

GREQAM

Groupement de Recherche en Economie
Quantitative d'Aix-Marseille - UMR-CNRS 6579
Ecole des Hautes Etudes en Sciences Sociales
Universités d'Aix-Marseille II et III

Document de Travail
n°2007-16

The influence of seller learning and time
constraints on sequential bargaining in an
artificial perishable goods market

Sonia MOULET
Juliette ROUCHIER

July 2007

DT-GREQAM

The influence of seller learning and time constraints on sequential bargaining in an artificial perishable goods market *

Sonia Moulet[†] Juliette Rouchier[†]

17th July 2007

Abstract

This paper investigates the formation of prices in a perishable goods market where agents bargain repeatedly through pair-wise interactions. After extensive field observations, we chose to focus on two aspects that seem important to actors of this market: the passage of time and update in judgement when gathering information. The main feature of the market is that a seller bargaining with a buyer has incomplete information about buyer's willingness to pay and is not sure how her trading partner will evaluate an offer or compare it with other options. On the other hand, buyers have limited time to look for goods and cannot meet all possible sellers before making a decision. Hence agents cannot calculate the best price to offer but receive information through limited interactions, and use this information to choose their actions.

An agent-based model was built to represent a framework that mimics the observed market institution and where agent's possible behaviors and learning was made as consistent as possible with gathered data. Simulations were run, first for sensitivity analysis concerning main parameters, then to test the dependance of agents' learning to (a) the time buyers can spend on the market and (b) the frequency of update in learning by sellers. To validate the model, features produced by the simulated market are compared to the stylized facts gathered for negotiation about four goods. We reproduce the main features of the data on the dynamics of offers, transaction prices and agents' behavior during the bargaining phases.

Key words: agent-based model, bargaining, perishable goods, pair-wise interaction, decentralized market.

JEL Classification: D40; D82; D83; C63; L14.

*We wish to thank Alan Kirman for his thoughtful remarks. Cars Hommes, anonymous reviewers for JEDC, WEHIA and AESCS conferences in 2005 made extremely useful comments on previous versions of this paper. Organizers of the Porquerolles summer school provided a rich framework for considering our research. Finally, Juliette Rouchier wishes to thank Toru Ishida and JSPS for enabling her to pursue her current research and Sonia Moulet wishes to thank the French Ministry of Research, which provides her grant

[†]GREQAM, EHESS Centre de la Vieille Charité, 2 rue de la Charité, F-13236 Marseille cedex 02, France. moulet@univmed.fr or rouchier@univmed.fr

Résumé

Cet article porte sur la formation des prix dans un marché de biens périssables où les agents négocient de gré à gré de façon répétée. Après avoir accompli un important travail sur le terrain, nous choisissons de nous concentrer sur deux aspects importants pour les acteurs : le temps et les changements de jugements des individus qui reçoivent de nouvelles informations. Ce marché se caractérise par une méconnaissance par le vendeur des critères de choix de son client (valeur de réservation, comment ce client évalue les offres reçues et les compare aux autres options sur le marché). D'autre part, les acheteurs ont une contrainte temporelle et ne peuvent pas rencontrer tous les vendeurs avant de prendre une décision. Ainsi, acheteurs et vendeurs ne peuvent pas calculer le meilleur prix à offrir ou à accepter mais, recevant des informations au cours de leurs rencontres, ils adaptent leur comportement à cette information.

Nous avons construit un modèle multi-agents dont l'environnement reproduit un marché empirique (le MIN, marché de gros en fruits et légumes de Marseille) et dont les agents ont des règles de décision et d'apprentissage cohérentes avec les données de terrain. Des simulations ont permis de tester l'impact marginal de chaque paramètre puis de deux variables précises : (a) le temps que les acheteurs passent sur le marché et (b) la fréquence de réévaluation de ses croyances par le vendeur. Pour valider le modèle, nous comparons les faits produits par le marché simulé et ceux provenant de données empiriques réelles concernant les négociations et transactions pour 4 types de biens. Nous reproduisons les faits principaux concernant la dynamique des offres, les prix des transactions et l'émergence de négociation.

Mots clés : modèle multi-agents, négociation, biens périssables, interactions bilatérales, marchés décentralisés.

Codes JEL : D40 ; D82 ; D83 ; C63 ; L14.

In this paper, we consider certain aspects of the issues related to learning on markets in a situation of imperfect information. For this we developed an agent-based model that represents an empirical market and the bargaining interactions that take place in it. We analyze through simulations the evolution of bargaining dynamics and the setting of prices. The data were gathered in the wholesale market for fruit and vegetables in Marseille.

Our agent-based model has been developed in the light of two previous studies. The first focused on the rationality of sellers as revealed by the way they set prices, and showed that long-term relationships help sellers to ascertain the willingness to pay of buyers (Rouchier & Mazaud, 2004). The other is a quantitative description of bargaining interactions as a succession of offers, counter-offers and prices for one seller over a number of days (Kirman *et al.*, 2005). The first data set contains assumptions about rationality, including descriptions by the actors themselves, that we used to write the model. The second data set is here used as a benchmark for simulation results.

The model structure mimics the observed market where a step is a time-sequence with a succession of bargaining and transaction interactions. Each participant is represented by an adaptive artificial agent that uses an algorithm to make decisions, on the basis of its past experience with the other agents. Rationality is represented by a reinforcement learning process. Simulations enable us to compare the characteristics of the artificial market with those highlighted by the empirical data described in Kirman *et al.* (2005). Two main parameters are studied : the time Buyers can stay in the market and the frequency of updating Sellers' beliefs. The main results are consistent with the dynamics of the empirical market : (1) the ex-post bargaining power index of the sellers is always higher than that of the buyers, (2) it tends to decrease over the course of the day, (3) the loyalty of the buyers does not influence the transaction prices and (4) the relative frequency of the different sorts of matches is reproduced.

The paper is organized as follows : in the first section we describe the market from which the data was obtained and we briefly survey the literature on bargaining processes in real markets. In the second section a description of the model is given. In the third section the main results are presented. Finally, in the last section, we compare our results with the empirical features of the real markets and we propose some topics for further research.

1 Representing bargaining processes in real markets

1.1 Agent-based representation of markets

Traditional game-theory studies bargaining under strong assumptions such as full rationality of the agents and common knowledge of beliefs and preferences. In real life these criteria are rarely met. Using computational models, we can model bargaining processes under weaker assumptions. We can introduce bargaining agents who have little a-priori knowledge and who gradually adapt and search for optimal solutions through trial and error.

Our methodology is usually described as agent-based modeling (Epstein & Axtell, 1996) or agent-based computational economics (ACE) (Testfatsion, 2005). Its aim is to

produce, generate and analyze emergent phenomena at the macro-level by describing decentralized interactions at micro-level. The computational approach does not need to make the restrictive assumptions of perfect rationality and common knowledge as is often the case in game-theoretical models. Agent-based modeling implies that agents are formalized as autonomous entities making decisions on the basis of their own goals, the information they have about the environment and their expectations regarding the future. There is a large literature focusing on the issues of network formation (e.g., Riolo, 1997; Ashlock *et al.*, 1996) and bargaining. About bargaining, Oliver (1996) was the first to demonstrate that a system of adaptive agents can learn effective negotiation strategies using evolutionary algorithms. This literature was then further developed by the use of different learning processes. For example, using the evolutionary approach, Matos *et al.* (1998) generate a model of offers and counter-offers by a linear combination of simple bargaining techniques. Using Q-learning, Rocha & Oliveira (2001) model virtual organizations in an e-commerce environment. Zeng & Sycara (1997) present a study of a practical application using Q-learning. They explore the hypothesis that this learning is beneficial in sequential negotiation and present experimental results. To represent bargaining, Brenner (2002) studies the dynamics of market prices under the assumption of behavioral learning by sellers and buyers. More precisely, he explores the question of whether bargaining on prices becomes common and whether the dynamics of prices converges to an equilibrium. He simulates a learning process and compares its results to the game theoretic predictions under the assumption of rational agents.

Following this literature, our aim is to produce a simple algorithmic behavioral process for artificial agents, which can match the history of intermediate offers and transaction prices we use as a benchmark. To reproduce the price sequences, we choose to consider only the bargaining process, the learning being based on offers and counter-offers without any additional circulation of information. This goes in line with results gathered by Rouchier (2004), including actors' analysis of their own actions, which inspired the development of our algorithm. She studies the same wholesale market as we do, but without reproducing any bargaining. The market is such that buyers are price-takers and switch between different sellers to get the commodities they need. The presence of loyal and opportunistic agents in unequal proportions is used to show the importance of both types of agents for maintaining a sufficient stock by the sellers. Agents possess even less information than in the situation we study. Her paper provides a description of the real market which serves as a basis for our model-building.

1.2 Fitting simulation data with real agent behavior

Agent-based modeling represents a complementary approach to the empirical understanding of markets. The empirical data allow an external and empirical validation of the ACE models. This can give some degree of confidence that the model might serve as an aid in understanding aspects of markets. The empirical understanding of markets can be deduced both from human subject experiments and historical data. Duffy (2006) has compared the empirical performance of a variety of learning models (e.g. zero-intelligence agents, reinforcement learning, evolutionary algorithm). He surveys literature in which agent-based models have been used to study findings from

human subject experiments and compares these findings with those from agent-based simulations. Janssen & Ahn (2006) compare the empirical performance of a variety of learning models and theories of social preferences in the context of experimental games involving the provision of public goods. They prove that it is important to evaluate model relevance at the local and at the global level when validating a learning model. Some other researchers have explored the empirical puzzles of stylized market facts. Andreoni & Miller (1995), Duffy & Uvner (2007) and Rouchier & Robin (2006) focus largely on auction mechanisms. Andreoni & Miller (1995) and Rouchier & Robin (2006) create and analyze a model of adaptive learning and demonstrate that such a model can capture the bidding patterns among human subjects in experimental auctions. Duffy & Uvner (2007) develop a model of dynamic internet auctions that reproduce both the macro and micro phenomena observed in laboratory experiments with human subjects. LeBaron (2006) replicates empirical facts from financial data. He argues that agent-based model approaches make more sense economically than their representative agent model counterparts. In the same spirit, Kirman & Vriend (2001) and Weisbuch *et al.* (2000) have worked on reproducing the behavior of agents in Marseilles fish markets. They build an ACE model and compare the macro-system with the stylized facts derived from empirical data.

1.3 The empirical wholesale fruit and vegetable market and the historical database

1.3.1 The empirical wholesale fruit and vegetable market

The wholesale market in Marseille is one of the markets created by the French government during the sixties in order to ensure the supply of fruit and vegetables to the city. It is an area controlled by the state where professional wholesalers and retailers can meet. All the sellers are located in a small area of less than two hectares and display all the goods on sale in the morning. The market opens at 3 :30 a.m. and closes at 9 :30 a.m. Professional buyers enter the market freely. In the main alley they can see the fruit and vegetables on display, and, from this, they get an idea of the qualities and quantities of all available products. All transaction prices are bilateral, and they are set through a private face-to-face bargaining process.

Prices are not posted and have to be requested by the buyers. The acquisition of the information used to decide on prices is a step-by-step process of inference from "private signals" (acceptance, rejection of offers and bargaining by agents). There is no centralized information on available quantities and almost no communication among agents of the same type. Neither buyers nor sellers communicate amongst themselves about prices or available quantities.

The buyers' search time is limited, because they have to return to their own shops. They cannot visit all the sellers to compare offers and choose the best price of the day. They therefore have to use some procedure other than systematic comparison on each market day, so as to limit their search time. To solve this problem, buyers use long-term information obtained from regular buyer-seller links. We note in this regard that buyers' loyalty to sellers is an important feature observed in wholesale markets (e.g. Rouchier & Mazaud, 2004; Kirman & Vriend, 2001)

1.3.2 The historical database

A quantitative data-set has already been analyzed by Kirman *et al.* (2005) and was recorded on this market for 50 market days in 1983. These observations are a description of the activity of a seller over one day with one record every 5 minutes. It gives a list of prices offered even if no transaction takes place (offers, counter-offers and transaction prices), as well as idle periods (no buyer). Only one seller's bargaining details were kept, but this consists of 3960 prices for four kinds of fruit and vegetables : leeks (an out-of season commodity), local tomatoes (a local in-season commodity), imported tomatoes (imported in-season commodity) and oranges (a commodity coming to the end of its season).

This is a very original data set in that it shows intermediate offer prices, whereas usually only final prices are observed. The data allows us to model the bargaining process (to be specified later on). Apart from the sequences of bargaining, Kirman *et al.* (2005) study the evolution of prices, and find an expression for the ex-post bargaining power index as follows :

- 0 if the transaction takes place at the buyer's counter-offer.
- 1 if the transaction occurs at the first offer proposed by the seller.
- $\frac{P-CO}{FO-CO}$ where FO is the seller's first offer, CO is the buyer's counter-offer, and P the transaction price.

Even though the data-set is based on the observation of only one seller, whose purchase price is unknown, and although it is not possible to discriminate among buyers when reading the prices, we use the data and the statistical analysis from that contribution as a benchmark in this paper. We do not present all the data that can be found in the paper (Kirman *et al.* , 2005), to which we refer readers who wish for more details.

2 The model

2.1 General framework of the model and agents' characteristics

As a convention we denote *Sellers* and *Buyers* to denote the artificial agents and *sellers* and *buyers* for the real agents. The model was implemented in Smalltalk using VisualWork 7.4nc. The pseudo codes are available online.

We make the following assumptions :

- (A1) There are two classes of agents : Buyers and Sellers. The two populations are static and repeatedly matched.
- (A2) Agents interact through private pair-wise interactions and there is no central market.
- (A3) Agents learn using a reinforcement learning model and, more precisely, classifier systems.
- (A4) If a good is not sold at the end of a market day it is lost (this represents perishability).
- (A5) There are 3 decisions on the market :
 1. The Sellers manage their amount of stock : This decision is based on a rule shown to be optimal in Abel (1985).

2. The Buyers choose which Sellers to visit. This decision uses discrete choice or multinomial logit model. Brock & Hommes (1997) implement this as an evolutionary selection of expectation rules based on realized profits, similar to this set-up. Heterogeneous agent models containing this kind of discrete choice selection rules have recently been surveyed by Hommes (2006). The selection is based on a choice where the Seller who gave the Buyer the highest payoff in the past is visited with a higher probability than the other Sellers. This represents a trade-off between the exploitation of old knowledge (past experiences) and the acquisition of new knowledge.

(A6) The bargaining process consists of 4 steps :

- Step 1.** The Seller makes an offer which can be accepted or rejected by the Buyer.
- Step 2.** The Buyer either accepts the offer, or rejects it and walks away, or decides to make a counter-offer.
- Step 3.** The Seller either accepts or makes a counter-counter-offer.
- Step 4.** The Buyer either accepts the counter-counter-offer or rejects it and walks away.

We can represent this sequence by the decision-tree in extensive form with 4 decision nodes, alternating between the Seller and the Buyer.

- The Seller has 2 decision nodes : S_1 and S_2 .
- The Buyer has 2 decision nodes : B_1 and B_2 .

At each decision node, the deciding agent has a number of options :

- S_1 : {First offer}
- B_1 : {Accept, Reject, Counter-offer}
- S_2 : {Accept, Counter-counter-offer}
- B_2 : {Accept, Reject}

The outcome of each decision is given in figure 1. The bargaining process can only end in two states : transaction or no transaction. We go back to the abbreviations used by Kirman *et al.* (2005) to describe the outcomes of the interactions. If the Buyer accepts the first offer, FO , then a Transaction No Bargaining (TNB) occurs. If the Buyer breaks the meeting just after FO without deal, the outcome Only Offer (OO) is observed. If the Buyer rejects the offer, makes a counter-offer CO and rejects the counter-counter-offer CCO then we observe Only Counter-offer (OC). Lastly, if the bargaining leads to a transaction, a Transaction With Bargaining occurs (TWB).

[FIG. 1 about here.]

(A7) (1) The Sellers' choice of offers and the Buyers' choice of counter-offer decisions are based on a classifier system without genetic algorithm (Holland, 1975). That is, it is not a learning classifier system, and it does not employ a genetic algorithm to generate new rules. (2) The Sellers' choice of counter-counter-offers are based on a simple rule as revealed by the way they set them.

(A8) There are two costs : a bargaining cost (c^B) and a search cost (c^S). We hence consider the total value of the time an Agent can spend in the market. More precisely, the two variables c^B and c^S represent the value of the time cost when a participant is either looking for a Seller or actually negotiating. This idea corresponds to the time

costs found in the search literature (e.g. Diamond, 1984; Mortensen, 1982; McMillan & Rotschild, 1994)

A simulation consists of the initialization of the artificial market and a succession of market days for which the model is run. At the beginning of each market day, Buyers enter simultaneously and Sellers have an initial stock of inventory and a limit value. Then, the time is divided into rounds. The maximum number of rounds is called the Buyers' time constraint from 1 to 4. A round is defined as follows :

- Order of Buyers randomly set at each time-step.
- Buyers choose one Seller to visit (using the discrete choice selection rule) and form a queue.
- Each Seller interacts with all the Buyers in her queue and updates the strength of the rules in the classifier system.

At the end of a round, if a Seller has not sold all her goods or if a Buyer has a non saturated time-constraint and has not yet made a transaction, then they will participate in the next round. Note that Sellers do not know how much time Buyers can stay in the market. They only know that the maximum number of rounds is four. The value of four was chosen to correspond to the real market.

Each Buyer j is defined by :

- **Initial parameters** : limit value (v^B) and time constraint for each market day t

$$c^B * S_t^B + c^S * S_t^S \leq C$$

where c^B is the cost of bargaining, c^S is the cost of searching, S_t^B is the number of counter-offers already made on this day, S_t^S is the number of Sellers already visited on this day and (C) is the upper bound of the buyers' time constraint.

- **Updated parameters** : probability of visiting each Seller k during the market day t ($p_{j,k}^t$)
- **Methods** : updating process for the classifier rule strengths and the discrete choice rule decision process for choosing Sellers and bargaining.

Each Seller k is defined by :

- **Initial parameters** : limit value (v^S), initial stock, frequency of Sellers' learning (ϕ)¹.
- **Updated parameters** : amount of stock.
- **Methods** : updating process for the classifier rule strengths, decision process before the market opens and during bargaining.

2.2 Decision rules

2.2.1 Payoff of matches

The payoff of a meeting depends on its outcome and the time Buyers have spent searching before this meeting. Suppose Buyer j matches with Seller k during day t . The Buyer obtains a payoff $\pi_{(j,k)}^t$ ² and the Seller obtains $\pi_{(k,j)}^t$

¹ ϕ can also be expressed as the inverse of Sellers' learning speed

²We can consider that Buyers receive the same utility by buying early at a high price as they do by buying at a slightly lower price but after searching or bargaining more. This expression allows us to introduce the trade-off between exploitation and exploration. The cost of bargaining and the cost of searching are seen as a global cost that can be interpreted as opportunity cost.

If a transaction occurs at price p , then the payoffs are :

$$\begin{aligned}\pi_{(j,k)}^t &= \pi_1^t + \pi_2^t = (v^B - p) + (C - c^B * S_t^B - c^S * S_t^S) \\ \pi_{(k,j)}^t &= p - v^S\end{aligned}$$

If no transaction occurs :

$$\begin{aligned}\pi_{(j,k)}^t &= \pi_1^t + \pi_2^t = (0) + (C - c^B * S_t^B - c^S * S_t^S) \\ \pi_{(k,j)}^t &= 0\end{aligned}$$

Where v^B is the limit value of the Buyer, S_t^B is the number of bargaining events in which she has taken part on this day and S_t^S is the number of Sellers she has met on this day.

2.2.2 Buyer's choice of Seller to visit

At least once per market day, each Buyer makes a decision about which Seller to visit. Due to the time constraint, she cannot visit all the Sellers. The probability of choosing a Seller depends on the δ -discounted past utility the Buyer acquired when visiting this Seller, denoted by $\Pi_{j,k}(t)$.

Here, we consider a parameter μ which measures the loyalty of Buyers towards Sellers with whom they make a high profit. This μ corresponds to the 'intensity of choice' parameter in Brock and Hommes (1997). The extreme cases are the situations : $\mu = 0$ where the Buyer randomly chooses the Seller she wants to visit without taking into account the past utility. The other extreme is the case where μ tends to ∞ . The Buyer stays completely loyal to the Seller with whom she makes a high profit. Formally, the δ -discounted past utility j acquired with k up until t is :

$$\Pi_{(j,k)}(t) = \begin{cases} 0 & \text{if } t = 1 \\ \delta * \Pi_{(j,k)}(t-1) + (1 - \delta) * \pi_{(j,k)}(t) & \text{otherwise} \end{cases} \quad (1)$$

where $\pi_{(j,k)}(t)$ is the payoff of the Buyer j in her meeting with k . Then the probability that j chooses to visit k at t is given by :

$$p_{(j,k)}(t) \equiv \frac{\exp(\mu * \Pi_{j,k}(t))}{\sum_l \exp(\mu * \Pi_{j,l}(t))}. \quad (3)$$

2.2.3 Bargaining over prices

Once matched, Buyers and Sellers make decisions about their matches. At each decision node, the deciding agent has a number of options :

- S_1 : {First offer}.
- B_1 : {Accept, Reject, Counter-offer}.
- S_2 : {Accept, Counter-counter-offer}.
- B_2 : {Accept, Reject}.

The decisions in the bargaining process S_1 and B_1 (see above) are modeled separately for each individual agent by means of a classifier system without genetic algorithm (Holland, 1975). Each rule is defined as a triplet c, a, s consisting of a condition 'c', an action 'a', and a measure of its strength 's' (if 'c' then 'a' with strength 's').

To choose a rule among those that are active, a bid b is made as follows :

$$b(t) = s(t) + \varepsilon. \quad \text{where } \varepsilon \sim U_{[0,a]} \quad (4)$$

Then, the rule with the highest bid gets activated.

When a rule has been used, the classifier system updates the strength as follows : the strength s of a rule that has been used and has generated a reward π at time $t - 1$ is :

$$s(t) = (1 - \zeta) * s(t - 1) + \zeta * \pi(t - 1) \quad \text{where } 0 < \zeta < 1 \quad (5)$$

All the conditional parts of the classifier rules concern the round

(*< if it is round n >*). The frequency of updating the rule strengths is not the same for Buyers and Sellers : Buyers update the classifier system at the end of each meeting, Sellers update only after every ϕ visits (we call ϕ the frequency of Sellers' learning).

Decision S_1 : The first decision consists in choosing an offer. For each conditional part, the rule set is :

$$\langle \text{offer } v^S \rangle, \langle \text{offer } v^S + 0.1 \rangle, \dots, \langle \text{offer } v^B \rangle$$

To evaluate the performance of each rule in the rule set, the Seller keeps track of the average payoff she obtained when using the rule in the past (her experience). The strength of each activated rule is updated every ϕ times using the payoff $\pi_{(k,j)}^t$ as reward from the external environment.

Decision B_1 : Suppose that Buyer j has just received the first offer FO but can stay on the market η rounds after the current round. j does not choose her action directly from her classifier system. She builds 3 other rules, one for each of 3 possible alternatives : *< accept FO >*, *< reject FO and make a counter-offer CO >*, *< reject FO and stop >*. The strength of these rules depends on the bids of the rules already in the classifier system. Thus, we denote by b_m^i the bid³ associated with the rule *< if it is round i then offer m >* and \bar{b}^i is the highest bid of all the rules with the conditional part *< if it is round i >*. Table 6 gives the expected payoffs associated with each possible rule. The rule that is activated by the classifier system is the rule with the highest expected payoff. If there is a tie between the rules, then there is a ranking between the 3 alternatives.

[TAB. 1 about here.]

Decision S_2 : The Seller can accept or reject CO . If she rejects it, she makes a counter-counter-offer CCO equal to or lower than FO . The decision depends on the number of Buyers in the current round, the expected number of Buyers in the following rounds and the Seller's remaining stock. The Seller calculates the value of CCO . If the Seller makes a counter-counter-offer, this counter-counter-offer is in the interval

³The term 'bid' refers to the weight of a rule in the classifier system, not to the economic meaning of a bid as an offer to buy.

between $[CO, FO]$ since we impose the following rules : \langle If $C\tilde{C}O \leq CO$ then the Seller accepts CO \rangle , \langle If $C\tilde{C}O \geq FO$ then the Seller again offers FO \rangle , \langle if $CO < C\tilde{C}O < FO$ then the Seller offers $CCO = C\tilde{C}O$ \rangle . In the first rule the Seller is in principle willing to sell at a lower price than CO , so she accepts CO . In the second rule, the Seller wants to sell at a higher price than FO but cannot do so, so she asks for FO again. In the last rule, the seller is willing to offer a price between FO and CO , she offers $C\tilde{C}O$. $C\tilde{C}O$ is calculated as follows :

$$C\tilde{C}O = CO + h(FO - CO), \text{ with :}$$

$$h = \frac{\Upsilon}{\text{remaining stock}},$$

$$\Upsilon = \frac{E[B^i(t)] + Q^i(t)}{Q_0^i(t)}.$$

Where :

$E[B^i(t)]$: Expected number of Buyers still to come in all the future rounds after the current round i but on the same market day. This is an extrapolation of the day before.

$Q^i(t)$: Number of Buyers in front of the Seller when she makes her decision.

$Q_0^i(t)$: initial number of Buyers in the current round but still on the same market day.

Decision B_2 : Suppose j has just received CCO as counter-counter-offer. At this stage in the bargaining, j has two alternatives : accepting or rejecting the counter-counter-offer CCO . Suppose that Buyer j can stay in the market for y rounds after the current one. As in decision 2, b_m^i denotes the bid associated with the rule \langle if it is round i then offer m \rangle and \bar{b}^i denotes the highest bid of all the rules with the conditional part \langle if it is round i \rangle . Table 2 gives the expected payoffs of each alternative. The rule to be activated is the rule with the highest expected payoff. If the payoffs are the same for the 2 rules, the Buyer accepts CCO .

[TAB. 2 about here.]

2.2.4 Decision rule : Sellers' stock re-initialization

At initialization, a Seller receives an amount of stock exogenously. Then, during the following market days, she manages her stock following the rule : \langle if I sold M units of the good at t , then I purchase $M + Q$ units at time $t + 1$, where Q is the number of Buyers that I did not bargain with because of a depleted stock \rangle . The depleted stock is in fact the effective excess demand for this Seller.

3 Simulations

3.1 Parameters and sensitivity analysis

All the simulations were run 10 times (varying the random seeds for the stochastic elements). The number of time steps varies and corresponds to the number of market

days the system takes to stabilize. The results presented here are those arising from a single simulation run. Table 3 displays the tested and fixed parameters.

[TAB. 3 about here.]

At initialization, the strengths of all the rules are set to the maximum possible payoff ($v^B - v^S = 2$). Before any comparison with empirical data, we ran a sensitivity analysis⁴ and we studied the evolution of the system. We performed a sensitivity analysis for the parameters with varying values: the frequency of Sellers' learning (ϕ) and the upper bound of Buyers' time constraint denoted by C .

- We test the frequency of the learning process by which Sellers update the strength of their rules in the classifier system. These rules are used to decide the first offer. We tested the values: $\{1, 20, 40, 90\}$. When $\phi = 1$, Sellers update the classifier systems after each meeting, when $\phi = 20$ (resp. $\phi = 40$), Sellers update it after every 20 (resp. 40) times. Finally, for $\phi = 90$, Sellers update the classifier system after every 90 matches.
- We test the Buyers' time-constraint. Our empirical observation is that buyers can visit at most four sellers before leaving the market. This empirical fact is common knowledge in the real market. We therefore tested 4 situations: (1) When $C=0.49$, we assume Buyers can visit at the most one Seller. $C=0.74$ (resp. $C=0.99$) signifies that Buyers can visit at the most 2 (resp. 3) Sellers. For $C = 1.44$, Buyers can visit 4 Sellers. In every case, Sellers only know that Buyers can visit at most 4 Sellers.

3.2 Observation

Different indicators show the evolution of the market. It is relevant to observe them both in the first steps in the learning process, during the very first market days when agents are discovering a situation, and in the long run. We hence insist on the different phases in the evolution of the system. In the observed market, individuals often have to learn in the face of new situations. We describe how the following indicators vary over the learning process: (1) Evolution of offers and counter-offers, (2) transaction prices, (3) The matches: emergence of bargaining and rounds of transaction and (4) the Sellers' ex-post bargaining power index (Sellers' EBP).

4 Impact of time constraint (C) and frequency of Sellers' learning (ϕ): results, support and rationale

We describe here the evolution of the indicators mentioned above depending on the two most significant parameters. To summarize, agents learn in different steps. At the beginning of the learning process, transactions occur systematically at the last opportunity of the Buyers and prices are very heterogeneous (Phase 1). As time goes by, the transaction prices become more concentrated around the Sellers' limit value (Phase 2). Then, agents turn to buying more and more often with the first Seller and

⁴The results of the sensitivity analysis can be obtained from the corresponding author

to bargaining (Phase 3). Eventually, the largest proportion of transactions occurs with the first Seller and the prices stabilize around the Sellers' limit value (Phase 4). Table 4 summarizes the impact of C and ϕ on the indicators. The impact can concern both the speed of convergence of the system and the final state.

[TAB. 4 about here.]

4.1 The intermediate offers depending on C and ϕ

4.1.1 Evolution of Sellers' first offers FO :

Result 1 : (1) If the frequency of Sellers' learning is lower than 50 ($\phi \leq 50$), offers converge but the speed of convergence decreases with ϕ . (2) If $C > 0.49$, the evolution of the distribution of first offers differs among the rounds.

Support for result 1 : If Buyers can visit only one Seller ($C = 0.49$) : during the first market days (Phase 1), Sellers try out all the rules. Then, in the long run, they always offer (as their first offer) prices lower than 4 (see left-hand graphs in Fig. 3, $\phi = 1, 20, 40$). If $\phi > 50$ offers do not converge and are made randomly (see the last left-hand graph in Fig. 3).

If Buyers can visit several Sellers ($C > 0.49$) : when $\phi = 90$, Sellers always randomly make their offers. The distribution is the same for all the rounds. In this case, the panel of offers in each round looks like the panel obtained for $C = 0.49$ and $\phi = 90$.

The interesting cases are when ($\phi \in 1, 20, 40$). For example, consider the case $C = 0.99$ and $\phi = 1$, Sellers can receive up to 3 rounds of Buyers and adapt the strength of the rules after each meeting. Fig. 2 shows that Sellers do not act in the same way for all the rounds. For the first round, after less than 100 matches, Sellers always offer prices lower than 3.4. For the second round, there is no dominant action (prices are offered randomly). For the third round, Sellers start by trying all the actions during the first 26 market days (Phase 1), and then, after this 26th market day, they always offer prices lower than 3.4 (Phases 2-3-4).

[FIG. 2 about here.]

Rationale for result 1 : The value $\phi = 50$ plays an important role. It is, in fact, equal to the ratio of the number of Buyers to the number of Sellers, each Seller seeing, on average, 50 Buyers per day. So, if $\phi > 50$, the Seller does not learn and makes random first offers. If $\phi \leq 50$, Sellers learn about which first offers to make. The payoff obtained by a Seller when using a rule depends on the transaction price, which, in turn, depends on the first offer and on the counter-counter-offer (in the event of bargaining). Since most of the time the transaction price is around the Sellers' limit value, the payoff is the same whatever the first offer. But making low offers increases the probability of selling. So low first offers are preferred to high ones. First offers are therefore concentrated around the Sellers' limit value.

4.1.2 Evolution of Buyers' counter-offers CO :

Result 2 : Whatever the value of ϕ , the distribution of counter-offers converges towards a distribution highly concentrated around the Sellers' limit value.

Support for result 2 : If Buyers can visit only one Seller ($C = 0.49$) : when $\phi = 1$ (right-hand graphs in Fig. 3), the convergence occurs after three days. For the other values of ϕ (right-hand graphs in Fig. 3, $\phi = 20, 40, 90$), the dynamics is almost the same : during the first 10 days (Phase 1), the counter-offers are very heterogeneous. Then there is an intermediary phase (Phases 2-3), where counter-offers are highly concentrated between 3 and 3.6 with some counter-offers dispersed between 3 and 5. The length of this phase increases with ϕ . For $\phi = 20$, the phase covers 10 market-days, for $\phi = 40$, it covers 20 market days and for $\phi = 90$, it covers the 40 days. Eventually, after this phase, the counter-offers become stabilized between 3 and 3.6 (Phases 3-4)

[FIG. 3 about here.]

If Buyers can visit several Sellers ($C > 0.49$) : the dynamics of counter-offers is linked to the dynamics of offers. When Sellers' offers are systematically low, Buyers' counter-offers are either very low (3 or 3.2) or non-existent. When Sellers do not always offers low prices, the dynamics is different. In the case $c = 0.99$ and $\phi = 90$, for the first round, two succeeding phases take place (Fig. 4). In the first phase (days 1-20), Buyers make counter-offers between 3 and 5. Then, in days 21-100, they either do not make counter-offers or offer prices below 3.4. For the second round, we cannot conclude because Buyers do not make enough counter-offers. For the third round, there are 2 phases. In the first phase (days 1 to 10), Buyers make counter-offers between 3 and 5. In phases 2-3, (the 30 following market days) their counter-offers are systematically lower than 3.4 to begin with, and then, eventually, become highly concentrated between 3 and 3.2 in the phases 3-4.

[FIG. 4 about here.]

Rationale for result 2 : Consider two cases : (1) In the first case, Sellers make low offers, and Buyers then either accept them or make lower counter-offers. The counter-offers are therefore systematically low. (2) In the second case, Sellers do not systematically make low offers, but Buyers still make low offers. This is rational since the counter-counter-offer is a linear combination of the two offers. Buyers learn by reinforcement that the lower the counter-offer, the lower the counter-counter-offer should be. Since Sellers do not learn about counter-counter-offers, they just apply the simple decision rule S_2 .

4.2 The transaction prices depending on C and ϕ

Result 3 : (1) If the frequency of learning is lower than 50, the transaction prices converge towards Sellers' limit values. (2) If the frequency of learning is strictly higher than 50, there is no convergence towards the Sellers' limit value, but the transaction offers are highly concentrated around this value.

Support for result 3 : If $\phi < 50$, transaction prices converge. Whatever the value of C , the dynamics is the same : there is first a phase (noted P1) where prices stay between 3 and 5. Then, prices become stable and converge towards the Sellers' limit value (Phase 2). The only impact of C concerns the length of P1. Fig. 5 shows that when $C = 1.44$, P1 lasts 3 days when $\phi = 1$ and 20 days when $\phi = 20$. If $\phi = 90$: Prices do not converge, there is just a higher concentration between 3 and 3.6 after the 28th day.

[FIG. 5 about here.]

Rationale for result 3 : (1) This result is linked to the fact that agents bargain : if Buyers did not bargain, Sellers would only make high offers and Buyers would accept them. Whatever the Sellers' counter-counter-offers, at their last opportunity, '*accept the counter-counter-offer*' gives Buyers a higher profit than '*reject the offer*'. On the basis of their reinforcement learning, Sellers would increase their prices during the last round. However, Buyers bargain and Sellers apply the decision rule S_2 to calculate the counter-counter-offer. This decision rule is not based on reinforcement learning : it is just a simple rule⁵ that expresses the counter-counter-offer as a linear combination of the first offer and the counter-offer. So, there is nothing in the model to ensure that Sellers sell at high prices. On the contrary, the fact that Buyers bargain enables them to benefit from low transaction prices, and we see Buyers take almost all the profit.(2) If the frequency of Sellers' learning is higher than 50, offers are randomly chosen between v^S and v^B . We have seen that Buyers make systematically lower counter-offers (result 2) but bargain in 60% of matches. The previous points (1) and (2) apply and lead to a distribution of prices between v^S and v^B characterized by a high concentration around v^S .

Result 4 : In the long run, **transaction prices** do not decrease over the rounds during one market day.

Support for result 4 : As we saw above, we can distinguish between 2 groups of market-days : (1) those with stable prices and (2) those with unstable prices. During the former, prices are stable within each day. During the latter, prices decrease, but simulations show that waiting until the last round does not necessarily result in a lower price. Indeed, when we do an 'ANalysis Of VAriances between groups' test to compare the average transaction prices the Buyers obtain if they make the transactions in the first round or in the last round, the results show that in both cases, the means are not equal and we cannot conclude anything about the round giving the lowest price. The only evidence is that when Buyers transact at the fourth round they cannot gain enough to cover their search cost.

Rationale for result 4 : The Buyers' time constraints are fixed for each day, and Sellers learn about this. Since agents bargain, we can explain the evolution of transaction offers in terms of the evolution of counter-counter-offers. After a period of

⁵This rule was built from empirical observations and the seller's explanations

learning, Sellers easily anticipate when Buyers will come, and they do not change their counter-counter-offer. By applying the decision rules S_2 , counter-counter-offers (and so, transaction prices) do not decrease over the rounds.

4.3 The nature of matches depending on C and ϕ

On the first market day, Buyers enter the market simultaneously and receive offers, which they always reject until their last opportunity. They try the different possible actions. As time goes by, Buyers start to accept offers earlier, instead of waiting until the last Seller they visit. Both parameters have an impact on the results : C has an impact on agents' strategies and ϕ has an impact on the emergence of bargaining.

Result 5 : If Buyers can visit more than 2 Sellers, in the long run, Buyers buy from the first Seller visited and bargain most of the time.

Support for result 5 : In all cases, the system evolves continuously to end up at the same state. Only the speed of convergence differs. To begin with (Phase 1), agents buy in their last round, then (Phases 2-3-4) they learn to buy in the first round. During phase 1, in intermediate rounds, Buyers simply ask for a price and break off the bargaining process. During the phases from 1 to 4, when Buyers purchase, they have always bargained before, whenever the first offer is higher than the Seller's limit value. Since the evolution of the distribution of first offers depends on ϕ , the evolution of the proportions of transactions with and without bargaining also depends on ϕ .

Rationale for result 5 : (1) Consider the fact that agents learn to buy from the first Seller visited. According to result 4, prices do not decrease within each market day. Using their reinforcement learning, Buyers learn that it is more profitable to buy early. The change from "buying from the last Seller visited" to "buying from the first visited" happens as follows. During the first steps in the learning process, Buyers always buy from the last Seller visited and never in the previous rounds. At the same time, Sellers update the strength of their rules for these previous rounds, and when prices become sufficiently low in those rounds, Buyers start to accept the offers made in the first round. There is almost no activity in the intermediate rounds. The decision to buy is based on the decision rules B_1 and B_2 . Using these rules, the Buyer compares the profit to be obtained from accepting the current offer with the profit expected from waiting for the next rounds. Prices being quite equal in all rounds, it is more profitable for Buyers to buy in the first round, so as to avoid search costs, rather than waiting to buy later on⁶. So, if there are more than 2 rounds, credible activity takes place almost exclusively in the last round, at the beginning of the learning process, and in the first round, later on in the process. (2) Consider the fact that agents bargain most of the time. If agents make a counter-offer, they are sure to pay a price equal to or lower than the initial price, and the negotiation cost appears to be low compared to this profit.

⁶'quite equal' because of the noise in the activation of the rules

4.4 The evolution of the ex-post bargaining power (EBP) index with variations in C and ϕ

Since these evolutions are similar for all Sellers, we limit the analysis to one Seller.

Result 6 : The bargaining power of Sellers either decreases or remains constant over the course of the day. The sample mean and the median are higher than 0.6. There are two mass points. The first is around 1 and it represents the highest concentration. The second is 0.

Support for result 6 : The simulations make it clear that C has an impact on the evolution of the Seller's EBP and that ϕ has no impact. If $C = 0.44$, the Seller's EBP increases over time and converges towards 1. When a transaction occurs, the price is much closer to the Seller's first offer than to the Buyer's counter-offer (Figure 6, Graph 1). If $C = 0.74$ or $C = 0.99$, after few days, the distribution of the Seller's EBP presents two mass points : one between 0.8 and 1 and one at 0. When a transaction occurs, the price is either almost equal to the Seller's first offer or equal to the Buyer's counter-offer (Figure 6, Graphs 2 and 3). If $C = 1.44$, the Seller's EBP remains dispersed between 0 and 1, with, in the long run, a high concentration between 0.8 and 1, around 0 and at 0.5 (Figure 6, graph 4).

[FIG. 6 about here.]

The Seller's EBP index either decreases or remains constant over the course of one day, but it always remains higher than 0.6. The decreases are observed most of the time at the beginning of the simulations. The EBP becomes constant when learning takes place.

Rationale for result 6 : The result means that when a transaction occurs, the price is either very close to the Seller's first offer (FO) or equal to the Buyer's counter-offer (CO). The mass point around 1 can be explained by the regularity of Buyers' behavior. Every day, almost the same number of Buyers visit a given Seller. Consequently, this Seller has no surplus in her stock. She is not surprised by some customers coming late (as time goes by, fewer and fewer Buyers visit the Seller in the last round). The Seller anticipates the number of Buyers well and does not reduce her first offer very much. The Seller's ex-post bargaining power is therefore close to 1.

5 Fitting the model with the empirical facts

5.1 Summary of the Stylized facts highlighted in Kirman *et al.* (2005)

(1) Concerning the seller's first offer, the data shows very little or no change during most days for in-season goods (oranges and domestic and imported tomatoes), while in the market for the out-of-season good (leeks), it fluctuates during all trading days

and moves downward over time. (2) Turning to the counter-counter-offers made by the seller, the data show a downward trend similar to the one observed for the first offer. (3) With regard to the matches, the number of negotiations for the out-of-season good decreases during the trading day. The decrease is even more pronounced for the local in-season good (local tomatoes). In contrast, it increases for the end-of-season good (oranges) and the imported in-season good (imported tomatoes). (4) For each product, the sample mean and median of the seller's bargaining power index are greater than 0.5. In every situation, whenever the bargaining takes place, the transaction price is closer to the first offer of the seller than to the buyers' counter-offers. The index distributions for each product have mass points at 0, 0.5 and 1. Hence, in most of the transactions, seller and buyer meet either exactly half-way between their offers FO and CO or at one of the two.

5.2 Comparison of the two markets and empirical validation of the model

Once the model is built, it produces data that can be validated by comparison with the historical data. We do not compare the data sets directly but rather the stylized facts. Our question is whether the behavior of this ACE model is consistent with the empirical stylized facts of the real market. We (1) calibrate the model empirically i.e. we choose the values of the parameters ϕ and C and we (2) identify the phases in the evolution of the system that maximize the correspondence of the model's behavior with the real-world system. The values which seem to us the most appropriate are : for the out-of-season good (leeks), the values $\phi = 90$, $C = 0.99$ and the learning phase around the 50th (Phases 2-3). Indeed, for these values both the empirical markets and the artificial market present the same characteristics. During bargaining, the Sellers' offers vary. The bargaining power index has two mass points at 0 and 1 and decreases over the market day. The nature of matching is quite well-reproduced (see table 5).

[TAB. 5 about here.]

In this case, the intuition is that when out of season, a good does not have a 'normal price' and individuals learn about this current value by making their offers. They therefore make heterogeneous offer prices, and individuals buy either early and late but not in the intermediate round.

For the in-season goods (oranges, imported tomatoes and domestic tomatoes), the value $\phi = 90$ is not appropriate, since we see that this value involves very heterogeneous first offers by the Seller, while on the real market, the seller's offers are almost constant over different days. In these 3 cases, the frequency of matches without transaction is between 4.5% and 5.30%. For the three products, the calibrated value of C is $C = 1.44$ and the state of the market-days is similar to the one produced after 22 market-days. Indeed, the cases $C = 0.44$, $C = 0.74$, $C = 0.99$ and $C = 1.44$ before day 22 (Phase 1-2) seem to be inappropriate. The case $C = 0.44$ requires that all the matches are concluded with a transaction. The other cases require a too-high percentage of matches without transactions (around 50% for $C = 1.44$ before the day 22, around 40% for $C = 0.74$ and at least 15% for $C = 0.99$). As far as the values of ϕ and C are concerned, the difference between the three products involves the frequencies of bargaining.

For all the products, the empirical stylized facts are well reproduced in the phase 2, around the market day 30, because after it, agents start to bargain less and we do not produce the very high frequency of *Transaction With Bargaining*. So, the value of ϕ seems to differentiate the cases. ϕ is equal to 40 for the end-of-season good (oranges), 20 for the imported in-season good (imported tomatoes) and 1 for the local in-season good (domestic tomatoes). With these values, we are able to reproduce the characteristics of the seller's bargaining power, the frequencies of the nature of matches and, finally, the fact that the offers are almost stable.

In accordance with our interpretation, we could assume that Sellers take into account their experiences with their customers to make their offers when the good we analyze is in season. The more we move away from the heart of the season, the less they take into account their experiences. Simulations seem to demonstrate that Buyers spend more time searching in the case of an out-of-season good than in the case of an in-season good. Finally, in all cases, individuals seem to be neither at the beginning of the simulation nor in the long-term situation. They have already acquired a certain level of experience in selecting some better rules.

6 Conclusion

Our model appears to reproduce the main stylized facts of the real markets for the four goods under consideration. Due to the simultaneous learning of Buyers and Sellers, the emerging system is sometimes surprising and difficult to explain. A variation of only one parameter sometimes completely changes the market macro-structure.

The model was fitted empirically and reproduced most of the empirically-observed stylized facts. The differences between goods can be explained as follows : in the market for the out-of-season good, buyers can stay in the market longer, but learn to buy earlier during the market day. The value of the good is unknown to the buyers and the sellers try to impose a price by not updating their beliefs after each meeting. In the market for in-season-goods, agents bargain much more, even if the first offer does not change. These situations can be explained by the fact that individuals bargain to ensure that the price they obtain for buying earlier is better than the one they could obtain by waiting until the end of the market. Indeed, in the artificial market, agents start by buying late and bargaining, then they bargain less, then they buy earlier (in round 1) and bargain. The differences between the markets for different goods can be explained by the frequency of Sellers' learning, the length of time that Buyers can stay in the market and the stage reached in the acquisition of experience. The agents only apply systematic rules and learn by reinforcement learning using classifier systems. There is no circulation of information, apart from that of intermediate offers. In the model, we are interested in reproducing the facts observed in one seller's shop. Although we have shown that in all the simulations, the transaction prices converge to the Sellers' limit value, we observed differences in the emergence of bargaining, the number of matches and the offer prices. The convergence toward the Sellers' limit value is as strange in the artificial market as in the empirical market. Indeed, in the empirical market, wholesalers declared that when a transaction occurs with a buyer, they do not make a high margin in comparison with their limit value. This does not entail the disappearance of

sellers, but rather the emergence of a new source of supply through the system of sales in consignment.

In this paper, we have built a model drawn from the empirical agents in order to test whether the behavior they say they follow generates the same long-run macro structure as that highlighted in Kirman *et al.* (2005). A next step should be to test whether the results are robust to changes in the decision process and in the assumptions. Thus, we could test how the results change when we add discrimination among agents, or when we allow agents to exchange information.

Références

- Abel, A.B. 1985. Inventories, stocks out and production smoothing. *Review of economics studies*, **42**.
- Andreoni, J., & Miller, J.H. 1995. Auctions with artificial adaptative agents. *Games and economic behavior*, 39–64.
- Ashlock, D., Smucker, M.D., Stanley, E.A., & Tesfatsion, L. 1996. Preferential partner selection in an evolutionary study of prisoner's dilemma. *Biosystems*, **37**(1–2).
- Brenner, T. 2002. A behavioural learning approach to the dynamics of prices. *Computational economics*, February, 67–94.
- Brock, W.A., & Hommes, C. 1997. A rational route to randomness. *Econometrica*, **65**(5), 1059–1096.
- Diamond, P.A. 1984. *A search equilibrium approach to the micro foundation of macroeconomics*. MIT Press.
- Duffy, J. 2006. Agent-based models and human subject experiments. *Chap. 19, pages 949–1011 of*: Tesfatsion, Leigh, & Judd, Kenneth L. (eds), *Handbook of computational economics*. Handbook of Computational Economics, vol. 2. Elsevier. available at <http://ideas.repec.org/h/eee/hechp/2-19.html>.
- Duffy, J., & Uvner, U. 2007. Internet auctions with artificial adaptative agents. *Journal of economic behavior and organization*.
- Epstein, J.M., & Axtell, R.L. 1996. *Growing artificial societies : Social science from the bottom up*. Washington, D.C. : Cambridge, Massachusetts, and Washington D.C. : MIT Press.
- Holland, J.H. 1975. *Adaptation in natural and artificial systems*. Ann Arbor : University of Michigan Press.
- Hommes, C. 2006. Heterogeneous agent models in economics and finance. *In* : Tesfatsion, Leigh, & Judd, Kenneth .L. (eds), *Agent-based computational economics*, vol. 2. Elsevier Science B.V.

- Janssen, M.A., & Ahn, T.K. 2006. Learning, signaling, and social preferences in public-good games. *Ecology and society*, **11**(2), art. 21.
- Kirman, A., & Vriend, N. 2001. Evolving market : 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.
- LeBaron, B. 2006. *Post walrassian macoeconomics*. Cambridge University Press. Chap. Agent-based financial markets : Matching stylized facts with style, pages 231–238.
- Matos, N., Sierra, C., & Jennings, N.R. 1998. Determining successful negotiation strategies : an evolutionary approach. *Pages 182–189 of : Demazeau, Y. (ed), Proceedings of the 3rd international conference on multi-agent systems (icmas-98)*. Paris, France : IEEE Press.
- McMillan, J., & Rothschild, M. 1994. Search. *Pages 905–927 of : Aumann, R.J., & Hart, S. (eds), Handbook of game theory*, vol. 2.
- Mortensen, D.T. 1982. The matching process as a non-cooperative/bargaining game. *Pages 233–258 of : McCall, J.J. (ed), The economics of information and uncertainty*. Chicago, Illinois : University of Chicago press.
- Oliver, J.R. 1996. A machine learning approach to automated negotiation and prospects for electronic commerce. *Journal of management information systems*, **13**(3), 83–112.
- Riolo, R. 1997. *The effects of tag-mediated selection of partners in evolving populations playing the iterated prisoner's dilemma*. Santa Fe Institute Working Paper, 97-02-016.
- Rocha, A.P., & Oliveira, E. 2001. Agents advanced features for negotiation in electronic commerce and virtual organisations formation process. *Pages 78–97 of : Sierra, C. (ed), Agent mediated electronic commerce : the european agentlink perspective*, vol. volume 1991 of Lecture Notes in Artificial Intelligence. Springer.
- Rouchier, J. 2004 (May). *Interaction routines and selfish behaviours in an artificial market*. WEHIA, Workshop of Economics with Heterogenous Interacting Agents, Kyoto.
- Rouchier, J., & Mazaud, J.P. 2004. *Trade relation and moral link in transactions among retailers and wholesale sellers on the arnavaux market*. 11th World Congress for Social Economics.
- Rouchier, J., & Robin, S. 2006. Double-auction for real and artificial agents : Information perception and price dynamics. *Simulation and gaming*, **37**(2), 195–208.

Testfatsion, L. 2005 (July). *Agent-based computational modeling and macroeconomics*. Economic Report 05023, Iowa State University.

Weisbuch, G., Kirman, A., & Heirreiner, D. 2000. Market organisation and trading relationships. *The economic journal*, April.

Zeng, D., & Sycara, K. 1997. How can an agent learn to negotiate. *Pages 233–244 of : Intelligent agents iii : Agent theories, architectures, and languages*, vol. 1193 of Lecture Notes in Computer Science (LNCS). Springer.

Table des figures

1	Representation of the bargaining process : a decision-tree in the bargaining process	24
2	For all rounds 1 – 3 : Offer prices. Parameters : $C = 0.99, \phi = 1$. . .	25
3	Offers and counter-offers made by agents over market days : $C = 0.49$ and $\phi \in \{1, 20, 40, 90\}$: Sellers' offers (left-hand column) and Buyers' counter-offers (right-hand column)	26
4	$C = 0.99, \phi = 90$: For all rounds 1 – 3 : Counter-Offer prices. . . .	27
5	$C = 1.44$ and $\phi \in \{1, 20, 90\}$: Evolution of transaction prices over the market-days	28
6	Evolution of Seller's ex-post bargaining power index over market days	29

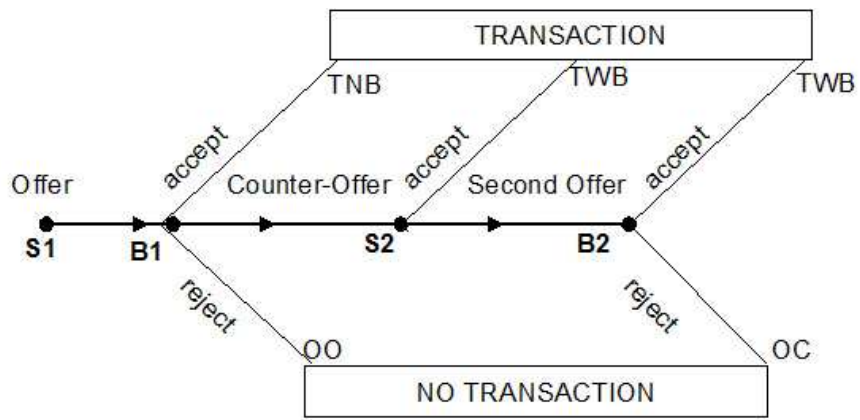


FIG. 1 – Representation of the bargaining process : a decision-tree in the bargaining process

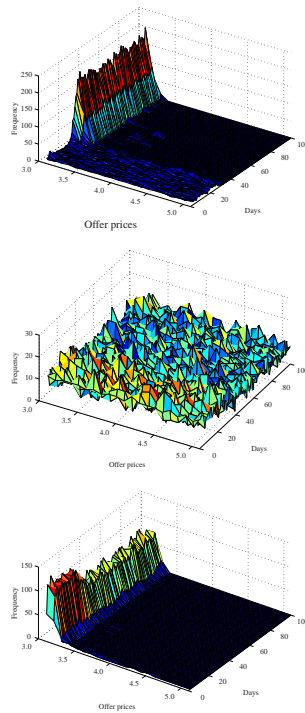


FIG. 2 – For all rounds 1 – 3 : Offer prices. Parameters : $C = 0.99$, $\phi = 1$.
 [The 3 graphs show the evolution of the first offer distribution over the market days. The first and the third graph show that after around 20 market days, a significant proportion of offers made in the first and the third round are lower than 3.5. The second graphs refers to the second round and show that the Sellers' first offers stay quite uniformly distributed along the learning. There is no dominant action]

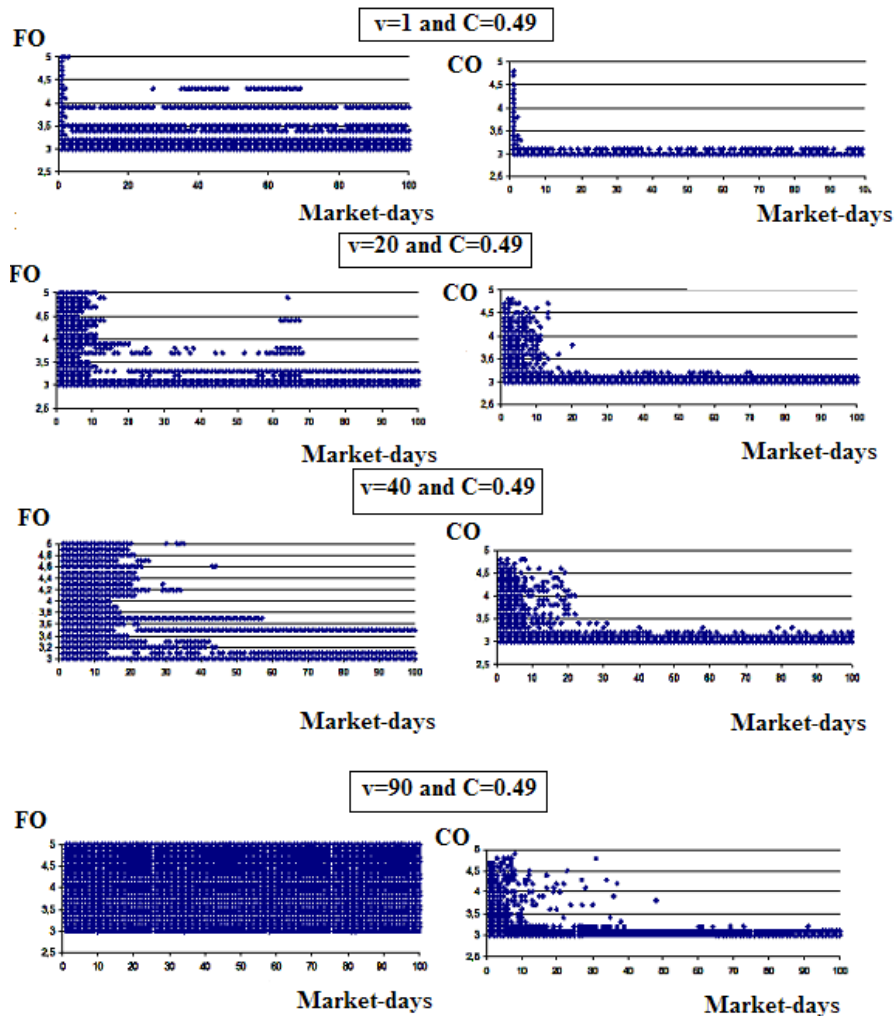


FIG. 3 – Offers and counter-offers made by agents over market days : $C = 0.49$ and $\phi \in \{1, 20, 40, 90\}$: Sellers' offers (left-hand column) and Buyers' counter-offers (right-hand column)

[The x-abcissa corresponds to market-days. The y-abcissa corresponds to the offers FO (on the left-hand column) and CO (on the right-hand column). The graphs on the same line refer to the same values of parameters. For example, if we consider the third line : The left-hand graph shows that the first offers are distributed between 3 and 5 during the 35 first market-days. They are then limited to values lower than 4. The right-hand graph shows that after 30 market-days, counter-offers stay lower than 3.5]

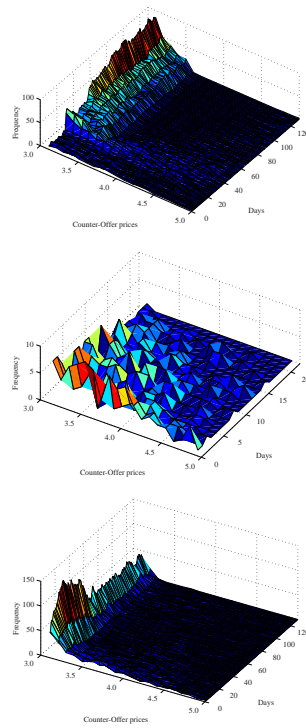


FIG. 4 – $C = 0.99$, $\phi = 90$: For all rounds 1 – 3 : Counter-Offer prices.

[The 3 graphs show the evolution of the counter-offers distribution over the market days. Each graph refers to a round. The first and the third graph show respectively that the first and the third round gathers two phases. For the first round, in the first phase (20 market days), Buyers make counter-offers between 3 and 5. Then, in the 80 following market days, if they make a counter-offer, the counter-offer is systematically lower than 3.5. The second graph refers to the second round but we cannot conclude because Buyers do not make enough counter-offers.]

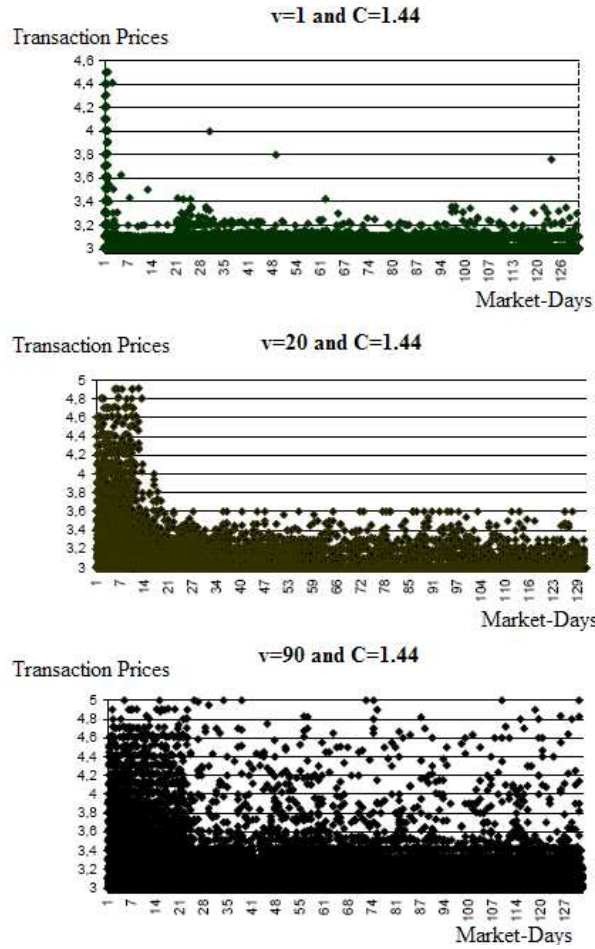


FIG. 5 – $C = 1.44$ and $\phi \in \{1, 20, 90\}$: Evolution of transaction prices over the market-days

[Each graph refers to a value of ϕ . The x-abcissa corresponds to market days and the y-abcissa corresponds to transaction prices. For example, graph 1 corresponds to the case $\phi = 1$. On the graph, we see that after less than 7 days, almost all transaction prices are lower than 3.2.]

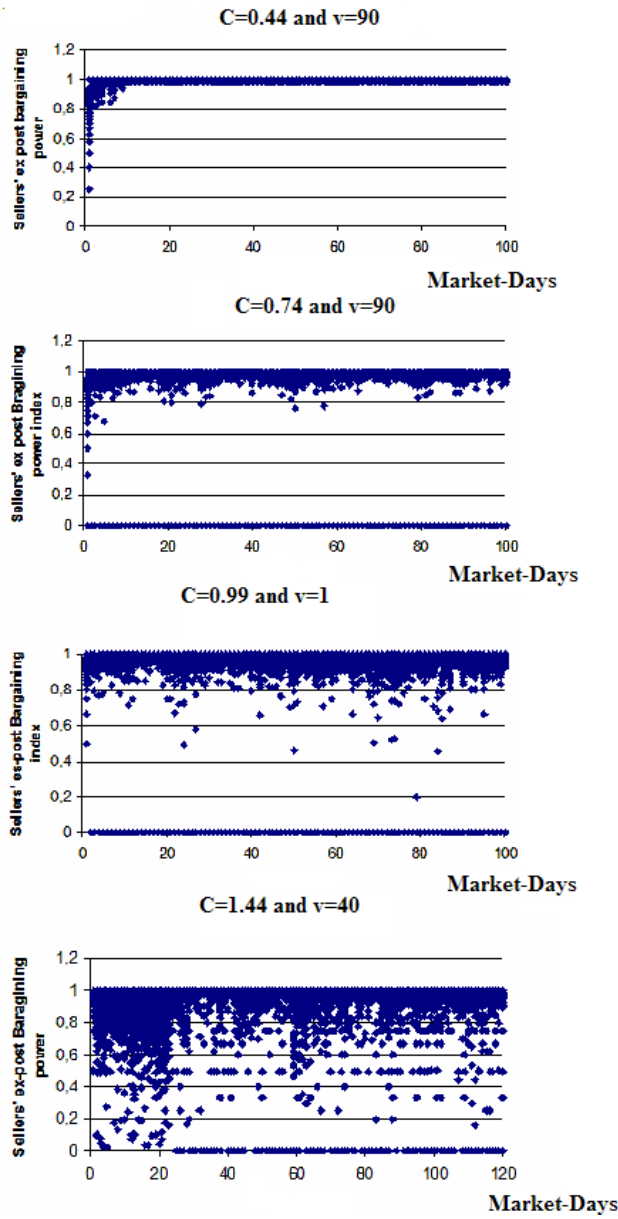


FIG. 6 – Evolution of Seller’s ex-post bargaining power index over market days [The x-abcissa corresponds to the market days and the y-abcissa corresponds to the Seller’s ex-post bargaining power index. Each graph corresponds to different values of C and ϕ . For example, the first graph corresponds to the values $C = 0.44$ and $v = 90$. It shows that after less than 10 days, the Seller’s bargaining power index converges toward 1.]

Liste des tableaux

1	Buyers : decision B_1	31
2	Buyers : decision B_2	32
3	General parameters	33
4	Impact of C and ϕ	34
5	Empirical and simulated nature of matches	35

Action	Expected Payoff
<accept FO >	$v^B - FO$
<reject and break off the bargaining >	$\frac{1}{y} \sum_{h=i+1}^{y+i+1} b^h$
<reject and make a counter-offer Co >	$\arg(\max_{x \leq FO} (b_x^i))$

TAB. 1 – The possible actions and their expected payoff for decision 2 : where FO denotes the first offer, i the current round, y the maximum number of Sellers that j can visit, b_m^i the bid associated with the rule < if it is the round i then offer m > and b^h the highest bid of all the rules with the conditional part < if it is the round i >

Action	Expected Payoff
< accept FO >	$v^B - FO$
< reject >	$\frac{1}{y} \sum_{h=i+1}^{y+i+1} b^h$

TAB. 2 – The actions and their expected payoff of the Buyer j concerning the decision 4 with b_m^i the bid associated with the rule < if it is round i then offer m >, \bar{b}^i the highest bid of all the rules with the conditional part < if it is round i >, i the current round, j the buyer and, y the maximum number of Sellers that j can visit

Notation	Parameter	values	interval
$\#B$	Number of Buyers	250	
$\#S$	Number of Sellers	5	
ϕ	Frequency of Sellers' learning		{1, 20, 40, 90}
v^S	Sellers' limit value	3	
v^B	Buyers' limit value	5	
C	The upper bound of the Buyers' time constraint		{0.49, 0.74, 0.99, 1.44}
c^B	Buyers' cost to bargain	0.1	
c^S	Buyers' cost to search	0.25	
δ	discount factor	0.95	
ζ	discount factor	0.95	
μ	measure of the Buyer's loyalty for the Seller with whom she makes a high profit	10	
ε	Noise used for the calculation of the bid	$\varepsilon \sim U_{[0,0.2]}$	

TAB. 3 – General parameters

Indicators	Impact of C	impact of ϕ
Evolution of offers	✓	✓
Counter-offers	✓	✓
Transaction prices		✓
Emergence of bargaining		✓
Rounds of transaction		✓
Ex-post bargaining power index (EBP)	✓	

TAB. 4 – Impact of time constraint (C) and frequency of Sellers' learning (ϕ)

Matches	Artificial market	Empirical market
OO	31%	19.73%
OC	14%	10.12%
TNB	11%	20.75%
TWB	42%	49.40%

TAB. 5 – Empirical and simulated data