Market Organisation and Trading Relationships

GÉRARD WEISBUCH [(1)]

ALAN KIRMAN [(2)]

and DOROTHEA HERREINER [(3)]

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

(2) GREQAM
EHESS et Université d'Aix-Marseille III
13002 Marseille, France

(3) Economic Theory III
University of Bonn, Adenauerallee 24 - 26
53113 Bonn, Germany

email:weisbuch@physique.ens.fr

November 18, 1997

Keywords: market organisation, reinforcement learning, customer loyalty, perishable goods.

1 Introduction

Many markets are characterised by trading relationships. Individuals systematically trade with particular partners in certain markets whilst in others no such stable links are observed. Some other markets exhibit a mixture of stable links and "searching" behaviour. Yet the way in which such organisation develops and its economic consequences are not considered in standard theoretical models. The notion of equilibrium in a market, in the Walrasian model for example, is simply characterised as a situation in which aggregate excess demand for the good in question is zero. This definition of equilibrium leaves two kinds of questions unanswered:

- How do agents get the information about who demands or supplies which good at what price? Who determines those prices?
- How is that information used to determine who will make which transaction with whom, thereby clearing the market at each stage and determining market organisation in the long run?

In standard theory, the process which adjusts prices is either left unspecified (Adam Smith's "invisible hand") or some fictitious player, such as Walras' "auctioneer", adjusts prices as a function of aggregate excess demand through some process such as a "tatonnement". The main idea behind this simplification is that prices are adjusted and supply and demand respond to these adjustments, during a preliminary phase, until equilibrium prices are reached. Transactions that clear the market are then performed in some unexplained way.

As is well known, even under strong assumptions the sort of price adjustment process envisaged by Walras is not, in general stable, and hence equilibrium would not necessarily be achieved (See e.g. Mas-Colell et al (1995)). Furthermore, even if equilibrium prices were to be announced, many markets are not organised centrally and the way in which those who wish to buy, at those prices, are matched with those who wish to sell has to be specified; this is one of the objectives of this paper.

A number of models have been developed to provide at least partial answers to these question. Such models examine situations in which sellers set prices individually and in which buyers choose which seller to buy from. The best known of these are "search models", (see, for example Diamond (1989)), which are usually for a market with a single good. More complete models with individuals setting prices and buyers searching have been developed for example by Fisher (1973) and Lesourne (1992). In standard search models, buyers sample sellers according to some rule and buy from the cheapest. All sellers are anonymous and are searched with equal probability. There is no memory of where favourable opportunities were found in the past. Such models seem to be plausible for transactions which take place infrequently, when sellers may have some knowledge of the distribution of prices but cannot be sure as to the prices charged by particular individual sellers. This is the case, for example, when an individual makes an infrequent purchase such as buying a car, is seeking a job, or when a firm invests in a large capital item.

Yet many markets are ones on which individuals trade frequently with each other. Of particular interest is the case of markets for perishable goods. Since sellers cannot hold inventories, they only supply the quantities they expect to sell during one session. A buyer who takes a considerable amount of time to search for the best price runs the risk of not finding anything to buy by the end of the session. Rather than gathering a lot of new information at each session, the best strategy for him is to use the experience gained from transactions made with different suppliers during previous sessions. We shall show that trading relationships develop because buyers learn about the value of trading with particular partners. Stable trading relationships are also profitable to sellers who can then predict with some accuracy the demand they will face in each session and determine their supply accordingly. The more loyal the customers, the better the prediction and the more likely the customer is to find the goods he is seeking. Thus the establishment of regular trading relationships may be mutually profitable. The basic aim of this paper is to suggest and test a simple search mechanism that would result in the establishment of stable trading relationships and to characterize the conditions under which this happens.

The standard game theoretic approach to the problem of trading relationships is to develop a game theoretic equilibrium notion for the network of trading links in the sense that no individual has any interest in adding or removing any of the links in which he is involved. This is the approach adopted by Jackson and Wolinsky (1996). Whilst such models provide a benchmark with which various trading structures can be compared, they do not explain how such structures might develop and, in addition, they assume that agents are perfectly capable of working out the consequences of changing links and of the reaction of other participants to such changes.

By contrast, our model falls into the class of adaptive economic models. In such models, agents are not endowed with perfect rationality, but behave according to some procedural rationality, using information obtained from other agents or from their own experience. Modeling economic agents as adaptive rather than perfectly rational makes sense in particular when they have incomplete information, which is the case for buyers in the Marseille wholesale fish market, where prices are not posted and may vary according to seller, time of the day and from day to day.

A typical example of the sort of procedural rationality that we have in mind is that of modifying one's behaviour by attributing greater weight to the use of rules that have proved to be profitable in the past. This is the approach developed by Arthur et al (1996) for example. Another example is the idea that one may, in the light of observation or experience, wish to imitate the behaviour of others. Such imitation may be motivated by the success of other agents or by inference about the information they possess and may be based on more or less sophisticated reasoning. A number of authors have adopted this approach to "social learning", in particular those who use discrete choice theory, (see e.g. Aoki (1996), Brock and Durlauf (1995), Durlauf (1990), Kirman (1993), Lesourne (1992). See Anderson et al (1992) for a recent review of the discrete choice theory literature).

In this paper, however, we shall focus on situations in which individuals have to rely on their own experience and do not observe that of others directly. We shall be interested here, in particular, in markets in which transactions are not made public, that is, there is no central market clearing mechanism and no prices are posted. In such markets agents have to rely on their own information. This is the case for many markets such as the Marseille wholesale fish market from which our empirical evidence is drawn. We will therefore develop a model which seeks to explain some of the phenomena that characterise this type of market and which will be based on learning from past experience.

We will adopt an approach which allows us to obtain analytical results for the simplest version of our model and we then use simulations to check that these results still hold in more complicated and realistic versions.

The structure of the paper is as follows. We start by proposing a very simple model of a market for a perishable good, in which at each time step buyers (retailers) meet sellers (wholesalers) and buy quantities of the homogeneous good to resell on their own local market. They do this in a shop which is chosen according to the information gathered during previous purchases. This model is analytically solved using the "mean field" approximation. The theory predicts that two distinct types of behaviour for the agents should be observed according to their learning and choice parameters: some agents should remain loyal to one selected shop, while others should keep on shopping around for ever. We then use multi-agent simulations to study more complex, and more realistic versions of the model, allowing for instance several purchases per buyer during the same day, varying prices, and more complicated adaptive behaviour of buyers and sellers. Our simulations show that the same patterns of dynamic behaviour persist. We finally verify that our theoretical predictions are consistent with the empirical data from the wholesale fish market in Marseille, while other theories are not.

2 The Simplest Model

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

2.1 Basic Assumptions

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

- Customers choose one shop every day according to their memory of previous transactions. As long as the shop has supply, a customer purchases a quantity $q_i(t)$ implying a profit $\pi_i(t)$. Whether the customer is served when he visits the shop depends on which shop j is visited at time t, how many people bought from that shop before, and how much endowment the shop had at the beginning of the day.

- Since the good is perishable and therefore cannot be stored between days, each day a seller supplies a quantity $Q_j(t)$ which he expects to sell on that day. In the simplest version of the model, this quantity is simply the quantity he sold yesterday.

- Every day the same market scenario is repeated.

These simplistic assumptions will be used in sections 2, 3 and 4. More realistic assumptions will be made in section 5.

2.2 Preference coefficients, Learning and Choice Probabilities

Our model seeks to explain trading relationships. Therefore, our assumptions about how buyers choose which shop to visit are crucial. These assumptions[1] are kept constant throughout the entire paper, i.e. they are the same for the basic model and its extensions.

A buyer has to choose one shop each day. The basic assumption of the present model is that his present choice of which shop to visit is based on his previous experience. His decision rule is therefore a mapping from the time series of the transactions he has had with different sellers and the profits associated with the transactions[2], $\mathbf{I}(t)$, to the unit simplex Δ_m, where m is the number of sellers:

$$P(t) : \mathbf{I}(t) \longrightarrow \Delta_m \tag{1}$$

A point in the simplex Δ_m represents the probabilities with which an individual chooses each of the m sellers.

The mapping $P(t)$ can be decomposed into two components, a mapping from $\mathbf{I}(t)$ to a vector $\mathbf{J}(t)$ of "preference coefficients" and a second mapping from $\mathbf{J}(t)$ to Δ_m. The first mapping is an encoding based on a learning process and the second mapping describes the probabilistic choice process.

Let us first specify the learning process of the buyers. By assumption the only information available comes from past transactions, so each buyer has a record for each shop. The profits buyer i made when buying from shop j are mapped into the preference coefficient J_{ij} by adding profits every period and discounting previous profits at a constant rate γ. Since we use discrete time for transactions, preferences are updated at each time step according to:

$$J_{ij}(t) = (1-\gamma) \cdot J_{ij}(t-1) + \pi_{ij}(t), \qquad \forall i, \forall j, \tag{2}$$

In other words, at each time step, all preference coefficients are discounted at a constant rate, and the preference coefficient for the shop with which a transaction occurs is increased by the profit made in that shop. Preference coefficients thus appear as the sum of discounted past profits. Discounting can be interpreted in different ways: it describes gradual forgetting of past events; it also serves to ensure that information is relevant to the current situation. In real life[3] shops do not necessarily have stationary characteristics in terms of the profits that they offer, because of possible changes in prices for many possible reasons or in the initial endowment relative to the number of customers they have to serve.

[1] The learning and probabilistic choice process described in this section was inspired by the formal neural networks approach to reinforcement learning as described for instance in Weisbuch (1990).

[2] Other models might give more information to the buyer: for instance, he might be aware of the transactions made by others; alternatively, individuals might have less information if, for example, their memory were limited.

[3] In our model we will use a stationarity hypothesis to facilitate calculations, see further in the next sections.

The decision rule used by the buyers then maps these preference coefficients into the choice of a shop. One deterministic way to do so is to choose the shop with the best record, that is the shop with the highest $J_{ij}(t)$. This would amount to mapping the $\mathbf{J}(t)$ into one of the apices of Δ_m. However, by doing this, the buyer would become a captive of the selected shop which would then be in a position to diminish the buyer's profit and to increase its own profit by changing prices. The shop could do this until the buyer's profit becomes negative before running any risk of losing that buyer. It is therefore in the buyer's interest to search from time to time among other sellers to check whether he could get a better profit elsewhere. In other words, a good strategy for the buyers would be a balance between the deterministic choice in favor of those shops which gave the best profits in the past and random search among other sellers. This raises the well known issue of the trade-off between exploitation of old knowledge and exploration to acquire new knowledge.

We use a probabilistic choice rule here, which characterizes this trade-off with a single parameter β. We suppose that the decision rule by which a buyer i assigns a probability P_{ij} of visiting seller j is proportional to the exponential of the preference coefficient for that seller. That is:

$$P_{ij} = \frac{\exp(\beta J_{ij})}{\sum_{j'} \exp(\beta J_{ij'})}, \qquad \forall i, \forall j, \tag{3}$$

where β, the discrimination rate, measures the non-linearity of the relationship between the probability P_{ij} and the preference coefficient J_{ij}. This specification[4] allows for any choice rule in the range of equal probabilities ($\beta = 0$) to best-reply ($\beta = \infty$).

In our case, the exponential rule can be derived directly (see Brock (1993) for a discussion). This is done by maximizing the weighted sum F_i of two terms; one of which favors immediate profit:

$$G_i = \sum_j P_{ij} J_{ij}.$$

G_i is approximately the expected discounted sum of profits. The other term favors search. To maximise the information gained during visits, buyers should maximize the Shannon entropy[5] of the distribution of search probabilities:

$$S_i = -\sum_j P_{ij} log P_{ij},$$

The function F_i to be maximized is then a linear combination of preferences and entropy terms:

$$F_i = \beta G_i + S_i. \tag{4}$$

The smaller β the stronger the weight given to disorder, i.e. to information gathering at different shops. The larger β the more important (short-run) payoff concerns.

[4]The exponential rule has been widely used in economics and elsewhere. Several justifications for its use are given in the discrete choice literature, see e.g. Anderson et al. (1992).

[5]Entropy is a measure of the disorder of a system; it is maximal (for each i) if all $P_{ij} = 1/m$, "the most random probability measure" as Brock (1993) calls it. Entropy is minimized if $P_{ij} = 1$ for one j and the other $P_{ij} = 0$.

Setting the derivatives of F_i with respect to P_{ij} equal to zero under the constraint that the sum of the probabilities is 1 gives equation (3).

3 Mean Field Approach

The simple model can be formally analysed within the framework of the Mean Field approach. This consists in replacing randomly fluctuating quantities by their average, thus neglecting fluctuations. It is only an approximation, but is often convenient to obtain at least a qualitative understanding of the behavior of the system.

The model is soluble in the continuous limit, when the changes of variables are small at each time step, i.e. $\gamma \to 0$. Equation (2) can be expressed as a difference equation in τ by multiplying γ and $\pi_{ij}(t)$ by τ and then rewriting it as:

$$\frac{J_{ij}(t+\tau) - J_{ij}(t)}{\tau} = -\gamma J_{ij}(t) + \pi_{ij}(t). \tag{5}$$

Taking the limit for $\tau \to 0$, leads to a stochastic differential equation

$$\frac{dJ_{ij}}{dt} = -\gamma J_{ij} + \pi_{ij} \tag{6}$$

in π_{ij}. The Mean Field approximation[6] consists in replacing the π_{ij} by its expected value $< \pi_{ij} >$, thereby transforming the stochastic differential equation into a deterministic differential equation.

3.1 The order/disorder transition

The time evolution of J_{ij} is thus approximated by the following equations:

$$\frac{dJ_{ij}}{dt} = -\gamma J_{ij} + < \pi_{ij} > \tag{7}$$

$$< \pi_{ij} > = Prob(q_i > 0) \cdot \pi_{ij} \frac{\exp(\beta J_{ij})}{\sum_{j'} \exp(\beta J_{ij'})}; \tag{8}$$

the fraction represents the probability that buyer i visits shops j; $Prob(q_i > 0)$ is the probability that the shop still has goods to sell when the buyer comes to shop j, in which case he gets a quantity q_i resulting in profit π_{ij}. Suppose the market converges to a stationary state in which buyers' preference coefficients do not change. Such a state is called an equilibrium in dynamical systems theory and it is obtained by setting the derivatives (equation (7)) equal to zero.

Let us consider the simplest case of two shops and to further simplify computation, let us suppose, for the moment being, that $Prob(q_i > 0) = 1$, which happens when

[6] For a discussion of this step see for instance Brout (1965).

buyers always find what they require at the shop they visit[7], and that profits in both shops are equal to π (see the next section for unequal profits). The equilibrium relations are in this case:

$$\gamma J_1 = \pi \frac{exp(\beta J_1)}{exp(\beta J_1) + exp(\beta J_2)}, \quad (9)$$

$$\gamma J_2 = \pi \frac{exp(\beta J_2)}{exp(\beta J_1) + exp(\beta J_2)}. \quad (10)$$

We dropped the index i refering to the buyer, and the remaining indices 1 and 2 refer to the sellers. Subtracting equation 10 from equation 9, we see that the difference between the two preference coefficients, $\Delta = J_1 - J_2$, obeys the following implicit equation:

$$\frac{\gamma \Delta}{\pi} = \frac{exp(\beta\Delta) - 1}{exp(\beta\Delta) + 1}. \quad (11)$$

The right hand side of the equation is 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. If [8]

$$\beta < \beta_c = \frac{2\gamma}{\pi} \quad (12)$$

there is only one stable solution $\Delta = 0$ and $J_1 = J_2 = \frac{\pi}{2\gamma}$. The average J_j are small and equal. A buyer visits both shops approximately half the time, switching at random between the shops. We call such a market *disordered* or *disorganized*.

In the opposite situation, if $\beta > \beta_c$, the zero solution is unstable and the other two solutions are stable and symmetric, with one preference coefficient large and the other one small[9]. At the stable solutions a buyer visits one shop with high probability and frequency (high preference coefficient) and the other shop with very low probability and therefore rarely (low preference coefficient). We call such a market *ordered* or *organized*; buyers are loyal.

The transition from the disordered to the ordered market is abrupt; the difference between the preference coefficients Δ stays 0 for $\beta < \beta_c$, it changes with infinite slope at $\beta = \beta_c$, and it increases approximately by the square root of the distance $(\beta - \beta_c)$ (close to $\Delta = 0$):

$$\Delta = \sqrt{\frac{12(\beta - \beta_c)}{\beta^3}} \quad (13)$$

as can be seen in figure 1.

---

[7]If this were always the case, there would be no rationale for the learning and choice algorithm described in section 2. We take here $Prob(q_i > 0) = 1$ only as a limiting case allowing analytical computations which are then checked against numerical simulations in section 4 with $Prob(q_i > 0) \leq 1$.

[8]By developing the hyperbolic tangent in series for small values of $\beta\Delta/2$. See appendix A for more details.

[9]The ratio between the two preference coefficients is exponential in $\beta\pi/\gamma$; see appendix A.

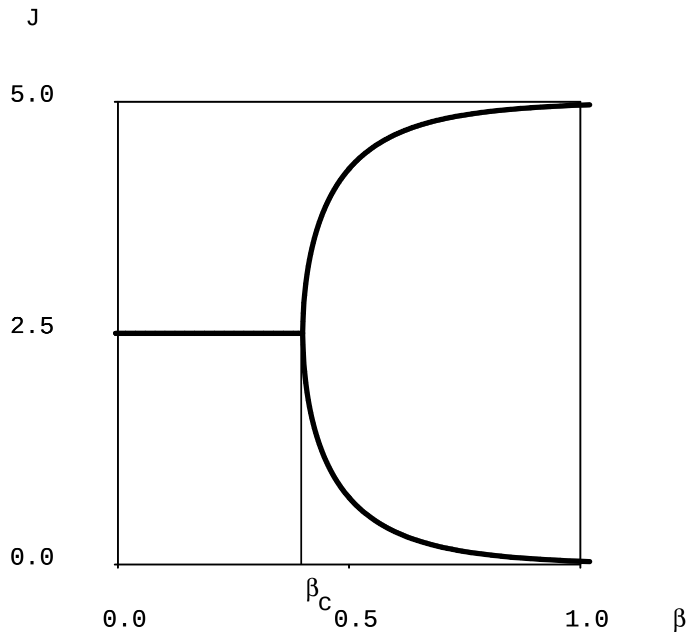

Figure 1: The order/disorder transition in β. Plot of both equilibirum preference coefficients versus the discrimination rate β. Below the transition rate β_c, preference coefficients are equal, but they rise or plummet sharply when the discrimination rate β increases above the transition. When profits in both shops are equal (as in this figure), either loyalty describes the upper branch, while the other describes the lower branch. (The figure is drawn for two shops with $\pi = 1$ and $\gamma = 0.2$ using GRIND software, De Boer 1983).

In the case of m shops, the fixed point equations are:

$$J_j = \frac{\pi}{\gamma} \frac{f(J_j)}{\sum_k f(J_k)} \tag{14}$$

where $f(x)$ is the exponential choice function $exp(\beta x)$. Summing over j the fixed point equations (14) one sees that any solution $\mathbf{J}$ satisfies

$$\sum_j J_j = \frac{\pi}{\gamma}. \tag{15}$$

Obviously, the symmetric fixed point

$$J_j = \frac{\pi}{m\gamma} \quad j = 1, ..., N \tag{16}$$

satisfies equation (14). This fixed point is an attractor *iff* the right hand side of equation (14) has a slope smaller than one. This condition is easily checked since the derivative of the denominator of the RHS of equation (14) is zero at the symmetric point, due to equation (15) and equality of derivatives of f. We thus obtain:

$$\beta_c = \frac{m\gamma}{\pi}. \tag{17}$$

In this case, there is either one stable stationary point (if $\beta < \beta_c$), where the customer visits all shops with equal likelihood, or there are m stable stationary points (if $\beta > \beta_c$), where a buyer is loyal to one of the m shops[10].

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 the ratio between β and β_c. As long as β/β_c is kept constant, changing the original parameters m, β, and π, only changes the scale of equilibrium variables such as actual profits of the buyers or the fraction of unsold endowments. The time scale of learning depends on γ: order, when achieved, is reached faster for larger values of γ.

Within the approximations made in this section, buyer dynamics are uncoupled: each buyer behaves independently of other buyers. As a result, if we now consider a set of buyers with a distribution of π, β and γ parameters, we expect to observe two distinct classes of buyers within the same market: loyal buyers with $\beta > \beta_c$, who visit the same shop most of the time, and searchers with $\beta < \beta_c$, who 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 and which will be discussed in section 6. Furthermore, because of the sharp transition in behavior when β goes across the transition, the distribution of behavior is expected to be bimodal even if the distribution of the characteristics π, β and γ is unimodal.

We can now compare the predictions of our model where agents learn individually from their past experience with those of models where agents imitate each others' behavior through social interactions (Föllmer (1974), Arthur/Lane (1993), Brock/Durlauf (1995), Orléan (1995)). Both types of models exhibit an abrupt phase transition between order for large β values and disorder for small β values. Two main differences exist.

- In the ordered regime, in the case of imitation, all agents make the same choice (at least when interactions among all agents are a priori possible[11]); in our model different agents are loyal to different shops. Imitation and positive social interactions favor uniformity, while decisions based on agents' memory favor diversity.

- In our model heterogeneity of buyer parameters results in having two classes of behavior, searchers and loyal buyers. Order is a property of buyers, not of the

[10] In total there are $2^m - 1$ equilibria for the differential equations associated with equation (17), however, $2^m - 1 - m$ of these equilibria are not stable.

[11] Imitation favors uniformity, but according to whether one uses a mean field approach (all interactions being possible) as in Arthur/Lane (1993), Brock/Durlauf (1995), Orléan (1995), or Markov random fields (interactions restricted to some neighborhood) as in Föllmer (1974), one observes global or local order. All agents make the same choice in the first case. Different choices can be made in the second case, with local patches of agents making the same choice.

market. In imitation models, the market as a whole is organised or disorganised, even in the presence of heterogeneity of agents[12].

3.2 Hysteresis

Up to this point we have considered a situation in which sellers propose the same prices, resulting in equal profits for buyers. However it is of some interest to examine what happens when profits differ. Let us come back once more to the case of two shops 1 and 2, and now suppose that they offer different prices and hence different profits π_1 and π_2. Replacing profit π in equations 9 and 10 by respectively π_1 and π_2, equation 11 becomes:

$$\frac{\gamma\Delta}{(\pi_1+\pi_2)/2} - \frac{\pi_1-\pi_2}{\pi_1+\pi_2} = \frac{\exp(\beta\Delta)-1}{\exp(\beta\Delta)+1.} \tag{18}$$

Here the critical β_c is $2\gamma/\bar{\pi}$ with $\bar{\pi} = (\pi_1+\pi_2)/2$. Let us assume without loss of generality that $\pi_1 > \pi_2$. Equation (18) amounts to shifting the left-hand side of equation (11) to the right by $(\pi_1-\pi_2)/(\pi_1+\pi_2)$.

If β is above[13] β_c, the three intersections remain as long as the difference in profits is not too large. Which of the two asymmetric intersections is actually reached by the learning dynamics depends on initial conditions.

Thus, as illustrated in figure 2, buyers can remain loyal 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 asked a lower price. When the most often frequented shop changes its price, the loyalty to that shop describes the upper branch of the loyalty versus profit curve (figure 1). The loyalty remains on the upper branch as long as it exists, i.e. until the point where the slope is vertical. When profit decreases beyond that level, a sudden and discontinuous transition to the lower branch occurs. This is the point when customers change their policy and visit the other shop. But, if the first shop reverses its high price/low buyer profit policy when loyalty 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.

Thus an 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. A consequence of this phenomenon, is that in order to attract customers who are loyal to another shop, a challenger has to

[12] Once more, this statement applies rigorously to the mean field approach. In the case of large heterogeneity of local interactions in Markov random fields, ordered and disordered regions might coexist.

[13] If $\beta < \beta_c$ then there remains only one stable solution, in which there is a small difference in preferences proportional to the difference in profits (if $\beta\Delta$ is small):

$$J_1 - J_2 \simeq \frac{2(\pi_1-\pi_2)}{(\beta_c-\beta)\bar{\pi}}. \tag{19}$$

Compare with footnote 9.

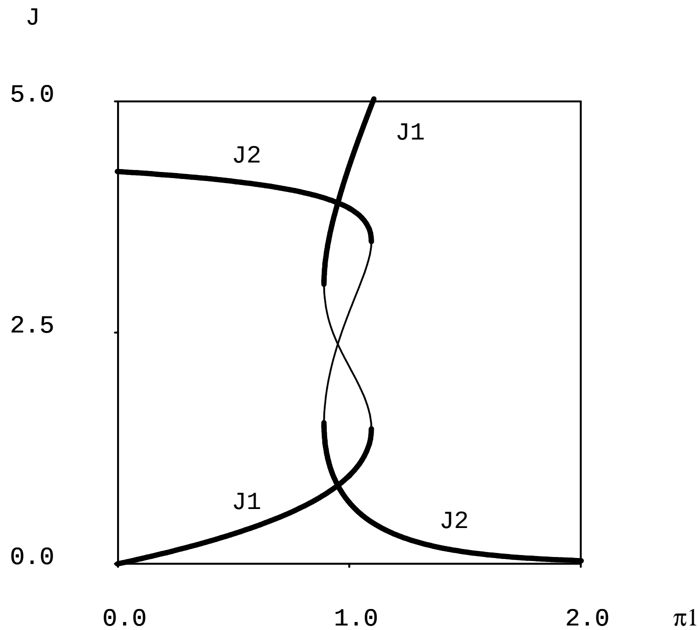

Figure 2: Hysteresis of preference coefficients. Plot of both preference coefficients versus π_1, the profit to be obtained from shop number 1 when π_2 the profit to be obtained from shop number 2 is held equal to 1. ($\beta = 0.5$ and $\gamma = 0.2$). The thick lines correspond to stable equilibria for both preference coefficients, J_1 and J_2, and the thin lines to unstable equilibria (if $\pi_1 \simeq \pi_2$). In the three solutions region, if the initial conditions are such that J_1 is large (and J_2 is small), J_1 remains large when π_1 is decreased, even when $\pi_1 < \pi_2$. The stability of this metastable attractor is lost when $\pi_1 = 0.89$. In a symmetrical manner, the high J_2 attractor existing at low π_1 can be maintained up to $\pi_1 = 1.095$. (the figure was drawn using GRIND software, De Boer 1983).

offer a profit significantly greater than the profit offered by the well established shop: once preference coefficients have reached equilibrium in the ordered regime, customers switch only for differences in profits corresponding to those where the slopes of the curves $J(\pi)$ in figure 1 are vertical (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.

4 Results

4.1 Indicators of order

We next proceed to run a number of numerical simulations of our model. This first enables us to check whether the theoretical results obtained from the mean field approximation are consistent with those obtained by running the discrete stochastic process as described by equation (2) and (3). Second, as discussed in the next section, it allows us to compare the simple model with more complicated, analytically intractable versions.

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, involving the dynamics of the buyers' J_{ij} vectors, is harder to monitor. We used two methods to do this.

Firstly, adapting a measure used in Derrida (1986) for instance, we define an order parameter y_i by

$$y_i = \frac{\sum_j J_{ij}^2}{(\sum_j J_{ij})^2}, \tag{20}$$

In the organized regime, when the customer is loyal to only one shop, y_i is close to 1 (all J_{ij} except one being close to zero). On the other hand, when a buyer visits m shops with equal probability, y_i is of order $1/m$. 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 loyalty of every single buyer. The first three graphs of figure 3 and 4 display simplex plots of a simulation at different steps. Each agent is represented by a small circle of a specific colour or shade, which represents the agent's probabilistic choice, i.e. the probability distribution over the 3 shops (corresponding to the 3 apices of the triangle). Proximity to one corner is an indication of loyalty to the shop corresponding to that apex. Agents represented by circles close to the center search all shops with equal probability.

4.2 A simple model

A simple model was run with 3 sellers and 30 buyers, for a large variety of parameter configurations and initial conditions. In the simulations, time is discrete and buyers receive equal profits when a transaction is made. Sellers' endowments at the begining of each session are finite, which implies that $Prob(q_i > 0)$ does not have to be one as in the simplest version solved analytically. The figures 3 and 4 correspond to a memory constant $\gamma = 0.1$. The critical non-linear parameter corresponding to a unitary profit is then $\beta_c = 0.3$ (equ. 13). Initial J_{ij} were all 0. Depending on the value of the non-linear parameter β, the two predicted dynamic regimes, order and disorder, are observed.

4.2.1 Disorganized behavior

For low values of the non-linear parameter β buyers never build up any loyalty. This is observed in figure 3, which describes the dynamics obtained with $\beta = 0.15\beta_c$. The daily profit of buyers averaged over all buyers and over 100 days after a transition period of 100 days, is only a fraction[14] of the buyer's profit per transaction. This is due to all those occasions on which 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 only a fraction of ten times the seller profit per transaction (the factor 10 corresponds to the average number of buyers per shop). This difference was also generated indirectly by buyers who visited empty shops since, at the same time some shops with supply were not visited resulting in in losses for their owners.

As seen in the simplex plots of figure 3, even at time 50, agents are still scattered around the barycenter of the triangle, an indication for a disordered regime without loyalty of any agent to any shop. Similarly, the order parameter y fluctuates well below 0.50 and thus corresponds to randomly distributed J_{ij}. Figure 3 shows that the performance of shop number 1 exhibits large fluctuations. The same is true for the two other shops.

4.2.2 Organized behavior

In sharp contrast, the same analysis performed with $\beta = 2\beta_c$ shows a great deal of organisation, see figure 4.

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 loyalty 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 very close to their profit per transaction. Because buyers have not changed shops during the last 100 days, sellers learnt to purchase the exact quantity needed to satisfy all their buyers; they incur no reduction in profit (as compared to the model predictions).

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. In that sense, the ordered regime is Pareto superior to the disordered regime.

4.2.3 Heterogeneity of buyers

Let us recall at this stage that in the case of real markets, we expect a mix of buyers with different β and γ parameters, such that some buyers will be loyal to certain sellers, while others will continue to search. Herreiner (1997) shows that buyer heterogeneity does not qualitatively change the above described results. Organized or disorganized behavior is here a property of buyers, not a property of markets.

---

[14]The exact percentage figures depend on the specific demand and supply functions, i.e. on the relationship between purchase and resale price for both, sellers and buyers. The simulations presented here were done with the specific functions discussed in section 5.1. However, the observed decrease in profit for buyers and sellers is generic.

4.3 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 small. Variables γ and π thus have to be small, which is true for the simulation results given in figures 2 and 3. 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 customers might change loyalty: they may "divorce" one shop for another one. Indeed increasing γ results in larger steps taken by customers on the simplex, which might make them move 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: if fewer steps are needed the probability that the process will generate such changes becomes higher. Because of the exponential growth of the time of the "divorce" process with respect to γ and π, a small change in those parameters 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 of a given shop as a function of time: "infidelities" appear as peaks (temporary changes of the number of customers) and "divorces" as steps (permanent changes).

5 More complicated models and results

In this section, we will discuss further refinements of the simple model and see what influence they have on the behaviour of the agents. All the variants to be discussed share the same fundamental mechanism by which buyers choose sellers and the same way of updating preference coefficients as defined in section 2.2.

These more realistic variants of the model are no longer analytically tractable and we therefore have to resort to computer simulations to compare their dynamical properties with those of the simple soluble model and with empirical data.

It is important at this stage to specify the type of comparison that we intend to make between the variants of the model and empirical evidence. We certainly expect some changes to occur at the global level when modifications are introduced in the way in which individual agents make their decisions. Nevertheless, the main point here is to check whether the *generic properties* of the dynamics are still preserved after these changes. The existence of two distinct, ordered and disordered regimes, separated by a transition, is such a generic property. On the other hand, we consider as non-generic the values of the parameters at the transition and the values of variables in the ordered or disordered regime. Since even the more elaborate versions of our model are so simplified in comparison with a very complex reality, a direct numerical fit of our model to empirical data would not be very satisfactory, if only because it would involve many parameters which are not directly observable. The search for genericity is based on the conjecture[15] that the large set of models which share the same generic properties also includes the "true" model of the real system itself.

5.1 Prices

We first need some assumption about the specific relationship between prices, purchased quantities and profits to run more realistic simulations. Let us suppose rationality at the level of a single transaction. Each buyer, being himself a retailer, faces a local demand function $p(q)$, which determines the relationship between the price and the quantity q that he brings to the local market. Let us suppose in order to simplify matters that $p(q)$ is known by the buyers, that it is the same for all buyers and that it is a simple function of q such as[16]:

$$p(q) = \frac{b}{q+c}. \tag{21}$$

The buyer's profit in this particular example is then:

$$\pi_b = q\left(\frac{b}{q+c} - p\right), \tag{22}$$

where p is the price asked by the seller. We then suppose 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. This quantity is:

$$q = \sqrt{\frac{bc}{p}} - c, \tag{23}$$

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

$$\pi_s = q\,(p - p_a) = \left(\sqrt{\frac{bc}{p}} - c\right)(p - p_a), \tag{24}$$

with respect to the price p that they charge to the buyers[17], where p_a is the price at which the sellers themselves purchase the fish.

[15]This general conjecture, which is basic in the dynamic modeling of complex systems, is proven rigorously for specific systems such as classes of universality in physics (see for instance Pfeuty/Toulouse (1977)) or structural stability in mathematics (see for instance Thom (1975)).

[16]The particular choice of the function $p(q)$ is of no importance, it allows to run simulations and to make comparisons between the different scenarios. For the model any monotonic decreasing function would do.

[17]The profit-maximizing price p is the solution to a cubic equation with first-order condition

$$p^3 - \frac{b}{4c}(p + p_a)^2 = 0, \tag{25}$$

which we calculate for the specific values used in the simulations.

The numerical simulations described in section 4 and 5 were done with $b = c = 1$ and $p_a = 0.3$. If these values are used, in the disorganized regime the buyer profits reduce to 88% and the seller profits to 74% of the possible profit per transaction.

5.2 Two sessions

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 regarded as the main motivation for refusal in standard search models. An alternative explanation is that customers refuse deals now in order to induce better offers in the future. In either case, to take this into account, we have to consider a model in which customers are given at least two occasions to purchase goods.

One further assumption to relax particularly 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 exactly the quantity they bring to the market. If they were able to predict precisely how many customers will visit their shop and accept their price, they would know exactly how much to supply. But, when their forecasts are not perfect they may not have the appropriate quantity for the number of buyers they actually face at the close of the market. It might therefore be better for them to sell at a lower price rather than to keep goods that they are not, by assumption, able to sell the next day. We ran the simulations with a constant afternoon price which is the morning price lowered by a factor $1-\epsilon$. A more intelligent choice for the sellers, namely monitoring previous fluctuations of the number of buyers and decreasing afternoon prices in proportion was also tested.

To summarise, we divide the day into two periods:

- During the morning, sellers maximize their profit and sell at a price p_{am} equal to p. Buyers visit one shop in the morning.
- During the afternoon they sell at a lower price $p_{pm} = (1 - \epsilon) \cdot p$ in order to reduce losses from unsold quantities. We assume that, because prices are lower in the afternoon, all buyers return for the afternoon session. Buyers visit one shop in the afternoon.

Sellers arrive in the morning with a quantity Q of the good corresponding to the number of customers they expect times q, plus some extra quantity of that good in case they have more customers than expected (see next section). The profit they expect from this additional amount is that obtained by satisfying new customers or unexpected former customers.

Buyers have to decide every morning whether to buy at the morning price or to wait for a better price in the afternoon. Of course waiting has a trade-off: they might not find anything to buy in the afternoon and thus make no profit. They choose an action according to their expectation of the average afternoon profit with respect to what they would get by buying in the morning, which they know from equation (22).

Average afternoon profit is estimated from their past history of afternoon profits. We used in the simulations a simple quadratic fit of the afternoon profit as a function of morning prices. For all reasonable choices of afternoon prices and extra supply by the sellers, expected afternoon profits for buyers are much smaller than morning profits, essentially because their chances of finding goods in the afternoon are smaller than in the morning. We discovered that even with their primitive prediction abilities, buyers soon (say after 50 time steps) realise that they would do better if they accepted the morning offers. Further investigations about the refusal issue can be found in Herreiner (1997).

All numerical simulations show that the introduction of a second session does not change the qualitative behaviour of the system: a low β disordered regime and a high β ordered regime still exist with the same characteristics as in the one session model. But the time to eventually reach the ordered regime and the width of the transition are increased. Estimated[18] β_c is at most 20% higher with two sessions than with one.

A change induced by the introduction of an afternoon session is that divorces are observed in the ordered regime for a wider range of the learning parameter γ: for $\gamma > 0.1$, as opposed to $\gamma > 0.3$ for the one-session model. On the occasion of an infidelity a buyer has a much better chance of making a higher afternoon profit with a new shop that has extra supply, she therefore takes larger steps across the simplex.

5.3 Sellers' initial endowment

We mentioned previously that the sellers may want to adjust their initial endowment to take into account the expected number of customers and possible fluctuations of that number. To do this sellers would need to know the probability distribution of the number of customers. Let us assume for the sake of comparison to results in search theory, that this distribution is continuous: $f(n_b)$ with $n_b \in [0, n]$. Define as $\hat{Q} = \hat{n} \cdot q$ the optimal endowment which maximizes a seller's expected profit if each buyer demands the same quantity q. Maximizing expected seller profit with respect to $\hat{n}$ yields the following condition, which determines $\hat{n}$ and therefore $\hat{Q}$:

$$1 - \int_0^{\hat{n}} f(n_b) dn_b = \frac{p_a}{p} \tag{26}$$

(p is the price at which they sell and p_a is the price at which they buy). The rule defined by this equation is optimal only for short-run considerations, if sellers assume that every market day is a one-shot game. It prevents strategic use of endowments, by which a seller tries to gain additional loyal customers by having extra units for unexpected customers.

In line with our general approach, we did not suppose for the simulations that sellers have a perfect knowledge of the probability distribution of visitors, but that they use a simple routine to add extra supply whenever they observe fluctuations in

[18] Since the transition is not abrupt as in the theoretical model, we have chosen a critical value for y, $y = 0.5$, to determine β_c, i.e. β such that $y = 0.5$.

the number of visits. The extra amount at time t is computed according to

$$\alpha(t) = (1-\epsilon)\cdot\alpha(t-1) + \epsilon\cdot var(n_b) \tag{27}$$

where ϵ is small and $var(n_b)$ is the variance of the number of buyers since the beginning of the simulation. The initial value of α is non zero at the beginning of the simulation. This equation simply describes the reduction of α in the absence of fluctuations. We checked by several numerical simulations with different choices of initial α and of ϵ that the only observable changes are variations of β_c, the critical threshold for order (in the 10 % range). The existence of two dynamic regimes persists.

Another possible refinement would consist in improving the predictive ability of the seller with respect to the number of customers. We tried a moving average prediction rather than the prediction based only on the preceding day but this only reduced performance (β_c increases).

5.4 Price fluctuations

The idea of a market with a uniform price is not realistic and we wanted to check the influence of price variations over time on the agents' behavior. In fact, the above section 3.2 on hysteresis already gives us a clue as to the possible results of price changes: price differences resulting in profit differences for the buyer lower than the width of the hysteresis curve do not change loyalty and therefore should not destroy order. For the parameters values of figure 2, one shop could increase its prices from equality with the other shop up to 19 % before losing its customers.

We ran simulations with morning price $p(t)$ fluctuating in each shop with an auto-regressive trend towards the price p (solution of equation (25)) which maximizes profits. The price is also decreased when potential buyers refuse the offer, a situation seldom encountered by the end of the simulations as mentioned earlier. The morning price of each shop is then varied in the simulations according to the following expression:

$$p_j(t+1) = \eta_j(t)\Big[p_j(t) - \lambda\big(p_j(t)-p\big) - \mu\frac{r_n}{n_j}\Big], \quad \eta_j(t) \overset{iid}{\sim} U[1-\epsilon, 1+\epsilon],\ \epsilon\in[0,1], \tag{28}$$

n_j and r_n are respectively the number of customers of the shop and the number of customers who refused the price of shop j during the last session.

The simulation results are remarkably close to the results obtained with constant morning price for both sessions: the transition is sharpened and order is obtained for slightly lower values of β.

6 Empirical Evidence

In order to see whether there was any empirical evidence of ordered or disordered behaviour of buyers in a market, we started from a data base for transactions on the wholesale fish market in Marseille (M.I.N Saumaty). The data base contains the following information:

- No. of buyers ca. 1400
- No of sellers 45
- For each individual transaction:
  - Name of buyer
  - Name of seller
  - Type of fish
  - Weight of fish
  - Price
  - Order in seller's transactions
- Dates: from 02 - 01 - 1988 to 29 - 06 - 1991
- Total number of transactions: 237162.

The market is organised as in our model, that is, no prices are posted, sellers start with a stock of fish which has to be disposed of rapidly because of its perishable nature. Buyers are either retailers or restaurant owners. Deals are made on a bilateral basis and the market closes at a fixed time. Of course the model is an extreme simplification of the real situation: there are different kinds of fish on the market, each species of fish is heterogeneous, buyers demand different quantities of fish. For a buyer the alternative to purchasing his optimal good is to purchase some inferior alternative.

Direct examination of the data file with the help of standard sorting facilities reveals a lot of organisation in terms of prices and buyer preferences for sellers. In particular, one observes that the most frequent buyers (those who visit the market more than once per week) with very few exceptions visit only one seller, while less frequent buyers visit several sellers, which is consistent with our model. The data will be analysed in this section only in terms of market organisation. Other aspects, such as data classification and price dynamics, which show persistent price dispersion, were analysed in Kirman and Vignes (1991) and Härdle and Kirman (1995).

6.1 Testing our model

A first step in comparing our theory with empirical data is to check whether individual buyers display ordered or disordered behaviour during those three and a half years. Since the classical approach to agent behaviour predicts search for the best price, and since searching behaviour implies visiting different shops, any manifestation of order would tend to support our theoretical prediction. If we find evidence of ordered behaviour for certain participants, a second step is then to relate the difference in the observed behaviour of these traders to some difference between their characteristics and those of other buyers.

	market shares of largest seller			monthly purchase share bought from one seller	
	1st	2nd	3rd	95%	80%
cod	43%	14%	12%	48%	
whiting	27%	8%	8%	24%	53%
sole	15%	14%	14%	33%	55%

Table 1: Loyalty in Cod, Whiting, and Sole Market

For the first step, to check for loyalty of buyers, we consider statistics for cod, whiting and sole transactions in 1989[19], see table 1.

Since we are interested in loyalty issues, we concentrated on the buyers who were present in the market for at least 8 months. As can be seen in the first three columns of table 1, the market for cod is much more concentrated than the market for whiting or sole. In the cod market almost half the buyers (86 of 178) buy more than 95% of their monthly purchases from one seller only, see the fourth colum of table. Also in the whiting and sole market buyers are loyal, but to a lesser degree: more than half of them[20] buy more than 80% from one seller. Hence, there are large fractions of loyal buyers in all three markets. Along the same line, because of the sharp transition between the two possible behaviours, our theory predicts a bimodal distribution of fidelities, which can be observed in figure 5.

For the second step, recall that our theory relates loyalty to the parameters β (discrimination rate) and π/γ (cumulated profit). β, the discrimination parameter is likely to vary from buyer to buyer, but we have no direct way to test it. However, π/γ is strongly and positively related to monthly purchases of buyers, and we therefore use the latter as a proxy variable. Figure 6 summarises loyalty of buyers in terms of relative frequency of visits to their favorite seller as a function of their monthly purchase of cod. One may observe that loyalty is high in general and that a number of buyers visit only one seller. A cubic fit shows that loyalty increases with monthly purchase.

All three features are consistent with our theory, and in contradiction with a random search behavior for all buyers. With standard statistical tests we check whether the population of buyers exhibits two different types of behaviour. We divide the buyers of cod into two naturally emerging groups according to the total size of their transactions. We choose as our dividing criterion a total purchase of two tons of cod over 36 monthes. The fraction of transactions with the most often visited seller was found to be 0.85 for big buyers and 0.56 for small buyers. If we consider, as in the model, that the two populations consist of individuals drawing their "favorite seller" with probability $P1$ in one population and $P2$ in the other one, we can test the hypothesis $P1{=}P2$. Given the two values for the tested data set, both the standard Maximum Likelihood test and Fisher's Exact test rejected the hypothesis $P1{=}P2$

[19]The statistics for other periods of comparable length are very similar.

[20]Whiting 124 of 229, and sole 154 of 280.

at all levels of confidence. Buyer behaviour is clearly linked to both measures of transaction volume and therefore to cumulated profits π/γ: the higher cumulated profits the more loyal a buyer.

6.2 Testing alternative models for order

The observed agreement between our model and empirical evidence does not "prove" that it is the only possible model. As most often with complex systems, several explanations at different levels of generality can be used to describe observed phenomena. Furthermore different models might not be mutually exclusive as we will discuss.

One alternative explanation to our model would ascribe a more strategic behaviour to the agents. If agents consider themselves in a repeated game, then they might establish implicit contracts which imply punishment for breaches of those contracts. Discussions[21] with sellers in Marseille reveal that they do not offer fish for specific customers, rather they claim to have "learnt" what "their" customers require; similarly, buyers do not order fish but assume that "their" sellers provides what they require. This is consistent with the mutual reinforcement mechanism suggested by our theory. If a particular buyer does not appear, this is not regarded as a breach of contract; if this happens repeatedly and some fish remains unsold, the seller will simply re-adjust his supply of fish accordingly.

At the same level of generality as our model, another alternative explanation could be based on the idea of "niches": a buyer would prefer a given seller because he provides him a product closer to his specific needs. Let us first note that the two hypotheses are not mutually exclusive: even if niches were an important factor, one would still have to explain why a seller chooses a niche stategy rather than selling a large variety of fish. Loyalty of buyers might be a condition for the profitability of "niches". Direct examination and surveys show that even though certain sellers specialise in serving supermarkets or institution cafeterias, all niches are occupied by several sellers. This is also consistent with the fact that many buyers are retailers who have to serve many different clients on their local markets. We check for the existence of niches via a cluster analysis according to average prices and quantities sold by sellers. Sellers are considered as members of the same cluster if their distribution of prices and quantities significantly overlaps. We find two clusters of cod sellers, low cost/large quantity sellers (5 sellers) and large cost/low quantity sellers (30 sellers). Since loyalty and search behaviour are observed in these two multi-member niches, the niche phenomenon cannot account by itself for the existence of loyalty; but according to our theory it facilitates loyalty by decreasing the number of sellers in competition, and thus lowering the critical transition parameter.

The model we used, including its variants, considers buyers as active agents and sellers as rather passive. Alternative and/or complementary explanations of the ob-

[21] It is perhaps worth emphasing that the basic theory of this paper was elaborated in the light of conversations with market participants who often were able to explain certain features of the data. Sellers comment on buyer behavior along the lines of "he (the buyer) comes here because he knows that he will find the sort of fish that he requires", whilst buyers explain their behaviour with statements such as "I go there because he has the fish that I want".

served organisation could be based on a more active role of sellers. A possible test for additional explanatory power of hypotheses which link loyalty to sellers' behaviour, is to check wether different sellers have different fractions of loyal buyers among their customers, and if so why. We did measure the fractions of loyal buyers of each seller and found them to be strongly[22] and positively correlated with the average quantity of fish per transaction sold by the seller (at least for all sellers making more than one transaction per day on average). We therefore conclude that the buyers' learning and search behaviour as described in our model is sufficient to explain the observed organisation without the necessity of further assumptions about seller behaviour.

7 Conclusions

We have examined a simple model of a market in order to see how the "order" that is observed on many markets for perishable goods develops. "Order" here means the establishment of stable trading relationships over the periods in which the market is open.

In the simplest model, we have shown analytically that an ordered regime appears whenever an agent's discrimination rate among shops divided by the number of shops is larger than the reciprocal of the discounted sum of their profit. When an individual's parameters place him in the organized regime, a buyer has strong preferences for one shop over all others. On the other hand, in the disordered regime, agents do not show any preference for a particular shop. The transition between the ordered and disordered regimes is continuous but very abrupt (at least for the simplest one session model) in terms of the order parameter.

Since individual properties of buyers govern the ratio of their discrimination rate β to the threshold rate $\beta_c = m\gamma/\pi$, a bimodal distribution of buyers, some with an ordered behavior some not, is to be expected in real markets. A comparison with empirical data from the Marseille fish market indeed shows the existence of a bimodal distribution of searchers and loyal buyers, and a positive correlation of the loyal behavior with the frequency of transactions.

When more realistic assumptions are introduced, such as adaptive behavior of sellers, fluctuations in prices, and a second session with a lower price to clear the market, simulations show that the critical value of the transition parameter is increased and the transition becomes somewhat less abrupt. However both regimes can still be observed. The simple model is thus robust with respect to changes that can be made to improve realism: its main qualitative property, namely the existence of two regimes of dynamical behavior is maintained.

Thus what we have shown within the context of an admittedly very simple model is that the presence of "order" and "organisation" in a market is very dependent on, and very sensitive to, the way in which agents react to their previous experience. As has been seen "order" in our model is more efficient in Pareto terms than disorder

[22]The correlation is stronger for sellers, with much less noise than for buyers. This is simply due to the fact that seller statistics involve more averaging than buyer statistics - on average there are more transactions per seller than per buyer.

and it is therefore of considerable economic interest to be able to identify under which conditions "order" emerges.

Acknowledgment

We thank Derek Smith for help with the Tk/Tcl interactive display of simulation results, Rob Deboer for the use of his GRIND software, Paul Pezanis for his help in the analysis of Marseille fishmarket data, Keneth Arrow, Larry Blume, Paul David, Steven Durlauf, Fang Fang Tang, Douglass North, Ariel Rubinstein and Nick Vriend for helpul comments, and Olivier Chenevez, Bernard Derrida, Jean Pierre Nadal and Jean Vannimenus for helpful discussions and important suggestions. The Laboratoire de Physique Statistique of ENS is associated with the CNRS (UA 1306). This work was started during a visit by GW and AK to the Santa Fe Institute which we thank for its hospitality.

8 References

Anderson S. P., A. de Palma and J.F. Thisse (1993), *Discrete Choice Theory of Product Differentiation*, Cambridge (MA): MIT Press.

Aoki M. (1996), *A New Approach to Macroeconomic Modelling*, New York: Cambridge University Press.

Arthur, B.W. and D.A. Lane (1993), "Information Contagion", *Structural Changes and Economic Dynamics*, 4, 81-104.

Arthur W.B., J.H. Holland, B. LeBaron, R.G. Palmer, and P. Tayler (1997), "Asset Pricing under Endogenous Expectations in an Artificial Stock Market", in: W.B. Arthur, D. Lane and S.N. Durlauf (eds), *The Economy as an Evolving, Complex System II*, 15-44.

Blume L. (1993), "The Statistical Mechanics of Strategic Interaction", *Games and Economic Behaviour*, 5, 387-424.

Brock W.A. and S.N. Durlauf (1995), "Discrete Choices with Social Interactions I: Theory", *Working Paper 95-10-084*, Santa Fe Institute, Santa Fe, NM.

Brout, E. (1965), *Phase Transitions*, New York (NY): Benjamin.

De Boer, R. J. (1983), *GRIND: Great Integrator Differential Equations*, Bioinformatics Group, University of Utrecht, The Nederlands.

Derrida B. (1986) "Phase Transitions in Random Networks of Automata" in: J. Souletie, J. Vannimenus and R. Stora (eds), *Chance and Matter*, New York (NY): North-Holland.

Durlauf S. (1990), "Locally Interacting Systems, Coordination Failure and the Behaviour of Aggregate Activity", *mimeo*, Stanford University, Stanford, CA.

Diamond P. (1989), "Search Theory", in: J. Eatwell, M. Milgate and P. Newman (eds), *The New Palgrave: A Dictionary of Economics*, London (UK): Macmillan, 273-79.

Fisher F. M. (1973), " Stability and Competitive Equilbrium in Two Models of Individual Price Adjustment", *Journal of Economic Theory*, 6, 446-470.

Föllmer H. (1974) "Random Economies with Many Interacting Agents", *Journal of Mathematical Economics*, 1/1, 51-62.

Härdle W. and A. Kirman, (1995), "Non Classical Demand: a Model-Free Examination of Price-Quantity Relation in the Marseilles fish market", *Journal of Econometrics*, 67, 227-257.

Herreiner, D. (1997), "Market Organization", *mimeo*, University of Bonn.

Jackson M. and A.Wolinsky (1996), "A Strategic Model of Social and Economic Networks", *Journal of Economic Theory*, 71, 44-74.

Kilani, K. and J. Lesourne (1995), "Endogeneous Preferences, Self-Organizing Systems and Consumer Theory", *Mimeo 95-6*, Laboratoire d'Econométrie, Conservatoire National des Arts et Métiers, Paris.

Kirman, A.P. and A. Vignes (1991), "Price Dispersion. Theoretical Considerations and Empirical Evidence from the Marseilles Fish Market", in K.G. Arrow (ed), *Issues in Contemporary Economics*, London (UK): Macmillan, 160-185.

Kirman A.P. (1993), "Ants, Rationality and Recruitment", *Quarterly Journal of Economics*, 108, 137-156

Kirman, A.P. and N. Vriend (1995), "Evolving Market Structure: A Model of Price Dispersion and Loyalty", *mimeo*, Department of Economics, Virginia Polytechnic Institute and State University, Blacksburg, VA.

Lesourne, J. (1992), *The Economics of Order and Disorder*, Oxford: Clarendon Press.

Luce, D.R. (1959), *Individual Choice Behaviour*, New York (NY): Wesley.

Marimon R. (1996), "Learning in Economics", *Working Paper 96/12*, Department of Economics, European University Institute, Florence, Italy.

Mas-Colell A, M. Whinston and J.R. Green, (1995), *Microeconomic Theory*, Oxford: Oxford University Press.

Orléan A. (1995), "Bayesian Interactions and Collective Dynamics of Opinions: Herd Behavior and Mimetic Contagion", *Journal of Economic Behavior and Organization*, 28, 257-274.

Pfeuty, P., Toulous, G. (1977), *Introduction to the Renomalization Group and to Critical Phenomena*, New York (NY): Wiley.

Thom, R. (1975), *Structural Stability and Morphogenesis*, New York (NY): Benjamin.

Vriend N. (1994), "Self-Organized Markets in a Decentralized Economy", *Working Paper 94-03-013*, Santa Fe Institute, Santa Fe, NM.

Weisbuch G. (1990), *Complex Systems Dynamics*, Redwood City (CA): Addison Wesley.

A Fixpoints of tanh()

Restate equation (11) as

$$\tanh(zx) = x, \quad z \in \mathcal{R}, \qquad \text{with} \quad x = \frac{\gamma\Delta}{\pi} \quad \text{and} \quad z = \frac{\beta\pi}{2\gamma} = \frac{\beta}{\beta_c}, \tag{29}$$

which is known to have one solution or fixpoint at $x = 0$ if $z \leq 1$; if $z > 1$ then there are three solutions.

The stability of these solutions can be analyzed by the slope of $\tanh(zx)$. When $\tanh(zx) > x$ then movement is to the right ($d\Delta/dt > 0$), and conversely:

$$\begin{aligned}
\tanh(zx) > x \quad &\Longleftrightarrow \quad \pi \cdot \frac{\exp(\beta J_1) - \exp(\beta J_2)}{\exp(\beta J_1) + \exp(\beta J_2)} \; > \; \gamma(J_1 - J_2) \\
&\Longleftrightarrow \quad \pi \frac{\exp(\beta J_1)}{\exp(\beta J_1) + \exp(\beta J_2)} - \gamma J_1 \; > \; \pi \frac{\exp(\beta J_2)}{\exp(\beta J_1) + \exp(\beta J_2)} - \gamma J_2 \\
&\Longleftrightarrow \quad \frac{dJ_1}{dt} > \frac{dJ_2}{dt} \Longleftrightarrow \frac{d\Delta}{dt} > 0.
\end{aligned}$$

If $\tanh(zx)$ has one solution, then the slope of $\tanh(zx)$ is flatter than the slope of x; the one solution is stable. If $\tanh(zx)$ has three solutions, then at $x = 0$ ($\Delta = 0$) the slope of $\tanh(zx)$ is steeper than the slope of x, i.e. the central solution is unstable, and then two other solutions are stable.

If $\tanh(zx)$ has three solutions, then the ratio of the preference coefficients at the outer stable solutions is approximately

$$\frac{J_1}{J_2} = \exp\left(\frac{\beta\pi}{\gamma}\right), \tag{30}$$

which can be obtained from equations (9) and (10) if $J_2 \approx 0$ and $J_1 \approx \pi/\gamma$.

To determine the speed of transition between the disordered and the ordered regime we calculate the third-order Taylor expansion of $\tanh(zx)$ at $x_0 = 0$ ($\Delta_0 = 0$):

$$\tanh(0) = 0, \qquad \tanh'(0) = z, \qquad \tanh''(0) = 0, \qquad \tanh'''(0) = -2z^3.$$

This yields

$$F(x_0 = 0) = zx - \frac{(zx)^3}{3} = x,$$

solving for x (Δ) leads to

$$x = \sqrt{\frac{3(z-1)}{z^3}} \qquad \text{and} \qquad \Delta = \sqrt{\frac{12(\beta - \beta_c)}{\beta^3.}} \tag{31}$$

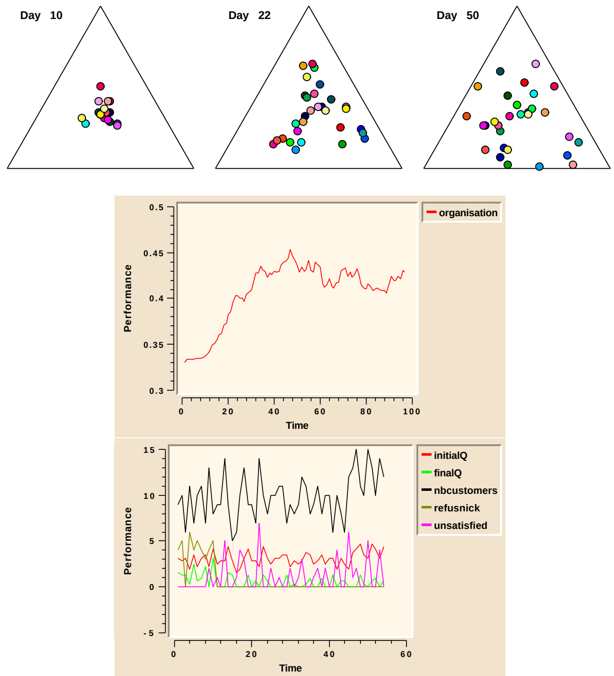

Figure 3: Charts for the disorganized regime. (30 agents visiting 3 shops, with $\gamma = 0.1$ and $\beta = 0.15\beta_c$) The first three graphs monitor market organization by simplex plots taken at time 10, 22 and 50. They show that no organization takes place. The fourth graph shows a time plot of the order parameter y (vertical axis: $[0.3, 0.5]$). The order parameter stays well below 1. The last graph gives a record of shop 1. The time charts display the initial and the final endowment, the number of customers, the number of customers refusing the proposed price (see section 5.2), and the number of unsatisfied customers who did not manage to buy anything. Fluctuations in the market do not decline over time.

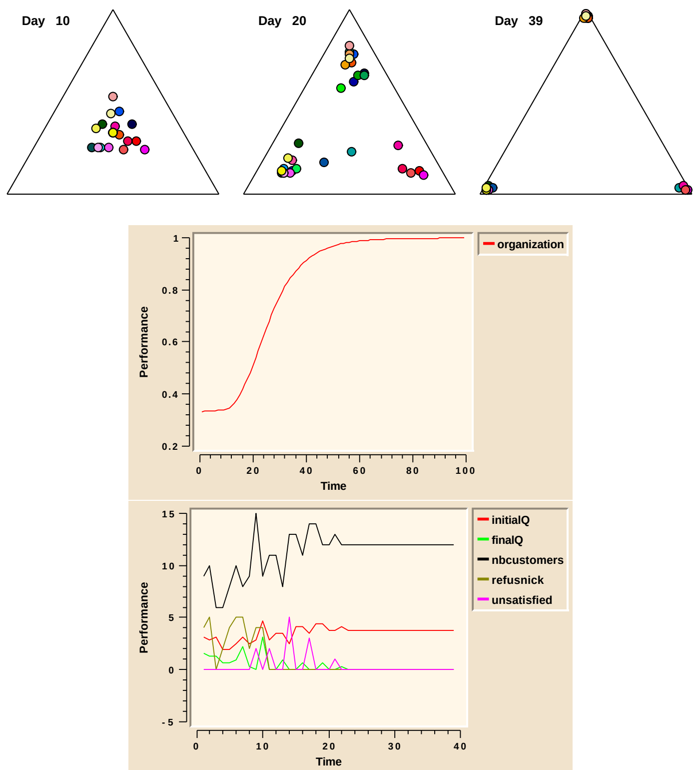

Figure 4: Charts for the organized regime (30 agents visiting 3 shops, with $\gamma = 0.1$ and $\beta = 2\beta_c$). All charts and notation are the same as for figure 3, except for the scale of the order parameter plot (y). In the three simplex plots, starting from indifference between all three shops, the circles move to the corners representing the preferred shops. Organization takes place. The order parameter y increases steadily from 0.33 to nearly one. The time charts show how fluctuations dimish quickly due to organization.

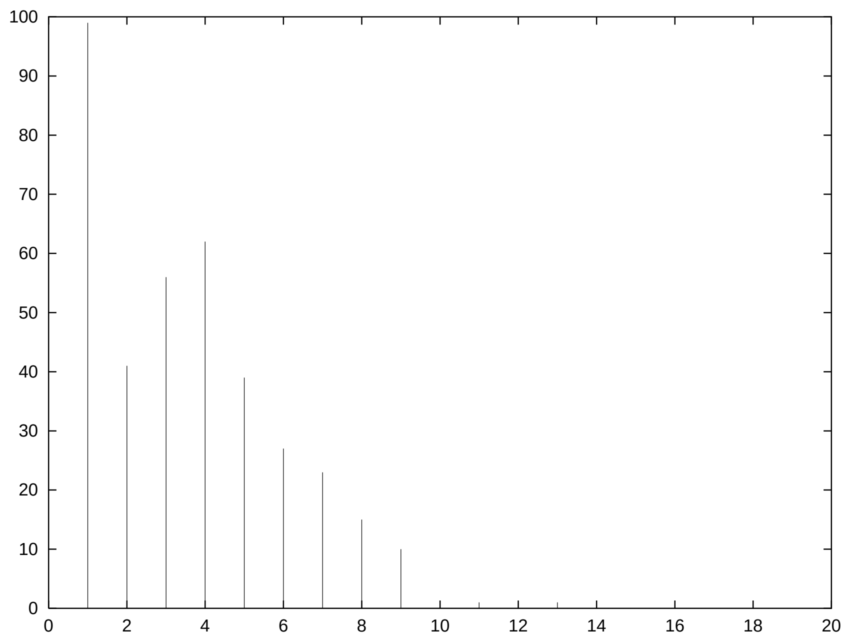

Figure 5: Histogram of the number of buyers of cod as a function of how many shops they visit in 1990. The sample of buyers includes only those visiting more than once a month, and who stay in the market for more than six months. The distribution is bimodal.

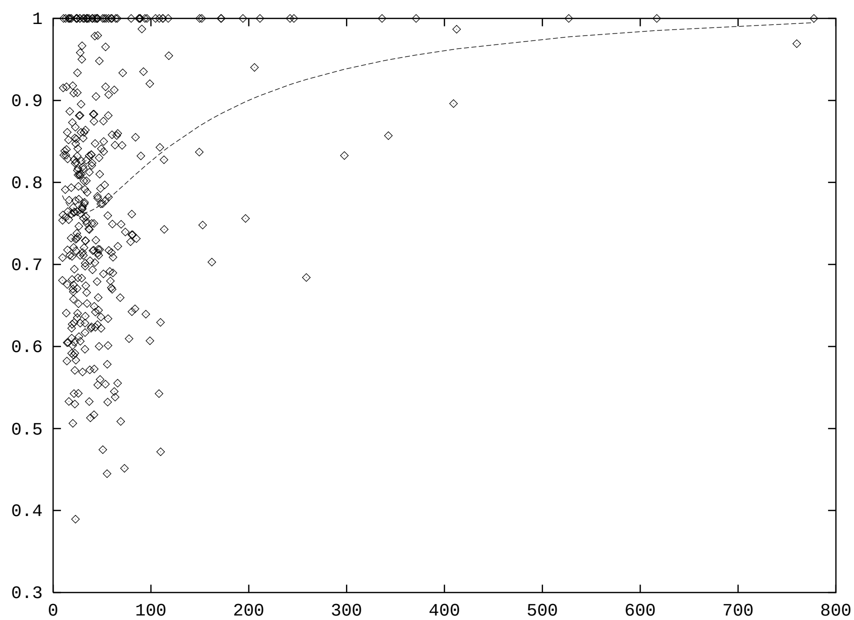

Figure 6: Each dot is an empirical evidence from the Marseilles fish market representing buyer loyalty to his favorite seller (relative frequency of visits) as a function of his monthly purchase of cod in kilograms. Low purchases correspond to unfrequent buyers, who generally visit once a week, while large purchase are those of buyers who visit nearly every day the market is open. The continuous line is a cubic fit which shows that loyalty increases with monthly purchase.