

Learning to Be Loyal. A Study of the Marseille Fish Market*

Alan P. Kirman¹ and Nicolaas J. Vriend²

¹ E.H.E.S.S., GREQAM, Marseille, France

² Queen Mary and Westfield College, University of London, Dept. Economics, Mile End Road, London E1 4NS, UK, <n.vriend@qmw.ac.uk>, <http://www.qmw.ac.uk/~ugte173/>.

Abstract. We study the wholesale fish market in Marseille. Two of the stylized facts of that market are high loyalty of buyers to sellers, and persistent price dispersion, although the same population of sellers and buyers meets in the same market hall on every day. We build a minimal model of adaptive agents. Sellers decide on quantities to supply, prices to ask, and how to treat loyal customers. Buyers decide which seller to visit, and which prices to accept. Learning takes place through reinforcement. We analyze the emergence of both stylized facts price dispersion and high loyalty. In a coevolutionary process, buyers learn to become loyal as sellers learn to offer higher utility to loyal buyers, while these sellers, in turn, learn to offer higher utility to loyal buyers as they happen to realize higher gross revenues from loyal buyers.

JEL classification: C7, D4, D8, L0, L1, L7

Keywords: Market structure, price dispersion, loyalty, adaptive behavior, coevolution

* An earlier version was first presented at the Economic ALife workshop at the Santa Fe Institute, May 1995, and subsequently circulated in several variants. We thank seminar and workshop participants at the Santa Fe Institute, the European University Institute, VPI&SU, UC-Los Angeles, UC-Irvine, and the Universities of Trento, Southern California, George Mason, Pompeu Fabra, London (QMW), Keele, and Ancona, participants at conferences of the SEDC in Barcelona, SCE in Geneva and Cambridge, ESEM and EEA in Istanbul, ASSET in Alicante and Marseille, and IIFET in Tromsø, and Antonio Cabrales, Art De Vany, and Johan Stennek for helpful comments and discussions. All errors and responsibilities are ours. Financial support through TMR grant ERBFMBICT950277 (NJV) from the Commission of the European Community, Brussels is also gratefully acknowledged.

1. Introduction

In this project we study the working of the wholesale fish market in Marseille (France). There are two reasons to study this specific market. First, we have a unique data set containing the records for all single transactions that have taken place in this fish market over a number of years. Rich as such a data set is, this is at the same time one of its limitations, as there are, for example, no data for transactions that did not take place. Second, this wholesale fish market is a relatively simple market. It is well-defined, and well-structured, compared with, for example, a market economy as such, or with a market like the one for second-hand cars. Also, fish is a perishable commodity, implying that there are no inventories carried over (cf. the early literature on market microstructures in finance).

Our real interest as economists is not in fish markets *per se*. But we do believe that some of the insights developed in this study of the fish market might be carried over to other real markets. In particular we believe that if we want to understand the dynamics of interactive market processes, and the emergent properties of the evolving market structures and outcomes, it might pay to analyze explicitly how agents interact with each other, how information spreads through the market, and how adjustments in disequilibrium take place.

In section 2 we will sketch the real fish market, and present two not easily explained stylized facts.¹ In section 3 we present a minimal model with interacting agents. Section 4 reports on some results, and section 5 concludes.

2. The Real Fish Market

The market is open every day from 2am to 6am. The real market has a fixed population of about registered 40 sellers. They buy their supply outside the market before the market opens. There is also a fixed population of buyers; about 400. These buyers are retail sellers or restaurant owners. During the market day, they shop around, visiting individual sellers. Standing face-to-face with a seller, the buyer tells the seller which type of fish and which quantity he is interested in. The seller then informs him about the price. Prices are not posted. And they are individual in a threefold sense. First, each individual seller decides upon his own prices. Second, each seller may have different prices for different buyers. Third, each seller may even ask a different price for a given type and quantity of fish if that is proposed by the same buyer at different times of the day. A price communicated by an individual seller to an individual buyer for a given

¹ For another recent paper about a fish market (Fulton), which has a structure similar to the Marseille market, see Graddy [1995].

transaction is not perceived by other buyers or sellers in the market. The prices are take-it-or-leave-it, and there is no bargaining. At the end of the day, unsold fish is thrown away, or sold to the European Union at bottom prices. It is forbidden to sell old fish.²

The most remarkable stylized facts of the real market are the following. First, there is a widespread high loyalty of buyers to sellers. Second, there is persistent price dispersion. Further documentation of these stylized facts can be found in Kirman & Vignes [1991], Kirman & McCarthy [1990], Vignes [1993], Härdle & Kirman [1995], and Weisbuch et al. [1996]. Since the same population of sellers and buyers is present every day in the same market hall, at first glance these stylized facts may seem somewhat puzzling. Of course, standard economic reasoning can explain these facts. If these agents behave in this way, they must have good reasons to do so, and the task for the theorist is simply to uncover these reasons. In the remainder of this paper, we will not focus on the empirical details concerning this fish market, but instead on the modeling approach we have chosen to explain the mentioned stylized facts.

For various reasons this turns out to be far from simple. As we will show now, some simple standard arguments do not have much explanatory power. Either they do not apply to this specific market, or they leave open too many possibilities. First, the stylized facts might seem to be related to the quality of the fish. It could be that the price dispersion simply correlates with quality differences of the fish. However, the classification scheme used to distinguish types of fish uses several quality measures. Clearly, there will be some residual quality differences within each type, but the price dispersion observed within each type seems excessive, e.g., when compared with the price dispersion between different types of fish.

Next, the effect of quality differences could seem to be related to the loyalty in case there would be an information asymmetry. Before discussing the contents of the information asymmetry argument, we would like to question whether an information asymmetry would be likely in this market. All buyers are professional traders themselves; retailers or restaurant owners. In any case, even if there were an information asymmetry concerning the quality of the fish, it is not clear what its effect would be. The standard textbook argument would run as follows. If the quality of the fish is not detectable *ex ante* for the buyers, then a seller has an incentive to deliver good quality with respect to repeat buyers, as an investment in his reputation. That is, the information asymmetry would explain the existence of loyalty of buyers to specific sellers. This seems a too loose application of the textbook argument. Notice that we have a fixed population of buyers in the Marseille fish market. That is, every buyer (loyal or not) is a potential repeat buyer; unlike the standard textbook example of tourists. Hence, a seller would have an incentive to deliver good quality to every single buyer. Clearly, a seller cannot offer higher quality to all buyers. A last attempt to save the causal link

² This is a simplification since some species can be offered for sale on the succeeding day. Buyers are however capable of recognizing the freshness of the fish.

between quality differences and loyalty would be to argue that non-loyal buyers are intrinsically non-loyal anyway, and that therefore there is no incentive to offer them good quality. But exactly the same argument would tell you that loyal buyers are intrinsically loyal anyway, and that therefore there would be no incentive to deliver them good quality. Hence, with respect to quality differences we must conclude that they may exist, but they are not observable in data, and it is unclear what they could explain, besides a minimal amount of price dispersion.

A second ready-to-use explanation economic theory might seem to offer for the stylized facts is the theory of implicit contracts. Because of the uncertainty related to variation in the daily supply to the fish market, the agents stipulate implicit contracts specifying the loyalty of buyers to specific sellers. This might well be, but the story cannot be that simple. First, on this fish market there are also variations in the daily demand. Given the variation in the daily supply, one might expect that buyers and sellers agree upon an implicit contract in which prices are somewhat higher than otherwise, while the supply to these buyers is guaranteed. With the variation of the demand one might expect an implicit contract in which prices are lower, while the demand to the seller is guaranteed (cf., newspaper subscriptions). With variation on both sides it is not clear what form the implicit contract should take. But also with variation on the supply side only, the form of the implicit contract is not obvious a priori. One contract might be the one mentioned above; with higher prices and better service to the buyers agreeing upon such a contract. Note that there is a trade-off between the price to be paid by those buyers and the service they get. This implies that those buyers might be equally well off under an implicit contract that gives them lower prices and worse service. In other words, there might exist an implicit contract specifying a long-term relationship between a buyer and a seller, in which the buyer is *not* served more frequently than buyers without such a contract. Many examples of explicit contracts of this type exist; for example in the natural gas market, or the market of power supply. In the cases in which such a buyer is not served, he will be forced to visit other sellers. Hence, in the transaction data such a seller might appear as one of the non-loyal buyers, even when there is this long-term relationship characterized by bad service. Hence, it might well be that there exist implicit contracts in the wholesale fish market in Marseille, but, besides the standard problems that it is not clear who should design the contracts and how they should be enforced, in our case the problem is that they might come in many forms, and it is difficult or impossible to indicate some of them in the transaction data. To the extent that implicit contracts might come in many forms, this might explain the existence of price dispersion. The problem, however, is that it leaves open many possibilities. Just as in some static models on price dispersion (Futia [1977], or Burdett & Judd [1983]), there might be multiple equilibria, and it is not clear whether the stable equilibria are characterized by price dispersion. Hopkins [1995] presents a dynamic analysis of such models.

A third standard argument is that it is known that in markets where suppliers first decide on their supply, and then enter a Bertrand game with given fixed

stocks, there need not exist a pure-strategy Nash equilibrium, whereas a mixed-strategy Nash equilibrium exists in general (see Stahl [1988], and the references therein). This might explain some of the price dispersion. Although, given that in our market there is price dispersion for a given seller even within one day, this means that we must imagine such a seller flipping coins for every single price asked. Moreover, it leaves open the issue of loyalty.

A more sophisticated approach than these simple arguments would be to design a multi-stage game with an equilibrium characterized by the stylized facts just mentioned. But designing and analyzing such a game is complicated. Moreover, it is easy to imagine such a game will have multiple equilibria; and not only because of the equilibria being asymmetric. Here is a very rudimentary example of the kind of games one could construct. Suppose there are two sellers, and two buyers. A seller's strategy S is a price proposed to buyer 1, and a price proposed to buyer 2. A buyer's strategy S is simply the choice of a seller to visit. Suppose that for both seller 1 and seller 2 we have $S=(a, b)$. And suppose for buyer 1 we have $S=(\text{seller 1})$, and for buyer 2 we have $S=(\text{seller 2})$. Now suppose, $a \neq b$, then we have a Nash equilibrium for the one-period game with loyalty and price dispersion. As we see immediately, many such equilibria may exist. Therefore, what we will do instead is something very simple. We will build a minimal model of an artificial fish market with adaptive agents. Besides being simple, an advantage of this approach is that the evolutionary outcomes of this model might offer some insights as to which of the possible equilibria of a corresponding multi-stage game might be more likely to occur.

With respect to quality, we will focus on one observable quality aspect, the seller's service rate, and analyze its role in the market. With respect to implicit contracts, we will keep these really *implicit*. That is, keeping aggregate demand and supply constant (although individual supply and demand faced may vary), we will analyze whether long lasting trading relationships do emerge in the artificial market, and we will analyze their characteristics. And with respect to the possibility of a mixed-strategy Nash equilibrium, we will not study how introspection may lead to equilibria, but we will instead look for the emergence of patterns in the behavior of adaptive agents.

Our comparison of the artificial and the real world will not be characterized by a quest for a complete model of the fish market. We will not try to build a model fitting all aspects of the real world for the following reasons. First, every model is by definition an abstraction. If enough data can be collected, statistical testing will reject any model. Second, when modeling by building artificial worlds, one might get a very good fit without gaining understanding. There exist economic simulation models with more than 10,000 variables. At some point it might be that one mainly succeeds in building a copy of the real world, about which we have the same degree of understanding as about the real world. Therefore, we will only consider specific questions concerning the stylized facts of the real market that appear remarkable or important. We will try to build a minimal model that generates, and with which to test those stylized facts. This might suggest ways to

understand, or not, those phenomena. This understanding is of the same type as with formal mathematical models. The question is whether we might consider the real world to be working 'as if' it were like our model.

3. The Artificial Fish Market

The place of action in our minimal model is the market hall. The actors are 10, initially identical sellers, and 100, initially identical buyers. They meet in the market hall for 5000 days, for a morning and afternoon session (see below). The commodity traded is indivisible, and perishable. On each day the sequence of events is the following.

In the morning, before the market opens, the sellers buy their supply outside the market for a given price that is identical to all sellers, and constant through time ($p^{\text{in}} = 9$). The market opens, and the buyers enter the market hall. Each buyer wants 1 unit of fish per day. All buyers simultaneously choose the queue of a seller. The sellers, then, handle these queues during the morning session (see below). When all queues have been handled, the end of the morning session is reached. In the afternoon, the market re-opens. All still unsatisfied buyers choose the queue of a seller. The sellers handle these queues, after which the end of the afternoon session is reached. All unsold stocks perish. The buyers re-sell their fish outside market at a given price that is identical for all buyers, and constant through time ($p^{\text{out}} = 15$). Notice that each buyer can visit at most one seller in the morning plus one seller in the afternoon. Figures 1 and 2 give a sketch of the market.

On each day, the buyers have to make the following decisions. First, they choose a seller for the morning session. Second, they decide which prices to accept or reject during the morning session. Third, if necessary, they choose a seller for the afternoon. And fourth, they decide which prices to accept or reject during the afternoon. The sellers face four decision problems as well. First, they decide the quantity to supply. Second, they decide how to handle their queues. Third, they decide which prices to ask during the morning. And fourth, they decide which prices to ask during the afternoon.

Each single decision problem is modeled separately for each individual agent by means of a Classifier System. Considering a Classifier System as a 'black box' for the moment, this implies that in our artificial fish market each single buyer, and each single seller has 4 decision boxes in his head. Hence, with 10 sellers and 100 buyers we model explicitly 440 decision boxes by separate Classifier Systems. Figure 3 presents one such stylized Classifier System.

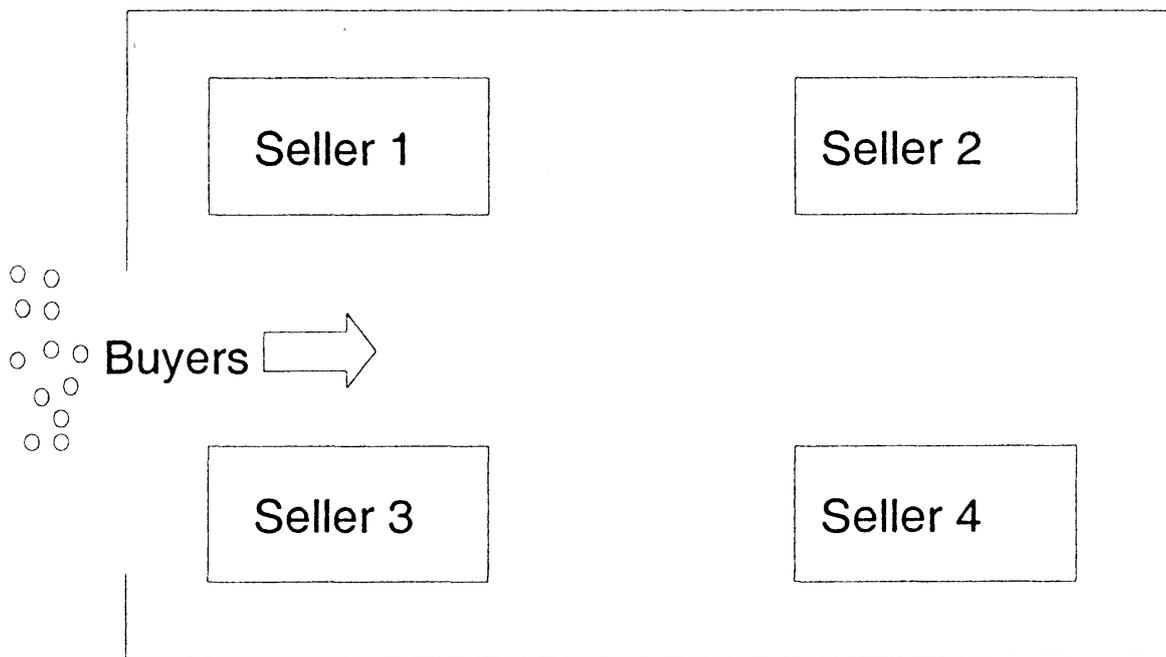


Fig. 1: The Market Hall with Buyers and Sellers

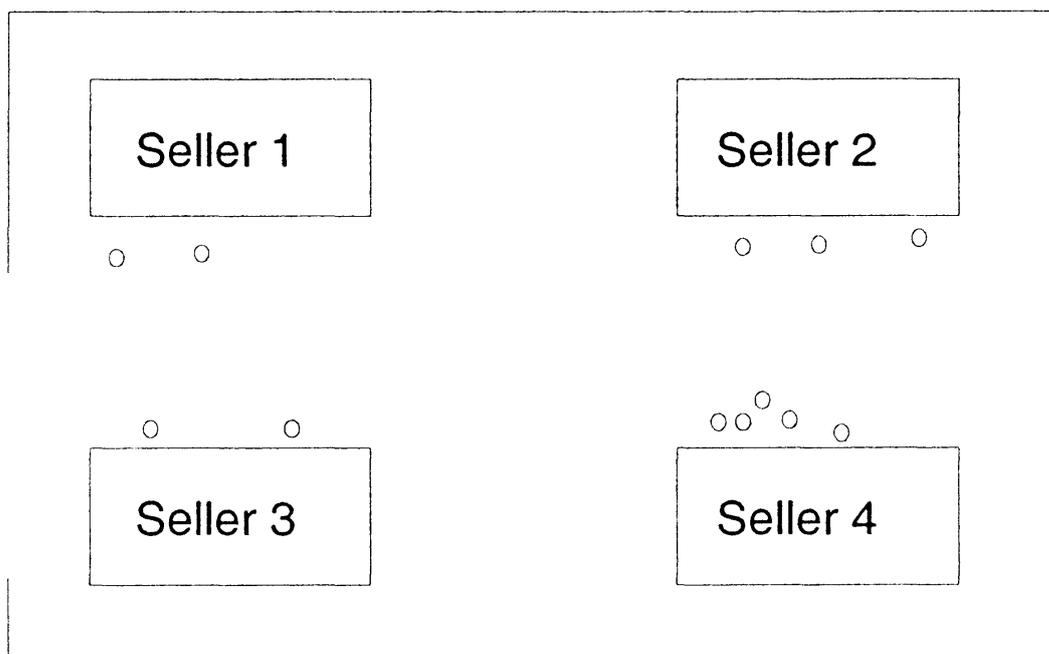


Fig. 2: Each Buyer Has Chosen a Seller

condition	action	strength
if	then
..
..

Fig. 3: A Classifier System

A Classifier System consists of a set of rules, each rule consisting of a condition 'if ...' part, and an action 'then ...' part, plus to each rule attached a measure of its strength. The Classifier System does two things. First, it decides which of the rules will be the active rule on a given day. Hence, it checks the condition part, and all rules satisfying the 'if ...' condition make a 'bid' as follows: $\text{bid} = \text{strength} + \varepsilon$, where ε is white noise. The rule with the highest 'bid' in this 'stochastic auction' wins the right to be active.³ Second, the Classifier System updates the strength s of a rule that has been active and has generated a reward from the environment on a given day $t-1$ as follows: $s_t = s_{t-1} - c \cdot s_{t-1} + c \cdot \text{reward}_{t-1}$, where $0 < c < 1$. Hence, $\Delta s_t = c \cdot (\text{reward}_{t-1} - s_{t-1})$. In other words, as long as the reward generated by the rule on day $t-1$ is greater than its strength at $t-1$, its strength will increase. Hence, the strength of each rule converges to the weighted average of the rewards from the environment generated by that rule.⁴ In our implementation of this model, the strengths of all rules are equal at the start.

Classifier Systems are a form of reinforcement learning. Reinforcement learning is related to multi-armed bandit problems, and is based on two principles. First, agents try actions. Second, actions that led to better outcomes in the past are

³ Besides through the white noise added to the 'bids', the agents experiment through some kind of 'trembling hand', disregarding any 'bid' with a given small probability.

⁴ We presented this specific learning model in Kirman & Vriend [1995]. Sarin & Vahid [1997] analyze the theoretical properties of this model, relating it also to evolutionary explanations. And in Sarin & Vahid [1998] they show its empirical relevance by explaining the data of some laboratory experiments.

more likely to be repeated in the future. Reinforcement learning is among the most basic forms of learning, for which there exist ample support in the psychological literature (see, e.g., Bush & Mosteller [1955], or Roth & Erev [1995]). Moreover, it is a minimal form of modeling learning, in the sense that we do not need to make many assumptions about the reasoning procedures followed by the agents.

We will now specify in a more detailed way the contents of each decision box for each of the types of decisions to be made by the buyers and the sellers.

Buyers

First, the buyer must make the choice of a seller for the morning session. This is the first decision related to the question 'who is going to interact with whom?'. The rules have no condition. The actions are simply: <choose seller 1>, <choose seller 2>, ..., up to <choose seller 10>. The strength of each activated rule is updated once a day using the following payoff as reward from the external environment: $\text{payoff} = \max \{p^{\text{out}} - p^{\text{morning}}, 0\}$ if a transaction takes place, and $\text{payoff} = 0$ if no transaction takes place during the morning session. For each seller visited the buyer keeps track of his average experienced payoff. Note that the payoff is greater or equal to zero because whichever seller he chose to visit, whatever price this seller asks him, he has always the fall-back option of rejecting the price. In other words, if a buyer accepts a price giving him a negative payoff, he should not blame the seller he chose to visit, but his decision to accept that price. No transaction takes place in case a buyer is not served, finding only empty shelves, or if he rejects a price.

Second, there is the choice of a seller for the afternoon session. This decision box is specified analogously to the one for the morning session. Note, however, that it is a separate decision box. Some sellers might be good sellers to visit in the morning, but not in the afternoon, or the other way round. Sellers that are sold out during the morning session simply close for the afternoon, and will not be visited by buyers. The strengths of the rules for the afternoon session are updated using the following payoffs: $\text{payoff} = \max \{p^{\text{out}} - p^{\text{afternoon}}, 0\}$.

Third, the prices to accept or reject during the morning session. There are 21 possible prices: 0, 1, 2, ..., 20. The rules are of the form: <if $p = 0$ then reject>, <if $p = 0$ then accept>, <if $p = 1$ then reject>, ..., <if $p = 20$ then accept>. Thus we have a set of $21 \times 2 = 42$ price acceptance/rejection rules. For the daily update of the strengths of the rules, the following payoff from the environment is used as a reward. If the price is accepted then we have $\text{payoff} = p^{\text{out}} - p^{\text{morning}}$. Note that this payoff will be negative if $p^{\text{morning}} > p^{\text{out}}$. If the price during the morning session is rejected then the reward for that rejection depends upon what happens during the afternoon session. If during the afternoon session this buyer does not transact, then the rejection during the morning led eventually to a zero payoff. If however the rejection during the morning is followed by a transaction during the afternoon, the reward for the rejection during the morning will be determined as follows: $\text{payoff} = \max \{p^{\text{out}} - p^{\text{afternoon}}, 0\}$. Analogous to the reward for the choice of a seller, this

payoff is never negative as the buyer has, after a rejection during the morning, always the fall-back option of not buying during the afternoon.

Fourth, there is the choice of which prices to accept or reject during the afternoon session. This is analogous to the morning acceptance/rejection decision, but modeled separately. After the morning session, buyers can retry during the afternoon session, but after the afternoon session the trading day is over. Hence, prices that are unacceptable during the morning session might be acceptable during the afternoon. The reward for the price acceptance/rejection decision during the afternoon is determined by the following payoff: $\text{payoff} = p^{\text{out}} - p^{\text{afternoon}}$ if the price is accepted, and $\text{payoff} = 0$ if no transaction takes place, i.e., if the buyer rejects or is not served.

Some readers might wonder, if the buyers know p^{out} , they should know which prices to accept or reject during the afternoon session, and which prices are definitely unacceptable during the morning session. Hence, why do we model buyers who still have to learn this? First, this may be considered as a step towards a more general setting in which p^{out} is not given, and in which the buyers do not know the price for which they can resell outside the market. Second, if the buyers would know which prices to accept/reject, and would actually follow that knowledge, the sellers could work out which prices they should propose. Note that when a seller proposes a price to a buyer during the afternoon session, the situation resembles closely that of an ultimatum game (see, e.g., Güth & Tietz [1990]), with pie size p^{out} . Although the subgame-perfect equilibrium of that game in extensive form is offering the minimum slice size, we know that in laboratory settings many human subjects deviate from the game-theoretically 'correct' actions, and that, in general, the outcomes do not converge to this equilibrium. One explanation offered in the literature (e.g., Gale et al. [1995], or Roth & Erev [1995]), is that the point of convergence depends crucially upon the relative speed of learning of the two sides. Therefore, we do not want to impose which actions the players choose. Instead, we will analyze which actions emerge as they happen to have led to better outcomes. We will come back to this point in the next section, when discussing the results.

Sellers

Now we turn to the sellers. First, they decide which quantity to supply. The rules do not have a condition, and specify simply a quantity to supply: <supply 0>, <supply 1>, ..., up to <supply 30>. The strengths are updated daily, using as a reward from the environment: $\text{payoff} = \text{net profit}$.

Second, the sellers have to decide how to handle the queues they face. This is an example of a decision related to the 'who is going to interact with whom?' question.⁵ Basically, facing a crowd of customers, the seller has to decide at any moment which potential buyer to serve. Note that we assume the queues to be

⁵ Two other papers dealing with this issue are Stanley et al. [1993], and Vriend [1995].

`Italian' rather than `British'. Besides physically `reshuffling' queues, sellers can obtain the same result by, e.g., by putting some of their fish aside for some customers; a common practice for many shop keepers, but not one which is practised in the Marseille fish market. We assume that the buyers differ among each other, as seen by the sellers, in the sense that their faces will have a different degree of familiarity for the sellers. More specifically, this degree of familiarity of the face of buyer i to seller j on day t is indicated by a variable $L_{ij}(t)$ as follows:

$$L_{ij}(t) = \sum_{x=1}^t \frac{r_{ij}(t-x)}{(1+\alpha)^{t-x}} 1$$

with $r_{ij}(t) = \alpha = 0.25$ if buyer i visits seller j , and $r_{ij}(t) = 0$ otherwise. Hence, $0 \leq L_{ij}(t) \leq 1$. This degree of familiarity of a face of a buyer for a seller is a weighted average of the past presences of this buyer in the seller's queue; assigning more weight to recent visits. The question, then, is, should a seller serve familiar faces first, later, or should he be indifferent? Assume that a seller uses a roulette wheel to decide which buyer to serve, and that the slot size for each buyer i in the queue of seller j is equal to $(1+L_{ij})^b$. If the seller were indifferent between serving loyal or casual customers, he would use a roulette wheel with equal slot sizes for all buyers, which is achieved by setting $b=0$. If the seller wanted to give advantage to loyal customers, he would assign larger slot sizes to more loyal customers, which is achieved by setting $b>0$. Finally, disadvantage to loyal customers can be given by assigning them smaller slot sizes, setting $b<0$. Hence, given the degree of familiarity of the buyers' faces to a seller, the only decision variable needed by a seller to determine whether to give advantage or disadvantage to a loyal buyer, is this variable b . The decision box to decide upon the value for b looks as follows. The rules have no condition, and the action part is simply a value for b : $\langle b = -25 \rangle$, $\langle b = -20 \rangle$, ..., $\langle b = 0 \rangle$, ..., up to $\langle b = 25 \rangle$. The strengths of these rules are updated daily by using the following reward from the environment: payoff = gross revenue. Note that once the supply decision has been made, supply costs are sunk costs, and the seller's objective is to maximize gross revenue. We do not want to pretend that real sellers actually carry out such calculations, but casual empiricism suggests that sellers are able to distinguish between their customers on the basis of their degree of familiarity. If their behavior coincides with that of the agents in our model, this might suggest that they are behaving `as if' they carry out such calculations.

Once a seller has picked his next customer to be served, his third decision to make is the price to ask. As mentioned above, there are 21 possible prices: 0, 1, 2, ..., 20. The condition part of the rules takes account of the following three state variables. The loyalty of the buyer (3 classes are distinguished, `low', `middle', `high', in order to make the number of rules not too high), the remaining stock of fish, and the remaining queue. In fact, we take the ratio of these last two variables (again distinguishing three classes as above in order to limit the number of rules). This leads to the following set of $3 \times 3 \times 21$ price rules: $\langle \text{if loyalty} = \text{'low' and}$

ratio='low' then $p^{\text{ask}} = 0$ >, <if loyalty='low' and ratio='low' then $p^{\text{ask}} = 1$ >, ..., <if loyalty='high' and ratio='high' then $p^{\text{ask}} = 20$ >. The strengths of these rules are updated using as a reward from the environment: payoff = revenue obtained using that price rule, where the revenue = price if accepted, and 0 if the price is rejected.

The fourth decision for a seller is the prices to ask during the afternoon. This decision box is analogous to the one for the morning session. But again, note that this is a separate decision box, since sellers may learn that it is profitable to charge different prices in the two sessions.

The model we have sketched is a simple model, abstracting from various aspects of the real fish market. The price p^{out} is assumed to be identical for all buyers, and constant through time. The same applies to the price p^{in} . All buyers, and all sellers are always present on the market. There is only one good, which is supposed to be homogeneous. Each buyer wants one unit per day only. There is a strict division of days into morning and afternoon sessions. And buyers can make only one visit per sub-period. These are all abstractions from reality. But even with these simplifications the model is already complicated. The first question is whether, and how, this minimal model can generate the stylized facts of the real market. Note that little of the agents' actions is predetermined. We do not assume a reservation price property. There are no adjustment rules of the form: if demand > supply then p^{ask} increases. The treatment of loyal customers is not predetermined. The adaptive agents have to find out themselves what actions are good ones, and reinforcement learning simply corresponds to the economic idea that agents are seeking to do the best they can, without imposing certain reasoning procedures or heuristics.

4. Results

The results presented here are for one run only. By experimentation and 'playing around' with the model we know that the results are representative, but a systematic reporting on this will be postponed till a later version of the paper.⁶ Since for each period all actions of all agents are recorded, for example the buyers' decision with respect to any of the 21 prices they might face, including all

⁶ The results are also robust in the sense that they do not seem to be very sensitive to the exact algorithm used. We have tried alternative algorithms like Genetic Algorithms, hill climbing algorithms, and annealing schedules for the sellers' decisions without observing qualitative changes in the results in any significant sense. This robustness conforms the experience reported by Roth & Erev [1995]. The reinforcement algorithm is based on the idea that actions that have led to better payoffs are more likely to be used in the future. This principle is common to many learning dynamics, like e.g., replicator dynamics and all generalizations thereof.

rejections, and all buyers that were not served, the generated data set is much richer than the one covering the real fish market.

Prices

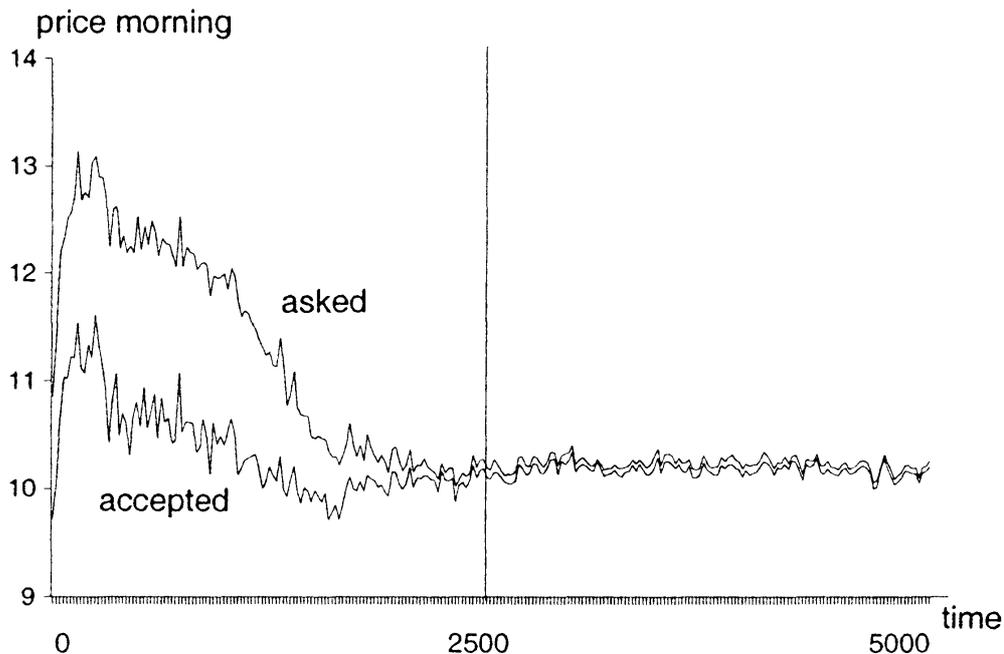


Fig. 4: Time Series Average Prices Asked and Accepted During Morning Sessions

Figure 4 presents the time series of all prices asked and all prices accepted during the morning sessions. For presentational reasons, each observation is the average of 20 days. Note, first, that the series start around 10, the average of the range of possible prices. Second, the prices asked increase faster than the prices accepted. Third, at some point, the latter starts 'pulling' down the prices asked. Fourth, prices accepted, then, decrease too much, and at some points are 'pulled' up again. Fifth, in the second half of the history there is almost perfect convergence of prices asked and prices accepted. Sixth, average prices are stable at a level of 10.2 during that period. A regression of prices asked in that half against time, gives a coefficient of -0.00001 that is statistically significant, but practically irrelevant (implying a price decrease of only 0.1 over a period of 10,000 days). These qualitative and quantitative features are representative for replications of the simulation.

Why do we not get a convergence of prices asked and prices accepted at a price equal to 14? That is at a price equal to $p^{\text{out}}-1$. As mentioned in the previous section, during the afternoon, when a buyer faces a seller, the situation resembles closely that of an ultimatum game. The subgame-perfect equilibrium of that game in extensive form is proposing the minimum slice size of the given pie to the

receiver. That is, $p^{\text{afternoon}} = p^{\text{out}} - 1$. Backward induction, then, leads to $p^{\text{morning}} = p^{\text{afternoon}} = p^{\text{out}} - 1$. A difference with the ultimatum game is that the interactions between the buyers and sellers are repeated. However, it is not simply a repeated ultimatum game, because the repetition is not enforced; both sellers and buyers can refuse to play the game again with each other. Hence, the effect of repetition is much weaker than in a standard repeated game.

Observe, in figure 4, that the sellers do try to drive the prices up towards this level of 14. But before the buyers learn to accept those prices, the sellers learn not to try them anymore, because they get rejected too often. This conforms to experience with the ultimatum game. Even with perfect information, there is in general no convergence to the subgame-perfect equilibrium; not in a laboratory setting with human subjects, and not in simulations with artificial players (see, e.g., Gale et al. [1995] with replicator dynamics, or Roth & Erev [1995] with reinforcement learning). As mentioned above, an explanation for this result is that the point of convergence depends crucially upon the relative speed of learning of the two sides. Hence, as far as the modeling of the individual players is concerned, there are two extreme possibilities. First, impose on all agents to play right from the start the game-theoretically 'correct' actions, although this would be at odds with much real world evidence. Second, do not tell the agents anything about which prices are good to ask, accept, or reject, let the agents learn about this all by themselves, and analyze which actions emerge. Any modeling option between these extremes would impose some of the agents' actions, taking away some of the emergent character of the actions, and breaking the symmetry between buyers and sellers, or between the morning and afternoon sessions. Clearly, if we had imposed on the buyers or sellers which prices to propose, accept, or reject, the emergent price patterns would have been very different.

Notice that a simple demand and supply analysis would have predicted a very different price than the game-theoretic analysis. Market supply is perfectly elastic at $p^{\text{in}} = 9$, whereas market demand is perfectly inelastic up to $p^{\text{out}} = 15$, leading to a market clearing equilibrium at $p = 9$.

For the afternoon, we get a similar graph for the prices, but as we see in figure 5 with some differences. There is slightly more convergence of prices asked and prices accepted. Almost no price is ever rejected during the afternoon. This convergence occurs at a stable, but somewhat higher price level than in the morning sessions; 11.2 against 10.2.⁷ This is in agreement with the real data.⁸ The precise pattern during the first half of the 5000 days is somewhat different from the morning sessions.

⁷ A regression of prices asked in that half against time, gives a coefficient of 0.00007 that is statistically significant (implying a price increase of only 0.7 over a period of 10,000 days).

⁸ In the real fish market there is some falling off at the very end of the day when there are mainly transaction of very small sizes. Notice that in our model all transaction have a unit size.

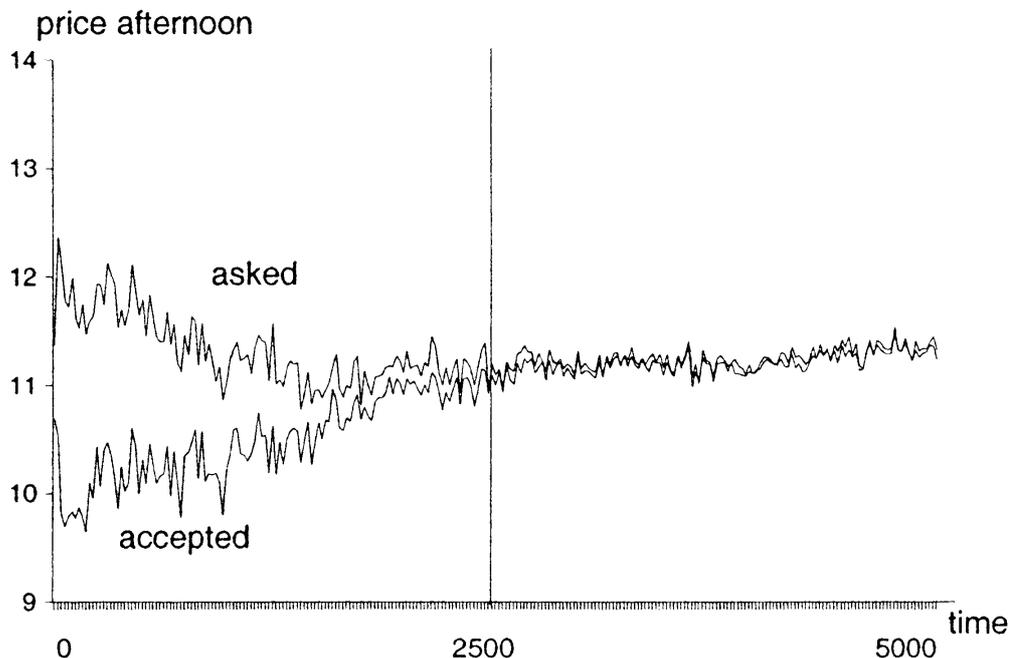


Fig. 5: Time Series Average Prices Asked and Accepted During Afternoon Sessions

Some readers might wonder, if prices are higher in the afternoon than in the morning, then buyers must be boundedly rational in an extreme way. But that would be a wrong observation. Suppose that there is exactly the same distribution of prices *asked* in the morning and in the afternoon, and suppose that those buyers that encounter prices in the upper tail of the distribution reject in the morning. The result will be that the average price *paid* in the morning will be lower than the average price *paid* in the afternoon. And it is only these prices actually *paid* that are available in the real market data. Now, suppose that we shift the morning price distribution somewhat to the left. The same argument used above implies that we could even have that the average price *asked* in the morning is smaller than the average price *asked* in the afternoon; and that with rational agents. The point is that it is not the average price that gets rejected. A sharper observation would seem that if we look at the simulation data, and if buyers are rational, then one should get that the average price *rejected* in the morning should be higher than the average price *asked* in the afternoon (neglecting for the moment the possibility of arriving 'too late'). But even that observation turns out to be incorrect; and that although our agents are so rational that the reinforcement of their acceptance/rejection rules in the morning are as follows: reinforcement $\text{accept}(\text{price}) = \text{payoff morning}$, and reinforcement $\text{reject}(\text{price}) = \text{payoff afternoon}$. How is this, then, possible? Suppose that the distribution of prices asked by the sellers adapts to the acceptance behavior of the buyers. The result will be that prices in the morning that are usually rejected will almost never be

asked, and the prices asked in the morning will usually be accepted. But the buyers experiment every now and then. Sometimes they reject a low price, although they know that on average accepting it had been better in the past. This brings the average price rejected in the morning down.⁹ Note that experimentation is not irrational. Its payoff is partly in the form of information. But the result of the adaptation by the sellers, plus the experimentation by the buyers is that there is a bias which means that the average price *rejected* in the morning may be lower than the average price *asked* in the afternoon.

What is the distribution of prices generated underlying these time series? Figure 6 presents the frequency distribution of the prices paid in the market during the last 2500 days of the artificial fish market. As we see, prices from 9 to 12 occur with frequencies greater than 8%, and prices from 9 to 11 with frequencies greater than 18%. Given the price for which the sellers buy outside the market ($p^{\text{in}}=9$), and the price for which the buyers resell outside the market ($p^{\text{out}}=15$), the 'effective' price range is 7.

Loyalty

The second stylized fact of the real fish market was the high loyalty of buyers to sellers. In the previous section we presented a measure of the loyalty of a given buyer i to a given seller j : $0 \leq L_{ij}(t) \leq 1$. Hence, for every buyer i , we can construct the following loyalty index: $\sum_{\text{seller } j} L_{ij}(t)^2 / \{\sum_{\text{seller } j} L_{ij}(t)\}^2$. This loyalty index will have a value equal to 1 if the buyer is perfectly loyal to one seller, and would have a value equal to $1/(\text{number of sellers})$ if the buyer has 'avoiding' behavior, visiting all sellers in a fixed sequence, and if there were no greater weight put on more recent visits.¹⁰

In figure 7 we present the time series for the morning sessions of this measure of loyalty averaged over all buyers, and the same series plus or minus two times the standard deviation. Each observation is the average of 20 days. We observe that loyalty does develop, but slowly, and with some variance over the buyers.

⁹ If buyers would sometimes accept a high price in the morning although they know that rejecting it had been better in the past, this will not influence the comparison given above.

¹⁰ Because of the discounting, the distribution of $L_{ij}(t)$ will not be uniform, but skewed, leading to a loyalty index of 0.14 in case of 'avoiding' behavior.

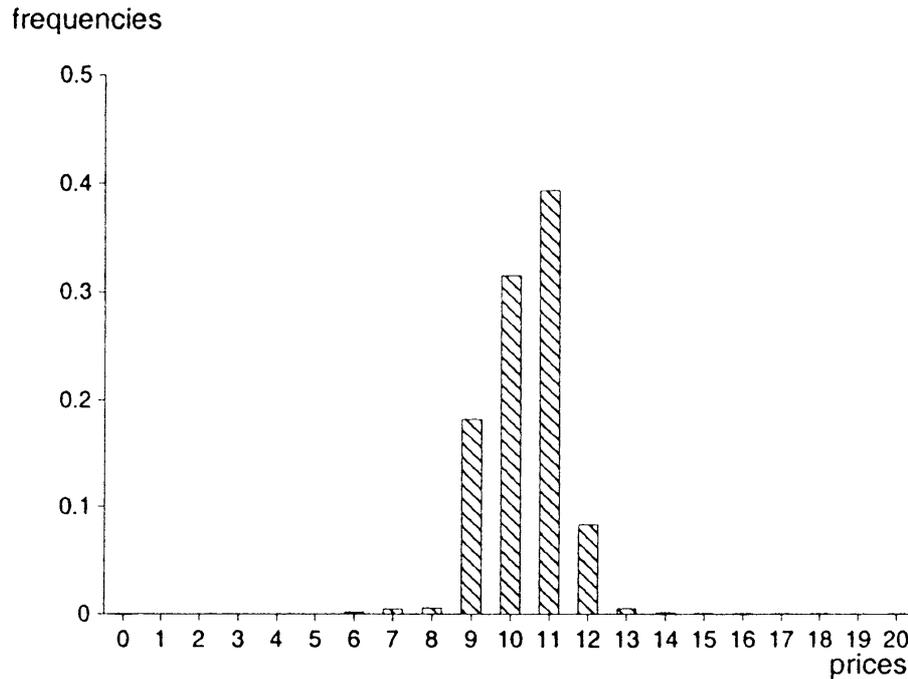


Fig. 6: Price Distribution Last 2500 Days

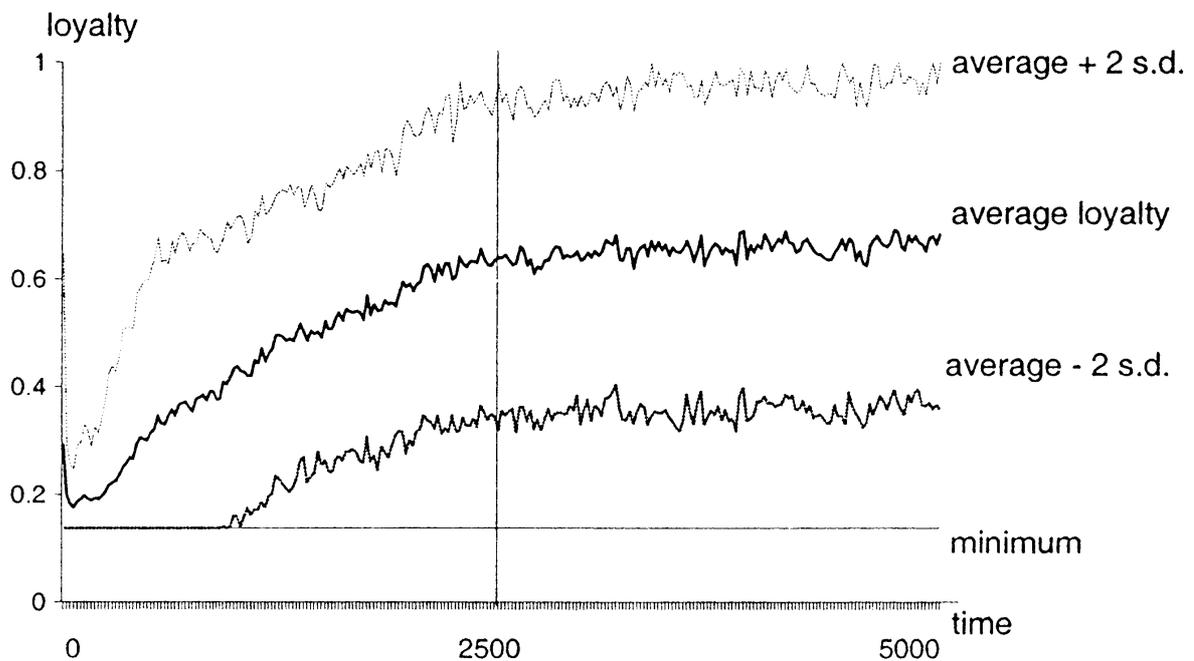


Fig. 7: Time Series Average Loyalty During Morning Sessions

The buyers do not know the concept of 'loyalty'. They simply choose a seller during every morning session. We did not impose assumptions which would favor the emergence of loyalty. For example, we could have imposed a 'forgetting'

mechanism; reducing the strength of the rule to visit a given seller with a fixed fraction every day that rule has not been chosen. Then, sellers not visited for some time will automatically gradually fade out from the buyers menu of potential sellers to be visited. Another assumption would have been to apply a so-called annealing schedule; gradually diminishing the noise added to the 'bids' by which sellers were chosen.

The sellers know the concept of 'loyalty'. They recognize which faces they have seen more often than others recently. But they are indifferent towards loyalty at the start. That is, they have no initial bias towards giving advantage or disadvantage to loyal customers in their queue, and no bias towards higher or lower prices to charge to loyal buyers.

Hence, the first question to ask is, why do buyers learn to become loyal? Figure 8 presents for all 100 buyers the difference in payoff they experience, averaged over the last 2500 mornings, when returning to the same seller as the previous morning, or when switching to a different seller.

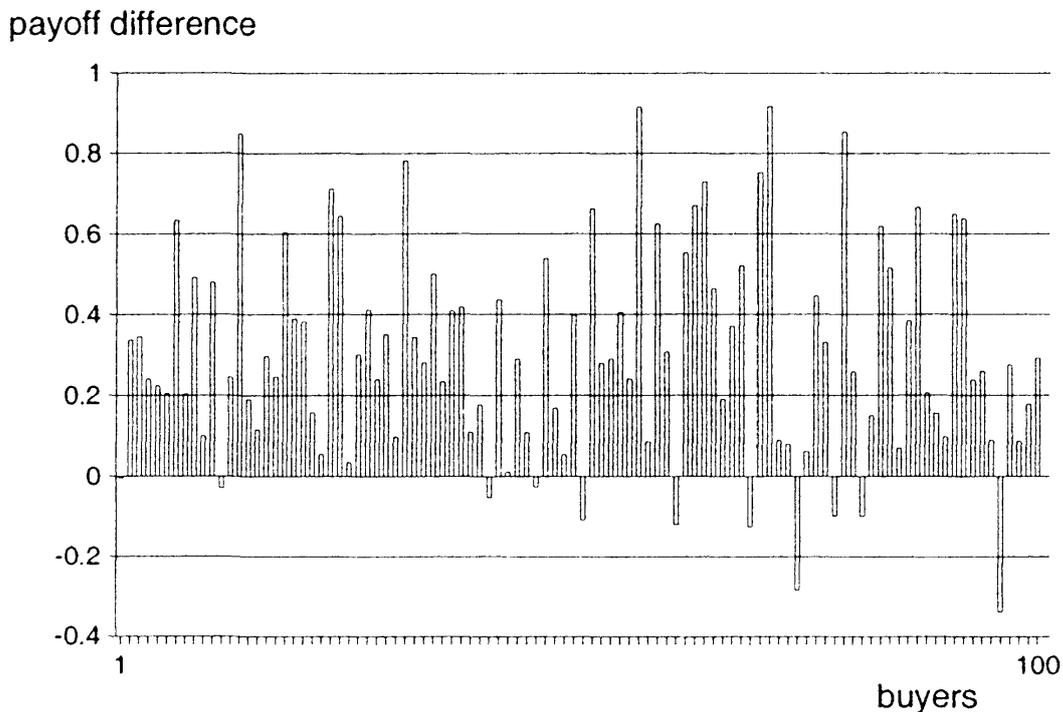


Fig. 8: Average Payoff Advantage Loyal Customers

As we see, 90% of the buyers experience a higher payoff when sticking to their seller than when switching, where the payoff to a buyer depends upon the prices

proposed to them and the service rate he gets.¹¹ Hence, we observe that visiting a seller where a buyer is loyal is simply more reinforced than visiting a different seller.

The next question is, why do sellers offer a higher payoff to loyal buyers than to other buyers? At the start, sellers were indifferent in this respect. Figure 9 presents for all 10 sellers the difference in payoff they experience, averaged over the past 2500 mornings, when dealing with a loyal customer or when dealing with a newcomer.

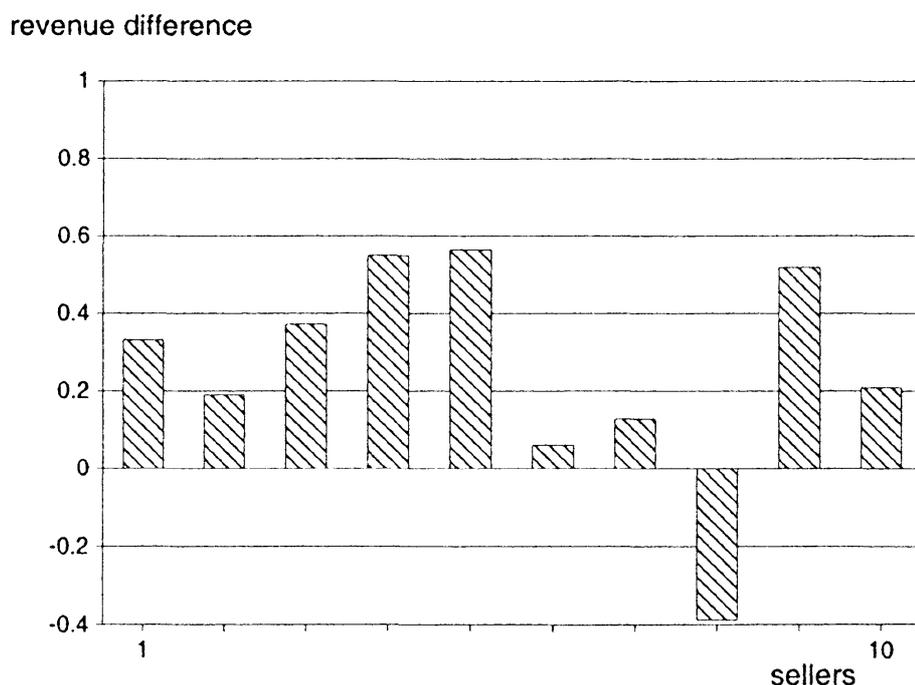


Fig. 9: Average Payoff Advantage Sellers Dealing with Loyal Customers

As we see, all but one sellers realize higher gross revenues when dealing with loyal buyers, where revenue depends upon the prices accepted and the rejection rate. Hence, offering a higher payoff to loyal customers makes sense.

In some sense what emerges here is remarkable. Both buyers and sellers are better off in loyal trading relationships. How is this possible? Note that this is not a constant-sum game. *If* two traders have decided to make a transaction then the determination of the price is a constant-sum game; given the pie size of $p^{\text{out}}=15$. The payoff for a buyer is $p^{\text{out}}-p^{\text{accepted}}$, whereas the payoff for a seller is p^{accepted} .

¹¹ Note that we made a simplification here. Loyalty is not a binary variable determined exclusively by the buyer's last visit, but a real-valued index between 0 and 1. It is a convenient approximation to present the results this way.

Hence, if, for a given the number of transactions, loyal buyers are better off on average than non-loyal buyers, then sellers dealing with loyal buyers are necessarily worse off on average. This is confirmed in table 1.

	Buyers	sellers	sum
Loyal	4.83	10.17	15.00
Switching	4.81	10.19	15.00

Table 1 Average Payoff per Transaction for Buyers and Sellers

But this is not the correct measure. In some cases a buyer chooses a seller, or a seller chooses a buyer with no transaction taking place. Sellers can sell out before serving a buyer, and buyers can refuse prices proposed by sellers. Hence, the game of choosing a potential trading partner is not a constant-sum game; the number of pies to be divided depends upon the players' actions. The average payoff offered to a buyer when choosing a seller can be written as follows, where the last factor on the right-hand side is simply the average service rate offered by the sellers chosen by the given buyer:

$$\frac{\sum (p^{out} - p^{ask})}{\#visits} = \frac{\sum (p^{out} - p^{ask})}{\#proposals} \cdot \frac{\#proposals}{\#visits}$$

Similarly, the average payoff for seller when proposing a price to a buyer in his queue can be written as follows, where the last factor on the right-hand side is the average acceptance rate of the buyers served by the given seller:

$$\frac{\sum (p^{accepted})}{\#proposals} = \frac{\sum (p^{accepted})}{\#transactions} \cdot \frac{\#transactions}{\#proposals}$$

Notice that a high service rate benefits the buyers without directly hurting the sellers, whereas a high acceptance rate benefits the sellers without directly hurting the buyers. Table 2 gives the average outcomes for the last 2500 mornings for the buyers, and for the sellers, distinguishing loyal and non-loyal interactions. As we see, loyal buyers tend to have a higher acceptance rate, and get a higher service rate than non-loyal buyers. In other words, the total surplus that can be realized is greater with loyal trading relationships, and therefore both sides can be better off through loyalty.

Thus, the important factor of the model is that both sides are learning about a situation which is changing precisely because their counterparts are also learning.

There have been few studies in which analytic results are obtained for such situations, although the reinforcement learning models used in game theory have this feature (see Roth & Erev [1995] for a discussion). Figure 10 presents the

coevolutionary process, graphing the simultaneous emergence of the payoff advantage from loyal trading relationships to both buyers and sellers, and the emergence of this loyalty itself.

	#proposals/#visits (service rate)	#transactions/#proposals (acceptance rate)
Loyal	0.83	0.86
Switching	0.79	0.83

Table 2 Average Service and Acceptance Rates for Two Types of Buyers

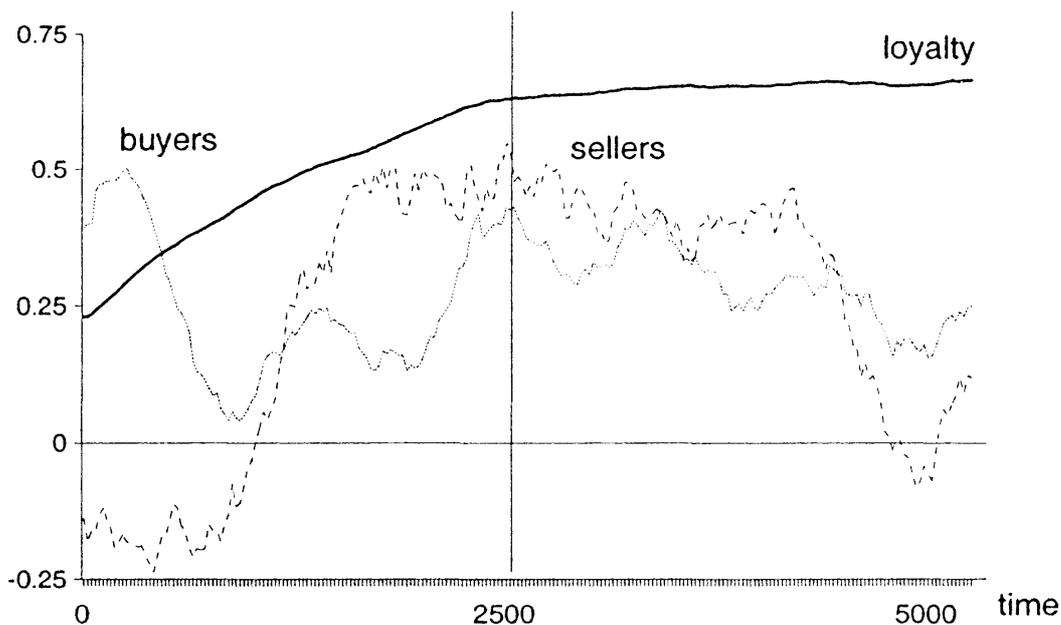


Fig. 10: Coevolution: Emergence Payoff Advantages of Loyalty to Buyers and Sellers, and Loyalty

Each observation is again the average of 20 mornings, and we smoothed the series by taking 25 period moving averages. Note, first, in the beginning sellers get a lower payoff from loyal customers than from new customers. Second, simultaneously, the payoff advantage to loyal buyers collapses. Third, only when a payoff advantage to sellers dealing with loyal customers emerges does the payoff advantage to loyal customers start to rise again. Fourth, in the meantime, loyal buyers had always experienced a positive payoff difference, and loyalty emerged.

Fifth, during the main part of the simulated history, the advantage of loyal relationships is positive to both buyers and sellers, and loyalty develops at a steady level around 0.66. Sixth, at the end, the advantage of loyalty to buyers and sellers collapses simultaneously, and re-emerges, whereas loyalty itself remains constant. Hence, one of the questions to be analyzed will be, what happens in the longer run with the variables in figure 9. For example, it might be that the actual advantage of loyalty to both sellers and buyers vanishes, but that loyalty itself continues simply because the buyers are locked in in their loyal shopping behavior.

5. Concluding Remarks

In our minimal model of the fish market both stylized facts price dispersion and loyalty do emerge during a coevolutionary process, in which both sides adapt to their changing environment. Buyers learn to become loyal, as sellers learn to offer a higher payoff to loyal buyers, while these sellers, in turn, learn to offer a higher payoff to loyal buyers, as they happen to realize a higher payoff from loyal buyers.

In some sense, this suggests that developing loyalty has some similarity with sending intrinsically meaningless signals like wearing a blue shirt. The similarity lies herein that with our assumptions there was nothing intrinsic in loyalty that makes it pay. There are no real costs of switching to another seller. Hence, at first sight one might conjecture that our conclusion could as well have been the other way round: 'Buyers learn to become non-loyal, as sellers learn to offer a higher payoff to non-loyal buyers, while these sellers, in turn, learn to offer a higher payoff to non-loyal buyers, as they happen to realize a higher payoff from non-loyal buyers'. However, matters are not this simple. Although this similarity with intrinsically meaningless signals seems true to some degree, there are also some important differences. Loyalty, whatever its degree, develops automatically as a result of market behavior, without explicit, additional non-market decisions like the color of your shirt. Moreover, being loyal or non-loyal has a direct economic meaning. Suppose, for example, that there appears to be serial correlation in a seller's decisions, and that a buyer is satisfied with that seller. Loyalty would pay to that buyer, but continuing to dress blue while shopping around randomly would not. Or suppose that a seller offers a bad service. One of his buyers becoming non-loyal would hurt, but that buyer merely changing the color of his shirt would not. The basic problem here is one of coordination. Loyalty means continuity, and allows buyers and sellers to avoid unproductive meetings.

We have shown how an explicit consideration of market interactions between the individual agents may lead to the explanation of stylized facts that were not easily explained by standard arguments neglecting this interaction. Already with this minimal model, starting with a population of identical buyers and identical sellers, the behavior displayed by the agents is rather rich, and much more has to be analyzed; for example, the exact pricing behavior of the sellers. Also, a cursory

examination of the data shows the emergence of heterogeneity among buyers and among sellers, i.e., of 'types' of agents, and relationships between 'types' of agents. For example, many sellers realize most of their sales in the morning, but some become typical afternoon sellers. And it appears that both types of implicit contracts do emerge. Some sellers offer their loyal customers good service with higher prices, whereas at the same time some other sellers combine lower prices with bad service for their loyal customers.

References

- Burdett, K., & Judd, K. (1983). Equilibrium Price Dispersion. *Econometrica*, 51, 955-969.
- Bush, R.R., & Mosteller, F. (1955). *Stochastic Models for Learning*. New York: Wiley.
- Futia, C.A. (1977). Excess Supply Equilibria. *Journal of Economic Theory*, 14, 200-220.
- Gale, J., Binmore, K., & Samuelson, L. (1995). Learning to Be Imperfect: The Ultimatum Game. *Games and Economic Behavior*, 8, 56-90.
- Graddy, K. (1995). Testing for Imperfect Competition at the Fulton Fish Market. *Rand Journal of Economics*, 26, 75-92.
- Güth, W., & Tietz, R. (1990). Ultimatum Bargaining Behavior: A Survey and Comparison of Experimental Results. *Journal of Economic Psychology*, 11, 417-440.
- Härdle, W., & Kirman, A.P. (1995). Nonclassical Demand: A Model-Free Examination of Price-Quantity Relations in the Marseille Fish Market. *Journal of Econometrics*, 67, 227-257.
- Hopkins, E. (1995). *Price Dispersion: an Evolutionary Approach* (mimeo).
- Kirman, A.P., & McCarthy, M. (1990). *Equilibrium Prices and Market Structure: The Marseille Fish Market*.
- Kirman, A.P., & Vignes, A. (1991). Price-Dispersion: Theoretical Considerations and Empirical Evidence from the Marseille Fish Market. In K.J. Arrow (Ed.), *Issues in Contemporary Economics: Proceedings of the Ninth World Congress of the International Economics Association, Athens, Greece*. (Vol. 1. Markets and Welfare, pp. 160-185). London: MacMillan.
- Kirman, A.P., & Vriend, N.J. (1995). *Evolving Market Structure: A Model of Price Dispersion and Loyalty* (mimeo) Paper presented at the Workshop Economic ALife, Santa Fe Institute.
- Roth, A.E., & Erev, I. (1995). Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term. *Games and Economic Behavior*, 8, 164-212.
- Sarin, R., & Vahid, F. (1997). *Payoff Assessments without Probabilities: A Simple Dynamic Model of Choice* (mimeo).

- Sarin, R., & Vahid, F. (1998). *Predicting How People Play Games: A Procedurally Rational Model of Choice* (mimeo).
- Stahl, D.O. (1988). Bertrand Competition for Inputs and Walrasian Outcomes. *American Economic Review*, 78, 189-201.
- Stanley, E.A., Ashlock, D., & Tesfatsion, L. (1993). *Iterated Prisoner's Dilemma with Choice and Refusal of Partners* (ISU Economic Report No. 30) Iowa State University.
- Vignes, A. (1993). *Dispersion de prix et marchés décentralisés: le cas du marché au poisson de Marseille* (Ph.D. thesis) European University Institute, Florence.
- Vriend, N.J. (1995). Self-Organization of Markets: An Example of a Computational Approach. *Computational Economics*, 8, 205-231.
- Weisbuch, G., Kirman, A., & Herreiner, D. (1996). *Market Organisation* (mimeo).