



Evolving market structure: An ACE model of price dispersion and loyalty[☆]

Alan P. Kirman^a, Nicolaas J. Vriend^{b,*}

^a*E.H.E.S.S., GREQAM, Marseille, France*

^b*Department of Economics, Queen Mary and Westfield College, University of London, Mile end Road, London, E1 4NS, UK*

Accepted 22 February 2000

Abstract

We present an agent-based computational economics (ACE) model of 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 it is every day the same population of sellers and buyers that meets in the same market hall. In our ACE model, sellers decide on quantities to supply, prices to ask, and how to treat loyal customers, while buyers decide which sellers to visit, and which prices to accept. Learning takes place through reinforcement. The model explains 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. The model also explains the effect of heterogeneity of the buyers. We analyze how this leads to subtle differences in the shopping patterns of the different types of buyers, and how this is

[☆]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), and Keele, participants at conferences of the SEDC in Barcelona, SCE in Geneva and Cambridge, ESEM and EEA in Istanbul, ASSET in Alicante and Marseille, IIFET in Tromsø, WEHIA in Ancona, and ASSA in NYC, and Antonio Cabrales, Art De Vany, Paul Pezanis, Willy Spanjaars, and Johan Stennek for helpful comments and discussions. The usual disclaimer applies.

* Corresponding author. <http://www.qmw.ac.uk/~ugte173/>.

E-mail address: n.vriend@qmw.ac.uk (N.J. Vriend).

related to the behavior of the sellers in the market. © 2001 Elsevier Science B.V. All rights reserved.

JEL classification: C70; D40; D80; L10; L70

Keywords: Price dispersion; Loyalty; Adaptive behavior

1. Introduction

We study the working of the wholesale fish market in Marseille (France), and in particular we focus on the following two stylized facts that characterize this market: a widespread high loyalty of buyers to sellers, and persistent price dispersion. Our real interest as economists is not in fish markets per se, and unlike the Scots we would not question the relevance of any story just heard with the standard expression ‘*but what’s it got to do with the price of fish?*’. But casual observations suggest that these stylized facts are common to many other markets,¹ we do believe that some of the insights developed in this study might be carried over to such markets.

More in general, 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. As we argued elsewhere (e.g., Kirman, 1994; Vriend, 1994, 1995, 1996, or Vriend, 1999), a natural way to do this is following an agent-based computational economics (ACE) approach.

Two additional reasons why we were attracted to studying this specific market are, first, that 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. This is a very rich data set, but at the same time it has some severe limitations, as there are, for example, no data concerning market interactions that did not lead to transactions. Second, this wholesale fish market is a relatively simple, well-defined, and well-structured market (cf. the ‘*market*’ for second-hand cars). Moreover, fish is a perishable commodity, implying that the issue of the strategic use of inventories does not arise (cf. the early literature on market microstructures in finance; see, e.g., O’Hara, 1994).

The paper is organized as follows. In Section 2, we will sketch the Marseille fish market, and discuss the two not easily explained stylized facts. In Section 3,

¹ For example, according to Royal Mail (1999) “*more than 25 percent of all people do not consider themselves loyal to any business*”, implying that about 75% is loyal to some business.

we present an ACE model of the market. Section 4 analyzes the model, and Section 5 concludes.

2. The Marseille fish market

2.1. Some stylized facts

The real market is open every day from 2 a.m. to 6 a.m.² The market has a fixed population of about 40 registered sellers. They buy their supply outside the market before the market opens. There is also a fixed population of approximately 400 buyers. 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 what 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 demanded by the same buyer at different times of the day. A price communicated by an individual seller to an individual buyer for a given transaction is not perceived by other buyers or sellers in the market. The prices are ‘take-it-or-leave-it’ prices, and there is essentially no bargaining. At the end of the day, unsold fish is thrown away, or disposed of in return for minimal compensation. Some species can be offered for sale on the succeeding day. Buyers are, however, capable of recognizing the freshness of the fish. Typically, on most days the market clears or gets very close to clearing.

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 and Vignes (1991), Kirman and McCarthy (1990), Vignes (1993), Härdle and Kirman (1995), and Weisbuch et al. (2000).³ In 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 stylized facts mentioned.

2.2. Explaining the stylized facts

Since the same population of sellers and buyers is present every day in the same market hall, and since there are no real search costs, switching costs, or

² Since the data were obtained, the market hours have changed and are now in the normal working day.

³ For an empirical analysis of some other fish markets, see Graddy (1995), or Wilson (1980).

product differentiation (see below), 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. But, for various reasons, explaining these stylized facts turns out to be far from simple. Some standard arguments do not have much explanatory power, in the sense that they either 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 is simply related to 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, for example, when compared with the price dispersion between different types of fish. Next, the effect of quality differences might 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 is 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. Hence, the information asymmetry would explain the existence of loyalty of buyers to specific sellers. But this seems a too loose application of the textbook argument. In the Marseille fish market, there is a fixed population of buyers. That is, every buyer (loyal or not) is a *potential* repeat buyer. Hence, a seller would have an incentive to deliver good quality to each buyer. Clearly, a seller cannot offer higher quality to all buyers. Another attempt to save the causal link 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. Apart from the fact that this seems a flawed argument anyway, following exactly the same argument would imply 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 the 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. But as there is this trade-off between the price to be paid by the 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 more frequently *not* served 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. In the transaction data, such a buyer might appear as one of the non-loyal buyers, even when there is this long-term relationship characterized by poor 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 for some of them it is difficult or impossible to indicate 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 of price dispersion (see, e.g., Futia, 1977; or Burdett and Judd, 1983), there might be multiple equilibria, and it is not clear whether the stable equilibria are characterized by price dispersion.⁴

A third standard argument in the literature is 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 standard arguments would be to design a multi-stage game with an equilibrium characterized by the stylized facts

⁴ Hopkins (1995) presents a dynamic analysis of such models.

just mentioned. But designing and analyzing such a game might be complicated. Moreover, it is easy to imagine such a game will have multiple equilibria; and not only because of the equilibria being asymmetric. De Francesco (1996) follows this game-theoretic approach to explain loyalty. In a choosing-the-seller game with exogenously given identical and constant prices, fixed and known capacity constraints, and fixed demand, De Francesco proves existence of an equilibrium characterized by the following rule for the buyers: stay if served, switch seller otherwise. Besides not examining price dispersion, and the restrictions listed, this result does not exclude that other equilibria exist as well, and it also remains an open question whether any plausible dynamics would lead to such an equilibrium.

Moreno and Wooders (1997) do not consider loyalty, but they analyze pricing behavior in a game where sellers deal sequentially with buyers with private values. Their analysis concerns a static game that would effectively correspond to one single market day of our fish market only. The essential issue arising when extending this kind of analysis to a repeated setup would be that private information, which plays a pivotal role in their analysis, becomes much less compelling, and that instead one would need to model exactly how and what the agents learn about these values over time.

Bergemann and Valimaki (1995, 1996) model price-setting behavior in markets as well as how individuals within those markets might learn which seller to choose.⁵ They consider a situation in which there are two products which are differentiated, and they are sold by duopolists to either a continuum or a finite number of buyers. In their case, the ‘standard’ product is known and the ‘value’ of the other product is unknown. Now the problem is that the firms will try and charge prices to make people experiment, while the buyers on the other hand will want to experiment to try and find out what the quality of the other product is. The value of the product is determined by the match of the buyer and the seller. In our case, the situation is somewhat different since what determines the value of a match from the buyer’s point of view is the price which the seller will charge. So in our case the seller controls it, whereas in Bergemann and Valimaki’s case it is exogenous. They consider their actors to be Bayesian and to behave accordingly, and their discussion centers around the problem of what will happen as the market evolves. In the Bergemann and Valimaki model, in each period all buyers receive a noisy signal concerning the true value of the product with an indeterminate quality, and in addition they receive an average or aggregate signal which tells them the aggregate assessment of the value of the new product. In the case where there are an infinite number of buyers, of course, it is sufficient for the buyers who are negligible in size to look only at the

⁵ Earlier papers in this spirit are those by Aghion et al. (1991) and Bolton and Harris (1993), together with Felli and Harris (1994).

aggregate signal. Firm prices in their model are adjusted over time as a result of experience. If the buyers have a strategic effect, that is, if their presence in the past as buyers at a particular firm has some influence over what the firm should do in the future, then the situation is very complicated.

The limitations of the case which Bergemann and Valimaki are able to analyze are indicated by the fact that they have only two products (sellers), one of which is of known value, and in most of their analysis there are very simple restrictions such as that if every consumer chooses one product then its quality turns out to be high whereas if fewer than a full measure of consumers choose that product then its value turns out to be low. These sort of extreme restrictions are imposed in order to obtain analytic results. Although their model allows for a high level of rationality on the part of individuals, it is very restrictive in terms of the cases it is able to consider. Thus, what we have here is a situation in which a complete model with completely rational individuals becomes analytically intractable as soon as we try to relax some of the restrictions. Nevertheless, even within their structure we see certain features which recur in our model. In particular, the idea that as the partners learn their prices ‘track’ each other occurs in both models. However, as we will see, our less rational individuals reverse direction as they learn, whereas in the Bergemann and Valimaki setup the learning is monotone.

As we will explain in detail in Section 3, what we will do instead is something very simple. We will build an ACE model of the Marseille fish market. With respect to the problems sketched above, an advantage of this ACE approach is that the analysis of the evolutionary aspects of the 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 things really *implicit*. That is, keeping aggregate demand and supply conditions constant (although individual supply and demand decisions may vary), we will analyze whether long-lasting trading relationships can be explained by the ACE model, 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.

3. An ACE model of the Marseille fish market

3.1. ACE modeling

Social theory is the explanation of social phenomena, where explanation is the modeling of the events to be explained. In contrast to the natural sciences, in the

social sciences “(e)xperimentation is impossible, and we have therefore no knowledge of definite regularities in the complex phenomena in the same sense as we have in the natural sciences” (Hayek, 1948b, p. 126). Although in the meantime experimental economics has been established as a well-developed research field, what still sets economics apart from, for example, physics is that, with or without experiments, economics yields no natural laws or universal constants.⁶ Therefore, as Hayek (1948a) put it, what we need to do is, “construct hypothetical models in an attempt to reproduce the patterns of social relationships which we know in the world around us” (p. 68). These theoretical models, then, can be ‘tested’ only in the sense that we check the consistency of our model with our knowledge of these social phenomena.

Such theoretical models can be presented in various forms. They could be either purely informal (verbal) or formal (mathematical or computational). In other words, a computer program as used in ACE constitutes in itself a model, i.e., an explanation of social phenomena. Notice that this implies that the following view of ACE research is incorrect. According to some people, reality produces *facts* (which need to be explained), and a computer program generates *artifacts* (which, in turn, need to be explained). The alleged objective, then, would be to show that the explanation for the facts could be the same as the explanation for the artifacts. As we explained, the output of an ACE computer program are *not* artifacts to be explained. The computer program *itself* is the model that explains the social facts. In fact, an ACE model stands to executing its program as a mathematical model stands to solving its equations. The only, and essential, reason to execute the computer program of an ACE model is to carry out the consistency checks; both with respect to reality and with respect to the expected behavior of the model. An advantage of quantitative models in general is that they can be analyzed more precisely. Such a consistency analysis can be a formal, mathematical analysis, or it could be a numerical analysis. A great advantage of ACE models, then, is that one can do very extensive and elaborate consistency checks.

The idea that it is not possible to test the truth of a social theory, and that the best we can aim for is doing consistency checks, resembles Friedman’s (1953) ‘*as if*’ argument. As Hayek (1967) put it: “We may not be able directly to confirm that the causal mechanism determining the phenomenon in question is the same as that of our model. But we know that, if the mechanism is the same, the observed structures must be capable of showing some kinds of action and unable to show others; and if, and so long as, the observed phenomena keep within the range of possibilities indicated as possible, that is so long as our

⁶ See the editorial in *The Economist* (1999) for a recent discussion of this issue. Another way to suggest a profound difference between physics and economics is the following comment attributed to Nobel Prize winner in physics Murray Gell-Mann: ‘*Imagine that electrons could think*’.

expectations derived from the model are not contradicted, there is good reason to regard the model as exhibiting the principle at work in the more complex phenomenon. () Our conclusions and predictions will also refer only to some properties of the resulting phenomenon, in other words, to a kind of phenomenon rather than to a particular event” (p. 15). Every model is by definition an abstraction, and if enough data can be collected, statistical testing will reject any model. The only way to avoid this would be to create a copy of the real market. Therefore, we will not try to build an ACE model fitting all aspects of the Marseille fish market. We will only consider specific questions concerning the stylized facts of the real market that appear remarkable or important. We will build an ACE model that explains those stylized facts, suggesting ways to understand those phenomena.⁷

3.2. *The setting*

The place of action in our ACE 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.⁸

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 one 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. Notice that each buyer can visit at most one seller in the morning plus one seller in the afternoon, but the model could be extended to allow for any finite number of sessions in the day. The sellers then serve the people in their queue in an order that they choose, after which the end of the afternoon session is reached. All unsold stocks perish. The buyers, then, resell their fish outside market at a given price that is identical for all buyers, and constant through time ($p^{\text{out}} = 15$). In a variant of the base model we introduce heterogeneity of the buyers. In particular, we consider a market with three types of buyers; 33 buyers with $p^{\text{out}} = 12$, 34 with $p^{\text{out}} = 15$,

⁷ See Baillie and Bollerslev (1997) for similar arguments. Since ACE models are *abstractions* from reality, and are not aimed at *replicating* reality, the term ‘*simulation*’ to describe ACE models might create confusion. ACE models do not try to simulate reality as such, but only to explain some general phenomena, the stylized facts.

⁸ See the appendix for the pseudo-code of the model.




Fig. 1. Sketch of the market: (a) Buyers enter the market hall; and (b) each buyer has chosen a seller.

and the remaining 33 buyers with $p^{out} = 18$. Fig. 1 presents a sketch of the market.

3.3. Adaptive behavior

On each day, the buyers need 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 (see Holland, 1975, for early ideas on this, or Holland et al., 1986, for a more elaborate treatment). Considering a classifier system as a ‘black box’ for the moment, this implies that in our ACE model each single buyer, and each single seller has four decision boxes in his head. Hence, with 10 sellers and 100 buyers we model explicitly 440 decision boxes by separate classifier systems. Fig. 2 presents one such stylized 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 s . The classifier system does two things. First, it decides which of the rules will be the active rule at a given time t . Hence, it checks the condition part, and each rule j satisfying the ‘if...’ condition makes a ‘bid’ b as follows:

$$b_j(t) = s_j(t) + \varepsilon,$$

where

$$\varepsilon \simeq N(0, \sigma). \tag{1}$$

condition	action	strength
if	then
..
..

Fig. 2. A classifier system.

The rule with the highest ‘bid’ in this ‘stochastic auction’ wins the right to be active.⁹ Second, the classifier system updates the strength s of a rule j that has been active and has generated a reward π from the environment on a given day $t - 1$ as follows:

$$s_j(t) = s_j(t - 1) - cs_j(t - 1) + c\pi(t - 1),$$

where

$$0 < c < 1. \tag{2}$$

Hence, $\Delta s_j(t) = c[\pi(t - 1) - s_j(t - 1)]$. In other words, as long as the payoff generated by a 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 some weighted average of the rewards from the environment generated by that rule. In our implementation of this model, the strengths of all rules are equal at the start.¹⁰

Classifier systems are a form of reinforcement learning (see, e.g., Bush and Mosteller, 1955; Sutton, 1992; or Roth and Erev, 1995). Reinforcement learning is based on two principles. First, agents try actions. Second, actions that led to better outcomes in the past are more likely to be repeated in the future. Reinforcement learning is among the most basic forms of learning. It is a minimal form of modeling learning, in the sense that one does not need to make many assumptions about the reasoning procedures followed by the agents. One would

⁹ Besides through the noise added to the ‘bids’, the agents in our model will experiment through some kind of ‘trembling hand’, effectively disregarding any ‘bid’ with a given small probability.

¹⁰ We presented this specific learning model in Kirman and Vriend (1995). Sarin and Vahid (1997) analyze some theoretical properties of this model, relating it also to evolutionary explanations.

expect more elaborate learning models to share many of the behavioral features of reinforcement learning models.

A closely related contribution to the literature on learning in economics is Easley and Rustichini (1999). Easley and Rustichini look at a situation in which individuals do not know the full structure of the model and are simply aware of the payoffs they receive from their own actions, which depends on the realized value of the state of the system. These states are identically and independently distributed but the individual does not know this distribution. They, then, impose a set of axioms which shows that the adaptive learning procedure of this type will lead to the Nash equilibrium which would have occurred had all individuals been aware of the full structure of the model.

The major difference between game-theoretic models and models such as ours, which is related to that of Easley and Rustichini, is the following. In a full blown game-theoretical setup we have a situation in which individuals contemplate their strategies in the light of the strategies of other people, and then work out which strategy is best to adopt. That is, game-theoretic models are based on the notion that one anticipates what opponents will do in the future given their rational behavior, and therefore one plays in function of this anticipation. That the individuals are aware of the strategies of the other players, and of the relationship between the strategies and the payoffs that they receive, is even true in dynamic approaches, such as ‘best-response’ behavior, or ‘fictitious play’. With ‘best response’ dynamics, players choose their best response to what they anticipate their opponents to be about to play or to their previous choices or strategies. With ‘fictitious play’, one plays a best response to the average or to the mixed strategy which is generated by the frequencies of the players’ strategies in the past. All of this relies on some sort of comprehensive knowledge of the underlying model. On the other hand, in the sort of reactive or primitive learning models such as Easley and Rustichini (1999) or the one which we use, no such strategic reflection takes place.

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.

3.4. Individual decisions: Buyers

As said above, a buyer needs to make four types of decisions. First, the buyer must make the choice of a seller for the morning session. The rules have no condition. The actions are simply: $\langle \text{choose seller } 1 \rangle$, $\langle \text{choose seller } 2 \rangle$, ..., up to $\langle \text{choose seller } 10 \rangle$. To evaluate the benefit of each of these rules, the buyer keeps track of the average payoff he experienced when visiting a given seller. The strength of each activated rule is updated every day using the following payoff as reward from the external environment: $\text{payoff} = \max\{p^{\text{out}} - p^{\text{found}}, 0\}$ if a transaction takes place or the price is rejected, and $\text{payoff} = 0$ if the buyer is not served, finding only empty shelves during the morning session. Note that the

payoff is greater or equal to zero because whichever seller he chose to visit, and whatever price this seller asks him, he always has 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.

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: $\text{payoff} = \max\{p^{\text{out}} - p^{\text{found}}, 0\}$ if a price offer has been received, and 0 if the buyer arrives late.

Third, a buyer needs to decide the prices to accept or reject during the morning session. We assume that there are 21 possible prices: 0, 1, 2, ..., 20. The rules are of the form: $\langle \text{if } p = 0 \text{ then reject} \rangle$, $\langle \text{if } p = 0 \text{ then accept} \rangle$, $\langle \text{if } p = 1 \text{ then reject} \rangle$, ..., $\langle \text{if } p = 20 \text{ then accept} \rangle$. 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 a buyer receives a 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 eventually led 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. In other words, if a buyer decides to purchase at a loss in the afternoon, he should not blame his rejection decision of the morning session for this.

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 may 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 that of an ultimatum game (see, e.g., Güth and Tietz, 1990), with cake 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 most human subjects deviate from the game-theoretically ‘correct’ actions, and that, in general, the outcomes do not converge to this equilibrium. 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 Section 4, when analyzing the model in detail.

3.5. Individual decisions: Sellers

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

Second, the sellers have to decide how to handle the queues they face.¹¹ Basically, facing a crowd of customers, the seller has to decide at any moment which potential buyer to serve next. Note that we assume the queues to be ‘Italian’ rather than ‘British’. Besides ‘reshuffling’ queues, sellers can obtain the same result, for example, by putting some of their fish aside for some customers; a common practice for many shop keepers, but not one which is practiced 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 modeled by a variable $L_{ij}(t)$ as follows:

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

¹¹ After the buyers’ choice of sellers, this is another example of a decision related to the ‘who is going to interact with whom?’ question. Some other papers dealing with this issue are Mailath et al. (1993), Stanley et al. (1993), Durlauf (1995), and Vriend (1995).

with

$$\begin{aligned} r_{ij}(t-x) &= \alpha && \text{if buyer } i \text{ visited seller } j \text{ on day } x, \\ r_{ij}(t-x) &= 0 && \text{otherwise,} \end{aligned} \quad (3)$$

where α is a parameter, with $0 < \alpha < 1$.

Hence, this degree of familiarity of a face of a buyer for a seller is a weighted average of the past appearances of this buyer in the seller's queue, assigning more weight to recent visits. As we see, $0 \leq L_{ij}(t) \leq 1$, with a value of one applying to perfectly loyal customers, and zero for a buyer never seen. 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 that he spins each time he needs to decide which buyer to serve next, 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 for a seller to determine whether to give advantage or disadvantage to loyal buyers, is this variable b , which we assume is chosen only once a day. 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: $\text{payoff} = (\text{gross revenue on the day}) / [(\text{highest possible price}) \cdot \text{supply}]$. 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 presume that real sellers actually carry out such calculations and spin roulette wheels, but casual empiricism suggests that sellers are able to distinguish between their customers on the basis of their degree of familiarity, and it is a well-known fact in social psychology that face recognition is a crucial determinant of how people treat each other in direct exchange contacts. If this behavior is consistent with that of the agents in our model, this might suggest that they are behaving 'as if' they carry out calculations similar to those modeled. The idea is not that a seller has an 'optimal' way of choosing buyers, but simply that his probability of choosing a buyer increases with familiarity if familiar buyers have turned out to be profitable in the past.

Once a seller has picked his next customer to be served, his third decision to make is the price to ask from that customer. 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 (distinguishing the

three classes ‘low’, ‘middle’, and ‘high’), the remaining stock of fish, and the remaining queue. In fact, to diminish the number of conditions, we take the ratio of these last two variables, again distinguishing three classes as above. This leads to the following set of $3 \times 3 \times 21 = 189$ price rules: $\langle \text{if loyalty} = \text{‘low’ and ratio} = \text{‘low’ then } p^{\text{ask}} = 0 \rangle$, $\langle \text{if loyalty} = \text{‘low’ and ratio} = \text{‘low’ then } p^{\text{ask}} = 1 \rangle$, ..., $\langle \text{if loyalty} = \text{‘high’ and ratio} = \text{‘high’ then } p^{\text{ask}} = 20 \rangle$. Notice that different customers may get different prices for the following two reasons. First, they may expect different prices since the conditions vary from customer to customer. Second, for given conditions prices may differ across customers because there is a stochastic element in the seller’s decision, and sellers decide upon the price for each customer separately. Just as with all other decisions, we assume that the strengths of these rules are updated once a day. This is done using the following reward from the environment: $\text{payoff} = (\text{total revenue obtained using that price rule}) / [\text{price} \cdot (\text{number of times rule used during that session})]$.

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 afternoon from those that they choose in the morning.

3.6. Discussion

The model we have sketched is a simple model, abstracting from various aspects of the real fish market. The price p^{in} for which the sellers buy their supply is assumed to be identical for all sellers, and constant through time. The same applies to the price p^{out} for which the buyers can resell outside the market in the base version, whereas in the variant we consider three different prices p^{out} . All buyers, and all sellers are always present on the market. There is only one homogeneous good, of which 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.

Besides these simplifications of the environment, the model is also simple as far as the specification of the behavior of the individual agents is concerned. The most basic assumption we made is that the adaptive agents have to find out themselves what actions are good ones, and the reinforcement learning algorithm used simply corresponds to the economic idea that agents are seeking to do the best they can, without imposing certain reasoning procedures or heuristics. Note that little of the agents’ actions is predetermined by standard conditions of economic rationality. We do not assume a reservation price property, and there are no adjustment rules of the form: if demand $>$ supply then p^{ask} increases. Also, the treatment of loyal customers is not predetermined.

The question, then, is whether the behavior of this simple ACE model is consistent with the stylized facts of the real market.

4. Analysis of the ACE model

In the analysis of the properties of our ACE model we will concentrate on one single, randomly selected, execution of the model. Extensive experimentation with the model showed that the main properties reported here are representative, and we will indicate below the relevant places where this is not the case. The model is also robust in the sense that it does not seem to be overly sensitive to the exact algorithm used. We analyzed alternative specifications of the model with genetic algorithms, hill climbing, and annealing schedules for the sellers' decisions, without observing any significant qualitative changes in the behavior of the model.

In Sections 4.1 and 4.2 we focus on the base model with one type of buyer with $p^{\text{out}} = 15$. The properties of the variant with multiple types of buyers will be analyzed in a separate section (Section 4.3).

4.1. Prices

Fig. 3 presents the time series of all prices asked and all prices accepted during the morning sessions. For presentational reasons, each observation in the graph is the average of 20 market days. Note, first, that the series start just above 10,




Fig. 3. Time-series average prices asked and accepted during morning sessions.

the middle of the range of possible prices. Second, the prices asked increase more than the prices accepted. Third, at some point, the latter start ‘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 presented there is almost perfect convergence of prices asked and prices accepted. Sixth, average prices are relatively stable at a level of 10.3 during that period, with a weak downward trend implying that it would take another 5000 market days to reach a price level of 10.

A simple demand and supply analysis shows that since market supply is perfectly elastic at $p^{\text{in}} = 9$, and market demand is perfectly inelastic up to $p^{\text{out}} = 15$, the market clearing equilibrium price equals 9. A game-theoretic analysis, on the other hand, would predict a very different price level, a price equal to $p^{\text{out}} - 1$. As mentioned above, during the afternoon, when a buyer faces a seller, the situation resembles that of an ultimatum game.¹² The subgame-perfect equilibrium of that game implies proposing the minimum slice size of the given cake to the receiver. This would correspond to $p^{\text{afternoon}} = p^{\text{out}} - 1 = 14$. Backward induction, then, leads to $p^{\text{morning}} = p^{\text{afternoon}} = 14$.¹³

Since we are not attributing strategic behavior to the individuals, one would not expect the model to determine such prices from the outset. However, it might be that the agents in our model learn to play such an equilibrium. Fig. 3 shows that the sellers do try to drive the prices up towards this level of 14. Although this disappears in the graph because of the averaging over 20 market days, there are occasional market days in which the average price asked is above 14, and in about 16% of the first 500 days the average price asked is above 13. However, 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 the general experience with laboratory experiments of the ultimatum game. Even with perfect information, there is in general no convergence to the subgame-perfect equilibrium. As both Gale et al. (1995) and Roth and Erev (1995) explain, this is related to the off-the-equilibrium-path incentive structure of the game. Suppose, for example, that an offer just above the subgame-perfect equilibrium is made in an ultimatum game, and the receiving player rejects. The result of this

¹² 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, we have a more complicated dynamic game in which the effect of repetition is much weaker than in a standard repeated ultimatum game. Moreover, the players in our model do not use dynamic strategies looking forward beyond the current market day.

¹³ According to Ramsauer (1998) it is the ‘take-it-or-leave-it’ aspect of the prices asked, just as in the ultimatum game, that is also the essential explanation of Diamond’s (1971) paradoxical result of prices converging to the monopoly level.

‘mistake’ is that he will himself forgo only a small payoff, whereas the other player, who made the proposal, loses almost the entire cake. In other words, the feedback received through these payoffs puts much more pressure on the proposer than on the receiver to change his behavior.

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 empirical evidence. Second, do not tell the agents anything about which prices to ask, accept, or reject, let the agents learn about this all by themselves, and analyze which actions emerge. In the spirit of ACE modeling, we choose this second option. Clearly, if we had imposed on the buyers or sellers which prices to propose, accept, or reject, the resulting price patterns of the model would have been very different.¹⁴

Fig. 4 presents a similar graph for the afternoon. The average prices asked and accepted slowly approach each other, although the pattern is much more erratic than for the morning sessions. This is related to the much smaller number of meetings during the afternoons, implying a larger influence of individual prices on the average price series, and also implying a slower learning process. At an average level of 11.2, the transaction prices in the periods 2501–5000 do not show a trend. Notice that the average prices paid in the afternoon sessions are




Fig. 4. Time-series average prices asked and accepted during afternoon sessions.

¹⁴Bergemann and Valimaki (1996) impose strategic behavior and Bayesian updating, and it would be worth comparing the evolution of the prices in their model with that in ours.

somewhat higher than in the morning sessions; 11.2 against 10.3. This is consistent with the real market.¹⁵

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 just used implies that we could even have that the average price *asked* in the morning is lower than the average price *asked* in the afternoon, with some rational agents rejecting prices in the morning. The point is that it is not the average price that gets rejected.¹⁶

What is the distribution of prices generated underlying the time series presented above? Fig. 5 presents the frequency distribution of the prices paid in the market during the last 2500 days of the ACE model. As we see, the most frequently realized price is 11 (49.4%), followed by 10 (28.8%). But even a price of 9 occurs with a frequency of 14.5%, whereas 3.7% of the prices are in the range from 12 to 15.¹⁷

4.2. Loyalty

The second stylized fact of the real fish market is the high loyalty of buyers to sellers. In Section 3, 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

¹⁵ Although in the real fish market there is some falling off of the prices at the very end of the day, when there are mainly very small transactions. In our model, all transactions have a unit size.

¹⁶ A sharper observation would be that 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 late). But even that observation turns out to be incorrect. Suppose that the distribution of prices asked by the sellers adapts to the acceptance behavior of the buyers such that prices in the morning that are typically rejected are no longer asked, and all 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, and sometimes they accept too high a price. In both cases, this brings the average price rejected in the morning down. 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.

¹⁷ Remember that the price for which the sellers buy outside the market, p^{in} , equals 9, while the price for which the buyers resell outside the market, p^{out} , is 15.




Fig. 5. Price distribution during the last 2500 days.

following loyalty index γ_i :

$$\gamma_i(t) = \frac{\sum_{\text{seller } j} [L_{ij}(t)]^2}{[\sum_{\text{seller } j} L_{ij}(t)]^2}. \tag{4}$$

This loyalty index $\gamma_i(t)$ will have a value equal to 1 if the buyer is perfectly loyal to one seller, and a value equal to $1/(\text{number of sellers})$ if the buyer has ‘avoiding’ behavior, visiting all sellers in a fixed sequence.¹⁸

In Fig. 6, we present the time series for the morning sessions of this measure of loyalty averaged over all buyers, and the same series for the buyers that make up the 5th and 95th percentile on that market day. Each observation is the average of 20 days. We observe that loyalty does develop, but slowly, and with considerable variance among the buyers, as indicated by the 5th and 95th percentiles. Whereas the first has an average γ of 0.45 in the last 2500 days, for the 95th percentile this is 1.00. Another way to measure the significance of this emergent loyalty is to look at the percentage of buyers that realize more than 95% of their monthly purchases from one single seller. In the last 25 days of our model, this is the case for 41% of the buyers.¹⁹

¹⁸ Because of the discounting of past experiences, putting a greater weight put on more recent visits, this avoiding behavior would lead to a loyalty index γ slightly above 0.10.

¹⁹ In the real market these percentages vary somewhat for the different types of fish, but typically they are between about 25 and 50.




Fig. 6. Time-series average loyalty during morning sessions.

The emergence of this loyalty occurs although the buyers do not even know what the concept of 'loyalty' means. During every morning session, they simply choose a seller, and they do not even remember which seller they visited the previous day. The sellers know the concept of 'loyalty', in the sense that they recognize which faces they have seen more often than others recently. But at the start they are indifferent towards loyalty. That is, they have no initial bias towards giving advantage or disadvantage to loyal customers in their queue, and no bias towards charging higher or lower prices to loyal buyers.

Hence, the first question to ask is, how come the buyers learn to become loyal? Fig. 7 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. As we see, 74% of the buyers experience, on average, that they are offered a more advantageous deal when returning to their seller than when switching to another seller, where the payoff to a buyer depends upon the prices proposed to him and the service rate he gets.²⁰ Hence, we observe that, typically, visiting a seller where a buyer is loyal simply happens to be more reinforced than visiting a different seller.²¹

²⁰Note that this graph is based on a simplification. Loyalty is not a binary variable determined exclusively by the buyer's last visit, but a real-valued index between 0 and 1 depending on a buyer's entire shopping history.

²¹The 74% shown is robust in the sense that it is always a large majority of buyers for whom this is the case, but it is relatively low as very often about 90% of the buyers are on average better off when returning to the same seller.

payoff advantage




Fig. 7. Average payoff advantage for buyers, returning to same seller compared to switching seller.

The next question is, how come the sellers offer a higher payoff to repeat customers than to other customers? Initially, sellers are indifferent in this respect. Fig. 8 presents for all 10 sellers the difference in payoff they experience, averaged over the last 2500 mornings, when dealing with a repeat customer or when dealing with a newcomer. As we see, all sellers realize higher gross revenues when dealing with repeat customers, where revenue depends upon the prices accepted and the rejection rate. Hence, offering a higher payoff to repeat customers makes sense.²²

In some sense what emerges here is remarkable. Both buyers and sellers are better off in loyal trading relationships. How is this possible? If two traders have decided to make a transaction, then the determination of the price is a constant-sum game. Given the cake size of $p^{\text{out}} = 15$, the gross payoff for a buyer is $p^{\text{out}} - p^{\text{accepted}}$, whereas the payoff for a seller is p^{accepted} . Hence, if, for a given number of transactions, a certain group of customers is better off on average than other customers, then sellers dealing with such customers are necessarily worse off, on average, than when dealing with other customers. This is

²² One seller offered such poor service that he did not get any repeat buyers in this period. The enormous disparity of the sellers' performance over a given period of time is a general finding, of which this one seller provides an extreme example.




Fig. 8. Average payoff advantage for sellers, dealing with repeat customers compared to switching customers.

Table 1
Average (gross) payoff *per transaction* to buyers and sellers

	Avg. payoff to sellers	Avg. payoff to buyers	Sum
Repeat customers	10.35	4.65	15.00
Switching customers	10.31	4.69	15.00

confirmed in Table 1, presenting the average gross payoff *per transaction* to buyers and sellers.²³

However, this is not the proper measure of the interest buyers and sellers may have in loyal trading relationships. The game of choosing a potential trading partner is *not* a constant-sum game, as the number of cakes to be divided depends upon the players' actions. In some cases, a buyer chooses a seller, or a seller chooses a buyer with no transaction taking place in the end, because

²³In this example, when a transaction takes place, switching customers pay slightly less than repeat customers. This is, however, neither a robust nor an essential feature of the model. The prices paid by switching and repeat customers are always very close on average, with sometimes repeat customers paying slightly more, and sometimes slightly less.

sellers may sell out before serving a buyer, and buyers may refuse prices proposed by sellers. Therefore, the correct measure of how wise it is to choose a certain seller needs to take account of the service rate offered by this seller, and the proper measure of the profitability of proposing a transaction to a certain buyer in a seller’s queue needs to include the acceptance rate.

Hence, 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:

$$\frac{\sum(p^{\text{out}} - p^{\text{ask}})}{\# \text{visits}} = \frac{\sum(p^{\text{out}} - p^{\text{ask}})}{\# \text{proposals}} \frac{\# \text{proposals}}{\# \text{visits}}. \tag{5}$$

Similarly, the *average payoff for a 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:

$$\frac{\sum(p^{\text{accepted}})}{\# \text{proposals}} = \frac{\sum(p^{\text{accepted}})}{\# \text{transactions}} \frac{\# \text{transactions}}{\# \text{proposals}}. \tag{6}$$

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. As it turns out, on average over the last 2500 mornings, sellers experience that repeat customers have a higher acceptance rate (0.92) than switching customers (0.88), whereas these repeat customers experience a higher service rate (0.97) than switching customers (0.93). The combined effect of this with the prices asked and accepted is summarized in Table 2. The table illustrates that the fact that the repeated interactions are characterized, on average, by higher acceptance rates and higher service rates implies that the total surplus that can be realized is indeed greater with loyal trading relationships, and that hence both sides can be better off through loyalty.

Thus, an 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 and Erev, 1995 for a discussion). Fig. 9

Table 2
Average (gross) payoff *per interaction* to buyers and sellers

	Avg. payoff to sellers	Avg. payoff to buyers
Repeat customers	9.50	4.50
Switching customers	9.09	4.25

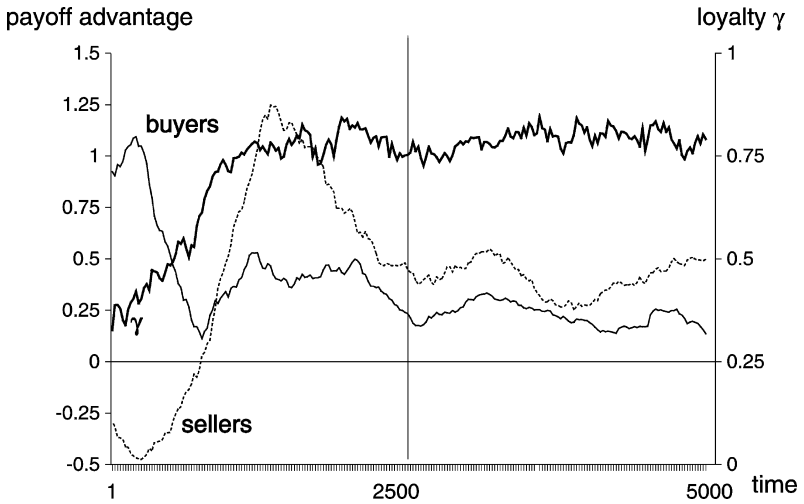


Fig. 9. Coevolution: emergence payoff advantages of loyalty to buyers and sellers, and loyalty.

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. Each observation is again the average of 20 mornings, and we smoothed the series of the payoff advantages by taking 25 period moving averages. Note that, first, in the beginning sellers get a lower payoff from repeat customers than from new customers. Second, simultaneously, the payoff advantage to repeat customers quickly disappears. Third, only when a payoff advantage to sellers dealing with repeat customers emerges does the payoff advantage to these repeat customers start to rise again. Fourth, in the meantime, repeat customers had always experienced a positive payoff difference, and loyalty emerged. Fifth, apart from the initial phase, the advantage of loyal relationships is positive to both buyers and sellers, and loyalty develops at a steady level around 0.79.

4.3. Model variant with multiple types of buyers

One of the simplifications we made in our base model was that all buyers could resell their fish outside the market for the same price $p^{\text{out}} = 15$. In the real market some obvious heterogeneity exists as some buyers are retail fish sellers, while others are restaurant owners.²⁴ In this section we analyze a variant of our

²⁴ In the market data all buyers are only identified by a code. Hence, we do not have information as to the various types of buyers.

ACE model, in which there are three types of buyers: 33 buyers can resell their fish outside the market for a price of $p^{\text{out}} = 12$, 34 can do so with $p^{\text{out}} = 15$, and another 33 for a price of 18. In line with the base model, neither the sellers nor the buyers are aware of this heterogeneity ex ante. The *only* difference with the base model is that the different types of buyers will experience different profit levels for given transactions. The question, then, is how this is going to influence the market outcomes, and in particular how this is related to loyalty and prices.

Table 3 presents the prices that are found, and those that are accepted by the three types of buyers during the last 2500 days. The first two lines show that the buyers with a higher p^{out} are prepared to accept higher prices in the market, both in the morning and in the afternoon, which is, in particular for the afternoon sessions, not surprising. More interesting is the fact that different types of buyers actually find different prices in the market, and this applies to the morning sessions as well as to the afternoon sessions, with buyers with a higher p^{out} on average finding higher prices. This is somewhat remarkable as the sellers do not recognize the various types of buyers as such. They only recognize the familiarity of a buyer's face. If different types of buyers find different prices, it must be that they develop distinct shopping patterns, which are (implicitly) recognized by the sellers. Before we analyze the effects of different shopping behaviors, we first have a look at these shopping patterns as such.

A first dimension in which shopping patterns might differ is the timing of purchases. In our model, all buyers enter the market simultaneously in the morning, but how many of them leave the market after the morning session, and how many return for the afternoon session is determined endogenously in the market. Table 4 shows the average experience for each type of buyer during the last 2500 morning sessions. Buyers with a higher p^{out} are more likely to realize their purchase during the morning sessions. Not only do buyers with a low p^{out} reject a price more often, but they are also served late more frequently. As a result of this endogenous shopping timing, the population mix of buyers in the

Table 3
Prices found and accepted by three types of buyers

	Type of buyer				
	'12'	'15'		'18'	
Price accepted (morning)	9.34	9.61	(+ 2.9%)	9.78	(+ 4.8%)
Price accepted (afternoon)	9.74	10.86	(+ 11.5%)	11.20	(+ 14.9%)
Price found (morning)	9.41	9.64	(+ 2.5%)	9.82	(+ 4.3%)
Price found (afternoon)	11.30	11.67	(+ 3.3%)	11.71	(+ 3.6%)

Table 4
Shopping experience of three types of buyers during morning sessions

	Type of buyer		
	'12'	'15'	'18'
Purchase	80.6%	90.1	91.9
Reject price	13.4	5.1	4.1
Late	6.0	4.8	4.0
Sum	100.0	100.0	100.0

Table 5
Shopping behavior by three types of buyers during morning sessions

	Type of buyer		
	'12'	'15'	'18'
Avg. loyalty L_{ij}	0.96	0.93	0.92
Rel. freq. switching seller	7.8%	8.2%	9.1%
Avg. loyalty γ_i	0.75	0.82	0.83
Avg. # sellers visited	7.4	5.7	6.0

afternoon sessions turns out to be very different from the morning sessions. In the afternoon, 51.5 % of the buyers in the market has a p^{out} of 12, 27.0% a p^{out} of 15, and only 21.5% a p^{out} of 18, whereas during the morning sessions the proportion of each type of buyer is one-third.

Since we saw above that even during the morning sessions the different types of buyers encounter different prices, we now focus on these morning sessions. Table 5 summarizes the shopping behavior of the three types of buyers during the morning sessions in the last 2500 days. The loyalty index L_{ij} reported in Table 5 measures the loyalty of a specific buyer to a specific seller, averaged over those visits that are actually made during the morning sessions. This loyalty index L_{ij} is used by the sellers when they decide whom to serve next, and which price to charge. As we see, buyers with a higher p^{out} tend to be *less* loyal with their visits. This is confirmed by the fact that they switch seller more frequently than those with a lower p^{out} . However, the last two lines of Table 5 indicate that we need to qualify this. The loyalty index γ_i , measuring the loyalty of a specific buyer i with respect to the population of sellers as a whole (see Eq. (4)), shows that the buyers with a higher p^{out} are not the less but the *more* loyal buyers. In other words, the shopping behaviors that the three types of

buyers develop differ in a subtle way. Some types are in some sense more loyal, but in another sense less loyal at the same time. What happens, among other things, is that the buyers with a higher p^{out} switch more frequently between sellers, but they do so within a more restricted group of sellers, as the last line of Table 5 also indicates.

The question, then, is how these differences in shopping behavior lead to different market outcomes for the three types of buyers. Apart from the buyers' own actions, this depends on how they are treated by the sellers. Sellers use the loyalty index L_{ij} to decide the order in which buyers are served, and to decide which prices to ask. The serving order may influence the prices to be asked as well, through the stock/queue ratio at the moment a price has to be announced. Moreover, both the serving order and the prices asked influence whether buyers conclude their transactions in the morning, or whether they have to go for a last chance in the afternoon session.

We first look at how sellers treat different types of buyers as far as the serving order is concerned. Seven out of ten sellers on average choose values of the treatment parameter that imply putting loyal customers more towards the end of the queue. However, since this function is non-linear, and since different sellers have different sales levels, and both treatment parameter values and sales levels vary over time, it is difficult to interpret the direct consequences of this for the three types of buyers. Therefore, we focus on the probability that the various types of buyers are served late, finding only empty shelves, during the morning sessions in the last 2500 days.

As we saw in Table 4, the buyers with a high p^{out} were less likely to be denied service. How is this related to the behavior of the sellers (who do not observe this p^{out})? A seller j observes the loyalty index L_{ij} of all buyers i in his queue. Computing the correlation coefficient between L_{ij} and the probability of being late shows that more loyal customers are somewhat *less* likely to be late (the coefficient is -0.27 , with a t -statistic of -2.8). This seems somewhat puzzling, because we had seen above that it was the group of buyers with a p^{out} of 12 that developed, on average, the highest values of this loyalty variable L_{ij} . Fig. 10 gives a first indication of what is going on. The graph shows how the three types of buyers form three separate clusters located from the top-right corner (where we find the buyers with $p^{\text{out}} = 12$) to the bottom-left corner (buyers with $p^{\text{out}} = 18$), within each cluster and in the three clusters combined a negative relationship between the probability of being late and average loyalty L_{ij} .

Hence, the question is what it is in the buyers' shopping behavior that separates them so neatly into these three clusters.²⁵ We saw already above that

²⁵ Remember that a seller j only knows L_{ij} and not γ_i . That is, a seller j knows how familiar a buyer i looks to him, but this seller j has no idea as to what his shopping behavior has been in all those periods in which this buyer i did not show up in his own queue.




Fig. 10. Average loyalty L_{ij} and service for all buyers.




Fig. 11. Average loyalty γ_i and service for all buyers.

the average loyalty L_{ij} of buyer i does not fully capture the pattern of his shopping behavior, and that L_{ij} and γ_i measure in fact different aspects of such behavior. How, then, does the service rate relate to this loyalty γ of the buyers? Fig. 11 shows that the correlation between γ and the probability of being late is clearly negative (the correlation coefficient is -0.76 , with a t -statistic of -11.5). Hence, the average value of a buyer's γ variable turns out to be a good predictor of his probability of being late.

Further illustration of the fact that γ captures some important aspects of the buyers' shopping behavior is provided by the following. Just as in the base

Table 6
Service for three types of buyers

	Type of buyer		
	'12'	'15'	'18'
Late repeat customers	5.8%	4.5	3.8
Late switching customers	10.9	8.7	7.8

model, the service rate for repeat customers is higher than for switching customers. The first are served late on average 4.6% of the time, and the latter 8.4%. Now, we had observed that buyers with a higher p^{out} tended to switch seller *more often*. How come, then, that they nevertheless experience better service? Again, this seems related to the buyers' loyalty as measured by γ . As Table 6 shows, for both repeat and switching customers the service rate increases with the p^{out} of the buyers (just as the loyalty γ). As a result, switching customers with a p^{out} of 12 are almost 3 times as likely to be denied service as repeat customers with a p^{out} of 18 (10.9% against 3.8%).

Next, we look at the prices. As indicated above, the different types of buyers find systematically different prices in the market. The question, then, is how this is related to their shopping behavior. When deciding upon prices, the sellers, besides distinguishing morning from afternoon sessions, look at the loyalty index L_{ij} , and at the actual stock/queue ratio, in both cases distinguishing three levels (high, middle, and low). Fig. 12 shows the pricing function learned on average by the sellers with respect to loyalty L . As we see, on average the sellers learn to charge lower prices to customers who demonstrate a higher loyalty L towards them. Hence, the buyers with a lower p^{out} are charged lower prices, since they develop a higher loyalty L (see Table 5).²⁶

However, the different experiences by the different types of buyers is not only due to the fact that individual sellers learn to distinguish these types implicitly. In part, these differences are due to the emergence of heterogeneity among the sellers, as different individual sellers learn to cater to different types of buyers. Fig. 13 shows the numbers of visits during the last 2500 mornings for each seller, distinguishing the three types of buyers. As we see, some sellers are rather small, focussing on one type of buyer, whereas the largest seller caters for all three types of buyers in approximately equal shares. If we focus, for example, on seller 3 on the one hand, and sellers 5 and 6 on the other hand, we see that for the first fewer than 8% of the customers have a p^{out} of 18, whereas for the latter two this is

²⁶ This is conformed by a positive correlation between the buyers' loyalty γ and the prices they find (the correlation coefficient is 0.26, with a t -statistic of 2.71).




Fig. 12. Average pricing function learned during the last 2500 days.

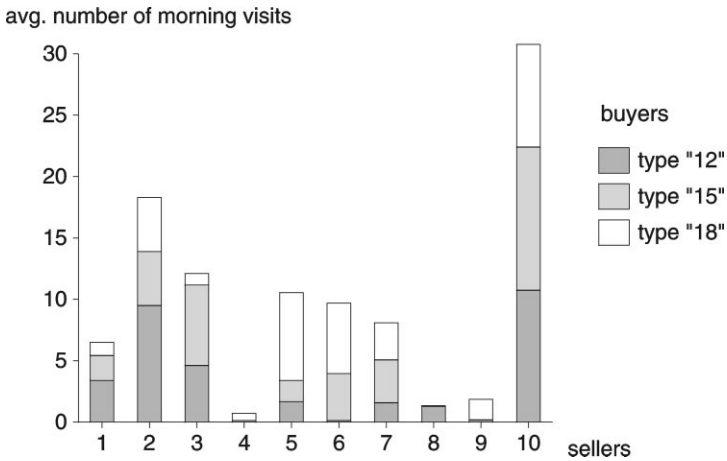


Fig. 13. Average number of visits for each seller during the last 2500 days.

around 60%, and for seller 6 only 1% of the customers has a p^{out} of 12. What distinguishes seller 3 from these other two, is that seller 3, the one basically without high-value customers, asks lower prices, has a lower supply to sales ratio, puts loyal customers towards the end of the queue, and offers a lower average service rate, while facing a higher rejection rate.

Hence, this variant of our ACE model explains how price dispersion and loyalty may be related to different types of buyers (who have different

opportunities to resell the fish outside the market). We observed that the buyers who can resell outside the market for a higher price develop a pattern of shopping behavior characterized by higher loyalty (as measured by γ) such that they find higher prices in the market, and experience a better service.

5. Concluding remarks

Our ACE model of the Marseille fish market, focussing on an explicit consideration of the market interactions between the individual agents, explains both stylized facts price dispersion and loyalty as the outcome of 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, the basic problem is one of coordination, and loyalty works as a coordination device. This suggests that developing loyalty has some similarity with sending intrinsically meaningless signals, say, like wearing a blue shirt. The similarity lies herein that there is nothing intrinsic in loyalty that makes it pay, as there are no real switching costs in our model (cf. Klemperer, 1995). Hence, at first sight one might conjecture that our explanation could as well have gone the other way around: '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, although there is some similarity with intrinsically meaningless signals, 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 one's shirt. Moreover, being loyal or non-loyal has a direct economic meaning. Suppose, for example, that there appears to be some serial correlation in a seller's decisions, and that a buyer is satisfied with that seller. Loyalty would benefit that buyer, but continuing to dress blue while shopping around randomly would not. Or suppose that a seller offers a poor service. One of his buyers becoming non-loyal would hurt, but that buyer merely changing the color of his shirt would not. Loyalty means continuity, and allows buyers and sellers to avoid unproductive meetings.

A variant of the model with multiple types of buyers explains how buyers who can resell outside the market for a higher price develop a pattern of shopping behavior characterized by higher loyalty such that they find higher prices in the market, and experience a better service. The sellers do not explicitly recognize different types of buyers (cf. Borenstein, 1991). But this price and service discrimination emerge as, in coevolution with the development of shopping patterns which differ in a subtle way among the buyers, the sellers on the one

hand learn to recognize these differences implicitly, whereas on the other, different types of sellers emerge. Hence, the underlying incentive, that missing out on a transaction is more costly for those buyers with a higher resell price, does not only lead those buyers to accept more prices, but it also leads to differences in shopping patterns and to differences among the sellers in ways that seem consistent with real markets.

To conclude, let us stress how simple our ACE model is, and which factors do not play a role in our model. We did not assume any real switching costs or search costs, and we show that to explain loyalty and price dispersion, there is no need to assume given social networks, direct personal relations, or cultural factors (cf. Richardson, 1960). In fact, in our model all sellers and all buyers are identical *ex ante*. In a variant of our base model, there are three types of buyers, but even there, our ACE model explains that for price and service discrimination to occur (plus the related loyalty), it is not necessary that the sellers explicitly recognize these different types of buyers. All agents in our model have very limited cognitive capabilities, and they cannot engage in sophisticated reasoning processes. The buyers do not even know what the concept of loyalty means, and sellers only recognize how familiar faces look to them, but do not recognize the identities of individual buyers as such. As the agents are not looking forward (and reasoning backward) beyond the end of the day, they do not play dynamic strategies. Each agent simply learns through reinforcement only. Hence, our ACE model of price dispersion and loyalty highlights the role of the evolving incentive structure of the market.²⁷

Appendix. Pseudo-code

```

program MAIN;
begin
  INITIALIZE_BUYERS; (see procedure below)
  INITIALIZE_SELLERS; (see procedure below)
  for all 10000 periods (5000 mornings + 5000 afternoons) do
  begin
    for all buyers i do for all sellers j do if morning then loyalty[i,j] :=
      loyalty[i,j]/1.25;
    for all buyers do
    begin

```

²⁷ Obviously, this is not to say that all these additional factors cannot or do not play a role in reality. But what our ACE model explains, is that they are not essential to explain phenomena like loyalty and price dispersion.

```
    CHOOSE_A_SELLER; (see procedure below)
    ACCEPT/REJECT_PRICES (see procedure below)
end;
for all sellers do
begin
    if morning then
    begin
        CHOOSE_SUPPLY_LEVEL; (see procedure below)
        CHOOSE_TREATMENT_PARAMETER_VALUE; (see procedure below)
    end;
    QUEUE_HANDLING; (see procedure below)
    if afternoon then REINFORCEMENT_SELLERS; (see procedure below)
end;
for all buyers i do for all sellers j do
    if i was in queue j then loyalty[i,j] := loyalty[i,j] + 0.25;
for all buyers do REINFORCEMENT_BUYERS; (see procedure below)
end;
end.
```

```
procedure INITIALIZE_BUYERS;
begin
    for all buyers do
    begin
        p_out := 12 (15, 18);
        with all sellers do
        begin
            fitness_choice_morning := 1.00;
            fitness_choice_afternoon := 1.00;
        end;
        with prices from 0 to 20 do
        begin
            fitness_reject_morning := 1.00;
            fitness_accept_morning := 1.00;
            fitness_reject_afternoon := 1.00;
            fitness_accept_afternoon := 1.00;
        end;
    end;
end;
end;
```

```
procedure INITIALIZE_SELLERS;
begin
    for all sellers do
    begin
```

```

p_in := 9;
with supply levels from 0 to 30 do fitness := 1.00;
with treatment parameter values from -25 to 25 (in steps of 5) do
  fitness := 1.00;
with all 189 morning price rules do fitness := 1.00;
with all 189 afternoon price rules do fitness := 1.00;
end;
end;

```

```

procedure CHOOSE_A_SELLER;

```

```

begin

```

```

  if morning then

```

```

    begin

```

```

      with all sellers do normalize actual fitness_choice_morning to [0, 1];

```

```

      for all sellers do

```

```

        begin

```

```

          bid := normalized_fitness +  $\varepsilon$  where  $\varepsilon \simeq N(0, 0.10)$ ;

```

```

          with probability 0.025 ignore bid;

```

```

        end;

```

```

        choose seller receiving highest bid;

```

```

      end

```

```

    else if afternoon and no transaction in morning then

```

```

      begin

```

```

        with all sellers do normalize actual fitness_choice_afternoon to [0, 1];

```

```

        if at least one seller open then

```

```

          begin

```

```

            for all open sellers do

```

```

              begin

```

```

                bid := normalized_fitness +  $\varepsilon$  where  $\varepsilon \simeq N(0, 0.10)$ ;

```

```

                with probability 0.025 ignore bid;

```

```

              end;

```

```

              choose seller receiving highest bid;

```

```

            end

```

```

          else if all sellers closed then no seller chosen;

```

```

        end

```

```

      else if afternoon and transaction in morning then no seller chosen;

```

```

    end;

```

```

procedure ACCEPT/REJECT_PRICES;

```

```

begin

```

```

  if morning then

```

```

    begin

```

```

for prices from 0 to 20 do
begin
  bid_ok := fitness_accept_this_price_in_morning;
  bid_reject := fitness_reject_this_price_in_morning +  $\varepsilon$  where  $\varepsilon \simeq N(0, 0.10)$ ;
  if bid_ok > bid_reject then accept this price else reject it;
  with probability 0.025 invert acceptance/rejection decision;
end;
end
if afternoon then
begin
  for prices from 0 to 20 do
  begin
    bid_ok := fitness_accept_this_price_in_afternoon;
    bid_reject := fitness_reject_this_price_in_afternoon +  $\varepsilon$  where  $\varepsilon \simeq N(0, 0.10)$ ;
    if bid_ok > bid_reject then accept this price else reject it;
    with probability 0.025 invert acceptance/rejection decision;
  end;
end;
end;

procedure CHOOSE_SUPPLY_LEVEL;
begin
  with all supply levels from 0 to 30 do normalize actual fitness to [0, 1]
  for all supply levels from 0 to 30 do
  begin
    bid := normalized_fitness +  $\varepsilon$  where  $\varepsilon \simeq N(0, 0.10)$ ;
    with probability 0.025 ignore bid;
  end;
  choose supply level receiving highest bid;
end;

procedure CHOOSE_TREATMENT_PARAMETER_VALUE;
begin
  with all treatment parameter values do normalize actual fitness to [0, 1]
  for all treatment parameter values do
  begin
    bid := normalized_fitness +  $\varepsilon$  where  $\varepsilon \simeq N(0, 0.10)$ ;
    with probability 0.025 ignore bid;
  end;
  choose treatment parameter value receiving highest bid;
end;

```

```

procedure QUEUE_HANDLING;
begin
  determine queue_length;
  check who is in queue;
  check their loyalty;
  give everybody in queue a weight := (1 + loyalty) ^ (treatment_
parameter_value);
  while customers in queue left and stock left do
  begin
    randomly pick a customer (each with probability of weight/(sum of
weights remaining queue));
    with served_customer do
    begin
      if loyalty < = 0.20 then loyalty_class := low
        else if loyalty < = 0.80 then loyalty_class := middle else
          loyalty_class := high;
    end;
    s/q_ratio := remaining_stock/remaining_queue_length;
    if s/q_ratio < = 0.75 then s/q_ratio_class := low
      else if ratio < = 1.25 then s/q_ratio_class := middle else
        s/q_ratio_class := high;
    if morning then for given loyalty_class and s/q_ratio_class do
      with all morning prices from 0 to 20 do normalize actual fitness to [0, 1];
    if afternoon then for given loyalty_class and s/q_ratio_class do
      with all afternoon prices from 0 to 20 do normalize actual fitness
to [0, 1];
    if morning then for given loyalty_class and s/q_ratio_class do
    begin
      for all morning prices from 0 to 20 do
      begin
        bid := normalized_fitness + ε where ε ≈ N(0, 0.10);
        with probability 0.025 ignore bid;
      end;
      choose price corresponding to rule receiving highest bid;
    end;
    if afternoon then for given loyalty_class and s/q_ratio_class do
    begin
      for all afternoon prices from 0 to 20 do
      begin
        bid := normalized_fitness + ε where ε ≈ N(0, 0.10);
        with probability 0.025 ignore bid;
      end;
    end;
  end;
end;

```

```
    choose price corresponding to rule receiving highest bid;
end;
update remaining queue information;
if price not rejected then update stock information;
end;
end;

procedure REINFORCEMENT_SELLERS;
begin
    compute gross revenue of the day;
    scaled_revenue := gross_revenue / (maximum_price * supply);
    net_profit of the day := gross_revenue - (supply * p_in);
    normalize net_profits last 200 days to [0, 1];
    with activated supply level rule do fitness := 0.95 * fitness + 0.05 *
    normalized_profit;
    with activated treatment parameter value rule do fitness := 0.95 * fitness
    + 0.05 * scaled_revenue;
    for all morning price rules do for number of times rule activated do
    begin
        reward := (#times accepted * price) / (times_active * maximum price);
        fitness := 0.95 * fitness + 0.05 * reward;
    end;
    for all afternoon price rules do for number of times rule activated do
    begin
        reward := (#times accepted * price) / (times_active * maximum price);
        fitness := 0.95 * fitness + 0.05 * reward;
    end;
end;

procedure REINFORCEMENT_BUYERS;
begin
    if transaction in morning then
    begin
        utility := (p_out - price_paid) / p_out;
        with seller visited in morning do
        begin
            reward_choice := max(0, utility);
            fitness_choice_morning := 0.95 * fitness_choice_morning + 0.05 *
            reward_choice;
        end;
        with price accepted do
        begin
            reward_accept := utility;
```

```

    fitness_accept_morning := 0.95*fitness_accept_morning + 0.05*
    reward_accept;
end;
end;
if reject in morning and transaction in afternoon then
begin
    utility_offered := (p_out-price_found_in_morning)/p_out;
    utility := (p_out-price_paid_in_afternoon)/p_out;
    with seller visited in morning do
    begin
        reward_choice := max(0, utility_offered);
        fitness_choice_morning := 0.95*fitness_choice_morning + 0.05*
        reward_choice;
    end;
    with price rejected in morning do
    begin
        reward_reject := max(0, utility);
        fitness_reject_morning := 0.95*fitness_reject_morning + 0.05*
        reward_reject;
    end;
    with seller visited in afternoon do
    begin
        reward_choice := max(0, utility);
        fitness_choice_afternoon := 0.95*fitness_choice_afternoon + 0.05*
        reward_choice;
    end;
    with price accepted in afternoon do
    begin
        reward_accept := utility;
        fitness_accept_afternoon := 0.95*fitness_accept_afternoon + 0.05*
        reward_accept;
    end;
end;
end;
if reject in morning and reject in afternoon then
begin
    utility_offered := (p_out-price_found_in_morning)/p_out;
    with seller visited in morning do
    begin
        reward_choice := max(0, utility_offered);
        fitness_choice_morning := 0.95*fitness_choice_morning + 0.05*
        reward_choice;
    end;
end;

```



```

with price rejected in morning do fitness_reject_morning:= 0.95*
fitness_reject_morning;
with seller visited in afternoon do
  fitness_choice_afternoon:= 0.95*fitness_choice_afternoon;
with price rejected in afternoon do
  fitness_reject_afternoon:= 0.95*fitness_reject_afternoon;
end;
if reject in morning and late in afternoon then
begin
  utility_offered:= (p_out-price_found_in_morning)/p_out;
  with seller visited in morning do
  begin
    reward_choice:= max(0, utility_offered);
    fitness_choice_morning:= 0.95*fitness_choice_morning+0.05*reward_
choice;
  end;
  with price rejected in morning do fitness_reject_morning:= 0.95*
fitness_reject_morning;
  with seller visited in afternoon do
    fitness_choice_afternoon:= 0.95*fitness_choice_afternoon;
  end;
if late in morning and transaction in afternoon then
begin
  utility:= (p_out-price_paid_in_afternoon)/p_out;
  with seller visited in morning do fitness_choice_morning:= 0.95*
fitness_choice_morning;
  with seller visited in afternoon do
  begin
    reward_choice:= max(0, utility);
    fitness_choice_afternoon:= 0.95*fitness_choice_afternoon+0.05*
reward_choice;
  end;
  with price accepted in afternoon do
  begin
    reward_accept:= utility;
    fitness_accept_afternoon:= 0.95*fitness_accept_afternoon+0.05*
reward_accept;
  end;
end;
if late in morning and reject in afternoon then
begin
  uility_offered:= (p_out-price_found_in_afternoon)/p_out;

```

```

with seller visited in morning do fitness_choice_morning:=
  0.95*fitness_choice_morning;
with seller visited in afternoon do
begin
  reward_choice:= max(0, utility_offered);
  fitness_choice_afternoon:= 0.95*fitness_choice_afternoon + 0.05*
  reward_choice;
end;
with price rejected in afternoon do
  fitness_reject_afternoon:= 0.95*fitness_reject_afternoon;
end;
if late in morning and late in afternoon then
begin
  with seller visited in morning do fitness_choice_morning:=
  0.95*fitness_choice_morning;
  with seller visited in afternoon do
    fitness_choice_afternoon:= 0.95*fitness_choice_afternoon;
  end;
end;
end;

```

References

- Aghion, P., Harris, C., Julien, B., 1991. Optimal learning by experimentation. *Review of Economic Studies* 58, 621–654.
- Baillie, R.T., Bollerslev, T., 1997. The forward premium anomaly is not as bad as you think. Mimeo.
- Bergemann, D., Valimaki, J., 1995. Learning and strategic pricing: further results. Mimeo, Northwestern University and Yale University.
- Bergemann, D., Valimaki, J., 1996. Learning and strategic pricing. *Econometrica* 64, 1125–1149.
- Bolton, P., Harris, C., 1993. Strategic experimentation. Discussion Paper TE/93/261, London School of Economics.
- Borenstein, S., 1991. Selling costs and switching costs: explaining retail gasoline margins. *RAND Journal of Economics* 22 (3), 354–369.
- Burdett, K., Judd, K., 1983. Equilibrium price dispersion. *Econometrica* 51, 955–969.
- Bush, R.R., Mosteller, F., 1955. *Stochastic Models for Learning*. Wiley, New York.
- De Francesco, M.A., 1996. Customer Markets as an Efficient Outcome of a Dynamic Choosing-the-Seller Game (Quaderni del dipartimento di economia politica 205), University of Siena.
- Diamond, P.A., 1971. A model of price adjustment. *Journal of Economic Theory* 3, 156–168.
- Durlauf, S.N., 1995. Neighborhood feedbacks, endogenous stratification, and income inequality. Working Paper 95-07-061, Santa Fe Institute.
- Easley, D., Rustichini, 1999. Choice without beliefs. *Econometrica* 67, 1157–1184.
- Felli, L., Harris, C., 1994. Job matching, learning and the distribution of surplus. Mimeo, London School of Economics.
- Friedman, M., 1953. The methodology of positive economics. In: Friedman, M. (Ed.), *Essays in Positive Economics*. University of Chicago Press, Chicago, pp. 3–43.
- 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.
- Hayek, F.A., 1948a. The facts of the social sciences (reprint). In: Hayek, F.A. (Ed.), *Individualism and Economic Order*. University of Chicago Press, Chicago, pp. 57–76.
- Hayek, F.A., 1948b. Socialist calculation I: the nature and history of the problem (reprint). In: Hayek, F.A. (Ed.), *Individualism and Economic Order*. University of Chicago Press, Chicago, pp. 119–147.
- Hayek, F.A., 1967. Degrees of explanation (reprint). In: Hayek, F.A. (Ed.), *Studies in Philosophy, Politics and Economics*. University of Chicago Press, Chicago, pp. 3–21.
- Holland, J.H., 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor.
- Holland, J.H., Holyoak, K.J., Nisbett, R.E., Thagard, P.R., 1986. *Induction: Processes of Inference, Learning, and Discovery*. MIT Press, Cambridge, MA.
- Hopkins, E., 1995. Price dispersion: an evolutionary approach. Mimeo.
- Kirman, A.P., 1994. *Economies with Interacting Agents*. Working Paper 94-05-030, Santa Fe Institute.
- Kirman, A.P., McCarthy, M., 1990. *Equilibrium Prices and Market Structure: The Marseille Fish Market*. Mimeo.
- Kirman, A.P., Vignes, A., 1991. Price-dispersion: theoretical considerations and empirical evidence from the Marseille fish market. In: Arrow, K.J. (Ed.), *Issues in Contemporary Economics: Proceedings of the Ninth World Congress of the International Economics Association*, Athens, Greece, Vol. 1. Markets and Welfare. MacMillan, London, pp. 160–185.
- 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.
- Klemperer, P., 1995. Competition when consumers have switching costs: an overview with applications to industrial organization, macroeconomics and international trade. *Review of Economic Studies* 62 (4), 515–539.
- Mailath, G., Samuelson, L., Shaked, A., 1993. Correlated equilibria as network equilibria. Discussion Paper No. B-244, University of Bonn.
- Moreno, D., Wooders, J., 1997. Price dynamics, trading patterns, and delay in a market with pairwise meetings and private information. Mimeo.
- O'Hara, M., 1994. *Market Microstructure Theory*. Blackwell, Oxford.
- Ramsauer, A., 1998. Non-Walrasian outcomes in a two-sided search model with ex ante heterogeneous agents and infinitesimal frictions. Mimeo.
- Richardson, G.B., 1960. *Information and Investment*. Oxford University Press, Oxford.
- 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.
- Royal Mail, 1999. (tv commercial), Fall.
- Sarin, R., Vahid, F., 1997. Payoff assessments without probabilities: a simple dynamic 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.
- Sutton, R.S., 1992. Introduction: the challenge of reinforcement learning. *Machine Learning* 8 (3/4), 225–227.

The Economist, 1999. News from the Lab, May 8.

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., 1994. A new perspective on decentralized trade. *Economie Appliquée* 47 (4), 5–22.

Vriend, N.J., 1995. Self-organization of markets: an example of a computational approach. *Computational Economics* 8, 205–231.

Vriend, N.J., 1996. Rational behavior and economic theory. *Journal of Economic Behavior and Organization* 29 (2), 263–285.

Vriend, N.J., 1999. Was Hayek an Ace? (Working Paper No. 403) Queen Mary and Westfield College, University of London.

Weisbuch, G., Kirman, A.P., Herreiner, D., 2000. Market organisation. *Economic Journal* 110 (465), 411–436.

Wilson, J.A., 1980. Adaptation to uncertainty and small numbers exchange: the new England fresh fish market. *Rand Journal of Economics* 11, 491–504.