The Economic Journal, **110** (*April*), 411–436. © Royal Economic Society 2000. Published by Blackwell Publishers, 108 Cowley Road, Oxford OX4 1JF, UK and 350 Main Street, Malden, MA 02148, USA.

MARKET ORGANISATION AND TRADING RELATIONSHIPS*

Gérard Weisbuch, Alan Kirman and Dorothea Herreiner

In this paper we give a theoretical model of buyers' behaviour on a market for a perishable good where no prices are posted. We show that if buyers learn from their own previous experience there is a sharp division between those who learn to be loyal to certain sellers and those who continue to 'shop around'. This feature remains in more general models which are simulated and is consistent with empirical data from the Marseille fish market.

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. 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 most standard theoretical models. In a Walrasian equilibrium, for example, the following questions are left 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?

One of the objectives of this paper is to examine a situation in which individuals set prices and the way in which those who wish to buy, at those prices, become matched with those who wish to sell. In standard search models (see e.g. Diamond (1989)),[1] 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. Yet many markets are ones on which individuals trade frequently with each other. This is particularly true for markets for perishable goods. Furthermore, in this case, since sellers cannot hold inventories, they only supply the quantities they expect to sell during one session. The essential risk, in our stylised context, for a buyer is not that of paying too higher a price but

* We thank Derek Smith for help with the Tk/Tcl interactive display of simulation results, Rob Deboer for the use of his GRIND software, Annick Vignes who originally collected the Marseille fish data, Paul Pezanis and Marlene Mueller for their help in its analysis, Kenneth Arrow, Larry Blume, Paul David, Steven Durlauf, Fang Fang Tang, Douglass North, Ariel Rubinstein and Nick Vriend for helpful 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 (with the help of a NATO Collaborative Research Grant CRG 95 1261) which we thank for its hospitality.

[1] For more elaborate models, see e.g. Fisher (1973) and Lesourne (1992).

rather of not being served at all. Hence, rather than gathering information about different sellers at each session, a buyer may find it worthwhile to use the experience gained from transactions with sellers during previous sessions.

In our model, trading relationships develop precisely because buyers learn, in this way, about the value of trading with particular partners. Such 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 examine the extent to which agents in a simple model of such markets will learn their way to establishing trading links and to characterise the conditions under which this happens.

We emphasise that the intrinsic uncertainty in our model is due to the behaviour of the buyers. There is no exogenous uncertainty about the parameters about which agents learn. Thus our model differs from those of (Felli and Harris, 1996; Bergemann and Valimaki, 1996; 1997), in which agents learn about some exogenous random variable and strategic behaviour influences such learning. Our model, on the other hand, uses a simple rule of thumb for learning with no strategic thinking. This would seem to be well adapted to a situation with no *a priori* stable distribution of the quantities available from each seller. This is the case for buyers in the Marseille wholesale fish market from which our empirical evidence is taken.

In this paper, we limit our attention to 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, and no prices are posted. In such markets agents have to rely on their own information. We will therefore develop a model which seeks to explain some of the phenomena that characterise this particular type of market in which learning is from individual past experience. Our emphasis is on the buyers' side of the market but the conclusions that we draw will be shown to hold even when different sellers behaviour is introduced. 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. The buyers choose their sellers according to the information gathered during previous purchases. This model is solved analytically using the 'mean field' approximation. The theory predicts that two distinct types of behaviour for the agents should be observed according to the values of their learning and choice parameters: some agents should remain loyal to one selected seller, 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. Finally, we verify that the theoretical predictions from our model are consistent with empirical data from the wholesale fish market in Marseille.

1. The Simplest Model

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

1.1. *Basic Assumptions*

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

- Customers choose one seller every day taking into account their memory of previous transactions. As long as the seller has a supply of the good, a customer purchases a quantity $q_i(t)$ and this, in turn, generates a profit of $\pi_i(t)$. Whether the customer is served when he visits the seller depends on which seller j is visited at time t, how many people bought from that seller before, and how much endowment the seller had at the beginning of the day.[2]
- Since the good is perishable and therefore cannot be stored from one day to the next, 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 1, 2 and 3. More realistic assumptions will be made in Section 4.

1.2. *Preference Coefficients, Learning and Choice Probabilities*

Our model seeks to explain trading relationships. Therefore, our assumptions about how buyers choose which seller to visit are crucial. These assumptions[3] 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 seller each day. The basic assumption of the model is that his present choice of which seller to visit is based on his previous experience. His decision rule is therefore a mapping from the time series of the transactions he had had with different sellers and the profits associated with the transactions,[4] $\mathbf{I}(t)$, to the unit simplex Δ_m, where m is the number of sellers:

[2] In the simplest model, customers visit only one seller and thus have only one chance to get served on each day.

[3] 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).

[4] In wholesale markets for perishable goods, profits are more pertinent to retailers than prices since retailers face uncertainties related to prices and available quantities. The relation between profits, prices and available quantities is discussed in Section 4.

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

A point is 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 seller. The profits buyer i made when buying from seller 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:

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

In other words, at each time step, all preference coefficients are discounted at a constant rate, and the preference coefficient for the seller with which a transaction occurs is increased by the profit made from that seller. 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 sellers do not necessarily have stationary characteristics in terms of the profits that they offer, because of possible changes in prices for many reasons, in their initial endowment and in their number of customers.

Buyers then map these preference coefficients into the choice of a seller. One deterministic way to do so is to choose the seller with the best record, that is the seller with the highest $J_{ij}(t)$. This would amount to mapping the $\mathbf{J}(t)$ into one of the apexes of Δ_m. However, by doing this, the buyer would become a captive of the selected seller who would then be in a position to diminish the buyer's profit and to increase his own profit by changing prices. The seller 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 reasonable rule-of-thumb for the buyers would be a balance between the deterministic choice in favour of those sellers who 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 characterises 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:

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

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[5] 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 maximising the weighted sum F_i of two terms; one of which favours immediate profit:

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

G_i is approximately the expected discounted sum of profits. The other term favours search. To maximise the information gained during visits, buyers should maximise the Shannon entropy[6] of the distribution of search probabilities:

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

The function F_i to be maximised 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 sellers. The larger β the more important (short-run) payoff concerns. Setting the derivatives of F_i with respect to P_i equal to zero under the constraint that the sum of the probabilities is 1 gives (3).

2. 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. This is only an approximation, but is often convenient to obtain at least a qualitative understanding of the behaviour of the system.[7]

The model can be solved 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

[5] 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).

[6] 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 minimised if $P_{ij} = 1$ for one j and the other $P_{ij} = 0$.

[7] The limits of applicability of Mean Field Theory are an advanced topic in Statistical Physics discussed for instance in Brout (1965). Some ideas are discussed in an economics context in Aoki (1996). Even without a full discussion (well beyond the scope of the present paper), we might expect that the Mean Field Theory applies reasonably well to the present case with no *a priori* connection structure. Even in those cases where its predictions are quantitatively inaccurate, the Mean Field Theory is still good enough to predict the transitions and some dynamical properties in their neighbourhood. This is sufficient for our purpose which is to solve a simple model to gain some insight into the structure of the market and to calibrate the simulation results.

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 consists in replacing the π_{ij} by its expected value $\langle \pi_{ij} \rangle$, thereby transforming the stochastic differential equation into a deterministic differential equation.

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

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

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

the fraction represents the probability that buyer i visits sellers j; $\text{Prob}(q_i > 0)$ is the probability that seller j still has goods to sell when the buyer comes to him, 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 (7) equal to zero.

Let us consider the simplest case of two sellers and to further simplify computation, let us suppose, for the time being, that $\text{Prob}(q_i > 0) = 1$, which means that buyers always find what they require at the seller they visit. (If this were always the case, there would be no rationale for the learning and choice algorithm described in Section 1. We take here $\text{Prob}(q_i > 0) = 1$ only as a limiting case which allows us to obtain analytical results which are then checked against numerical simulations in Section 3 where a buyer may choose a seller who has no more stock, i.e. with $\text{Prob}(q_i > 0) \leqslant 1$.

2.1. *The Order/disorder Transition for the Symmetric Case*

To make computations easier, let us first suppose for the moment that profits from both sellers are equal to π (see the next Section for unequal profits). The equilibrium relations are in this case:

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

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

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

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

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]

$$\beta < \beta_c = \frac{2\gamma}{\pi} \tag{12}$$

there is only one stable solution $\Delta = 0$ and $J_1 = J_2 = \pi/2\gamma$. The average cumulated profit, or preference coefficients, J_j are small and equal. A buyer visits both sellers approximately half the time, switching at random between the sellers. We call such a regime *disordered* or *disorganised*.

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 seller with high probability and frequency (high preference coefficient) and the other seller with low probability and therefore rarely (low preference coefficient). We call such a regime *ordered* or *organised*; buyers are loyal.

The transition from the disordered to the ordered regime 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 $\Delta = 0$):

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

as can be seen in Fig. 1 obtained by solving (11).

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

$$J_j = \frac{\pi}{\gamma} \frac{\exp(\beta J_i)}{\sum_k \exp(\beta J_k)}. \tag{14}$$

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

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

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

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

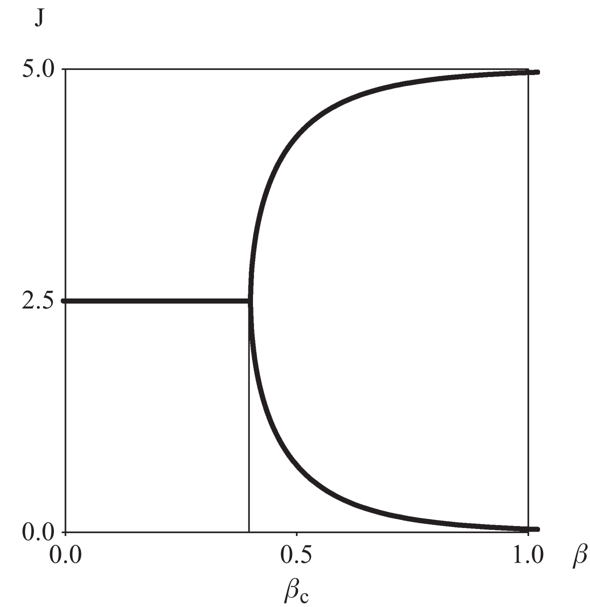

Fig. 1. *The order/disorder transition in* β.
Plot of both equilibrium preference coefficients versus the discrimination rate β. Below the transition rate β_c, preference coefficients are equal, but they rise or plummet sharply when the discrimination rate β increases above the transition. When profits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

satisfies (14). The symmetric fixed point is an attractor if and only if the right hand side of (14) has a slope smaller than one. This condition is easily checked since the derivative of the denominator of the RHS of (14) is zero at the symmetric point, due to (15) and equality of the derivatives with respect to the J_k. 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 sellers with equal likelihood, or there are m stable stationary points (if $\beta > \beta_c$), where a buyer is loyal to one of the m sellers.[10]

The above analysis shows that as long as the mean field approximation remains valid, the qualitative behaviour 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, the numbers of sellers m, the non linear parameter β and the profit from a transaction π, only changes the scale of equilibrium variables such as the current profits of the buyers or the fraction

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

of unsold endowments. The time scale of learning depends on γ: order, when achieved, is reached faster for larges 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 seller most of the time, and searchers with $\beta < \beta_c$, who wander from seller to seller. 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 5. Furthermore, because of the sharp change in behaviour when β goes across the transition, the distribution of behaviour 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' behaviour through social interactions (Föllmer, 1974; Arthur and Lane, 1993; Brock and 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 sellers. Imitation and positive social interactions favour uniformity, while decisions based on agents' memory favour diversity.
- In our model heterogeneity of buyer parameters results in having two classes of behaviour, 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 in the presence of heterogeneity of agents.[12]

The most important point to emerge from this analysis is that, even when profits from different sellers are the same, there is an abrupt transition from disloyal to loyal behaviour as the non linear parameter β passes a critical value. This is due to the fact that we have chosen a non linear learning mechanisms which means that the self reinforcing process suddenly becomes much stronger. This sort of phase transition is well known in physics in such non linear systems.[13]

[11] Imitation favours 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), Orléan (1995), or Markov random fields (interactions restricted to some neighbourhood) as in Föllmer (1974), one observes global or local order. All agents made 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.

[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 may coexist.

[13] For a discussion of other learning processes, see Nadal *et al.* (1998).

2.2. *Sellers Offering Different Profits and Hysteresis*

Up to this point we have considered a situation in which sellers propose the same prices, resulting in equal profits for buyers and we now have to check whether similar results apply when profits differ. Let us come back once more to the simple case of two sellers 1 and 2, and now suppose that they offer different prices and hence different profits π_1 and π_2. This is possible since in our market no prices are posted. Replacing profit π in (9) and (10) by π_1 and π_2 respectively, (11) becomes

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

The equation in Δ has once more either one or three solutions, depending now on the value of two reduced parameters, the relative profit difference $r=(\pi_1-\pi_2)/(\pi_1+\pi_2)$ and the ratio between β and $\beta_c=4\gamma/(\pi_1+\pi_2)$. The geometric interpretation of (18) in Δ is still the intersection of a straight line (the left hand-side of (18)) with the hyperbolic tangent (its right hand-side). Equation (18) amounts to shifting the left-hand side of (11) by $(\pi_1-\pi_2)/(\pi_1+\pi_2)$ and replacing the uniform profit by its average for the two sellers.

If β is above[14] β_c, the three intersections remain as long as the difference in relative profits r is not too large. The frontier between the two regimes, three versus one solution, in the β, r parameter space is described by the inequality:

$$\left(\frac{3r}{2}\right)^2<\left(1-\frac{\beta_c}{\beta}\right)^3 \tag{20}$$

in the neighbourhood of $\beta=\beta_c$.[15]

As we could have easily guessed, the existence of two regimes separated by an abrupt phase transition does not depend upon the simplifying assumption of symmetry between the two sellers. The above analysis can be generalised to a larger number of sellers and exhibit transitions in the number of solutions with respect to relative profit differences among sellers and average profit.

In the organised regime, which of the two asymmetric intersections is actually reached by the learning dynamics depends on initial conditions. The following analysis of the hysteresis effect can give us some clues about the consequences of changes in the profit offered by the different sellers.[16]

Fig. 2 is represents the preference coefficients J_1 and J_2 obtained by solving (18) by numerical methods (De Boer, 1983) for $\pi_2=1$, $\beta=0.5$, $\gamma=0.2$ and π_1 varying from 0 to 2. By following the evolution of preferences on the two

[14] 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):

$$J_1-J_2\simeq\frac{4(\pi_1-\pi_2)}{(\beta_c-\beta)(\pi_1+\pi_2)}. \tag{19}$$

Compare with footnote 9.

[15] Expression obtained by developing the hyperbolic tangent in series for small values of $\beta\Delta/2$.

[16] And hence about possible strategies for the sellers.

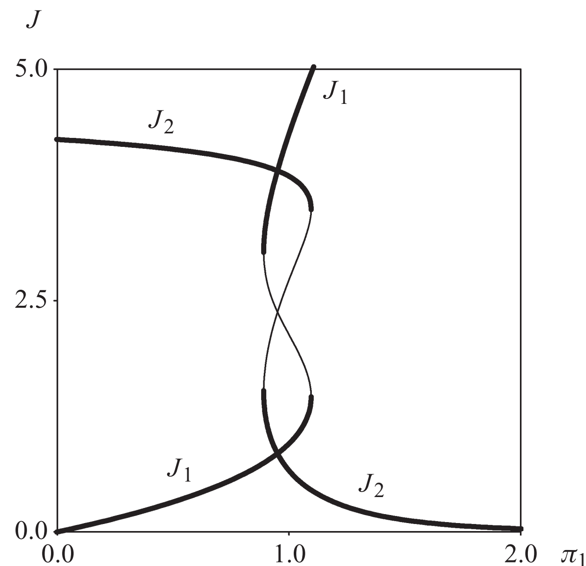

Fig. 2. *Hysteresis of Preference Coefficients.*
Plot of both preference coefficients J_1 and J_2 versus π_1, the profit to be obtained from seller number 1 when π_2 the profit to be obtained from seller number 2 is held equal to 1. ($\beta = 0.5$ and $\gamma = 0.2$). The thick lines correspond to stable equilibria for both preference coefficients, J_1 and J_2, and the thin lines to unstable equilibria (when $\pi_1 \simeq \pi_2$). 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 $\pi_1 = 0.89$. Similarly, the high J_2 attractor which exists for low values of π_1 can be maintained up to $\pi_1 = 1.095$. (The figure was drawn using GRIND software, De Boer (1983)).

branches of Fig. 2, we can see that buyers can remain loyal to a seller asking for a higher price (which results in a lower profit for the buyer), provided that they became attached to this seller when he asked a lower price. When the most often frequented seller changes his price, the loyalty to that seller describes the upper branch of the loyalty versus profit curve (Fig. 2). 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 fidelity and visit the other seller. But, if the first seller reverses its high price/low buyer profit practice 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 seller that offered good deals in the past, even though the current deals they offer are less interesting than those now offered by other sellers. A consequence of this phenomenon, is that in order to attract customers who are loyal to another seller, a challenger has to offer a profit significantly greater than the profit offered by the well established seller: 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 Fig. 1 are vertical (i.e. not when profits are equalised!). In other words, simple myopic economic rationality (i.e. choosing the seller offering the best deal) is not ensured in the region where hysteresis occurs.

3. Results

3.1. *Indicators of Order*

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

Our principal results concern the extent to which the market organises itself.

Simulations generate a large amount of data about individual transactions such as which seller was visited, how much was purchased, and agents' profits. However, the organisation process itself, involving the dynamics of the buyers' J_{ij} vectors of cumulated profit from each seller, is harder to monitor. We used two methods to do this.

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

$$y_i = \frac{\sum_j J_{ij}^2}{(\sum_j J_{ij})^2}. \tag{21}$$

In the organised regime, when the customer is loyal to only one seller, y_i is close to 1 (all J_{ij} except one being close to zero). On the other hand, when a buyer visits m sellers with equal probability, y_i is of order $1/m$. More generally, y_i can be interpreted as the inverse number of sellers visited. We were particularly interested in y, the average of y_i over all buyers.

Secondly, when the number of sellers is small, 2 or 3, a simplex plot can be used to portray the evolution of the loyalty of every single buyer. The first graphs of Figs 3 and 4 display simplex plots of a simulation at different steps. Each agent is represented by a small circle of a specific shade. This represents the agent's probabilistic choice, i.e. the probability distribution over the 3 sellers (corresponding to the 3 apexes of the triangle). Proximity to a corner is an indication of loyalty to the seller corresponding to that apex. Agents represented by circles close to the centre search all sellers with almost equal probability.

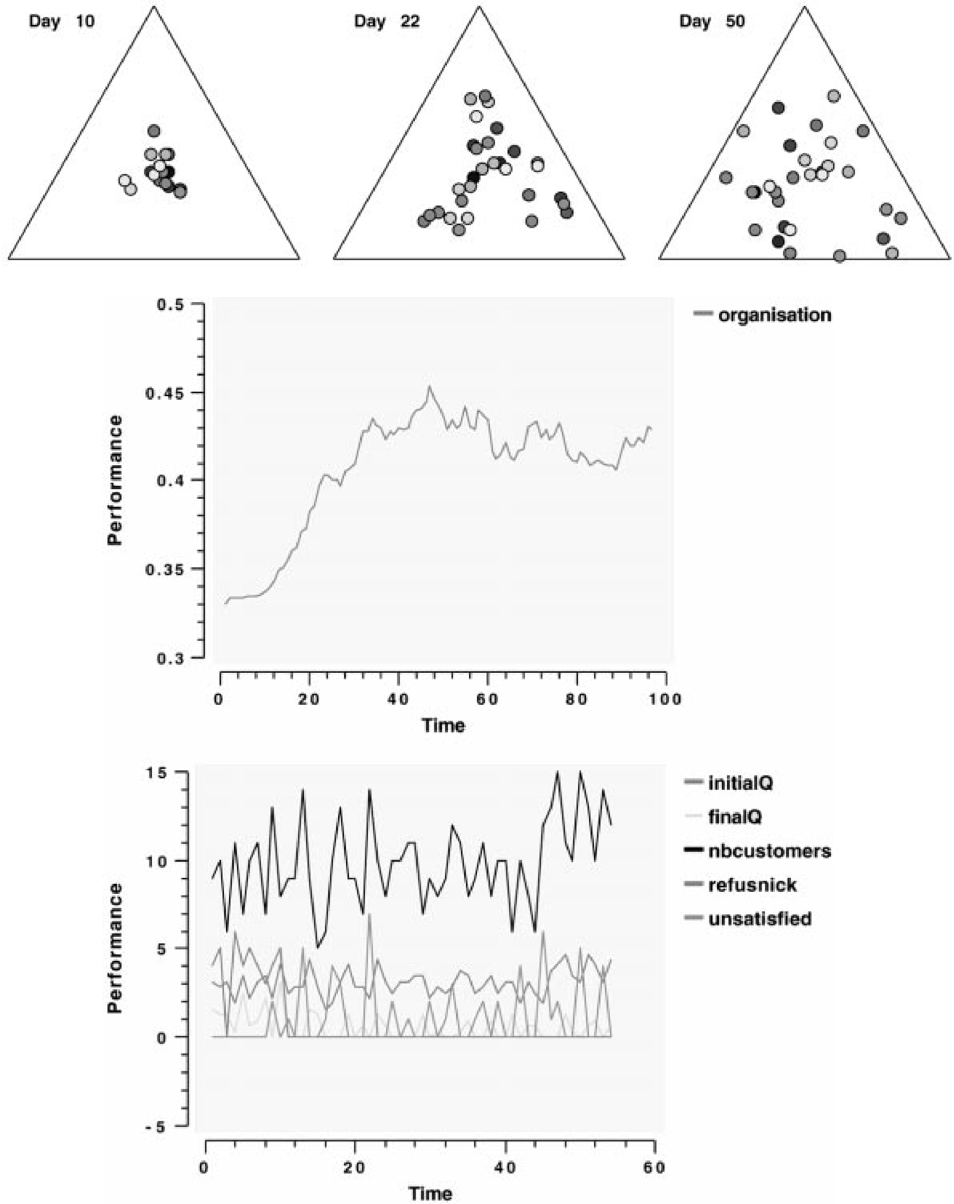

Fig. 3. *Charts for the Disorganised Regime.*
(30 agents visiting 3 sellers, with the discount parameter $\gamma = 0.1$ and the non-linear parameter $\beta = 0.15\beta_c$). The first three graphs show market organisation by simplex plots at times 10, 22 and 50. They show that no organisation takes place. 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 seller 1. The time charts dislay the initial and the final endowment, the number of customers, the number of customers refusing the proposed price (see section 3.2), and the number of unsatisfied customers who did not manage to buy anything. Fluctuations in the market do not decline over time.

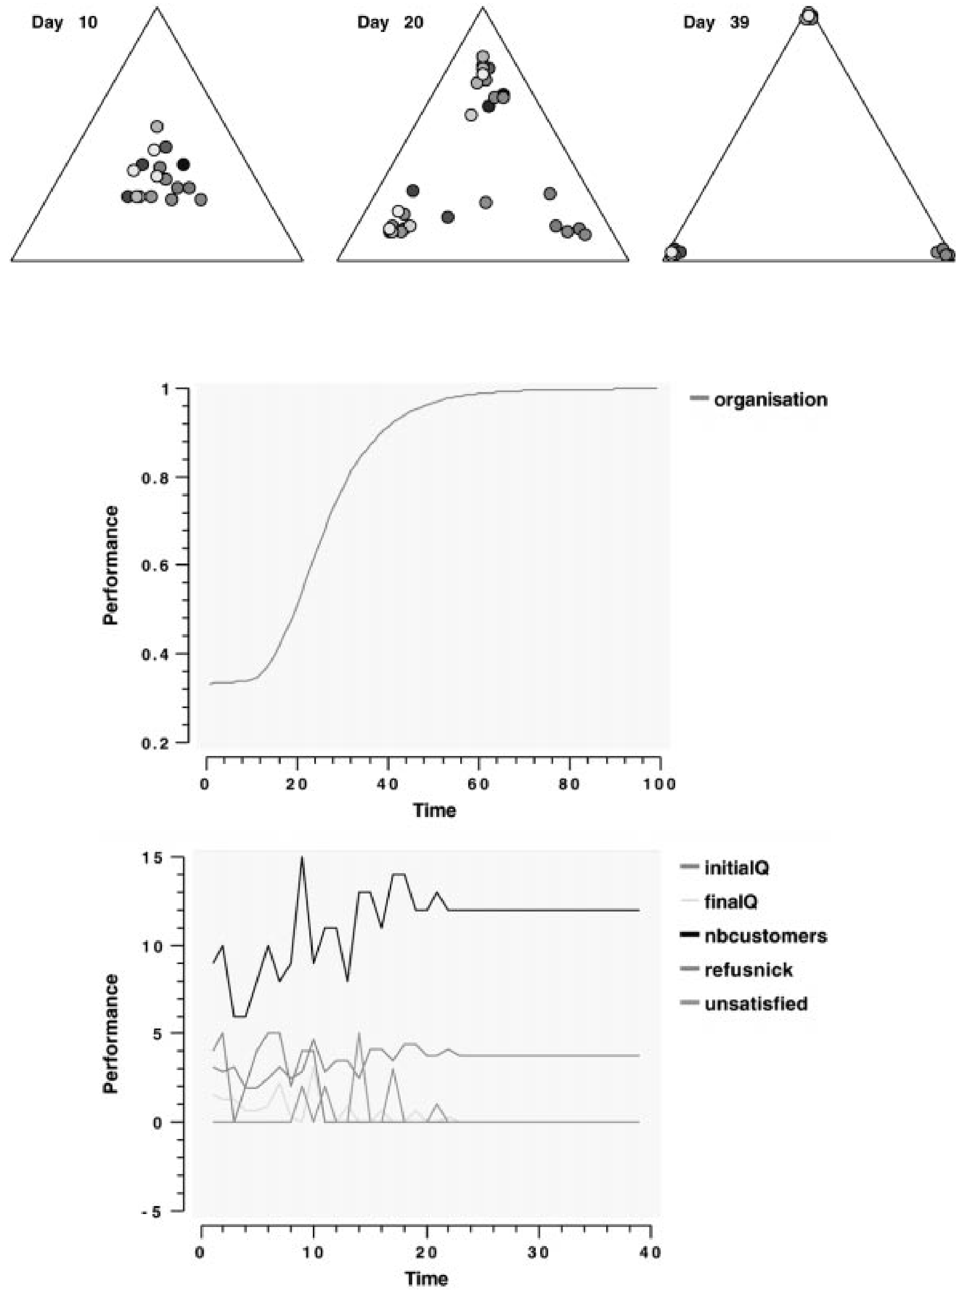

Fig. 4. *Charts for the Organised Regime.*
(30 agents visting 3 sellers, with $\gamma = 0.1$ and $\beta = 2\beta_c$). All charts and notation are the same as for Fig. 3, except for the scale of the order parameter plot (y). In the three simplex plots, starting from indifference between all three sellers, the circles move to the corners representing the preferred sellers. Organisation takes place. The order parameter y increases steadily from 0.33 to nearly one. The time charts show how fluctuations decrease quickly due to organisation.

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 the same profit π whenever they succeed in making a transaction. Sellers choose their 'endowments' at the beginning of each session using a simple rule of thumb. The rule used in the basic simulation was to choose resources to match yesterday's demand. This means that the probability of being served by the seller that one visits, i.e. that $\text{Prob}(q_i > 0)$ need not be one as in the simplest version solved analytically. Figs 3 and 4 correspond to a discount factor $\gamma = 0.1$. We set the value of the profit obtained from a transaction equal to 1 and the critical value of the non-linear parameter is then $\beta_c = 0.3$ (13). The initial value of each cumulated profit variable J_{ij} was set equal to 0. As in the analytic case, depending on the value of the non-linear parameter β, the two predicted dynamic regimes, order and disorder, are observed.

3.2.1. *Disorganised Behaviour*

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[17] of the buyer's profit per transaction. This is due to all those occasions on which a buyer visited a seller with no stocks. 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 seller). This difference was also generated indirectly by buyers who visited sellers with no stocks since, at the same time some sellers with supply were not visited resulting in losses for their owners.

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

3.2.2. *Organised Behaviour*

In sharp contrast, the same analysis performed with a higher value for the non linear parameter $\beta = 2\beta_c$ shows a great deal of organisation, see Fig. 4.

The average order parameter, y, steadily increases to 1 in 200 time steps. As

[17] 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 4. However, the observed decrease in profit for buyers and sellers is generic.

seen on the simplex plot at time 50, each customer has built up loyalty to one seller. Performance of seller number one also stabilises 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 the profit per realised transaction. Because buyers have not changed sellers during the last 100 days, sellers learn to purchase the exact quantity needed to satisfy all their buyers and incur no losses as foreseen by the model.

By avoiding daily fluctuations in the number of customers visiting a seller, the ordered regime is beneficial to both customers and sellers, that is both obtain higher profits than in the disorganised situation. In this sense, the ordered regime is Pareto superior to the disordered regime.

3.2.3. *Heterogeneity of Buyers and Sellers*

In real markets, we expect a mix of buyers with different non linear β and discount γ parameters. Thus, some buyers will be loyal to certain sellers, while others will continue to search. The same is true for sellers who might offer different prices and thus different profits to buyers. The generality of the theoretically derived results was tested by numerous simulations.

Table 1 gives a series obtained from a uniform distribution of γ_i coefficients varying from 0.1 to 0.5 among 30 buyers; three sellers offer different profits varying from -20% to $+20\%$ of the mean profit. $\beta_c = 1.56$ for the median seller ($\gamma = 0.03$) and the median buyer ($\pi = 0.057$). All data were taken after the simulations had been run for 500 time steps. Different lines of the table correspond to different values of β (first column). The next 7 columns represent the histograms of buyers order parameters: the leftmost bin correspond to $0.3 < y_i < 0.4$, the rightmost to $0.9 < y_i < 1.0$. The two rightmost

Table 1

Histograms and Averages of Order Parameters Across the Transitions

	Histograms of y_i							Order parameters	
β	0.3	0.4	0.5	0.6	0.7	0.8	0.9	y	y_s
0.3	29	1	0	0	0	0	0	0.35	0.35
0.6	26	1	2	0	1	0	0	0.38	0.35
0.9	24	3	1	1	1	0	0	0.39	0.35
1.2	18	5	2	1	2	1	1	0.45	0.34
1.5	8	5	2	2	1	6	6	0.63	0.34
1.8	1	3	5	1	1	7	12	0.79	0.38
2.1	1	0	0	1	2	6	20	0.90	0.36
2.4	0	0	0	0	0	6	24	0.94	0.36
2.7	0	0	0	0	0	1	29	0.98	0.47

columns are average buyers order parameter y and y_s an order parameter for sellers defined by

$$y_s = \frac{\sum_j n_j^2}{(\sum_j n_j)^2} \tag{22}$$

where n_j is the number of customers of seller j. This series of simulation shows that when β increases:

- more and more buyers move across the transition towards ordered behaviour, filling the rightmost bins of the histogram; such buyers become loyal to certain sellers, while others continue to search;
- individual transitions remain sharp; buyers maintain either ordered or disordered behaviour according to the position of their $\beta_c = m\pi/\gamma_i$ with respect to β; note that the intermediate bins have a rather small population with respect to the extremal bins, even for intermediate values of β;
- Since $y_s < 0.5$, it is clear that all three sellers are visited, even though one is offering a profit 40% above the lowest offered profit.

As also shown in Herreiner (1997), organised or disorganised behaviour is here a property of buyers, not a property of markets.

3.4. *Beyond the Mean Field Approximation*

The results of the mean field approach were obtained from a differential equation approximating a discrete time algorithm. They are valid when the changes at each step of the algorithm can be considered small. γ thus has to be small, which is true for the simulation results given in Figs 2 and 3.

One of the features noticed by observing on-line the motion of individual buyers in the ordered regime on the simplex plots is that agents sometimes move 'backward' towards sellers which are not the sellers that they 'prefer', i.e. those whose with highest preference coefficient. But since for most of the time they move towards preferred sellers, these moves only very infrequently make them change sellers and preferences. When the variable γ is increased, these moves have more important consequences.

- Customers might change loyalty from one seller to another one. Increasing γ results in larger steps taken by customers on the simplex, which might make them move from one corner 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 probabilty that the process will generate such changes becomes higher and they are indeed sometimes observable on-line on the simplex plots.
- Sellers offering higher profits are favoured on average by these changes. The stability of fidelity coefficients for sellers offering lower profits predicted

by the mean field continuous approach becomes metastability in discrete dynamics.[18]

4. More Complicated Models and Results

4.1. *Generality*

We now briefly discuss some generalisations of the model. Since these more realistic variants are no longer analytically tractable we had to resort to computer simulations to compare their dynamic properties with those of the simple soluble model and with empirical data. Full details of these simulations can be found in Weisbuch *et al.* (1998). All the variants share the same fundamental mechanism by which buyers choose sellers and the same way of updating preference coefficients as defined in Section 1.2.

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 to 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 in which individuals will find themselves, 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 transiton 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 generality is based on the conjecture[19] that the large set of models which share the same generic properties also includes the 'true' model of the real system itself.

The search for generic properties frees us from the necessity of having a comprehensive model for all the aspects of the market such as endogenous price dynamics or a fully developed theory of sellers' behaviour: as long as various exogenous price dynamics and the different sellers' behaviours that we have tested yield the same properties for buyers' behaviour, these properties would also be observed in fully comprehensive models.

[18] Metastability means that according to initial conditions and probabilistic events, buyers might be attracted towards sellers offering lower profits on a short time scale of order $1/\gamma$, but that sellers offering the best deals are selected in the ultra long term (this distinction is similar to the distinction between the long run and the ultra long run introduced in economics by Gale *et al.* (1995). But the times to reach an equilibrium distribution of fidelities can be extremely long depending on γ (see Herreiner (1997) for more simulation results). Naturally, in real life, sellers losing customers because they offer lower profit have time to readjust and to improve their offers!

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

4.2. *Profits*

For our theoretical discussion it was enough to know the profit realised from a transaction with each firm. A first step towards a more complete model would be to specify explicitly the origin of this profit. For example, one could define the demand faced by the retailers who are the buyers in our model. This in turn would give the optimal price to be set by the seller and the quantity that would be purchased by a buyer. This was done by using a simple hyperbolic specification for the demand curves faced by the retailers and assuming that these retailers were local monopolists in order to make the simulations discussed in Section 3.

4.3. *Several Sessions*

The one-session model described in Section 1 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 would have to consider a model in which customers are given at least two occasions to purchase goods, and for this we assume the existence of a 'morning' and an 'afternoon' session.

This would necessitate specifying the rules by which agents make their decisions taking into account the existence of a second session. Thus sellers would have to decide on the prices to charge in the morning in the light of the transactions they expect to make in the afternoon. They also have to decide on prices in the afternoon given the quantities of goods left on their hands after the morning transactions. Buyers on the other hand, would make a decision as to which prices to accept in the morning given the prices and quantities they expect to be available in the afternoon.

Simulations with an afternoon session do not reveal qualitative changes in the observed behaviour of buyers; we simply observed that when a buyer does not make a transaction in the morning she has a much better chance of making a higher afternoon profit with a new seller who still has stocks. This leads to a more rapid change in the probability of switching to another buyer and to a shorter duration of loyalty.

4.4. *Sellers' Behaviour*

Another feature that we mentioned previously that would enhance the model is that of the calculation of the amount that sellers wish to supply. As we have pointed out, the simple model simulated in Section 3 used a rule of thumb for sellers which simply forecast today's sales as being the same as yesterday's. We tried two other rules, neither of which changed the qualitative results. One rule was to use a weighted average of past sales and the other was to use that quantity plus a small amount to take account of what would be optimal behaviour. To make the actual optimal calculation would require that the sellers know the probability distribution of the number of customers with which they are faced. This distribution depends on the behaviour of the buyers and is difficult to calculate. However, if it is known by the seller, then the optimal quantity to supply is easily computed and is given in Weisbuch *et al.* (1998).

The last feature of our model which needs to be generalised is that of uniform behaviour by all sellers at any point in time. The ultimate answer to this would be to develop a complete model of sellers' behaviour. However, since our aim in this paper was to explain certain empirical regularities in the behaviour of buyers we confined ourselves to examining the case in which the distribution of prices varied over time. The fact that we obtained similar results would suggest that the introduction of strategic pricing behaviour as in the models of Bergemann and Valimaki (1996; 1997) would not affect the qualitative behaviour of the buyers in our model.

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 concerns 237,162 individual transactions between ca. 1,400 buyers and 45 sellers which occurred between 2nd January 1988 and 29th June 1991. For each transaction the following information is recorded: the name of the buyer, the name of the seller, the type of fish, the weight of fish, the price, the order in the seller's transactions.

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 our 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 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 one per week) with very few exceptions visit only one seller, while less frequent buyers visit several sellers, which is consistent with the results from 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).

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 sellers, 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.

For the first step, to check for loyalty of buyers, we consider statistics for cod, whiting and sole transactions in 1989,[20] see Table 2.

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 2, 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 column of table. Also in the whiting and sole market buyers are loyal, but to a lesser degree: more than half of them[21] buy more than 80% from one seller. Hence, there are large fractions of loyal buyers in all three markets.

Furthermore, one key conclusion of our model which distinguishes it from a number of other possible explanations of fidelity is that, due to the phase transition, there will be two separate types of behaviour in the market regardless of the underlying distribution of the exogenous variables describing individual characteristics. In other words, even if the exogenous explanatory variables such as frequency of visits or the amounts purchased have an unimodal distribution, the distribution of fidelity will be bimodal.

We then tested the distributions of the dependent variable of our theory, fidelity and of the two exogenous variables, frequency of visits and monthly volume of transactions for cod. The histograms are shown in Figs 5, 6 and 7. Using a test developed by Scott (1992)[22] we estimated the probability of different modes. In other words we estimated the probability that a particular

Table 2

Loyalty in the Market for Cod, Whiting, and Sole

	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%

[20] The statistics for other periods of comparable length are very similar.
[21] Whiting 124 of 229, and sole 154 or 280.
[22] The test developed by Scott involves the smoothed frequency polygons and is more difficult than

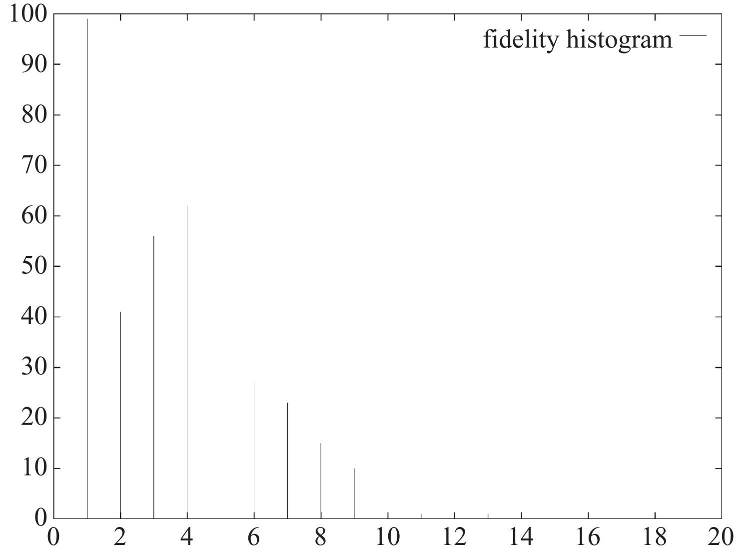

Fig. 5. *Histogram of the Number of Buyers of Cod as a Function of How Many Sellers They Visit on Average During One Month in 1990.*
The sample of buyers include those visiting more than once per month, and present in the market for more than six months. The distribution is clearly bimodal, with one peak corresponding to fidelity to one seller, and another peak centred on visting 4 sellers on average.

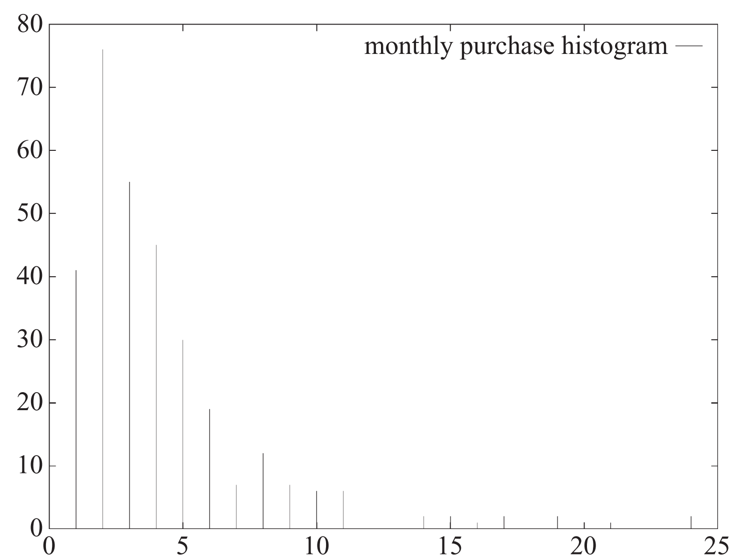

Fig. 6. *Histogram of the Number of Buyers of Cod as a Function of Their Monthly Purchase of Cod.* (Each bin corresponds to ten extra kilograms purchased). The distribution presents only one peak.

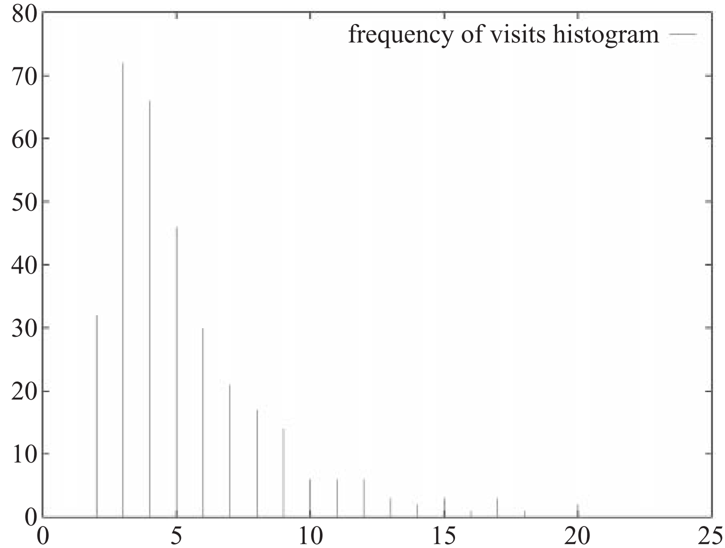

Fig. 7. *Histogram of the Number of Buyers of Cod as a Function of Their Average Number of Monthly Visits to the Fish Market.*
The distribution presents only one peak.

value was a local maximum of the density. For fidelity we found that there were two modes with probabilities 0.99 and 0.83 respectively whereas the other two distributions had single modes with 0.99 probabilities. This shows clearly that the two distinct groups of buyers found by examining their fidelity could not be explained by there being two distinct groups of buyers as far as quantities purchased or frequency of visits are concerned.

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. Fig. 8 summarises loyalty of buyers in terms of relative frequency of visits to their favourite seller as a function of their monthly purchase of cod on a logarithmic scale. One may observe that loyalty is high in general and that a number of buyers visit only one seller. The fit was done using a non-parametric adaptive smoothing algorithm called LOWESS originally developed by Cleveland in 1979. It is an extension of the usual k nearest neighbour smoothers since it allows k to vary as a function of the variance in the part of the data being smoothed. It shows that loyalty increases with monthly purchase.

All three features are consistent with our theory, and in contradiction with a random search behaviour for all buyers.

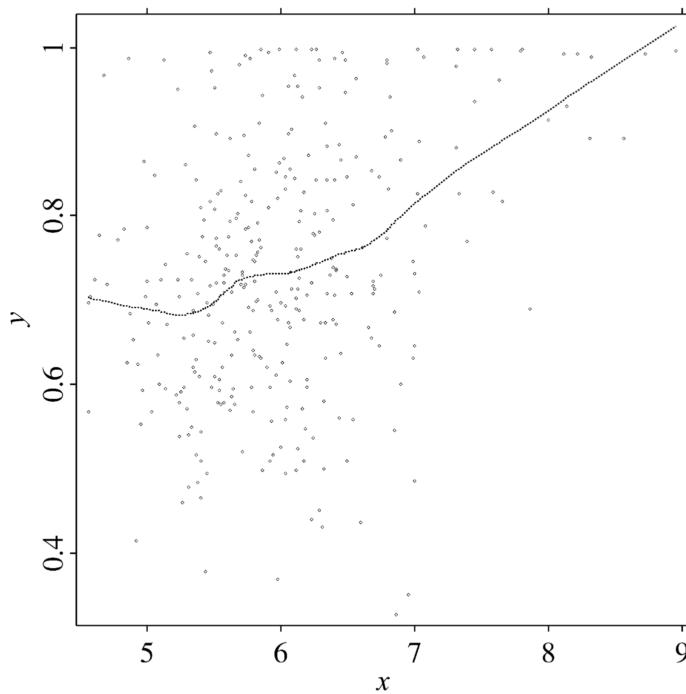

Fig. 8. *Buyer Loyalty on the Marseille Fish Market.*
Each dot is an empirical observation from the Marseille fish market representing buyer loyalty to his favourite seller (relative frequency of visits) as a function of the logarithm of his monthly purchase of cod in tens of kilograms. Low purchases correspond to infrequent buyers, who generally visit once a week, while large purchases are those made by buyers who visit nearly every day the market is open. The fit is a non-parametric adaptive smoothing algorithm which shows that loyalty increases with monthly purchase.

6. Conclusions

We have examined a simple model of a market in order to see how 'order' develops. 'Order' here means the establishment of stable trading relationships over the periods in which the market is open. Such order is observed on many markets for perishable goods. We focused on the behaviour of buyers to explain how the sharp division into those who are loyal to a particular seller and those who always search could arise. A simple theoretical model yielded this division which is observed on the Marseille fish market. The basic feature of our model is that agents learn by reinforcing those actions, in our case choice of sellers, which proved to be more profitable in the past. Simulations of models which incorporated more general features but which retained the same learning mechanism, showed the same qualitative features as the basic model. 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 'ordered behaviour' in our model is more efficient in Pareto terms than 'disorder' and it is therefore of considerable interest to be able to identify under which conditions such 'order' emerges.

Laboratoire de Physique Statistique, ENS, Paris

EHESS, Institut Universitaire de France and Université d'Aix-Marseille III

University of Bonn

Date of receipt of first submission: March 1998
Date of receipt of final typescript: July 1999

Appendix. Fixed Points of tanh( )

Restate (11) as

$$\tanh(zx) = x,\ z \in \mathscr{R},\ \text{ with } x = \frac{\gamma\Delta}{\pi} \text{ and } z = \frac{\beta\pi}{2\gamma} = \frac{\beta}{\beta_c}, \tag{23}$$

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

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

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

If $\tanh(zx)$ has one solution, then the slope of $\tanh(zx)$ is flatter than the slope of x; and the unique 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 the 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

$$\frac{J_1}{J_2} = \exp\left(\frac{\beta\pi}{\gamma}\right), \tag{24}$$

which can be obtained from (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 $x_0 = 0$ ($\Delta_0 = 0$):

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

This yields

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

solving for $x(\Delta)$ leads to

$$x = \sqrt{\frac{3(z-1)}{z^3}} \quad \text{and} \quad \Delta = \sqrt{\frac{12(\beta - \beta_c)}{\beta^3}}. \tag{25}$$

References

Anderson, S. P., de Palma, A. and Thisse, J. F. (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 Lane, D. A. (1993), 'Information contagion', *Structural Changes and Economic Dynamics*, vol. 4, pp. 81–104.

Bergemann, D. and Valimaki, J. (1996), 'Market experimentation and pricing', Cowles Foundation discussion Paper 1122.

Bergemann, D. and Valimaki, J. (1997), 'Market diffusion with two-sided learning', *RAND Journal of Economics*, vol. 28, pp. 773–95.

Brock, W. (1993), 'Pathways to randomness in the economy: emergent nonlinearity and chaos in economics and finance', *Estudios Economicos*, vol. 8, pp. 3–55.

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.

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

Cleveland, W. S. (1979) 'Robust locally weighted regression and smoothing scatterplots', *Journal of the American Statistical Association*, vol. 74, pp. 829–36.

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

Derrida, B. (1986) 'Phase transitions in random networks of automata' in (J. Souletie, J. Vannimenus and R. Stora, eds.), *Chance and Matter*, New York: North-Holland.

Diamond, P. (1989), 'Search theory', in (J. Eatwell, M. Milgate and P. Newman, eds.), *The New Palgrave: A Dictionary of Economics*, London: Macmillan, pp. 273–9.

Felli, L. and Harris, C. (1996), 'Learning, wage dynamics, and firm-specific human capital', *Journal of Political Economy*, vol. 104, pp. 838–68.

Fisher, F. M. (1973), 'Stability and competitive equilibrium 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*, vol. 1, pp. 51–62.

Gale, J., Binmore, K. and Samuelson, L. (1995), 'Learning to be imperfect: the ultimatum game', *Games and Economics Behavior*, vol. 8, pp. 56–90.

Härdle, W. and Kirman, A. (1995), 'Non classical demand: a model-free examination of price-quantity relation in the Marseille fish market', *Journal of Econometrics*, vol. 67.

Herreiner, D. (1997), 'Market Organization', mimeo, G and W III, University of Bonn.

Jackson, M. and Wolinsky, A. (1996), 'A strategic model of social and economic networks', *Journal of Economic Theory*, vol. 71, pp. 44–74.

Kirman, A. P. and Vignes, A. (1991), 'Price dispersion. Theoretical considerations and empirical evidence from the Marseilles fish market', in (K. G. Arrow, ed.), *Issues in Contemporary Economics*, London: Macmillan.

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

Nadal, J-P., Weisbuch, G., Chenevez, O. and Kirman, A. (1998) 'A formal approach to market organization: choice functions, mean field approximation and maximum entropy principle', in (J. Lesourne and A. Orlean, eds), *Advances in Self-Organization and Evolutionary Economics*, Paris: Economica, pp. 149–59.

Orléan, A. (1995), 'Bayesian interactions and collective dynamics of opinions: herd behavior and mimetic contagion', *Journal of Economic Behavior and Organization*, vol. 28, pp. 257–74.

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

Scott, D. W. (1992) 'Multivariate density estimation', New York: John Wiley.

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

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

Weisbuch, G., Kirman, A. and Herreiner, D. (1998), Internal report available on the web, http://www.lps.ens.fr/weisbuch/market.ps