MARKET ORGANISATION AND TRADING RELATIONSHIPS

G. WEISBUCH, J.-P. NADAL, O. CHENEVEZ
Laboratoire de Physique Statistique de l'Ecole Normale Supérieure[†]
ENS, 24 rue Lhomond
75231 Paris Cedex 05, France

AND

A. KIRMAN
GREQAM, EHESS and *Université d'Aix-Marseille III*
13002 Marseille, France
and
Economic Theory III
University of Bonn, Adenauerallee 24 - 26
53113 Bonn, Germany

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

1. Introduction

Many markets are characterised by trading relationships. In certain markets individuals systematically trade with particular partners, whilst in others no such stable links are observed. Yet the way in which such organisation develops and its economic consequences are not considered in standard theoretical models.

A number of models have been developed to provide at least partial answers to these questions. 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 commodity. More complete models with individuals setting prices and buyers searching

[†]Laboratory associated with CNRS (URA 1306), Ecole Normale Supérieure and Universities Paris 6 and Paris 7

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, seeking a job, or when a firm invests in a large capital item.

Yet on many markets individuals trade frequently with each other. Of particular interest are 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 searching 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 can exactly predict 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. Modelling 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 fish market from which our empirical evidence is drawn. An important aspect of this particular market, and of other markets for perishable goods, is that agents face a trade-off between finding the best possible transaction and being sure that they can actually make a transaction. They are aware of the possibility of short supply by the end of the session: since sellers only bring to the market the quantities they expect to sell during one session, a buyer who searches until he finds the best price may not find anything to buy by the end of the session. Similarly, a seller first offering too high a price and re-adjusting only by the end of the session would realise too late that buyers have been served elsewhere. 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 commodity, 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. We then discuss the dynamics obtained according to how choices are made with respect to information. These models are analytically solved using the "mean field" approximation. For the case of exponential choice functions, which we use in the rest of the paper, 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 forever. 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 on the part 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

2.1. BASIC ASSUMPTIONS

Let us consider a set of n buyers i and a set of m sellers j. In order to simplify assumptions as much as possible, let us suppose that:

(i) Customers choose one shop each day according to their memory of previous transactions. As long as the shop has supplies, 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 customers that shop had before, and what quantity of stocks the shop had at the beginning of the day.

(ii) Since the good is perishable, and therefore cannot be stored overnight, each day a seller supplies some 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.

(iii) Each day the same market scenario is repeated.

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

2.2. A GENERAL FRAMEWORK FOR STUDYING BUYERS' DYNAMICS

We are interested in modelling buyers' behaviour. In this paper we assume that each buyer makes use of previous experience to select a seller. Since we want to emphasise the role of the individual buyers' choice functions, we assume that there is *no direct interaction between buyers*. We also assume that information on a seller is only obtained through transactions with him (there are no "posted" prices; see (Weisbuch *et al.*, 1997) for a real instance

$$\begin{matrix}\text{information}\\ \text{(observations)}\end{matrix} \rightarrow \begin{bmatrix}\text{coding}\\ \ \end{bmatrix} \rightarrow \begin{bmatrix}\text{choice}\\ \text{function}\end{bmatrix} \rightarrow \begin{matrix}\text{decision}\\ \text{(seller selection)}\end{matrix}$$

Figure 1. The general model

of such a condition). The general framework, illustrated in Fig. 1, is thus the following. Each time a buyer makes a transaction with a seller, he acquires some information about what he can expect from this particular seller (quality of goods, profit,...). This information will be encoded in some way which updates the previously acquired information about sellers. This stored information is the input to the (possibly probabilistic) choice (or decision) rule used by the buyer to select a seller for the next transaction.

Let us illustrate this general model with simple specific examples. Considering a given buyer, we will denote by $\mathbf{J} = \{J_j, j = 1, ..., N\}$ the stored information, J_j being the information concerning the jth seller. In the simplest case, J_j is a scalar. For instance, J_j may be the profit obtained the last time the buyer dealt with the jth seller; or it may be some moving average value of past profits from seller j, e.g.,

$$J_j(t) = (1 - \gamma)J_j(t - 1) + \pi_j(t) \tag{1}$$

where $\pi_j(t)$ is the actual profit at time t if j is the seller visited at time t, and $\pi_j(t) = 0$ otherwise. The parameter γ is smaller than 1: events far in the past are progressively forgotten. An updating rule such as (1) is an example of a coding scheme. One may consider more involved rules, taking into account not only the mean profit obtained from each seller, but also some information on the frequency of visits to each seller. In the following, we will only consider the case of a single variable J_j stored for each seller. We note, however, that our approach can be easily generalised to more complicated situations.

The choice rule, to be denoted by $P(j|\mathbf{J})$, is the probability that the buyer will choose the jth seller based on the stored information $\mathbf{J}$. A large class of model encountered in the literature corresponds to a decision rule defined by

$$P(j|\mathbf{J}) = \frac{f(J_j)}{\sum_{k=1}^{N} f(J_k)} \tag{2}$$

where $f(.) \geq 0$ is some function chosen a priori — to be called below the *choice function*. If, as above, J_j is a scalar, f is a real valued function of a single variable. Typical choices for f are a linear or affine function (Kilani and Lesourne, 1995), or an exponential function(Anderson *et al.*, 1993; Blume, 1993), in which case the choice rule is called the *logit* rule.

2.3. CHOICE FUNCTIONS AND PHASE TRANSITION

2.3.1. *Mean field approximation*

In this section we consider, for a general choice function, the mean field approximation as used in (Weisbuch *et al.*, 1997) for the logit case (the mean field approach has been also applied to other economic problems, see e.g. (Aoki, 1996; Brock and Durlauf, 1995)). We consider a buyer whose choice rule is as defined in (2), and the coding rule as in (1) (the discussion can be easily generalised to other coding rules). Moreover, for simplicity we assume a constant, seller independent profit π from each transaction (this implies in particular that the transaction is always possible and realized between the buyer and the chosen seller). Hence we have:

$$\begin{aligned} \pi_j(t) &= \pi \ \text{ if } j \text{ is chosen,} \\ &= 0 \ \text{ otherwise.} \end{aligned} \qquad (3)$$

The mean field approach (Derrida, 1986) consists in replacing randomly fluctuating quantities by their expectations, thus neglecting fluctuations. Averaging (1), one gets

$$J_j(t) = (1-\gamma)J_j(t-1) + \pi P(j|\mathbf{J}(t-1))\,. \qquad (4)$$

In the large time limit, one gets the fixed point (mean field) equations:

$$J_j = \frac{\pi}{\gamma}\frac{f(J_j)}{\sum_k f(J_k)}\,. \qquad (5)$$

In the above equation we have replaced $P(j|\mathbf{J})$ by its expression (2).

Let us study now the solutions of the mean field equations. More precisely, the equations (5) are fixed point equations of a dynamical process: among the solutions only the stable ones are meaningful, so we will have to study the stability of the solutions.

As a preliminary remark, summing over j the fixed point equations (5) one sees that any solution $\mathbf{J}$ satisfies

$$\sum_j J_j = \frac{\pi}{\gamma}\,. \qquad (6)$$

Obviously,

$$J_j = \frac{\pi}{N\gamma} \quad j = 1, ..., N \qquad (7)$$

is always a solution. Developing (4) in the vicinity of this symmetric fixed point (7)[1], one finds that it is stable if the quantity α, defined by

$$\alpha \equiv \left. \frac{d \ln f(x)}{d \ln x} \right|_{x=\frac{\pi}{N}}, \tag{8}$$

is smaller than 1. Otherwise, that is if

$$\alpha \geq 1 \tag{9}$$

the symmetric solution (7) is unstable: there must exist other, stable solutions.

To simplify the discussion, let us consider the simplest case of two sellers, $N = 2$. In that case, we can work with the single variable J_1, since according to (6) the other one, J_2, is equal to $\frac{\pi}{\gamma} - J_1$. Then the mean field equations become simply

$$J_1 = \frac{\pi}{\gamma} \frac{f(J_1)}{f(J_1) + f(\frac{\pi}{\gamma} - J_1)} \equiv g(J_1)\,. \tag{10}$$

In fact, it is clear that if J_1 is a solution, then $\frac{\pi}{\gamma} - J_1$ is also a solution. Hence we have at least two stable solutions. Since we have $J_2 = \frac{\pi}{\gamma} - J_1$, each pair of solutions can be written $\{J_1, J_2\}$. To keep the discussion simple, we will restrict the discussion below to the simplest case of a unique stable pair of solutions (hence one unstable and two stable solutions). Geometrically, a solutions J_1 of (10) is given in the plane $\{x, y\}$ by the intersection of the straight line $y = x$ with the curve $y = g(x)$. One can show that the parameter α defined above is here equal to the slope of g at that value $\frac{\pi}{2\gamma}$ of J_1. Hence the condition for having the symmetric point unstable is

$$\alpha = \left. \frac{dg(x)}{dx} \right|_{x=\frac{\pi}{2\gamma}} \geq 1\,. \tag{11}$$

Remark: if $f(0) = 0$, it is easily seen that there are always (at least) three solutions, the symmetric point (7), and the pair $\{0, \frac{\pi}{\gamma}\}$. Performing the stability analysis one finds that the non symmetric solutions $\{0, \frac{\pi}{\gamma}\}$ are stable if

$$\frac{\pi f'(0)}{\gamma f(\frac{\pi}{\gamma})} < 1\,. \tag{12}$$

2.3.2. *Interpretation*

If the only stable solution of the equilibrium equations is $J_j = \frac{\pi}{N\gamma}$, the frequencies of visiting any seller are equal. The probability of visiting a

[1]We use the fact that the derivative of the denominator of Eq. (2) is zero at the symmetric fixed point because of Eq. (6)

seller simply fluctuates without any stable preference for one particular seller emerging.

If there are other stable solutions $J_j \neq \frac{\pi}{N\gamma}$, the frequency of visiting some seller is larger than for the others. The buyer has a stable preference for one seller. According to the above discussion, the qualitative behaviour of the buyer depends on the choice function $f(.)$, the number of sellers N, the memory parameter γ, and the profit π only through the quantity α defined in equation (8). If the buyer modifies his choice strategy, or if his profit varies in such a way that his α changes, an abrupt change of behaviour will be observed when α crosses the critical value 1. This is analogous to a second order phase transition in physical systems, where the parameter α has the meaning of the inverse of the temperature. If one starts with a small value of α, the stable solution $\{J_j = \frac{\pi}{N\gamma},\ j = 1, ..., N\}$ remains valid until α reaches 1. Just above the transition, $\mathbf{J}$ starts to depart from the symmetric solution, with

$$\left| J_j - \frac{\pi}{N\gamma} \right| \sim \sqrt{\alpha - 1}\,. \tag{13}$$

Now it is reasonable to assume that the buyers on a market have different choice strategies, and/or make different profits, so that they have different values of α. When there exists a wide range of α values distributed around the critical value 1, one will observe two categories of buyers: the ones who choose the seller they will visit randomly and the others, who have strong preferences. We say that the distribution is bimodal.

2.3.3. *Specific choice functions*

The linear and affine cases:

Let us consider the simplest case, that is an affine choice function. As can be seen from the definition (2) of the choice rule, $f(x)$ and $af(x)$, for any $a > 0$, give the same choice rule. Without loss of generality, an affine choice function can thus be defined by

$$f(J_j) = \beta J_j + 1\,, \tag{14}$$

with $\beta \geq 0$. For that case the quantity α is

$$\alpha = \frac{1}{1 + \frac{N\gamma}{\beta\pi}}\,. \tag{15}$$

The purely linear case $f(J_j) = J_j$ (studied in (Kilani and Lesourne, 1995)) is obtained for $\beta \to \infty$. For this case, $\frac{1}{\beta} = 0$, the number of solutions is infinite: every $\mathbf{J} = \{J_j, j = 1, ..., N\}$ such that $\sum_j J_j = \frac{\pi}{\gamma}$ is a stable solution. This is analogous to what happens in the classical model of Blackwell's urns. If $\beta < \infty$, this degeneracy does not subsist. There is only one solution, the symmetric one $\{J_j = \frac{\pi}{N\gamma}, j = 1, ..., N\}$ (which is indeed stable: $\alpha < 1$). In any case, that is whatever the value of β is, there will be no transition.

The power law case:
Let us consider the power law case, a simple generalization of the affine case:

$$f(J_j) = (\beta J_j)^n + 1 \tag{16}$$

with $n > 0$ and $\beta \geq 0$. For this choice function α is given by

$$\alpha = \frac{n}{1 + \left(\frac{N\gamma}{\beta\pi}\right)^n} \,. \tag{17}$$

For $n = 1$ one recovers the results for the linear and affine cases: α is always smaller than 1 for $\beta < \infty$, and equal to 1 if $\frac{1}{\beta} = 0$. For $n < 1$ α is always smaller than 1 and there is no transition as found in (Kilani and Lesourne, 1995).

For $n > 1$, there exists the possibility of observing a transition, hence a bimodal situation: α is larger than 1 for $\frac{\beta\pi}{N\gamma} > (n-1)^{-\frac{1}{n}}$.
Remark: in the particular case $\beta \to \infty$, that is for $f(J_j) = (J_j)^n$, one has $f(0) = 0$. Since $n > 1$, $f'(0) = 0$, so that according to (12), the non symmetric solutions $\{0, \frac{\pi}{\gamma}\}$ are stable.

Finally in this case (16) the convexity of f is a necessary condition for observing a bimodal behaviour.

The exponential case:
The standard logit case corresponds to an exponential choice function:

$$f(J_j) = \exp(\beta J_j) \,. \tag{18}$$

In that case α is simply given by

$$\alpha = \frac{\beta\pi}{N\gamma} \,. \tag{19}$$

The symmetric point is unstable if $\alpha = \frac{\beta\pi}{N\gamma} > 1$.

The exponential choice function will be used consistently in the next sections. Presently our main conclusion is that according to the value of β with respect to a critical point defined by $\beta_c = \frac{N\gamma}{\pi}$, two behaviors are possible for buyers: fidelity to one shop for $\beta > \beta_c$ or random search among all shops for $\beta < \beta_c$.

2.4. DERIVATION OF THE LOGIT FUNCTION FROM AN OPTIMIZATION PRINCIPLE

2.4.1. *An exploration-exploitation compromise*

In the previous section we studied the qualitative behaviour that we can expect for a general choice function. The next question is then: in what sense is a given choice function efficient? One attractive feature of a choice function is that it may represent some sort of "best behaviour" with respect to some criterion. Here we will show that the logit function can be derived from an optimization strategy. In particular, we argue below that, for modeling the buyers' strategy, one can define a maximization principle, formally identical to the so-called *maximum entropy principle* considered in statistical physics (Balian, 1992) — to be briefly presented later for comparison.

Let us assume that the buyer wants to find a compromise between getting the best profit at the next transaction and keeping the best possible knowledge of a market in order to be able to make good choices in the future: the market can vary in time because of external events or because the sellers' strategies change. This requires that he visits every seller as frequently as possible (he can only get information about a seller by making transactions with this seller). If p_j is the probability of visiting seller j, $p_j = p_j^0 \equiv 1/N$ would correspond to maximum information. The proper measure of the similarity between this uniform distribution $\{p_j^0\}_{j=1}^N$ and the actual distribution $\{p_j\}_{j=1}^N$ is the entropy $\mathcal{S}$,

$$\mathcal{S} = -\sum_j p_j \ln p_j \,. \tag{20}$$

The entropy is a measure of the uncertainty in the occurrence of the events $j = 1, ..., N$. In the context of Information Theory (Blahut, 1988), it is the minimal *amount of information* (measured in bits if the logarithm in (20) is taken in base 2) required in order to code the set of events.

One may thus want to choose the p_j's by a compromise between the maximization of entropy and a maximization of immediate profit. Taking the (moving) average J_j as an estimate of the profit to be obtained from seller j, we thus maximise

$$\mathcal{C} \equiv \mathcal{S} + \beta \sum_j p_j J_j \tag{21}$$

over all possible p_j's. The quantity $\frac{\ln 2}{\beta}$ is the amount of profit considered to be equivalent to one bit of information.

Introducing a Lagrangian multiplier λ in order to impose the normalization constraint $\sum_j p_j = 1$, one finally maximises

$$\mathcal{C} = \mathcal{S} + \beta \sum_j p_j J_j - \lambda \left(\sum_j p_j - 1 \right) \tag{22}$$

Taking the derivative of $\mathcal{C}$ with respect to p_j, one gets

$$-1 - \ln p_j + \beta J_j - \lambda = 0 \tag{23}$$

which gives precisely

$$p_j = \frac{1}{Z} \exp(\beta J_j) \tag{24}$$

with $Z = \sum_j \exp(\beta J_j)$.

The logit strategy is thus obtained as a consequence of the optimization of a cost function which expresses the compromise between short term profit and preservation of information for long term profits.

2.4.2. *Link with physics and inference theory*

The exponential family of probability distributions plays a central role in statistical physics, where it is derived from the *maximum entropy principle* (Balian, 1992). More generally, the maximum entropy principle is a tool for making inferences. In fact, it has already been used in economics in order to justify the choice of an exponential distribution - see e.g. (de Palma *et al.*, 1996; Williams, 1977).

For completeness, we restate here this inference principle. One constructs a probability distribution $\{p_j, j = 1, ..., N\}$ based on some prior knowledge, in such a way that the resulting probability law does not contain more information than what can be gained from this prior knowledge. The measure of uncertainty in the occurrence of the events is given by the entropy $\mathcal{S}$ of the probability distribution, as defined in (20). If we know some mean value E of an observable quantity E_j, we estimate the p_j's by maximizing the entropy $\mathcal{S}$ under the constraint that E is given. This leads to

$$p_j = \frac{1}{Z} \exp(-\beta E_j) \tag{25}$$

where Z is the normalization constant (the "partition function"). For a physical system $T \equiv \frac{1}{\beta}$ is the temperature, and E is the energy. If one works at a given value of β (instead of a given value of E), one sees that as T goes to zero (β goes to ∞) the system will choose the states with the smallest possible values of energy. In our model of buyers' strategy, the quantity which plays the role of the energy is thus *minus* the mean profit (since the profit has to be maximised). With the maximum entropy principle one predicts the probability distribution without making any hypotheses on the dynamics. The resulting probability distribution is the best guess based on the knowledge we have about the system: the logit function can be understood as the best description of the buyer's strategy based on the knowledge of the mean profit he obtains.

The specificity of statistical physics is that the application of this *inference* principle leads precisely to the correct *physical* description - the law of thermodynamics. Clearly, there is no reason *a priori* for expecting such success in the context of economics. Nevertheless, there are several approaches tending to show that the exponential family may play a fundamental role in economics, too, as discussed in particular in (de Palma *et al.*, 1996). What we have shown in this paper is that the maximum entropy principle has an appealing "physical" interpretation in the context of the search for an exploitation/exploration compromise.

A final remark is in order. One should note that to derive a choice function from an optimization principle does *not* imply that one assumes the buyer to be aware of optimizing some criterion. An analogy can be made with living systems evolving according to past experiences. One of the main approach to modelling evolution in nature assumes the optimization of some cost function, the survival fitness. Clearly, no genetic system is aware of what is really going on, and only mutation rules can be observed at the level of individuals. Similarly, it is commonly believed that the organization of the brain is *optimally* fitted to the tasks it has to solve, through evolution and adaptation. It is not unreasonable to expect that a buyer follows some empirical rule, the rule itself being chosen according to some kind of cultural knowledge based on past experiences, possibly including those of previous generations, in such a way that, implicitly, the rule implements the optimization of some cost function.

2.4.3. *Interpretation*

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 the actual profits of the buyers or the fraction of unsold goods. 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 models where agents imitate each others' behavior through social interactions (Föllmer, 1974; Arthur and Lane, 1993; Brock and Durlauf, 1995; Orlean, 1995). Both type of models exhibit an abrupt phase transition between order for the large β values and disorder for small β's. There are two main differences:

(i) 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[2]); in our model different agents are loyal to different shops. Imitation and positive social interactions favor uniformity, while decisions based on agents' memory favour diversity.
(ii) 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 market. In imitation models, the market as a whole is organised or disorganised, even for heterogenous agents (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).

2.4.4. *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 (we can assume without loss of generality that $\pi_1 > \pi_2$). Replacing profit π in equation (5)

[2]Imitation favors uniformity, but according to whether one uses a mean field approach (all interactions being possible) as in (Arthur and Lane, 1993; Brock and Durlauf, 1995; Orlean, 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.

by π_j, the computation of $\Delta \equiv J_1 - J_2$ gives

$$\Delta - \frac{\pi_1 - \pi_2}{2\gamma} = \frac{\bar{\pi}}{\gamma} \frac{\exp(\beta\Delta) - 1}{\exp(\beta\Delta) + 1.} \equiv G(\Delta) \tag{26}$$

with

$$\bar{\pi} \equiv (\pi_1 + \pi_2)/2.$$

It is the equation $\Delta = G(\Delta)$, as given above in (26), that one has to study instead of equation (10). One finds that the critical β_c is $2\gamma/\bar{\pi}$.

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):

$$J_1 - J_2 \simeq \frac{2(\pi_1 - \pi_2)}{(\beta_c - \beta)\bar{\pi}}. \tag{27}$$

If β is above 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 Fig. 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 charged a lower price. When the most often frequented shop changes its prices, the loyalty to that shop describes the upper branch of the loyalty versus profit curve (Fig. 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 start visiting 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 attractive 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 offer a profit significantly greater than the profit offered by the well-established shop: when 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 Fig. 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.

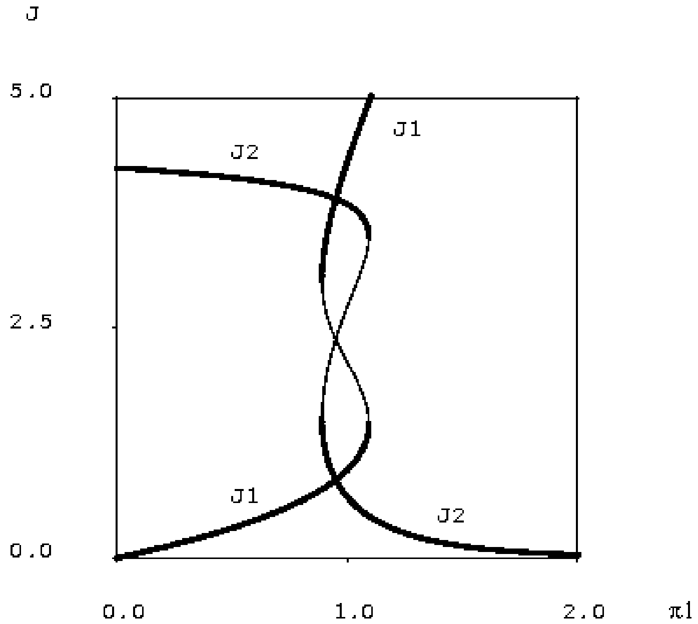

Figure 2. Hysteresis of preference coefficients. Plot of both preference coefficients versus π_1, the profit to be obtained from shop number 1 when π_2, the profit to be obtained from shop number 2 is held equal to 1 ($\beta = 0.5$ and $\gamma = 0.2$). The thick lines correspond to stable equilibria for both preference coefficients, J_1 and J_2, and the thin lines around $\pi_1 = \pi_2 = 1$ to unstable equilibria. In the region of three solutions, the larger value of J_1 is reached from initial conditions when J_1 is already large. Thus if π_1 is decreased from above one, J_1 is kept large (and J_2 is kept small), even when π_1 becomes less than π_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).

3. Results

3.1. INDICATORS OF ORDER

We next proceeded to run a number of numerical simulations of our model. This first enabled us to check whether the theoretical results obtained from the mean field approximation were consistent with those obtained by running the discrete stochastic process as described by equation 2 and 3. Secondly, as discussed in the next section, it allowed 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 vectors of buyers J_{ij}'s is harder to monitor. We used two methods to do this.

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

$$y_i = \frac{\sum_j J_{ij}^2}{(\sum_j J_{ij})^2}. \tag{28}$$

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 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 loyalty of every single buyer. Figures 2a and 3a, for instance, 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.

3.2. 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' stocks 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 following Figs. 3 and 4 correspond to a memory constant $\gamma = 0.1$. The critical non-linear parameter corresponding to a unitary profit is then $\beta_c = 0.3$ (Eq. 13). Initial J_{ij} were all 0. Depending on the value of the non-linear parameter β, the two predicted behaviours, order and disorder, are observed.

3.2.1. *Disorganized behavior*

For low values of the non-linear parameter β, buyers never build up any loyalty. This is observed in Fig. 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 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 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 some shops with supplies were not visited, and this resulted in losses for their owners. (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 Sec. 5.1. However, the observed decrease in profit for buyers and sellers is generic.)

As seen on the simplex plot, even at time 50, agents are still scattered around the barycenter of the triangle, an indication of 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.

3.2.2. *Organized behavior*

In sharp contrast to the above, the same analysis performed with $\beta = 2\beta_c$ shows a great deal of organisation (see Figure 3).

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 particular 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 times the number of daily transactions. Because buyers have not changed shops during the last 100 days, sellers have learnt to purchase the exact quantity needed to satisfy all their buyers, and they had no losses.

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.

3.2.3. *Heterogeneity of buyers*

Let us recall at this stage that in the case of real markets, we expect a mixture 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 results described above. Organized or disorganized behavior is here a property of buyers, not a property of markets.

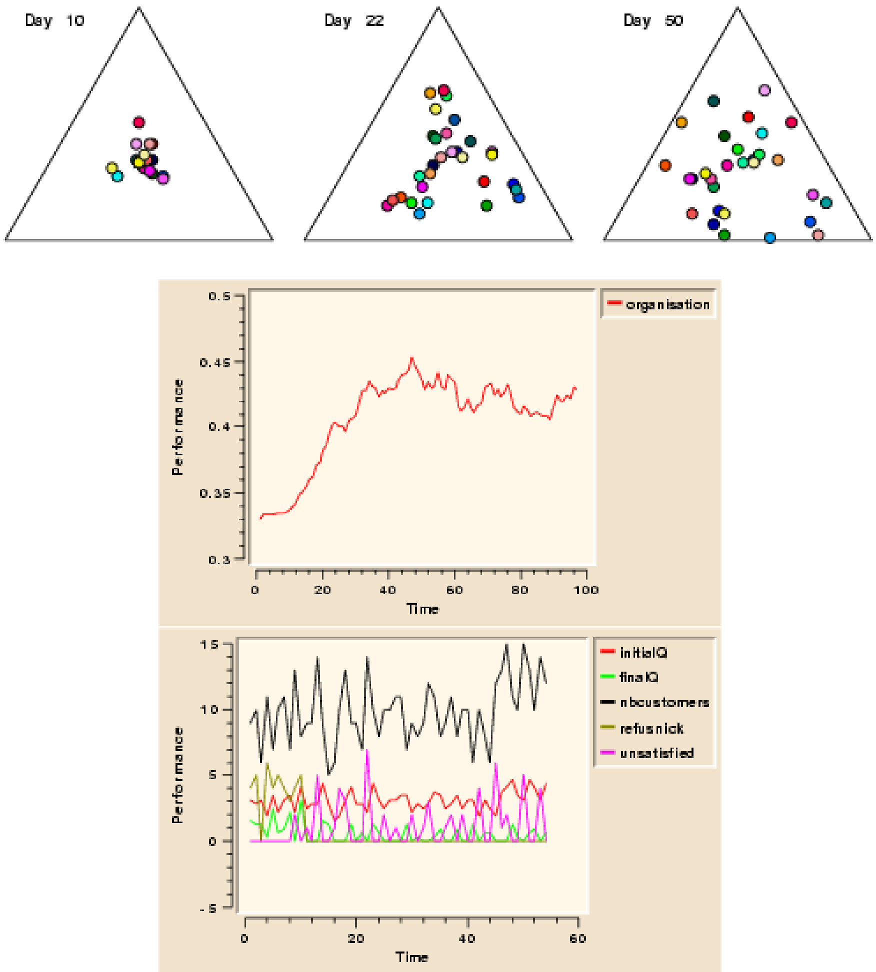

Figure 3. Charts for the disorganized regime (30 agents visiting 3 shops, when the learning parameter $\gamma = 0.1$ and $\beta = 0.15\beta_c$). The first three graphs monitor market organization by simplex plots taken at time 10, 22 and 50. The fourth graph shows a time plot of the order parameter y (vertical axis: $[0.3, 0.5]$). The last graph gives a record of shop 1. The time charts display the initial and the final stocks, the number of customers, the number of customers refusing the proposed price (see Sec. 5.2), and the number of unsatisfied customers who did not manage to buy anything.

3.3. BEYOND THE MEAN FIELD APPROXIMATION

The results of the mean field approach have been 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. The variables γ and π thus have to be small, which is true for the simulation results given in Figs. 2 and 3. One of the features noticed by observing

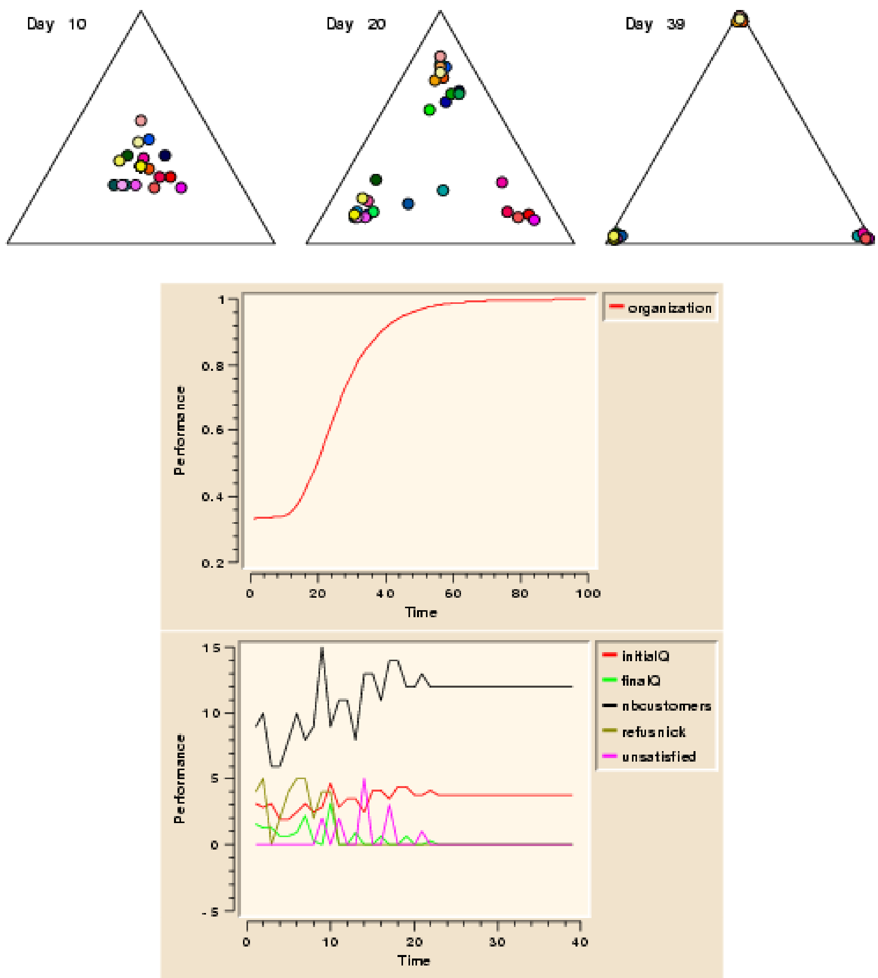

Figure 4. Charts for the organized regime (30 agents visiting 3 shops, when the learning parameter $\gamma = 0.1$ and $\beta = 2\beta_c$). All charts and notation are the same as for Fig. 2, except for the scale of the order parameter plot (y). Starting from all J_{ij} equal to 1, all representative circles move to the triangle corners representing the preferred shops. The order parameter y varies from 0.33 (equal interest for all shops) to nearly 1 (strong preference for only one shop). Due to organization, fluctuations of performance diminish in time.

the motion of individual buyers on the simplex plots on-line is that agents sometimes move "backward" towards shops which are not the shops that they prefer in the ordered regime. But since most of the time they move towards their 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 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 γ and π, a small change in the relevant parameters π or γ results in a switch from a no-divorce regime to a divorce regime. Divorces are observable on the simplex plots on-line and also by examining the evolution of the number of customers of a given shop as a function of time: "infidelities" appear as peaks and "divorces" as steps.

4. 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. 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 Sec. 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 so many parameters which are not directly observable. But the search for genericity is based on the conjecture[3] that the large set of models which share the same generic

[3]This conjecture, which is basic in the dynamic modeling of complex systems, rests

properties also includes the "true" model of the real system itself.

4.1. 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 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[4]:

$$p(q) = \frac{b}{q+c} \,. \tag{29}$$

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

$$\pi_b = q\left(\frac{b}{q+c} - p\right), \tag{30}$$

where p is the price asked for 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:

$$q = \sqrt{\frac{bc}{p}} - c \,. \tag{31}$$

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

$$\pi_s = q\,(p - p_a) = \left(\sqrt{\frac{bc}{p}} - c\right)(p - p_a)\;, \tag{32}$$

with respect to the price p that they charge to the buyers[5], where p_a is the price at which the sellers themselves purchase the goods.

on rigorous proofs about specific systems, such as classes of universality in physics or structural stability in mathematics.

[4]This particular choice of the function $p(q)$ is of no importance, it allows to run simulations and to make comparisons between the different scenarios. Any monotonic decreasing function would do for our model.

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

$$p^3 - \frac{b}{4c}(p + p_a)^2 = 0, \tag{33}$$

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

4.2. TWO SESSIONS

The one-session model described in Sec. 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 main motivation for refusal in standard search models. An alternative explanation is that customers refuse deals 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 all the quantity they bring to the market. If they were able to predict precisely how many customers will visit their shop and accept this price, they would know exactly how much to supply. But when their forecasts are not perfect, they may not have excess stocks at the close of the market. It might therefore be better for them to sell at a lower price 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 of $1-\epsilon$. 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:

1 During the morning, sellers maximize their profit and sell at a price p_{am} equal to p. Buyers visit one shop in the morning.
2 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 commodity corresponding to the number of customers they expect times q, plus some extra, in case they have more customers than expected. The profit they expect from this additional amount is that obtained by satisfying new customers or unexpected former customers who might appear.

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 according to their expectations of the average afternoon profit with respect to what they would get by buying in the morning, which they know from equation 30. The average afternoon profit is estimated from their past history of afternoon profits. In the simulations we used a simple quadratic fit of the afternoon profits as a function of morning prices. But 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 were smaller than in the morning. We discovered that even with their primitive prediction abilities, buyers soon (say after 50 time steps) realised that they would do better to accept the morning offers. Further investigations of 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. The estimated[6] β_c is at most 20 percent 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 instance as soon as γ is larger than 0.1, as opposed to γ larger than 0.3 for the one session model. This is because when an infidelity occurs, since a buyer has a much better chance of making a higher afternoon profit with a new shop that has extra supplies, she then takes larger steps across the simplex.

4.3. SELLERS' INITIAL STOCK

We mentioned earlier that the sellers may want to adjust their initial stock 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 $n_b \in [0, n]$. By maximizing the expected payoff $(E(\pi_s))$ with respect to $\hat{n}$, the sellers determine the optimal initial stock $\hat{Q} = \hat{n} \cdot q$ by:

$$1 - \int_0^{\hat{n}} f(n_b) dn_b = \frac{p_a}{p} \qquad (34)$$

The rule defined by Eq. (34) is optimal only for short-run considerations, when sellers assume that every market day is a one-shot game. It prevents the strategic use of stock, when a seller tries to gain additional loyal customers by having extra units for unexpected customers.

[6]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 $y = 0.5$.

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 stock whenever they observe fluctuations in the number of visits. The extra stock at time t is computed according to

$$\alpha(t) = (1-\epsilon)\cdot\alpha(t-1) + \epsilon \cdot var(n_b)\,, \tag{35}$$

where ϵ is small and $var(n_b)$ is the variance of the number of buyers computed from 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 were variations of β_c, the critical threshold for order, in the 10 % range. The existence of two dynamic regimes persists.

Another possible refinement would be to improve the predictive ability of the seller of 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).

4.4. 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, Sec. 3.2 above on hysteresis already gives us a clue 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 then should not distroy order. For the parameter values of Fig. 2, one shop could increase its prices to 19 % more than the other shop before loosing its customers.

We ran simulations with the morning price $p(t)$ fluctuating in each shop with an auto-regressive trend towards the morning price computed to maximize profits p. 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:

$$p(t{+}1) = \eta(t)\left[p(t){-}\lambda\Big(p(t){-}p\Big){-}\mu\frac{r_n}{n_j}\right],\ \text{ with } \eta(t) \overset{iid}{\sim} U[1{-}\epsilon, 1{+}\epsilon],\ \epsilon \in [0,1], \tag{36}$$

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

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

5. 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 700
No of sellers 40
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 is 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 and the alternative for a buyer to purchasing his optimal good is, in fact, to purchase some, in his view, 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 buyers preferences for sellers. In particular, one immediately 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 would visit several sellers, which is consistent with our model. The data will be analysed in this section only in terms of the organisation issue. Other aspects, such as price dynamics showing persistent dispersion, were analysed in Kirman and Vignes (1991) and Härdle and Kirman (1995).

5.1. TESTING OUR MODEL

A first step in comparing our theory with the empirical data is to check whether individual buyers displayed ordered or disordered behaviour during those three years. Since the classical approach to agent behaviour predicts searching 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 behaviours of these traders to some difference between their characteristics and those of other buyers.

TABLE 1. Loyalty in Cod, Whiting, and Sole Market

	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%

For the first step,we consider statistics for cod, whiting and sole transactions in 1989, see table 1.

Since we are interested in loyalty issues, we have 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. In the whiting and sole market buyers are also loyal, but to a lesser degree: more than one half of them[7] buy more than 80% from one seller. Hence, there are large fractions of loyal buyers in all three markets.

For the second step, recall that our theory relates loyalty to the parameters β (discrimination rate) and π/γ (cumulated profit). β probably varies from buyer to buyer, but we have no direct way to test it *apriori*. However, π/γ is strongly and positively related to the monthly purchases of buyers, and we therefore use the latter as a proxy variable.

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

Figure 4 summarises the loyalty of buyers in terms of the relative frequency of visits to their favorite seller as a function of their monthly purchase of cod . One may observe that loyalty is in general is high, that a number of buyers visit only one seller, and that a cubic fit shows that loyalty increases with monthly purchase. All three features are consistent with our theory, and contradict a random search behavior for all buyers. We also used standard statistical tests to check the idea that the population of buyers should exhibit two types of behaviour. We divided the buyers of cod into two groups. We have chosen as our dividing criterion a total purchase of two tons of cod over 36 months. We calculated the fraction of transactions with the most often visited seller and found 0.85 for the big buyers and 0.56 for the 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 $P1{=}P2$ at all levels of confidence.

5.2. 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 is often the case 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 shall see.

One alternative explanation that has been offered is that contractual arrangements develop between buyers and sellers. Discussion with Marseille sellers reveals that they do not offer fish for specific customers but that "he (the buyer) comes here because he knows that he will find the sort of fish that he requires". Similarly, the buyers do not order fish; they make the statement such as "I go there because he has the fish that I want". 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 and if this happens over a period and some quantity of fish remains unsold, the seller will simply readjust his supply of fish accordingly. In connection with the points now discussed, it is peharps worth emphasizing that the basic theory of this paper has been elaborated in the light of conversations with market participants who often were able to explain certain features of the data.

At the same level of generality, 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

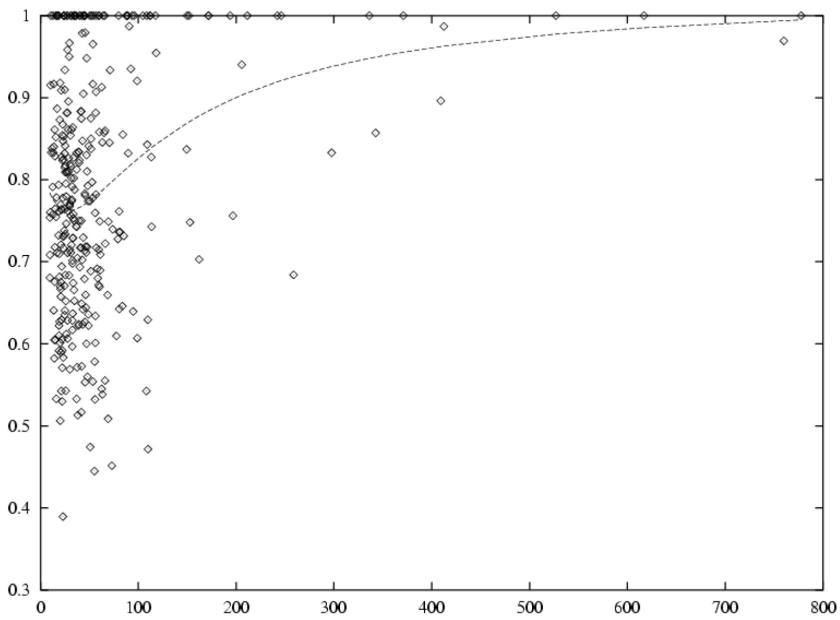

Figure 5. Each dot is an empirical evidence from Marseilles fishmarket representing a buyer's loyalty to his favorite seller (relative frequency of visits) as a function of his monthly purchase of cod in kilograms. Low purchases correspond to unfrequent buyers, who generally visit once a week, while large purchase are those of buyers who visit nearly everyday the market opens. The continuous line is a cubic fit which shows that loyalty increases with monthly purchase.

note that the two hypotheses are not mutually exclusive: even if niches were an important factor, one would still have to explain why sellers choose a niche stategy rather than selling a large choice of fish. Loyalty of buyers might be a pre-condition for the profitability of "niches". Anyway, direct examination and surveys show that even though certain sellers specialise in serving supermarkets or institution cafeterias, almost 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. Another check for the existence of niches is clustering analysis according to average prices and quantities sold by sellers. Sellers are considered as members of the same cluster, when their distribution of prices and quantities significantly overlap. We did find two clusters of cod sellers, low cost bulk sellers (5 sellers) and expensive 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 observed organisation could be based on a more active role of sellers. A possible test of the necessity of extra hypotheses implying that loyalty is due to sellers' behaviour is to check whether 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[8] 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.

6. 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. Here "Order" means the establishment of stable trading relationships over many periods in which the market is open.

In the simplest model, we have shown analytically that an ordered regime appears whenever the agents' discrimination rate between shops divided by the number of shops is larger than the reciprocal of the discounted sum of their profit. When an individual's parameters put him into 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. 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 $\beta_c = n\gamma/\pi$, 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 fishmarket indeed shows the existence of a bimodal distribution of searchers and loyal

[8]The fact that the correlation is stronger for sellers, with much less noise than for buyers, is due to the fact that sellers' statistics involve more averaging than buyers statistics.

buyers, and the positive correlation of loyal behavior with the frequency of transactions.

When more realistic assumptions are introduced, such as the adaptive behavior of sellers, fluctuations in prices, and later sessions with lower prices 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 strongly 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 terms than disorder, and it is therefore of considerable economic interest 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, Kenneth Arrow, Larry Blume, Paul David, Steven Durlauf, Fang Fang Tang, Douglass North, Ariel Rubinstein and Nick Vriend for helpul comments, and Bernard Derrida and Jean Vannimenus for helpful discussions and important suggestions. This work was started during a visit by GW and AK to the Santa Fe Institute which we thank for its hospitality.

References

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

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

Arthur, Brian W. and Lane, David .A, (1993), Information contagion, *Structural Changes and Economic Dynamics*, **4**, 81-104.

Arthur W.B., Holland J.H., LeBaron B., Palmer R.G., and Tayler P.(1996), Asset Pricing Under Endogenous Expectations in an Artificial Stock Market, Santa Fe Institute Working Paper 96-12-093, to appear in *The Economy as an Evolving, Complex System II*, W.B. Arthur, D. Lane and S.N. Durlauf. eds.

Balian R., (1992), *From Microphysics to Macrophysics*, Springer.

Blahut R. E., (1988), *Principles and Practice of Information Theory* Addison-Wesley, Cambridge MA.

Blume L. (1993), The statistical mechanics of strategic interaction, *Games and Economic Behaviour*, **5** pp. 387-424.

Brock W.A. and Durlauf S.N. (1995), Discrete choices with social interactions I: Theory, Working Paper 95-10-084, Santa Fe Institute, Santa Fe, NM.

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 *Chance and Matter* , Ed. by Souletie J. Vannimenus J. and Stora R., 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 *The New Palgrave: A Dictionary of Economics*, J. Eatwell, M. Milgate and P. Newman eds., Macmillan, London, pp. 273-79.

Fisher F. M. (1973), Stability and competitive equilbrium in two models of individual price adjustment, *Journal of Economic Theory*, Vol 6.

Föllmer H. (1974), Random economies with many interacting agents, *Journal of Mathematical Economics* **1**, 1, March, pp. 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*, vol. 67.

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*, Vol. 71., pp.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, Alan P. and A. Vignes, (1991), Price dispersion. Theoretical considerations and empirical evidence from the Marseilles fish market, in *Issues in contemporary economics*, ed. K. G. Arrow, Macmillan, London.

Kirman Alan P., (1993), Ants , rationality and recruitment, *Quarterly Journal of Economics* vol. 108, pp. 137-156

Kirman, Alan 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 Jacques (1992), *The Economics of Order and Disorder*, Clarendon Press, Oxford.

Orlean A. (1995), "Bayesian interactions and collective dynamics of opinions: herd behavior and mimetic contagion", *JEBO*, 28, pp. 257-274.

de Palma A., Kilani K. and Lesourne J., (1996), Network externalities and the Polya-Logit model, Preprint of the Laboratoire d'Econométrie, Conservatoire National des Arts et Métiers, Paris.

Weisbuch G., Kirman A. and Herreiner D., (1997), Market organisation, in the proceedings of the workshop *Self-Organization and Evolutionary Economics: New Development* (Paris, 1996), Economica (to appear).

Williams H.C.W.L., (1977), On the formation of travel demand models and economics evaluation measures of user benefit, Environment and Planning A **9**, 283-344