Proceedings of the 2013 Winter Simulation Conference
R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, eds.

AN AGENT-BASED MODEL FOR SEQUENTIAL DUTCH AUCTIONS

Eric Guerci

GREDEG (UMR 6227),
University of Nice Sophia-Antipolis
250 rue Albert Einstein,
06560 Valbonne, FRANCE.

Alan Kirman

GREQAM (UMR 7316), EHESS,
Aix Marseille University
2, rue de la vieille charité,
13236 Marseille Cedex 02, FRANCE.

Sonia Moulet

GREQAM (UMR 7316)
2, rue de la vieille charité,
13236 Marseille Cedex 02, FRANCE.

ABSTRACT

We propose an agent-based computational mode to investigate sequential Dutch auctions with particular emphasis on markets for perishable goods and we take as an example wholesale fish markets. Buyers in these markets sell the fish they purchase on a retail market. The paper provides an original model of boundedly rational behavior for wholesale buyers' behavior incorporating inter-temporal profit maximization, conjectures on opponents' behavior and fictive learning. We analyse the dynamics of the aggregate price under different market conditions in order to explain the emergence of market price patterns such as the well-known declining price paradox. The proposed behavioral model provides alternative explanations for market price dynamics to those which depend on standard hypotheses such as diminishing marginal profits.

1 INTRODUCTION

This paper uses and agent-based computational model to analyse sequential Dutch/descending auctions in order to explain the emergence of certain well-known "stylized facts" such as the declining price paradox. Starting from the Ashenfelter's (1989) observation of the declining price phenomenon at wine auctions, a considerable theoretical and empirical literature on this topic has developed. This paper pursues this strand of research with a specific focus on the nature of sequential descending auction in the context of fish markets. A related literature on this specific subject exists and ranges from theoretical/computational approaches to empirical investigations (Kirman and Vignes 1991; Weisbuch et al. 2000; Kirman and Vriend 2001; Kirman et al. 2005; Graddy 2006). Fish markets are particularly suited to the analysis of price formation since they have a number of "good" properties. Fish is a perishable good and therefore there is no possibility for the sellers to postpone the sale of the good (fresh fish cannot be stocked for more than a day and still be characterised as "fresh"), so the total quantity on must be traded. This greatly simplifies the economic analysis and makes the conclusions easier to analyse and to see to what extent they are in line with the empirical evidence.

The most common agent-based modeling approach is to exploit a well-established behavioral model which determines certain minimal behavioral requirements for the agents' decision making process capable of replicating "styled facts" such as the declining price paradox at the aggregate level. Such models deliberately

978-1-4799-2076-1/13/$31.00 ©2013 IEEE

avoid to carefully addressing context-dependent modeling assumptions which necessitate departing from standard learning models.Too often, the same algorithms are adopted for different economic context where the level of information available varies or where the repeated nature of the interaction among the same market actors is neglected.

We adopt a different approach in this paper and propose to identify a "rich" decision-making model which incorporates behavioral assumptions which are consistent with the empirical evidence. The authors had the opportunity to visit fish markets organised as sequential Dutch auctions and to discuss with market participants. In the economic context which we analyses, we aim to give an accurate reflection of the available level of information available to market operators about the market mechanisms and conditions, these conjectures about opponent's behavior and the inter-temporal arbitrage opportunities inherent in the daily sequential descending auctions. Thus we propose so to implement a non standard computational learning algorithm which captures behavioral aspects which can help us to reproduce the emergence of certain well-known "stylized facts". One outcome of this research work is thus to show how behavioral aspects which are frequently neglected can give rise to relevant economic phenomena.

2 COMPUTATIONAL MODEL

The goods, which are assumed here to be indivisible units of a perishable and homogenous good, are traded daily in the early morning wholesale market and are then sold on retail markets. The former market is modeled as a sequential Dutch/descending auction market comprising T identical rounds in each of which only one unit of the good is traded. The latter is represented by assuming that each buyer on the wholesale market is a monopolistic seller facing a linear demand on his own retail market. Every period/day d the same sequential trading procedure comprising T rounds is repeated. Furthermore, we assume, as is in fact the case, that each market participant meets all the other participants every day and the same population of agents is always present on the computational market model. Indeed, the wholesale buyers, on the markets that we have studied, meet repeatedly among themselves on a daily basis. The major focus of this research is on the functioning of the wholesale market. In what follows, we sketch its characteristics.

The market actors are one auctioneer, several sellers and buyers. In each round of the auction the auctioneer is responsible for choosing the starting price $\overline{P}$. The starting price is seldom varied by the auctioneer. This behavior is motivated by the attempt to avoid buyers' strategic reasoning. One auctioneer that we interviewed mentioned an analogy with the card game of bridge saying that auctioneers do not want buyers to implement the similar strategy of counting and memorizing the cards which have already been played so as to work out which remain in the hands of their opponents.

Several sellers or equivalently a representative seller are modeled in order to determine the daily supply of the good which is assumed to be exogenous. Sellers are characterized by a reservation price $\underline{P}$ corresponding to an opportunity cost related to a sale in another market. In our model, both sellers and the auctioneer are assumed to be "zero-intelligence" actors, the rationale for this is that this enables us to focus on the buyers who exhibit the most interesting and relevant behavior among market actors with respect to price formation. The n who are assumed to learn generate the daily demand. Both demand and supply may vary on a daily basis. The profit function of the buyer i for day d is:

$$\Pi_d^i(p_k^{i,m}, q_k^i) = Q^i \frac{a_i - Q^i}{b_i} - \sum_{k=1}^{K} p_k^{i,m} \tag{1}$$

where $Q^i = \sum_{k=1}^{K} q_k^i$ is the number of units bought in the wholesale market in K auction rounds where she won, which are then sold in the retail market. We assume that she can resell the total amount in the retail market where the buyer i is supposed to face a linear demand $D^{i,r}(p^{i,r}) = a_i - b_i p^{i,r}$. On the other hand, costs correspond to the daily sum of wholesale market prices $p_k^{i,m}$ paid in order to buy Q^i units of the good.

At round t of the auction of the day d, the winning buyer i, having already bought Q^i units of good, obtains a reward equal to the marginal profit of buying one more unit of good Q^i+1:

$$\pi_t^i(Q^i+1) = \Pi_d^i(Q^i+1) - \Pi_d^i(Q^i) = \frac{a_i - (2Q^i+1)}{b_i} - p_t^{i,m}, \qquad t \in 1,\ldots,T \tag{2}$$

since the profit is a quadratic function:

$$\Pi_d^i(Q^i) = (p_t^{i,r}(Q^i) - p_t^{i,m})Q^i = \frac{(a_i - b_i p_t^{i,m})Q^i - Q^{i2}}{b_i} \tag{3}$$

Buyers rationally exit the current daily market when the marginal profit at price 0 is negative, in other words when there is no positive price at which it would be profitable to trade and this in turn determines the maximum amount of quantity to buy in current day $Q^i > (a_i-1)/2$.

Figure 1 shows two distinct market scenarios derived of two different retail demand models, the linear model that we have described and a perfectly inelastic retail demand model. In the left-hand side of the figure, the linear retail demand model is sketched and both the continuous profit curve and discrete marginal profit values are drawn. The diminishing marginal profit effect is evident and the profit region where marginal revenue is greater or equal to marginal costs is shown as a grey area. We propose a second market scenario, in the right-hand side of the figure, where the retail demand is perfectly inelastic. The resulting marginal profits are constant up to the amount of goods bought in the retail market and after they decay linearly. This assumption might appear rather extreme, but it captures an interesting effect. When we refer to demand on the retail market we are referring to a short term (daily) demand. In the long run the demand is never perfectly inelastic, but in our model the retail demand is determined by the daily behavior of the customers. The seller has to daily forecast the retail demand and to buy the fish in advance on the same day in the auction-based wholesale market. The buyers who are the sellers on the

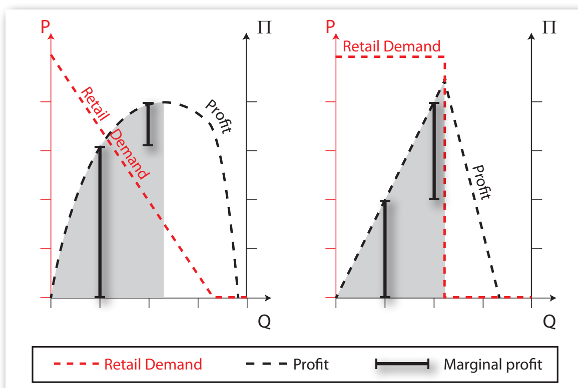

Figure 1: Retail demand models.

retail market, learn in the long run to maximize their long-term profits by means of an original learning algorithm that we refer to as the qEWA-learning model which is based on several behavioral assumptions. In particular, we adopt an action-state representation (condition-action rule), a reinforcement learning rule (explore and exploit mechanisms (Roth and Erev (1995)), the possibility of using counterfactual reasoning (EWA based, Camerer (1999)) and finally an intertemporal decision-making (q-learning (Sutton and Barto (1998), Watkins (1989)).

However in our market model, there is no room for belief-based reasoning (belief-based models assume that agents explicitly model opponents' behavior) because of the large number of buyers and the fast decision-making process. Indeed, as previously mentioned, in reality the auctioneer speeds up the trade by occasionally diminishing the starting price.

The action space $\mathscr{P}$ is a discrete set of bid prices $p \in \{\underline{P}, ..., \hat{p}, ...\overline{P}\}$ which is common to all buyers. We consider both prices below and above the true reservation price $\hat{p}$ for each buyer. The state space $\mathscr{S}^i = \mathscr{S}^i_{int} \times \mathscr{S}^i_{ext}$ is given by the cartesian product of an internal (private) state set $\mathscr{S}^i_{int}$ and an external (common) state $\mathscr{S}^i_{ext}$ set. The internal state is represented by the number of units of good that the buyer i has already bought $q^i \in 0, ..., Q^i$. The buyer can modify her reservation price conditional on the amount of units she has already obtained. The external state enables each buyer to determine the number of the current auction $t \in \{0, ..., T\}$. So, the buyer modifies her reservation price conditional on the time lasting until the daily market close.

We incorporate the action-state space representation in the classical EWA-learning model (Camerer (1999)) as follows.

The key element of the EWA algorithm are attractions $A^i : \mathscr{P}^i \rightarrow \mathbb{R}$. They represent a numerical measure of how much it is worth to play an action. In the classical EWA model, attractions are defined only over actions, because no state space is considered. We extend the original formulation by defining attractions over both actions and states $A^i : \mathscr{P}^i \rightarrow \mathbb{R} \times \mathscr{S}^i$.

Attractions values evolve over time on the basis of current experience. Buyers are assumed to have a complete representation of the state space when participating in the market, thus they update attractions values conditioned on the state they are experimenting. In particular, the rule for updating attractions is the following:

$$A^i_t(s^i_t, p^i) = \frac{\phi N_{t-1} A^i_{t-1}(s^i_t, p^i) + [\delta + (1-\delta) I(p^i_t, p^i)] R^i(p^i, p^{-i}, s^i_t)}{N_t}, \tag{4}$$

where p^i_t is the last action played, s^i_t is the current state, $I(\cdot,\cdot)$ is the indicator function and N_t stands for number of "observations equivalents" of past experience which are updated according to:

$$N_t(s^i_t) = \phi(1-\kappa) N_{t-1} s^i_t + 1. \tag{5}$$

For a detailed explanation of the role played by the three parameters δ, ϕ, κ, see the original paper by Camerer (1999). The reward term $R^i(p^i, p^{-i}, s^i_t)$ depends on the current state, in particular on the internal state because of the marginal profit of the current unit. Furthermore, we adopt a different reward function than the standard instantaneous reward/profit $R^i_t(p^i_t, p^{-i}_t, s^i_t) = \pi^{i,q+1}_t$. We use as a basis the Q-learning reward formula.

$$R^i_t(p^i_t, p^{-i}_t, s^i_t) \quad = \quad \pi^i_t(q+1) + \gamma A^{*,i}_t(s^i_t), \quad if\, p^i_t > \forall p^{-i}_t \tag{6}$$

$$\cdot \quad = \quad \gamma A^{*,i}_t(s^i_t), \quad if\, p^i_t < \forall p^{-i}_t \tag{7}$$

where:

$$A^{*,i}_t(s^i_t) = \max_{p^i \in \mathscr{P}^i} A^i_t(s^i_{t+1}, p^i). \tag{8}$$

The $A^{*,i}_t(s^i_t)$ term enables the optimizing agent to be forward-looking. She can anticipate the potential optimal profits in next state given past experience. The agent can thus incorporate in her current utility the expectation of the maximum future rewards in the next round/state. The probability of choosing an action is then monotonically related to the attractions by implementing a classical logit probabilistic choice model:

$$\pi^i_t(a^i) = \frac{e^{\lambda_t A^i_t(a^i)}}{\sum_{a^i} e^{\lambda_t A^i_t(a^i)})}. \tag{9}$$

It is worth noting that, the original EWA learning model assumes that the buyer reinforces not only the most recently played action but all other actions as well. Our qEWA learning model allows for almost full updating, because buyers need to be able to determine with certainty and by means of counterfactual reasoning the resulting reward and the next state. This is easy to understand with the following analysis of the consequences of a bid. In particular, we can infer the new state if the buyer lost in the following cases:

- she can easily understand that any other private reservation prices that she would have had below the market price would have yielded an identical profit equal to zero and the same future state.
- she can also infer what would have happened if her reservation prices had been above the market price.

and if the buyer won:

- she can easily understand that had she had any other private reservation prices above the second best market price (where the first was her own winning bid price) would have yielded a similar outcome of the market game (but different profit) and the same future state.
- she can also infer that for prices below the second best market price she would have had zero profits, previous internal state and next external state (one auction round more).
- finally, we assume that she is not capable of guessing what would have happened in the case where her bid at the same price of the second best market price.

In line with our modeling assumption, such as the reinforcement of the actions that were not taken, we adopt the following set of parameter of the original EWA model which corresponds in the EWA cube representation to the fictitious play corner: $\phi = 1$, $\delta = 1$, $\kappa = 0$ and $\lambda = 10^{-6}$.

3 SIMULATION RESULTS

In Figure 2 the experimental setup is illustrated and the parameters which are varied for the different scenarios are as follows. by providing some notations used for defining the different scenarios. T represents the amount available, or equivalently the number pf auction rounds, on a particular day, for the daily amount of auction, and D is the number of days for which one simulation is repeated. To analyse what happens once the learning process has had time to take place we select, F a number of days at the end of the simulation to examine the behaviour to which the buyers have converged. Finally, R is the number of independent simulations or repetitions of each market scenario. A scenario is also characterized by several parameters

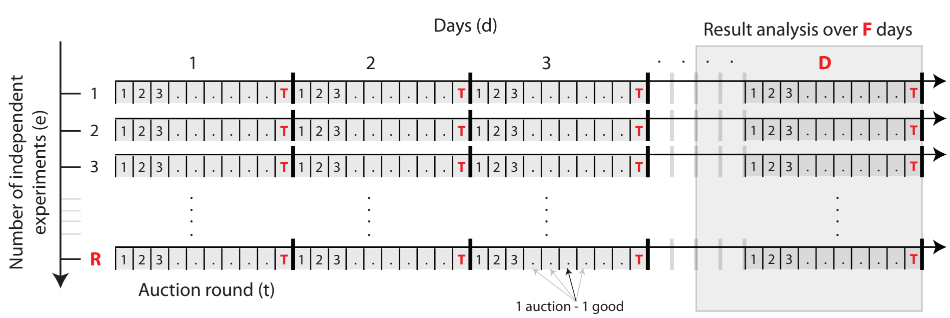

Figure 2: A daily market is composed of T identical auction rounds, on each of which only one unit of good is traded. One simulation is composed of D days. A number R of simulations, giving the same market and experimental condition, are run independently for each scenario studied in order to compute the results at the equilibrium as an average over the last F days of each simulation.

specifying the supply and the demand on the market. This is shown in Figure 3 and Table 1 for the first scenario we simulated.

3.1 Scenarios 1: one learning buyer plus one random buyer

Table 1 gives the simulation settings for the two scenarios, 1A and 1B, investigated in this session. Table 1 is composed of three parts: the left part lists the simulated scenarios, the central part reports the values of the

Table 1: Simulation settings for scenarios 1A and 1B.

Scenarios	*Simulation*			*Supply* ‖ *Demand*		
	R	D	F	T	P_1	$\mathscr{P}_r$
1A	20	10000	50	8	10, 6	1
1B	20	10000	50	8	10, 10	1

simulation settings defined previously in Figure 2, whereas the right part reports the values for characterizing supply and demand in each of the simulated scenarios. Let us look at the left part. The number of days that we have chosen before we consider that learning has taken place might appear excessive. However, this ensures that we can be reasonably confident of any evidence of statistical regularities that we find. By repeatedly exploring all possible states and actions given opponents' behavior over several days we ensure convergence. It is worth remembering that only the last F=50 days are then considered for studying market results. As far as supply is concerned, the parameter T (daily auction rounds) fully describes the daily overall supply given our assumption that one unit is traded in each round. Finally, demand is generated by n learning agents $P_1, P_2, \dots P_n$ or (and) a "zero-intelligence" buyer who chooses uniformly at random from the action space $\mathscr{P}$. The presence of the latter agent guarantees that there is always positive demand for each unit of the good. The action space is the set of 11 prices $p \in \{0, 1 = \underline{P}, 2, 3, ..., 10 = \overline{P}\}$. We introduce the action corresponding to a bid price at 0 which is below the reservation price of the seller $\underline{P}$, in order to give buyers the option to simply exit the market.

The first scenario (1A) is thus characterized by 8 auction rounds and one learning agent willing to buy two out of the eight traded units of good with diminishing marginal profits (10 and 6) and by the presence at each auction round of one uniformly random buyer. The internal states are given by the set of possible goods bought $\mathscr{S}^i_{int} \in \{0, 1, 2\}$ and the external states, by the set of auction rounds $\mathscr{S}^i_{ext}$ $t \in \{0, 1, 2, ..., 8\}$.

Figure 3 presents one plot representing the averaged action-state matrix of attractions (estimated over the final S days) for the learning buyer. The action-state attraction matrix reports what a buyer has learned about the desirability of an action/bid price (y-axis) given a specific state (x-axis). The figure shows with dots the best action for each state given the information that the individual has obtained while learning. It can be thought as the mental map adopted by the buyer to play on the market. Indeed, the dot correspond to the bid price played at each state, because after learning the buyer always select the action with the highest attraction given the state. The states (y-axis) are represented by first fixing the internal state and then running over the external states. Thus, states from 1 to 8 refer to the eight auction rounds, respectively, when the buyer has bought 0 unit of good out of the 2 units she is willing to buy. States from 9 to 16 refer to the sequence of the eight auction rounds when the buyer has bought only one unit of good (she still needs one more). It is worth noting that in Figure 3 there is one state where all actions are indicated with

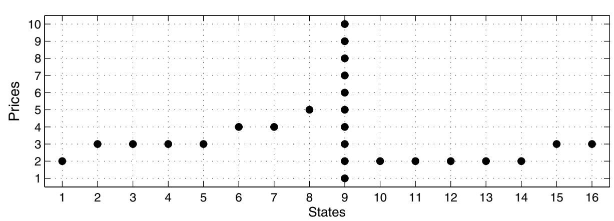

Figure 3: Action-state matrices for scenario 1A (linear demand). The plot highlights the best action for each state.

a dot, that is, state 9. This reflects the fact that such a state is impossible since it would be one in which the buyer at the outset of the auction has already purchased one unit. Since this state is never explored its attractions values always remain at zero and thus equivalent. On the y-axis only the prices between $\underline{P}$ and $\overline{P}$ are shown, but, as previously mentioned, we have introduced a further action corresponding to a bid price at 0 (which we do not plot) providing the buyer with the possibility to exit the market. This exit strategy is never chosen since the buyer is always willing to buy until she obtains the second unit. Indeed, one more state is always added in the simulation which we never mention and shows in the plots. This state is a final absorbing state, where buyers end when they have bought all the units they need. In the absorbing state, buyers bid at 0. When a new trading day occur buyers exit from this absorbing state and they restart from state 1.

The "mental map" of the unique learning buyer (reflected in Figure 3) represents the learned behavior of the buyer after several days of "practice". The first finding is that a significant difference exists between the set of states (1 to 8) when the buyer has not yet bought any good and the set of states from 9 to 16 when one good has been bought. The former set exhibits bid prices higher than the latter set for each corresponding auction round. This behavior simply reflects the fact that, in this scenario, the marginal profit yielded by the second unit is lower than that of the first unit. It is worth noting that all the best actions are significantly below the corresponding marginal profit, namely, 10 and 6. A second and original feature of the learning algorithm is that the buyer learns to recognise time pressure. As time passes without success the buyer learns to increase her bids. Thus, as auction rounds go by the bid price increases for both internal states. In the last auction rounds, the buyer still willing to buy some units of good increases her chance to win by raising the bid price. Given this and the fact that the disturbance term (the random buyer or "noise trader") has mean 5, we can understand how the level of the price evolves. For states from 1 to 8, bid prices are significantly below the corresponding marginal profit (10), but progressively in the later auction rounds bid prices increase very close to the average disturbance price. The buyer learns, in particular in earlier auction rounds, to make a risky bet (bidding at a price significantly lower than the average disturbance price), because she knows that she has still in the future a certain number of opportunities in the subsequent auction rounds to increase her chance of winning the first unit of good by raising her bid price. Similarly, the reasoning applies to states from 10 to 16, but it is also obvious that the bid for the last unit will be lower in the case where the individual has already obtained a unit. In effect, the buyer learns t0 bid against the distribution of bids with which she is actually faced.

Now we give the results for the inelastic retail demand market scenario (1B) where the marginal profits for the learning buyer are constant and equal to 10. One might expect the first finding in the previous scenario to vanish, and that a constant behavior (price strategy) would appear irrespective of the internal state, that is, how many units of good she has already bought. Figure 4 plots the action-state attraction matrix and the best action in each case. It is evident that the bidding behavior does not change significantly, in the sense that both previous findings are confirmed, while the second one is now obvious (given the experimental evidence of previous scenario), the first is not trivial. The buyer learns a different bid price for trying to obtain the first unit of good compared to the second one. This phenomenon cannot be due to different profitability of the units (marginal profits are constant), it is more the outcome of a betting behavior faced with time pressure, that is, the buyer can take more risk (diminishing bid prices) after having bought at least the first unit of good (states from 9 to 16). In other words with one unit in hand and the option of bidding in later rounds the buyer can bid lower.

Finally, Figure 5 reports the price dynamics for four averaged generic days for both scenarios. The comparison shows clearly and the previous discussion explains why. It is worth noting that the noise trader imposes an exogenous reference price level for trading, for the purchase of both units of the good. This scenario is extreme in one sense, the market environment with which the learning buyer is faced is stationary. We will see that the next scenario is extreme in another sense.

Scenarios 2: four learning homogeneous buyers

Table 2 shows the simulation settings for two new scenarios, 2A and 2B. In these scenarios we propose to investigate market outcomes with four learning buyers plus one uniformly random buyer or noise trader. The four buyers co-evolve determining a highly non-stationary environment where four interdependent learning dynamics occur. In a certain sense this scenario is at the opposite extreme to the previous one.

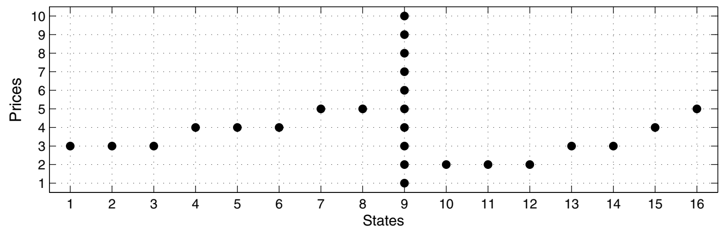

Figure 4: Action-state matrix for scenario 1B (inelastic demand).

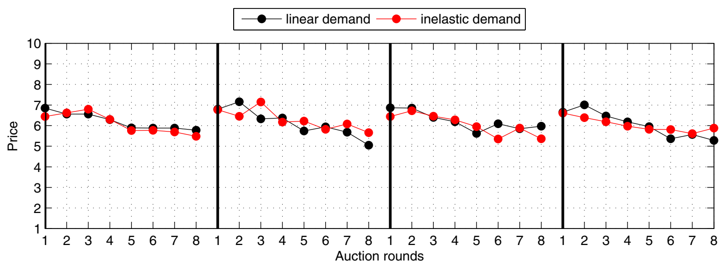

Figure 5: Transaction prices for four generic days (eight auction rounds per day) after l earning for both scenarios 1A (black curve) and 1B (red curve).

Table 2: Simulation settings for scenarios 2A and 2B.

Scenarios	*Simulation*			*Supply* ‖ *Demand*					
	R	D	F	T	P_1	P_2	P_3	P_4	$\mathscr{P}_r$
2A	20	10000	50	8	10, 6	10, 6	10, 6	10, 6	1
2B	20	10000	50	8	10, 10	10, 10	10, 10	10, 10	1

Now, we are assuming that the learning dynamic for each buyer starts from scratch and simultaneously, indeed, they have no priors at all over the market environment. In reality, of course, except in rare cases, buyers do not learn simultaneously from scratch, they do have priors based on their previous experience. For instance a market entrant faces expert players who have already developed a more stable bidding behavior.

Figure 6 and Figure 7 report the action-state attraction matrices after the learning period for both scenarios 2A and 2B, respectively. We report the action-state attraction matrix for only one of the four learning buyers, since, as one would expect, the four matrices are extremely similar in their structure, as all buyers are homogeneous and the situation is therefore symmetric.

We now observe significant difference between the bidding behaviors in the elastic and non elastic demand situations emerges. In the first scenario (see Figure 6), both the diminishing marginal profit and the time pressure effects are present as in previous scenarios. However, the level of the "optimal" bidding price with respect to each state, is now considerably higher. Obviously, we have generated a market condition with more demand. If the noise trader wins in any round the players can no longer all purchase two units. There are three more learning buyers competing for the same amount of scarce resource as in previous scenarios. Given this, figure 7 reports that there is now no difference between the bidding price for all states (both internal and external), the only effect which is still slightly present is the time pressure.
Figure 8 reports the price dynamics for four averaged generic days, after learning, for both scenarios. The comparison now reveals that the declining price phenomenon is not common to both scenarios. This effect is present only in the 2A (linear demand) model. The diminishing marginal profit condition is now the

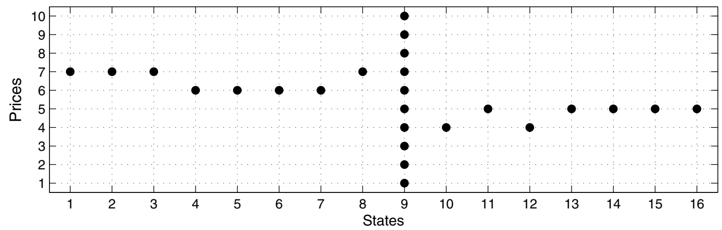

Figure 6: Action-state matrix for scenario 2A (linear demand).

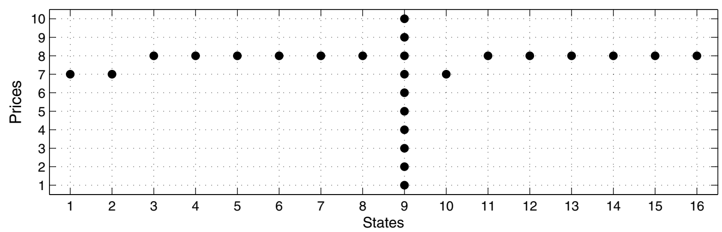

Figure 7: Action-state matrix for scenario 2B (inelastic demand).

only cause for the declining price paradox. In earlier auction rounds, the units of good with the highest value to the buyers are bought, in later rounds the remaining ones. This effect vanishes in the 2B (inelastic demand) model. The buyers have learned to compete strongly to get the desired amount of goods and the prices bid remain at the same level for both units of the good. Indeed, the opposite effect, a slight increase in the price levels as rounds go by, now emerges, because the time pressure effect now plays a prominent role. It is worth noting that the purchases of both units of good occur at prices significantly higher than the average disturbance price or noise trader bid.

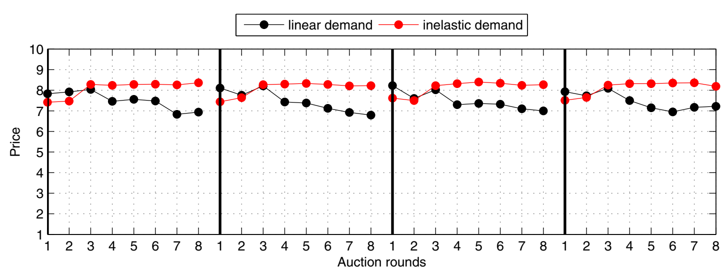

Figure 8: Transaction prices for four generic days (eight auction rounds per day) after learning.

Scenario 3: four learning homogeneous buyers plus random buyers

Table 3: Simulation settings for scenarios 3A and 3B.

Scenarios	*Simulation*			*Supply* ‖ *Demand*					
	R	D	F	T	P_1	P_2	P_3	P_4	$\mathscr{P}_r$
1A	20	10000	50	8	10, 6	10, 6	10, 6	10, 6	1
1B	20	10000	50	8	10, 10	10, 10	10, 10	10, 10	1

Finally, Table 3 shows the simulation settings for the last two scenarios considered, 3A and 3B. These scenarios are identical to scenarios 2A and 2B except for the fact that one day is composed of 12 auction rounds instead of eight. The aim is to explore what happens when the demand of the learning buyers is kept the same whilst more units of the good are available. The competition amongst the learning buyers is thus, in effect, diminished. Figure 9 and Figure 10 report the action-state attraction matrices for both scenarios. Comparing the previous scenario 2A and the new one 3A, similar bidding behavior emerges. The bids for the first unit of good are again around 6, whereas those for the second one are slightly lower 4 instead of 5, just below the average disturbance price of 5. In the new market context, buyers can take more risk and the presence of the random buyer is more strongly felt.

But now a big difference emerges between scenarios 3A and 3B. The decline in competitive pressure allows learning buyers in later rounds become more aware of the presence of the random buyer. Thus they start to decrease their bids for the second unit of the good. Again a different valuation is present in the context of constant marginal profits. This effect is evident also in Figure 11. Figure 11 shows the price dynamics for four averaged generic days after learning for both 3A and 3B scenarios. The declining price phenomenon is now present again in both scenarios.

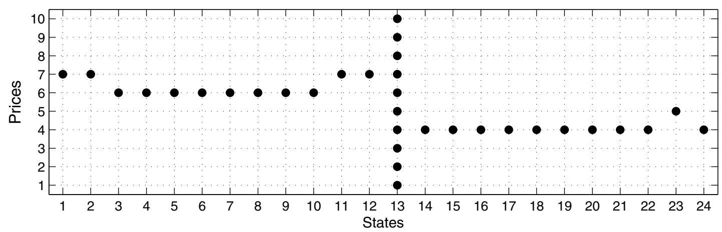

Figure 9: Action-state matrix for scenario 3A (linear demand).

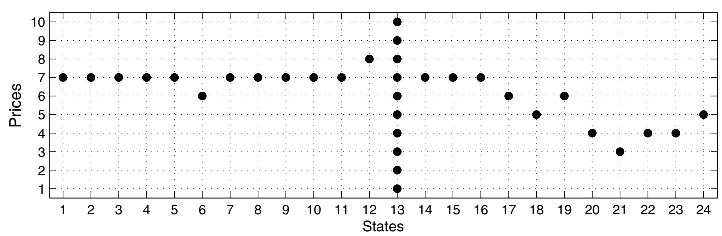

Figure 10: Action-state matrix for scenario 3B (inelastic demand).

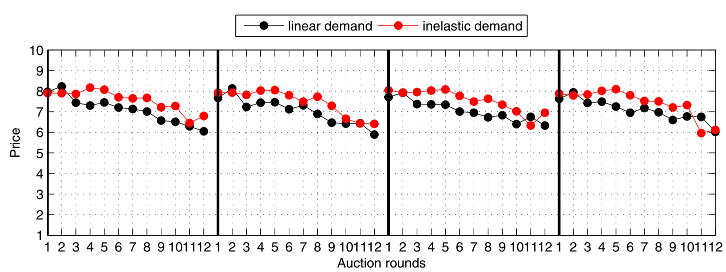

Figure 11: Transaction prices for four generic days (twelve auction rounds per day) at convergence.

4 CONCLUSIONS

This paper has proposed an agent-based computational model to investigate the market price dynamics of sequential Dutch auctions in particular in the context of fish markets. Our approach was to see to what extent learning by market participants who reinforce on their previous experience and can do counterfactual reasoning will generate the sort of phenomena observed in empirical perishable good markets. The particular market on which we based our model, that in Ancona is one on which wholesale buyers meet each other repeatedly and participate in a sequence of Dutch auctions at which the catch of the previous night is sold. The model focuses on the behavior of buyers, and, in particular incorporates inter-temporal profit maximization, conjectures on opponents' behavior and fictive learning. Nonetheless, some behavioral aspects have been neglected. We implicitly neglect the emergence of loyalty among buyers and sellers. In reality, auctioneers try to mitigate this inefficient market outcome by randomly ordering the presentation of the goods of the various sellers throughout the auction rounds. Obviously, this action is not sufficient, because buyers might wait, taking some risk, given that the identity of the vessel that caught the fish that is currently being sold is posted. However this is offset by buyers not knowing when and if fish from the vessels that they favour will appear. Indeed, some empirical studies on auction-based fish markets (Giulioni and Bucciarelli (2011), Gallegati et al. (2009)) have statistically confirmed the presence of loyalty. However, this aspect does not seem to be important in explaining the declining price phenomenon, since there is no obvious reason for the random sorting of the vessels to induce a decline in price as auction rounds go by.

A further behavioral component which has not been incorporated in our buyer's behavioral model concerns the formation of expectations about daily supply. In reality, when buyers come to the market in the early morning, they start to form expectations about the daily overall market supply. They can collect information by looking at the amount of boxes or even speaking with other market actors, they also will have an idea as to how much impact the weather will have had on the catch. Buyers thus form priors about the expected daily supply and bid accordingly. Conversely, our behavioral model, which is a first step to modelling the results of learning in this complex context, has been trained over different scenarios which have been kept stable throughout each independent simulation (number of buyers, number of auction rounds/unit of goods sold).

In conclusion then, our agent-based computational model has successfully replicated the price declining phenomenon which has long been thought of as paradoxical in economics. Our approach involves three potential sources for this feature. They are diminishing marginal profits (linear retail demand model), time pressure (perfectly inelastic retail demand model) and more abundant supply. All these aspects in a more complex and realistic market settings can jointly contribute to the emergence of such price formation phenomenon. Finally, it is worth remembering that all the simulated scenarios have considered an homogeneous population of learning buyers. The heterogeneity of buyers (in particular in terms of marginal profits) is an aspect that has not been addressed here and has been left for future research. Obviously, this may further contribute to "exogenously" generate the declining price effect.

The other important line of research on this topic will be to undertake experimental and empirical analysis in order to validate the behavioral model. Statistical investigation of human-subjects experiments and the bidding data of participants on real wholesale fish markets will enable us to confront our theoretical/computational findings to the empirical and experimental evidence for some of the apparently anomalous features of market prices.

REFERENCES

Ashenfelter, O. (1989), How auctions work for wine and art, *Journal of Economic Perspectives*, 3(3), 23-36.

Camerer, C., Ho, T., (1999). Experience-weighted attraction learning in normal-form games. *Econometrica*, 67, 827-74.

Erev, I., Roth, A. K., (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review* 88 (4), 848-881.

Graddy, K. (2006), Markets: The Fulton Fish Market. *The Journal of Economic Perspectives*, 20(2), 207-220.

Gallegati, M., Gianfranco Giulioni, G., Kirman, A. and Palestrini (2011). A. What's that got to do with the price of fish? Buyers structure behavior on the Ancona fish market. *Journal of Economic Behavior and Organization*, 80, (1), 20-33.

Giulioni, G, Bucciarelli, E, (2011). Agents' ability to manage information in centralized markets: Comparing two wholesale fish markets. *Journal of Economic Behavior and Organization* 80, 34-49.

Kirman, A., & Vignes, A. (1991), Price Dispersion: Theoretical Considerations and Empirical Evidence from the Marseilles Fish Market. In Arrow, Kenneth J. (ed), *Issues in Contemporary Economics: Proceedings of the Ninth World Congress of the International Economic Association*, Athens, Greece. New York: New York University Press.

Weisbuch, G., Kirman, A., & Heirreiner, D. (2000). Market organisation and trading relationships. *The Economic Journal*, April.

Kirman, A.P., & Vriend, N.J. (2001). Evolving Market.s An ACE Model of Price Dispersion and Loyalty. *Journal of Economics Dynamics and Control*, 25(3/4), 459-502.

Kirman, A., Schulz, R., Hardle, W., & Werwatz, A. (2005). Transactions That Did Not Happen and Their Influence on Prices. *Journal of Economic Behaviour and Organization*, 56, 567-591.

Sutton, R. S., Barto, A. G., (1998). *Reinforcement Learning: An introduction*. The MIT press Cambridge.

Watkins, C., Dayan, P., (1992). Q-learning. *Machine Learning* 8 (3-4), 279-292.

AUTHOR BIOGRAPHIES

ERIC GUERCI is Associate Professor at the University of Nice Sophia Antipolis and member of the (CNRS) GREDEG - Nice and also associate member of the (CNRS) GREQAM - Marseille. His research interests lie broadly in agent-based modelling and simulation and behavioral economics. His email address is eric.guerci@gredeg.cnrs.fr.

ALAN KIRMAN is Professor emeritus at the Aix Marseille University, Director of Studies at the Ecole des Hautes Etudes en Sciences Sociales, (EHESS) and member of the Institut Universitaire de France. He studied at Oxford University and obtained his Ph.D in economics from Princeton. He has been a professor at Johns Hopkins University, Universitè Libre de Bruxelles, the University of Warwick and at the European University Institute in Florence. He is a Fellow of the Econometric Society, of the European Economics Association, and has held a Houblon Norman Fellowship at the Bank of England. He was awarded the Humboldt Prize and has been president of the Association of Southern European Economic Theorists, (ASSET) and was awarded the Fondacion Urrutia prize in 2010 for diversity in economic research. He has been a member of the Institute for Advanced Study in Princeton. He is honorary editor of the Journal of Economic Behavior and Organisation, and associate editor of several others. He has published more than 120 articles in refereed journals and is the author of three books, the most recent "Complex Economics" was published in 2010 and editor of 12 others. His main interest are in the way in which economies, markets and their participants actually function and the link between micro and macro behaviour. His e-mail and web addresses are alan.kirman@univ-amu.fr and http://www.vcharite.univ-mrs.fr/PP/kirman/index.html respectively.

SONIA MOULET is phD. She is affiliated to the GREQAM-Marseille. She holds a M.S. and Ph.D. in Mathematical EConomics from the EHESS in France. Her research interests concern the learning on decentralized markets. Her email address is sonia.moulet@gmail.com.