best deal) is not ensured in the region where hysteresis occurs.

5 More complicated models and results

We will discuss, in this section, further refinements of the simple model and see what influence they have on the behaviour of the agents.

5.1 Beyond the mean field approximation

The results of the mean field approach were obtained from a differential equation modeling a discrete time algorithm. They are valid when the changes at each step of the algorithm can be considered as small. Variables γ and π thus have to be small, which is true for the simulation results given in figures 1 and 2. One of the features noticed by observing on-line the motion of individual buyers on the simplex plots is that agents sometimes move "backward" towards shops which are not the shops that they prefer in the ordered regime. But since for most of the time they move towards preferred shops, these "infidelities" never make them change shops and preferences permanently. They commit "adultery", but do not "divorce".

When variables γ and π are increased, infidelities have more important consequences, and customer might change fidelity: they may "divorce" one shop for another one. Indeed increasing π results in larger steps taken by customers on the simplex, which might make them go from one corner neighborhood to another one in a few time steps. In fact the probability of a given path on the simplex varies as the product of probabilities of individual time steps: when fewer steps are needed the probability that the process will generate such changes becomes higher. Because of the exponential growth of time of the "divorce" process with respect to π , a small change in relevant parameters, π or γ results in a switch from a no-divorce regime to a divorce regime. Divorces are observable on-line on the simplex plots and also by examining the evolution of the number of customers as a function of time. In figure 6, "infidelities" appears as peaks and "divorces" as steps. The bias towards an increase in the number of customers is due to the supply available at the beginning of each day: shops with a large number of customers have a large supply and a new customer has more chance of being served there. Hence the inherent bias in the number of customers per shop: large shops tend to increase in size, small shops tend to decrease.

5.2 Morning and afternoon

The one-session model described in section 2 is a considerable simplification of the way buyers search for sellers. As is commonly observed in several markets with the sort of structure we are modelling here, customers that refuse a deal with one seller, usually shop around to find other offers. Indeed this is generally regarded as the principle motivation for refusal. An alternative explanation is that customers refuse now in order to induce better offers in the future. In either case we should then consider a model in which customers are given at least two occasions to purchase goods.

One further assumption to relax in the case of perishable goods is the idea of a constant price for all sessions. In fact p is the price sellers would charge at each transaction if they were sure to sell all the quantity they bring to the market. If they were able to exactly predict how many customers would visit their shop and accept their price, they would bring to the market the exact quantity necessary to satisfy the buyers at that price. But, when their predictions are