JASSS

Marta Posada and Adolfo López-Paredes (2008)

How to Choose the Bidding Strategy in Continuous Double Auctions: Imitation Versus Take-The-Best Heuristics

Journal of Artificial Societies and Social Simulation vol. 11, no. 1 6
<http://jasss.soc.surrey.ac.uk/11/1/6.html>

For information about citing this article, click here

Received: 05-Feb-2007 Accepted: 25-Nov-2007 Published: 31-Jan-2008

Abstract

Human-subject market experiments have established in a wide variety of environments that the Continuous Double Auction (CDA) guarantees the maximum efficiency (100 percent) and the transaction prices converge quickly to the competitive equilibrium price. Since in human-subject experiments we can not control the agents' behaviour, one would like to know if these properties (quick price convergence and high market efficiency) hold for alternative agents' bidding strategies. We go a step farther: we substitute human agents by artificial agents to calibrate the agents' behaviour . In this paper we demonstrate that price convergence and allocative market efficiency in CDA markets depend on the proportion of the bidding strategies (Kaplan, Zero-Intelligence Plus, and GD) that agents have on both market sides. As a result, price convergence may not be achieved. The interesting question to ask is: can convergence be assured if the agents choose their bidding strategies? Since humans are frugal we explore two fast & frugal heuristics (imitation versus take-the-best) to choose one of three bidding strategies in order to answer this question. We find that the take-the-best choice performs much better than the imitation heuristic in the three market environments analyzed. Our experiment can be interpreted as a test to see whether an individual learning outperforms social learning or individual rationality (take-the-best) outperforms ecological rationality (imitation), for a given relevant institution (the CDA) in alternative environments.

Keywords:
Agent Based Models, Double Auction, Individual and Social Learning, Computational Organization, Bounded Rationality

Introduction

1.1

The Continuous Double Auction (CDA) is the dominant institution for the real-world trading of equities, CO2 emissions permits, derivatives, etc. In the CDA, buyers and sellers announce and accept bids and asks at any time. The information is held separately by many market participants (in the form of privately known reservation values and marginal costs). How is it that we observe quickly and accurately coordination through the trading process to reach the competitive equilibrium price? Does this coordination ability of the CDA relate to the exchange rules of the market and/or to the agents' behaviour? How does the agents' behaviour influence on the market dynamics? These questions have inspired a great deal of both experimental and artificial research as well as this paper.

1.2

Human traders have been used to examine some of these issues by experimental economists beginning with the seminal work of the Nobel Prize Vernon Smith (Smith 1962). Following Smith (1989), there are three dimensions that are essential in the design of any market experiment (I × E × A): the institution (I, it is both the exchange rules and the way the contracts are closed), the environment (E includes the agent endowments and values (reserve prices and marginal cost), resources, and knowledge), and the agents' behaviour (A). Alternative ontologies could be used to define the artificial model and the experiment. But there are many advantages from adopting this standard because it is common to all the works and results from experimental economics with humans. In terms of the triple I × E × A, one can describe our paper as an experiment to test if for a given I (CDA) and alternative Es, the agents' behaviour matters in terms of allocative efficiency and price convergence.

1.3

As Marks (2005) points out, possible criteria for judging how good a design auction market is might include: maximizing sellers revenue, maximizing market efficiency, discouraging collusion or predatory behaviour, discouraging entry-deterrent behaviour, or, as in financial markets, assuring that market prices reflect all valuable dated information and the true value of the company equities. We are interested in two design criteria: price convergence towards a stable equilibrium and market efficiency in allocative terms.

1.4

Some concepts such as efficiency are not univocally defined and accepted in the social science community, we make some comments about efficiency measures which will help the readers to situate the paper. Allocative efficiency is the criteria we use which is different from the efficiency measure in finance[1]. We deal with market efficiency as understood in microeconomics and industrial organization, and of course, following all the experimental economics research on market and policy design. Following Smith (1962) and the related literature in experimental economics, we define allocative efficiency as the total profit actually earned by all the traders divided by the maximum total profit that could have been earned by all the traders (i.e., the sum of producer and consumer surplus).

1.5

We choose the CDA because there is extensive experimental economics evidence on its properties. The main conclusion from the experimental economics approach is that the institution matters in terms of market efficiency and price convergence (Smith 1976). Market human-subject experiments have established for a wide variety of environments that the CDA guarantees the maximum efficiency (100 percent) and that the transaction prices converge quickly to the competitive equilibrium price (Holt 1995). They have also disclosed the features of the path of convergence to equilibrium. Although the human-subject experiments directly reveal the agents' aggregated behaviour, their individual decision rules and the impact of these rules on individual and aggregate performance are not directly revealed.

1.6

As Kirman and Vriend (2001) argued, if we want to understand the dynamics of interactive market processes and the emergent properties of the evolving market structures and outcomes, it might pay to analyze explicitly how agents interact with each other and how information spreads through the market. Muchnick and Solomon (2007) conclude with a desideratum of proposed extensions of their NatLab model: to compare the efficiency of different trading strategies, to isolate the influence of traders' bidding strategies on the market, to study the co-evolution of traders behaviour, to find ways to improve market efficiency and stability, and to study how people depart from rationality. Some of these questions have been widely described in Posada et al. (2005, 2006, 2007a, 2007b) and Posada (2006).

1.7

In this paper each soft-agent chooses endogenously his bidding strategy from a portfolio of three alternatives using a heuristic choice which has been called 'a frugal and fast' heuristic by Gigerenzer and Goldstein (1996). A heuristic is not good or bad nor rational or irrational *per se*. It is 'rational' to the degree that it can exploit environmental structures. We have explored two alternative heuristic choices: imitation and take-the-best. Imitation heuristic, as a proxy for social learning, is based on the other agents' past experience. Take-the-best heuristic, as a proxy for individual rational learning, is based on the own agent's individual reasoning from past experience.

1.8

In this paper we study whether a social heuristic (imitation) can compete in terms of allocative efficiency and price convergence with a rational, albeit frugal, take-the-best heuristic. And to what extent the results depend on different environments. Our goal is to analyze the effect that the agents' behaviour has on price convergence, market efficiency and individual surplus. Although the take-the-best heuristic requires less information (private information) than the imitation one, we have discovered through simulation that it works much better. Price convergence to the competitive equilibrium price, under the take-the-best heuristic, is always achieved. It provides higher market efficiency than the imitation choosing heuristic. In particular, we examine in heterogeneous populations the adverse effects on the results of the parasitic behaviour.

1.9

This work is motivated by Gode & Sunder's (1993) paper. Their experiment showed that for a symmetric environment market populated by Zero-Intelligent agents (the lowest possible individual learning skill), allocative efficiency was achieved in a CDA, thus indicating that efficiency was robust against individual learning, and it was an ecological property.

1.10

There is an ecological property (in our case in terms of allocative efficiency and price convergence) when social interaction leads to an outcome that is surprisingly different from or robust against individual behaviour. Adam Smith taught us that efforts of entrepreneurs for maximization of individual profit lead to minimization of profit for all. Thus imitation, the simplest heuristic for social knowledge, might turn out to be a good agent's heuristic in a CDA market and more generally in impersonal exchanges.

1.11

This idea has a long history in economics but it has been brought to a forefront because of the results from experimental economics first, and recently from artificial economics. Nevertheless the term ecological rationality was "rediscovered" and popularized in economics by Vernon Smith in his Nobel Prize lecture (Smith 2003): *Constructivism and ecological rationality in economics*. Accordingly we can define ecological rationality, related with social learning (Conte et al. 2001), as the match between the structure of a heuristic and the structure of an environment (see several chapters of Gigerenzer et al. 2001). Vriend (2000) gives a more technical issue of individual and social learning, but very related to our two agent's strategic choices (imitation versus take-the-best). For example, reinforcement learning as in Roth and Erev (1995) takes place at the individual level whereas replicator dynamics (e.g. Weibull 1995) is a form of social learning.

1.12

Our experiment can be interpreted as a test to see whether an individual learning outperforms social learning or individual rationality outperforms ecological rationality for a given relevant institution: the CDA.

1.13

The rest of this paper is organized as follows. In Section 2 we present the antecedents and motivations to analyse the dynamics market in (E × A) settings. In Section 3 we describe the agent-based model built to handle the two fast and frugal heuristics to choose the bidding strategy. In Section 4 we calibrate the model and describe the experiments. In Section 5 we analyze the emergent results in terms of price convergence and individual surplus when the agents have fixed bidding strategies, and in Section 6 we compare the results when the agents choose their bidding strategies using imitation versus take-the-best choice. In Section 7 we consider the main conclusions of the paper.

Market experiments with artificial agents in a CDA market.

2.1

To understand the dynamics of the CDA market, we must answer the following questions: *How much should the agents bid or ask? When should they place a bid or an ask? And when should they accept an outstanding order of some other trader?* But in Agent-Computational Economics approach there is a previous decision which is the core of this paper: *which bidding strategy should the agents choose to obtain higher profit?* This decision has usually been taken by the modeller and has not been an endogenous choice of the agents.

2.2

Different agents' behaviours have been proposed in previous models of homogeneous bidding behaviour to understand the interaction of the CDA exchange rules and individual behaviour (Easley and Ledyard 1993, Gode and Sunder 1993, Cliff and Bruten 1997, Preist and van Tol 1998, Gjerstad and Dickhaut 1998). The first shocking result from Agent-Computational Economics approach to CDA markets was that the Zero- Intelligence (ZI) agents developed by Gode and Sunder (1993) may lead to market efficiency near 100 percent. The ZI agents submit random orders between their own valuation and the best order in the market. This result proves that the CDA exchange rules are robust with respect to individual learning. It is in agreement with the eighteenth century classical philosophers, Adam Smith and David Hume, and later on with Hayek, that claimed that spontaneous order may be an outcome of the individual interactions. Although market efficiency is an ecological property, thus robust against individual learning, the price convergence depends on individual learning as Cliff and Bruten (1997) first demonstrated.

2.3

The interaction between different agents' behaviours has been analyzed in previous models of heterogeneous bidding behaviour to extend the understanding of the interaction of the CDA exchange rules and individual behaviour, including Tesauro and Das (2001), Das et al. (2001), Walsh et al. (2002), Chen and Tai (2003), Posada et al. (2005) and Posada (2006). In markets where ZI agents compete with intelligent agents like ZIP agents who use an elementary form of machine learning (Cliff and Bruten 1997) or GD agents who place orders which maximize the expected surplus (Gjerstad and Dickhaut 1998), the ZI agents perform poorly in terms of individual surplus. If the agents were not limited by a static analysis in which they have fixed bidding strategies, and they could choose the bidding strategy, agents will choose the ZIP or GD bidding strategy. This decision will be different if ZIP and/or GD agents compete in the market with Kaplan (K) parasitic agents, who only accept the orders submitted by other agents and never submit orders (Rust et al. 1993). The reason is that ZIP and GD agents achieve lower profits than the K parasitic agents. But if all the agents choose to be parasitic then no trade will take place (Chen 2000). The K parasitic agents must be parasites on the intelligent agents to trade and to obtain profit.

2.4

We simulate different scenarios with heterogeneous agents to answer the following questions: How many K agents will accept the market before it collapses? In case that the market collapses, is there a way to prevent it? To answer the first question, we extend previous works analyzing the whole bidding strategy space instead of the reduced analysis of two heterogeneous settings[2] (Tesauro and Das 2001). To answer the second question, we allow the soft-agents to choose between three alternative bidding strategies from the menu {ZIP, GD or K} using a fast and frugal heuristic, in an adaptive approach to rationality, instead of assuming fixed bidding strategies.

2.5

We explore two different heuristics from the tool-box (Gigerenzer and Selten 2001): imitation and take-the-best to study whether a social learning heuristic can compete in terms of allocative efficiency and price convergence with an individual learning heuristic.

2.6

We have considered the take-the-best heuristic to account for individual level learning. It is based on the own agent's individual reasoning from past experience as it happens in the CDA market laboratory experiments where human traders only know their own valuations. In take-the-best, cues that guide search have precomputed validities based on past experience. The heuristic picks up the cue with the highest validity and bases the decision on that, if possible. If not, it takes up the next cue. If it runs out of cues, it guesses randomly.

2.7

We have considered the imitation heuristic to account for social learning. It is based on other agents' past experience. The imitation of the successful behaviour has played an important role in the recent literature on learning and adaptive behaviour in economics, for example, in evolutionary economics (Nelson and Winter 1982) and in evolutionary game theory (Weibull 1995). It seems likely in situations with little information or understanding as it happens in CDA markets. In informal terms, imitation seems to understate the fact that markets can perform a good information gathering reflected in prices. It seems to be beyond human capacities, and to channel social (ecological) learning. One is tempted to dismiss this fact initially. However, such dismissal is not in accordance with social behaviour. The capacities to infer underlying information from the behaviour of others is so commonplace that we hardly notice the complex process that it is taking place. When searching for a restaurant in an unfamiliar motorway, people look to see where many others are eating. If there are people looking into a store window, those passing on the sidewalk will frequently stop to look also. These common examples illustrate the ability of humans to infer subtle information from systemic properties and suggest that such inferences might be common place not only among humans.

Design and modelling of the experiment: CDA-Environment-Agents' behaviour.

3.1

Our aim is the study of a market organized by a CDA and populated by artificial traders who are endowed with a learning heuristics which allows them to choose their bidding strategy from a portfolio of three alternatives (ZIP, GD, and K). We are particularly interested in the effects that the K parasitic behaviour has on the market results: market efficiency, price convergence, and individual surplus.

3.2

The Figure 1 summarises the structure of the model we have implemented for simulations[3]. We outline that this model follows the guidelines adopted from human-subject market experiments: institution (CDA), environment (supply and demand scenarios) and agents' behaviour (to choose the bidding strategy (ZIP, GD, K) using a frugal and fast heuristic -imitation or take-the-best-, to calculate an order, to submit an order, to accept an order).

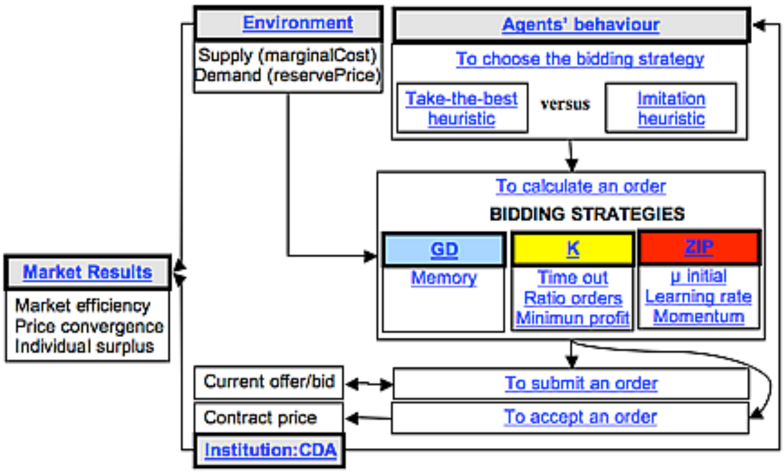

Figure 1. Model structure

The Institution: CDA

3.3

We consider a Persistent CDA with a bid-ask spread reduction (a new bid/ask has to provide better terms than previous outstanding-bid/outstanding-ask[4]). A trade takes place when a new ask (a) made is less than a pre-existing bid (bc) or when a new bid (b) made is greater than a pre-existing ask (ac). The pre-existing order is used to set the trade price. Otherwise, the order is appended to a queue of current open orders. Although the CDA protocol is well known, it is detailed in Table 1.

Table 1: Pseudo-code of CDA protocol (He et al. 2003)

```
1.{Initialize the initial bid/ask : ac=2*maximum reserve price, bc= 0
   when a period starts or a deal takes place}
2.{Several situations might arise during a round}
  Repeat
    {When a seller-agent submits an ask a}
      if a≥ac then a is an invalid ask (spread reduction rule)
      if bc<a<ac then ac is update to a (spread reduction rule)
      if a≤bc then
            this seller-agent makes a deal at bc
            goto 1
    {When a buyer-agent submits a bid b}
      if b≤bc then b is an invalid bid (spread reduction rule)
      if bc<b<ac then bc is update to b (spread reduction rule)
      if b≥ac then
            this buyer-agent makes a deal at ac
            goto 1
  until no bids/asks are submitted
```

The Environment

3.4

We consider that each trader plays either as seller or as buyer. The assumption of fixed roles conforms to extensive prior studies of the CDA related to this work, and cited in Section 2, including experiments involving human (Smith 1962, 1976) and artificial agents (Gode and Sunder 1993, Cliff and Bruten 1997, Preist and van Tol 1998, Gjerstad and Dickhaut 1998, Tesauro and Das 2001, Das et al. 2001, Walsh et al. 2002)[5].

3.5

Each agent is endowed with a finite number of units to trade. Seller i has n_i units to trade and he has a vector of marginal costs (MC_{i1}, MC_{i2},…, MC_{ini}) for the corresponding units. Here MC_{i1} is the marginal cost to seller *i* of the first unit, MC_{i2} is the cost of the second unit, and so on. Similarly, Buyer *j* has n_j units to trade and he has a vector of reserve prices (RP_{j1}, RP_{j2},…, RP_{jmj}) for the corresponding units. Here $RP_{j1}$is the reserve price to buyer *j* of the first unit, RP_{j2} is the reserve price of the second unit, and so on. The valuations are private, i.e., each trader only knows his own valuations.

3.6

This conceptualization of the environment enables to simulate different scenarios of market supply and demand (see Table 6).

Agents' Behaviour

3.7

The artificial traders play four decisions: (1) Which bidding strategy should he choose to obtain a higher profit? Once this decision has been taken, each agent deliberates the following questions: (2) How much should he bid or ask? (3) When should he place a bid or an ask? (4) When should he accept an outstanding order?

3.8

In the following subsections we describe widely the deliberative processes that agents use to play the market. The three bidding strategies we consider (ZIP, GD, and K), have been widely described in previous literature.

The bidding strategy choice: imitation versus take-the-best heuristics

3.9

In the beginning of each trading period, each agent chooses a bidding strategy from a portfolio of three alternatives (ZIP, GD, and K), with the exception of the two first trading periods where agent plays the same way. The agent can modify the bidding strategy by means of one of these heuristics: *imitation* and *take-the-best*.

a. **Imitation**: Each agent changes his bidding strategy imitating the best bidding strategy in the following way: *IF my profit is lower than the market average profit in the previous period THEN I will change my bidding strategy with a probability* ϵ. In case that the agent decides to change the bidding strategy, he imitates the bidding strategy with the highest average profit from the previous period (see Table 2 for pseudo-code). To take the decision each trader knows his own profit, the market average profit, and the average profit for each bidding strategy.

Table 2: Pseudo-code of the imitation heuristic

```
If profit (t) ≥ market average profit (t-1)
Then No change the bidding strategy
Else if randomNumber < changing probability
          Then No change the bidding strategy
          Else Imitate the bidding strategy with the highest average profit
```

The changing probability is calculated as the ratio of the difference between the market average profit and the agent's profit with respect to the market average profit.

b. **Take-the-best**: Each agent changes his bidding strategy looking for the best in the following way: *IF my profit in this period is not lower than it was in the previous one or my profit is greater than the minimum profit I expected THEN I don´t change my bidding strategy. Alternatively, I choose the bidding strategy which returns the highest profits based on past experience* (see Table 3 for pseudo-code). The information each trader knows is his reservation prices and his memory (the transaction prices he made in the past, and both the maximum and minimum transaction price in each period). The agent does not know others' bidding strategies or the profit achieved by other agents.

Table 3: Pseudo-code of the take-the-best heuristic

```
If profit (t) ≥ profit (t-1)
Then No change the bidding strategy
Else if profit ≠ 0
          Then if profit > profit belief
                    Then No change the bidding strategy
                    Else Choose the bidding strategy with the highest profit belief
          Else Choose randomly the bidding strategy
```

3.10

To form his beliefs, each agent compares if the orders from an alternative bidding strategy could have had lower or greater profits than the profits that have been obtained using the current bidding strategy in the following way (see Table 4 for pseudo-code):

1. If the bid (ask) from an alternative bidding strategy was greater (lower) than the realized bid (ask), it would has been accepted but he could have obtained lower profits.
2.
3. If the bid (ask) from an alternative bidding strategy was lower (higher) than the realized

bid (ask), but greater (lower) than the minimum (maximum) transaction price for that period, then the buyer (seller) considers that the bid (ask) would have been accepted and he could have obtained greater profits.

4. If the bid (ask) from an alternative bidding strategy was lower (higher) than the minimum (maximum) price of exchange of this period, the buyer (seller) assumes that no seller (buyer) would have accepted it.
5.
6. If the bid (ask) from an alternative bidding strategy was lower (higher) than the accepted ask (bid), the buyer (seller) would have rejected the ask (bid) with no profit.
7.
8. If the bid (ask) from an alternative bidding strategy was greater (lower) than the accepted ask (bid), he could have obtained the same profits whatever the value of the bid (ask) was.

Table 4: Pseudo-code to form agent's beliefs

BUYER's beliefs

```
If bid from an alternative bidding strategy < realized bid
Then if bid from an alternative bidding strategy > minimum transaction price for that period
      then greater profits
      Else profits =0
Else lower profits
If bid from an alternative bidding strategy < accepted ask
Then profits=0
Else same profits
```

SELLER's beliefs

```
If ask from an alternative bidding strategy > realized ask
Then if ask from an alternative bidding strategy < maximum transaction price for that period
      then greater profits
      Else profits =0
Else lower profits
If ask from an alternative bidding strategy > accepted bid
Then profits=0
Else same profits
```

How much bid or ask?

3.11

The **GD agent** (Gjerstad and Dickhaut 1998) has belief-learning skills. Each agent chooses the order that maximizes his expected surplus, defined as the product of the gain from trade and the probability for an order to be accepted. The GD agents use the history HM of the recent market activity to calculate a belief function q(x) estimating the probability (F or G) for an order at price x to be accepted.

$$\text{Buyers: } \mathit{max}\ F_x\ \mathit{(reserve\ price\text{-}price)}\ \text{with } \hat{q}(x) = \frac{ABL(x)+AL(x)}{ABL(x)+AL(x)+RBG(x)}, \quad (1)$$

$$\text{Seller: } \mathit{max}\ G_x\ \mathit{(price\text{-}marginal\ cost)}\ \text{with } \hat{q}(x) = \frac{AAG(x)+BG(x)}{AAG(x)+BG(x)+RAL(x)}, \quad (2)$$

where ABL is the number of accepted bids in HM with price lower than x, AL is the number of accepted and rejected asks with price lower than x, RBG is the number of rejected bids with price greater than x, AAG is the number of accepted asks in HM with price greater than x, BG is the number of accepted or rejected bids with price greater than x, and RAL is the number of rejected asks with price lower than x. Interpolation is used for prices at which no orders are registered in HM

3.12

The **K agent** (Rust et al. 1993) exhibits a parasitic behaviour. The basic idea behind this bidding strategy is: "wait in the background and let others negotiate and accept an order when it is interesting". So the buyers place bids equal to the current ask only when one of the following three conditions are met: (i) the fraction of time remaining in the period is less than a time out factor; (ii) the best ask is lower than the minimum trade price in the previous trade, and (iii) the best ask is lower than the maximum trade price in the previous period and the ratio of the bid-ask spread and the best ask is lower than a ratio order factor, and the expected profit is greater than a minimum profit factor.

3.13

Each **ZIP agent** (Cliff and Bruten 1997) has a mark-up μ that determines the price at which he is willing to buy or sell in an adaptative way. The agents learn to modify the profit margin using the information generated in the last trading cycle. The profit margin of a buyer is:

$$\text{Buyers: } \mu = 1 - \frac{order_{t-1} + \Delta t_t}{Reserve\,Price}, \quad (3)$$

Sellers: $\mu = 1 + \frac{order_{t-1} + \Delta t_t}{MarginalCost}$, (4)

where Δt is calculated using the individual trader's learning rate (β), the momentum learning coefficient (γ), and the difference between the bid target and the bid in the last round in the following way:

$$\Delta t_t = \gamma * \Delta t_{t-1} + (1-\gamma) * \beta * (target\ order - order_{t-1}), \quad (5)$$

where the target order is calculated following Preist and van Tol (1998), increasing the margin if the last order was accepted and decreasing it if the last order was rejected.

When to submit a bid or an ask?

3.14
When an agent is active at a given time step, he may submit an order (a new one or the replacement of an open order). The agents have a constant activation probability of 25 percent. Of course, the bid/ask must be also in agreement with spread reduction rule of the institution. So the K agents never submit a bid/ask in the market.

When to accept an outstanding order?

3.15
A buyer accepts the current ask if his bid (submitted or not) is equal to or lower than the current ask. Similarly, a seller accepts the current bid if his ask (submitted or not) is equal to or greater than the current bid.

Market Results

3.16
We analyze the market performance in terms of price convergence, market efficiency, and individual surplus. While the individual surplus is measured in each period step, the price convergence and the market efficiency are measured at the end of each period.

3.17
To measure the price convergence, we use the α convergence coefficient that was defined by Smith (1962) as the ratio of the standard deviation of exchange prices to the competitive equilibrium price. Hence it is a measure of the exchange price variation in relation to the competitive equilibrium price.

3.18
We also follow Smith (1962) to measure the market efficiency who defined the allocative efficiency as a measure of the actual total profit achieved in relation to the maximum total profit that can be achieved. To have a graphical idea of the profit achieved by agents with the same bidding strategy, we represent the individual profit achieved by each agent in each exchange using a colour code (ZIP: red, GD: blue, K: yellow).

3.19
The objective of the analysis is a full exploration of the bidding strategy space. A model with n agents, each one with 3 bidding strategies (ZIP, GD, and K) requires the computation of n2/2+3n/2+1 populations[6].

Parameters and scenarios of the simulation.

4.1
The model parameters are related to the decision: How much should an agent bid or ask. We have used the same parameters which have been proposed in previous works by Cliff and Bruten (1997) for ZIP agents and by Tesauro and Das (2001) for both GD agents and K agents (see Table 5). The ZIP agents' parameters were chosen following some trial-and-error experimentation. Later Cliff (2001) demonstrated that similar parameters values were obtained using genetic algorithm in parameter-optimization.

Table 5: List of parameters and their default values

ZIP parameters	*Value*
Initial profit margin (μ)	[0.05 , 0.35]

Learning rate coefficient (β)	[0.1 , 0.5]
momentum learning coefficient (γ)	[0 , 0.1]
GD parameter	
Memory	8
K parameters	
Minimum profit	[0.01 , 0.03]
Time out	[0.05 , 0.15]
Ratio orders	[0.0125 , 0.0375]

4.2

The number of agents is twenty (10 buyers and 10 sellers). In each experiment all the agents are either imitators or taker-the-best. We do this to isolate *ceteris paribus* the effect of each bidding behaviour on the price dynamics and on the market efficiency, which is the core of the paper.

4.3

We have simulated three environments (E1, E2, E3) and two choosing heuristics: take-the-best versus imitation (Table 6). Each run consists of a sequence of fifteen consecutive trading periods, each lasting 100 time steps.

4.4

To prevent the agents from having relative initial advantage, each trader has the same number of trading units (10 units each agent) and their valuations for each unit are equal for all the agents that are in the same market side.

4.5

In the symmetric environment (E1), the market demand and supply are drawn from the following individual demand function for the buyer *j*: $RP_{j\ k}=270-10q_k$, and from the following individual supply function for the seller *i*: $MC_{i\ k}=150+10q_k$, where k=1,…,10 units, i=1,…,10 sellers, and j=1,…,10 buyers. The corresponding competitive equilibrium price is P_0 = 210, and the quantity Q_0 =60. In this case, theoretically, the agent's maximum individual surplus (Smax) per period is 50, the agent's total individual surplus per period is 150, the buyers' surplus per period is 1500, the sellers' surplus per period is 1500, and the market surplus per period is 3000.

Table 6: Experiments for three environments (E) and two choosing heuristics (A)

	Take-the-best heuristic	Imitation heuristic	(P_0,Q_0)	Buyers' surplus //sellers' surplus
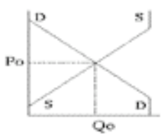 E1 Environment	SE1C	SE1I	(210, 60)	1500 // 1500
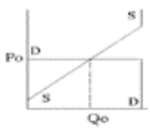 E2 Environment	SE2C	SE2I	(210, 60)	0 // 1500
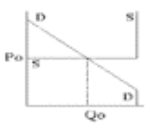 E3 Environment	SE3C	SE3I	(210, 60)	1500 // 0

4.6

Both E2 and E3 asymmetric environments correspond to a modification of E1: we introduce alternatively an elastic demand and an elastic supply at the competitive price. The total market surplus is the same in both cases and it is 1500 but lower than in the E1. In E2, theoretically, the buyers' surplus per period is 0, and the sellers' surplus per period is 1500. While in E3, theoretically, the buyers' surplus per period is 1500, and the sellers' surplus per period is 0. P_0 and Q_0 are identical in the three environments.

Price convergence when the agents have fixed bidding strategies.

5.1

The main issue of this paper is the analysis of the market performance, in terms of efficiency, price convergence, and individual surplus, when we endow each agent with a mechanism to choose endogenously the bidding strategy from a portfolio with the individual goal of getting higher profits.

5.2

Previous to such analysis, we show simulation results when the agents can not update their bidding strategies. Although this is a simple setting, it is a convenient starting point to get an idea of the different paths of convergence to the competitive equilibrium price and to validate our findings with previous computational works (Cliff and Bruten 1997, Gjerstad and Dickhaut 1998, Tesauro and Das 2001) and with experimental works (Smith 1976).

5.3

From these analyses we obtain that, in K homogeneous populations and in some heterogeneous populations, price convergence to the competitive equilibrium is not achieved, and market efficiency is lower than expected by evidence from human-subject experiments.

5.4

The main result that we obtain with markets populated by homogeneous populations (all GD agents or all ZIP agents) is to confirm that the path of convergence depends on individual learning (compare Figure 2.a with Figure 2.b, where the discontinue line is the competitive price and x-axis is the time steps). We obtain that the GD agents take less time than the ZIP agents both to learn and to trade at prices closely to the competitive equilibrium in any environment. The GD sellers' surplus is similar to GD buyers' surplus (compare Figure 2.c with Figure 2.e). When in one market side the agents' surplus is greater than the agent's maximum individual surplus per period (Smax), it is lower in the other market side.

E1 **(GD buyers, GD sellers)**	**E1** **(ZIP buyers, ZIP sellers)**

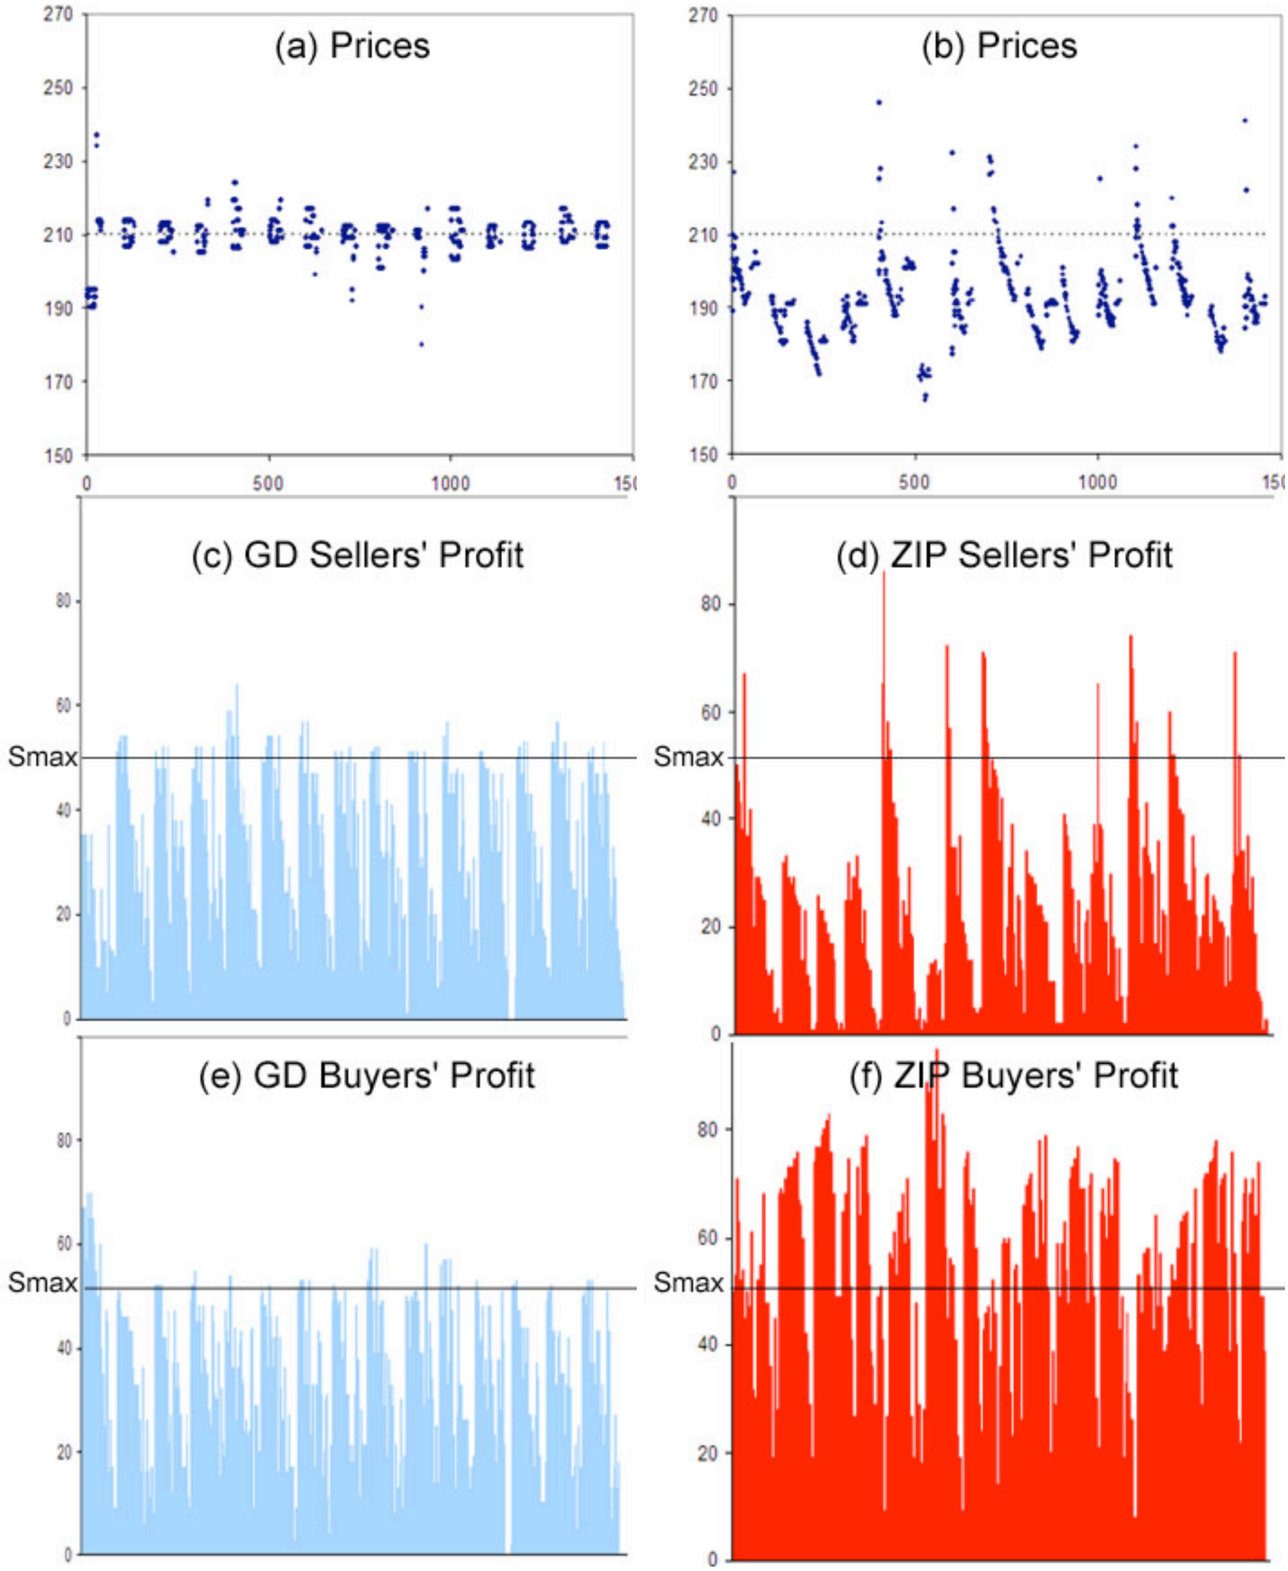

Figure 2. Path of price convergence and agents' surplus in a run of a market populated by a homogeneous population (all GD or all ZIP) in E1 environment

5.5

In heterogeneous populations we obtain that the K parasitic agents achieved higher profits than both the ZIP and the GD agents in all populations of the bidding strategy space. We also obtain that the presence of K parasitic agents has consequences on the price convergence. Price convergence depends, not only on the intersection of the supply and demand schedules, but also upon the agents' parasitic behaviour and their distribution between both sides of the market:

i. (i) The effect of the proportion of K agents: If the parasitic bidding strategy is used by less than 50 percent of the agents, there is price convergence to the competitive equilibrium price. The path of convergence is the path of the dominant bidding strategy.
ii.
iii. (ii) The effect of the location of K agents. Not only the proportion of parasitic agents in the population is relevant, but it is also important their distribution between both sides of the market. If the parasitic bidding strategy is used by 50 percent or more of the agents, price convergence is not achieved and the market collapses except in the singular situation when the K agents are equally distributed on both sides of the market.

5.6 We remark that E2 and E3 asymmetric environments are not mirror image because of the effect of the location of K agents. In Figure 3 we report the path of price convergence and the surplus of a population where buyers are K and sellers are GD in both E2 and E3 environments. When demand is perfectly elastic (E2), the K buyers' surplus should theoretically be zero (Smax=0 in Figure 3.e) because the K buyers' valuations are equal to the competitive equilibrium price. Surprisingly, it is not what we observe in simulations. The profits achieved by the K buyers (yellow bars in Figure 3.e) are greater than the profits of GD agents (blue bar in Figure 3.c) and the price contracts are far from the competitive equilibrium price (Figure 3.a). We observe the same surplus results when supply is perfectly elastic (E3) but now the contract prices are closely to the competitive equilibrium price (Figure 3.b). K buyers steal all the sellers' surplus (compare Figures 3.e and 3.f with Figures 3.c and 3.d), and accordingly, influence the price convergence. It does not depend on the elasticity of the market side they are.

5.7 We find equivalent results in experimental economics: contracts tend to be executed to the disadvantage of the market side having the price initiative (Smith 1962). As in experimental laboratory it is not possible to control the subject behaviour, the parasitic agents' behaviour is implemented by the rules of institution. The institutions in which one market side is not permitted to make an order are labelled offer auction (if buyers are not permitted to make bids) and bid auction (if sellers are not permitted to make asks).

5.8 Since agent's bidding strategy matters and the profit depends on the proportion and the location of agents' behaviour, in the next section we try to answer two questions: Would the change of bidding strategies assure the price convergence? Does a social heuristic learning process beat an individual one?

E2 (demand is perfectly elastic)
(GD sellers, K buyers)

E3(supply is perfectly elastic)
(GD sellers, K buyers)

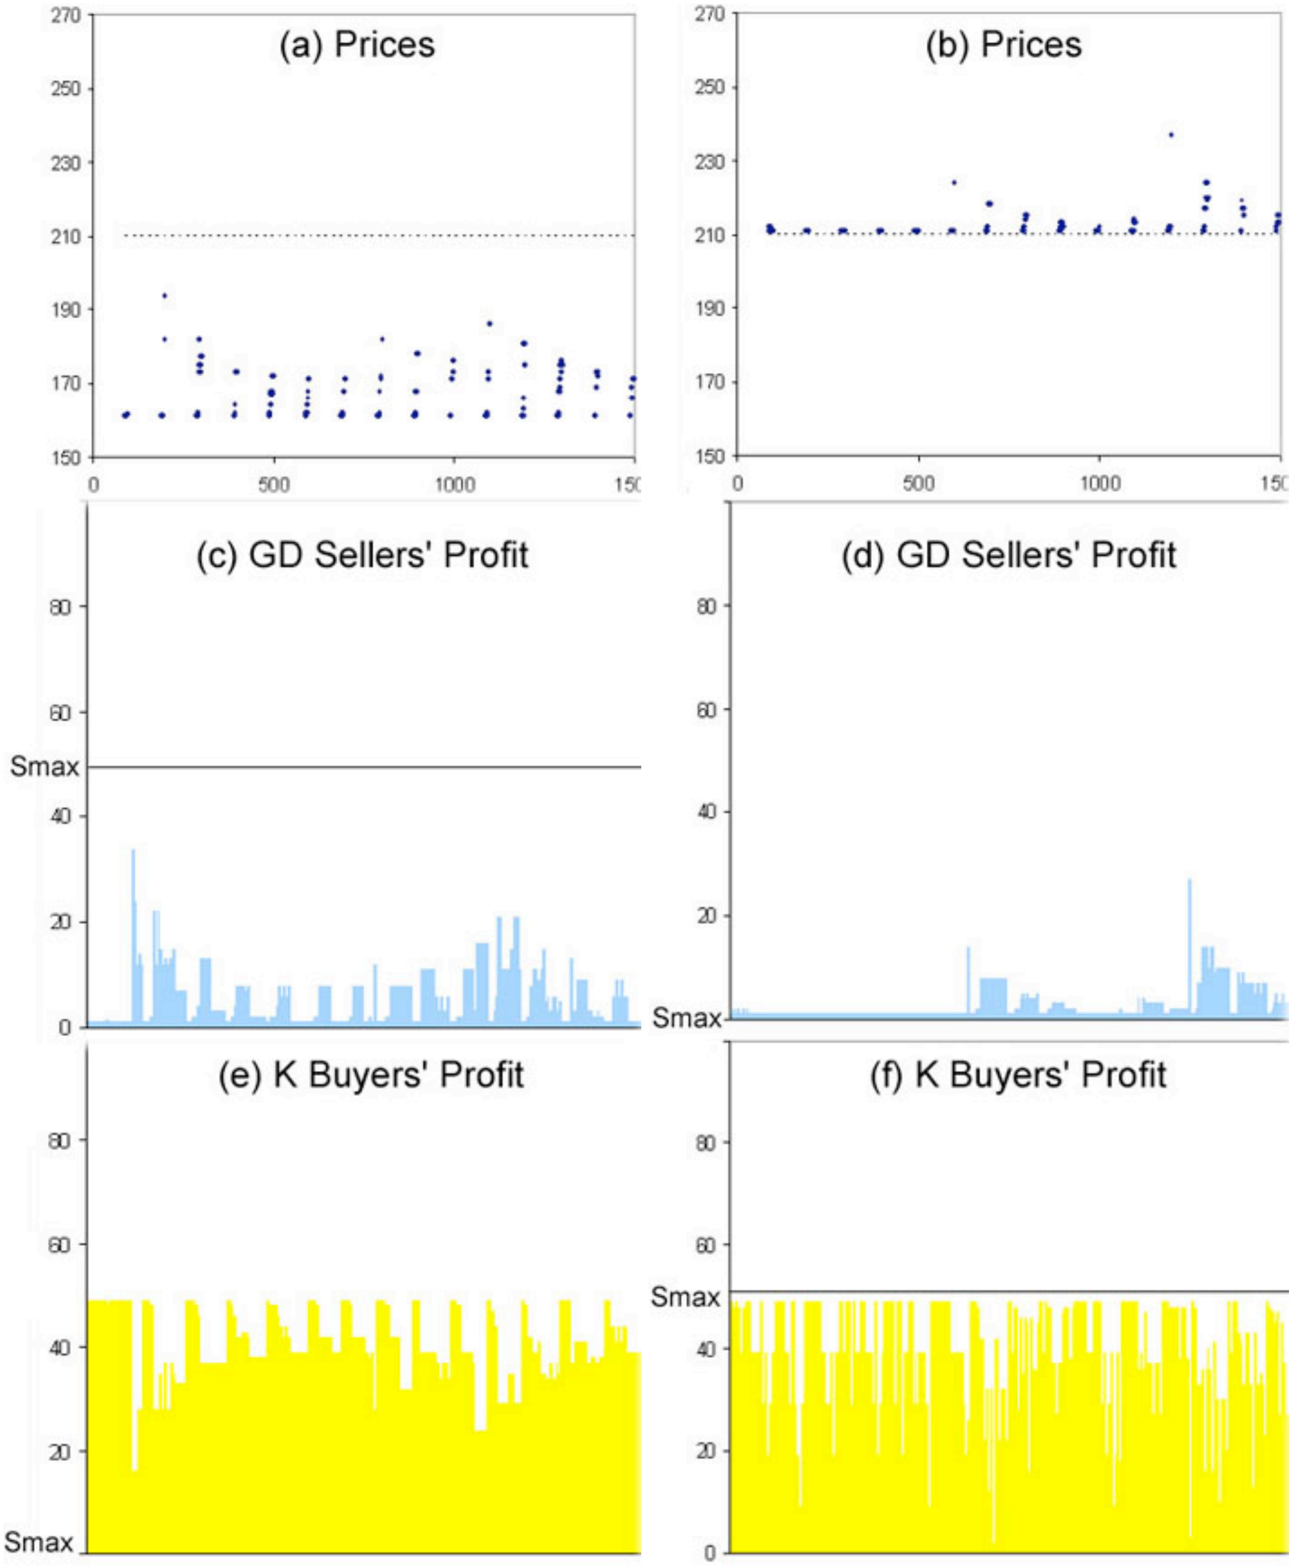

Figure 3. Path of price convergence and surplus in a run of a market populated by a heterogeneous population (10 GD sellers and 10 K buyers) in E2 and E3 asymmetric environments. K buyers steal all the sellers' surplus. It does not depend on the elasticity of the market side they are.

Simulation evidence of market performance: imitation versus take-the best heuristics

6.1

Price convergence can be improved if we allow the agents to choose their bidding strategy with one of the two alternative heuristics. The success of a heuristic is mostly explained by its ability to exploit the structure of the environment (ecological rationality). We find that although the take-the-best heuristic requires less information than the imitation one, it works better.

6.2

In Table 7 we show the α convergence coefficient computed for each trading period and for each heuristic as a measure of price convergence. In every scenario, the take-the-best heuristic beats the imitation one. We observe the α convergence coefficient decreases with time run indicating price convergence. After some periods, it is lower than 5 in E1 (first column: SE1C) and E2 environments (third column: SE2C), and it is always lower than 3 in E3 environment (fifth column: SE3C). While the α measure for imitation heuristic indicates that prices do not converge. It is higher than 5 in E1, E2, and E3 environments (second column: SE1I, fourth column: SE2I, and sixth column: SE3I).

Table 7: Average α convergence coefficient for 30 runs

	E1		E2 Perfectly elastic demand		E3 Perfectly elastic supply	
Period	**SE1C** Take-the-best heuristic	**SE1I** Imitation heuristic	**SE2C** Take-the-best heuristic	**SE2I** Imitation heuristic	**SE3C** Take-the-best heuristic	**SE3I** Imitation heuristic
1	22,8	23,0	23,0	23,2	0,8	1,0
2	17,1	18,2	20,2	19,9	2,6	3,1
3	13,7	15,6	16,5	17,7	2,7	3,5
4	12,2	14,2	13,6	16,5	2,6	4,8
5	9,4	10,8	9,5	13,9	2,2	4,4
6	7,0	6,7	7,9	10,6	1,8	4,6
7	5,9	6,1	6,5	7,7	1,8	5,7
8	4,8	7,2	5,3	7,6	1,6	7,8
9	4,8	8,0	4,4	8,6	2,7	8,4
10	4,4	6,9	3,6	10,9	1,9	9,7
11	4,5	8,0	2,8	13,3	1,9	9,3
12	3,6	6,6	2,0	14,5	1,4	10,3
13	2,7	7,0	1,6	12,5	2,1	10,8
14	3,9	8,1	1,8	10,6	1,6	9,7
15	2,2	5,7	1,8	7,5	2,2	9,3

6.3

In the following subsections we analyze in detail price dynamics, surplus, and market efficiency in the three environments: the E1 symmetric environment, the E2 asymmetric environment with a perfect elastic demand and the E3 asymmetric environment with a perfect elastic supply.

Symmetric environment E1

6.4

When the agents use the take-the-best heuristic to choose the bidding strategy, the price convergence is achieved (Figure 4.a). The α convergence coefficient is lower than 5 after some trading periods (see first column of Table 7: SE1C). There are K agents in both market sides. The surplus achieved by them in both market sides is similar (compare yellow colour in Figure 4.c and Figure 4.e) and it is similar to the agents' surplus achieved using other bidding strategy, GD or ZIP (compare yellow, blue, and red bars in Figure 4.c and Figure 4.e). No bidding strategy seems to dominate.

6.5

When the agents use the imitation heuristic to choose the bidding strategy, there is no evidence of price convergence (Figure 4.b). The α convergence coefficient does not settle down and it is higher than it is for the take-the-best heuristic (compare first and second column of Table 7: SE1C and SE1I). The imitation heuristic induces in the three environments that agents adopt the K bidding strategy. The traders imitate the parasitic K bidding strategy which is the bidding strategy with the highest profit. The market is mainly populated by K parasitic agents. The surplus achieved by them is similar in both market sides (compare yellow colour in Figure 4.d and Figure 4.f).

Take-the-best heuristic SE1C	**Imitation heuristic SE1I**

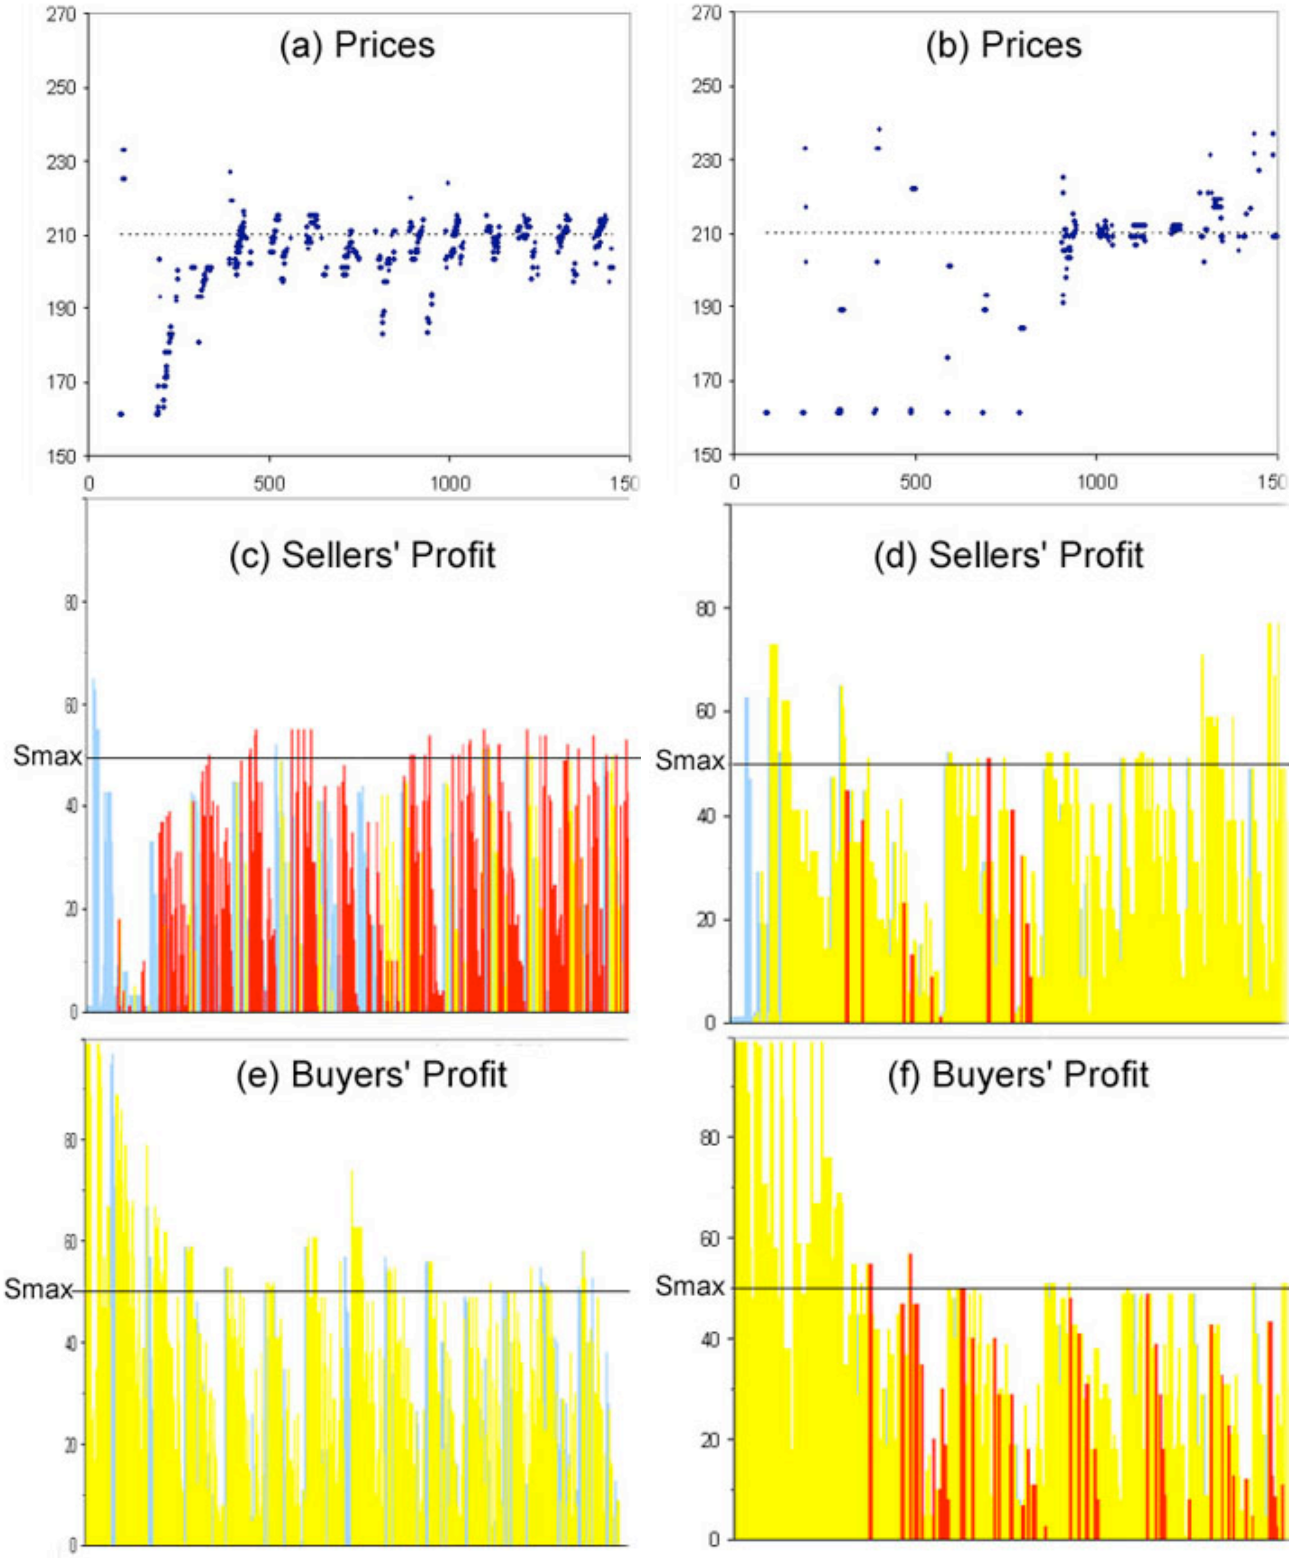

Figure 4. Path of price convergence and agents' surplus in a run of a market populated by an initial population of 10K buyers and 10 GD sellers in E1 environment. In E1, demand and supply are symmetric. There is no evidence of price convergence when agents use "imitation" to update their bidding strategies. On the other hand, take-the-best leads to price convergence and the maximum surplus of both buyers and sellers.

Asymmetric environments E2

6.6

In asymmetric environments the prices converge from the market side with lower elasticity. In the E2 environment, where the demand is perfectly elastic, the prices converge from below. We find equivalent results on CDA market human-subject experiments (Smith 1962).

6.7

Except for this difference in the path of price convergence, the market performance under the E2 asymmetric environment with perfect elastic demand (Figure 5, SE2I and SE2C) is similar to the market performance under the E1 symmetric environment (Figure 4, SE1I and SE1C) for the two heuristics. The price convergence is achieved when the agents use the take-the-best heuristic to choose the bidding strategy but not when the agents use the imitation one. If we compare the α convergence coefficient in both cases, after some trading periods, it is lower than 5 for the take-the-best heuristic (see third column of Table 7: SE2C) but higher than 7 for the imitation one (see fourth column of Table 7:SE2I).

6.8

When the agents use the take-the-best heuristic to choose the bidding strategy, the K bidding strategy is not the dominant. The K parasitic agents are mainly located in the market side with lower elasticity (in E2: the supply side) while GD agents are mainly located in the market side with higher elasticity (in E2: the demand side). The buyers' surplus is not zero but it is less than the sellers' surplus (compare Figure 5.e with Figure 3.e).

6.9

When the agents use the imitation one, again, the market is mainly populated by K parasitic agents because the parasitic K strategy is the bidding strategy with the highest profit. As a result, the transactions volume decreases.

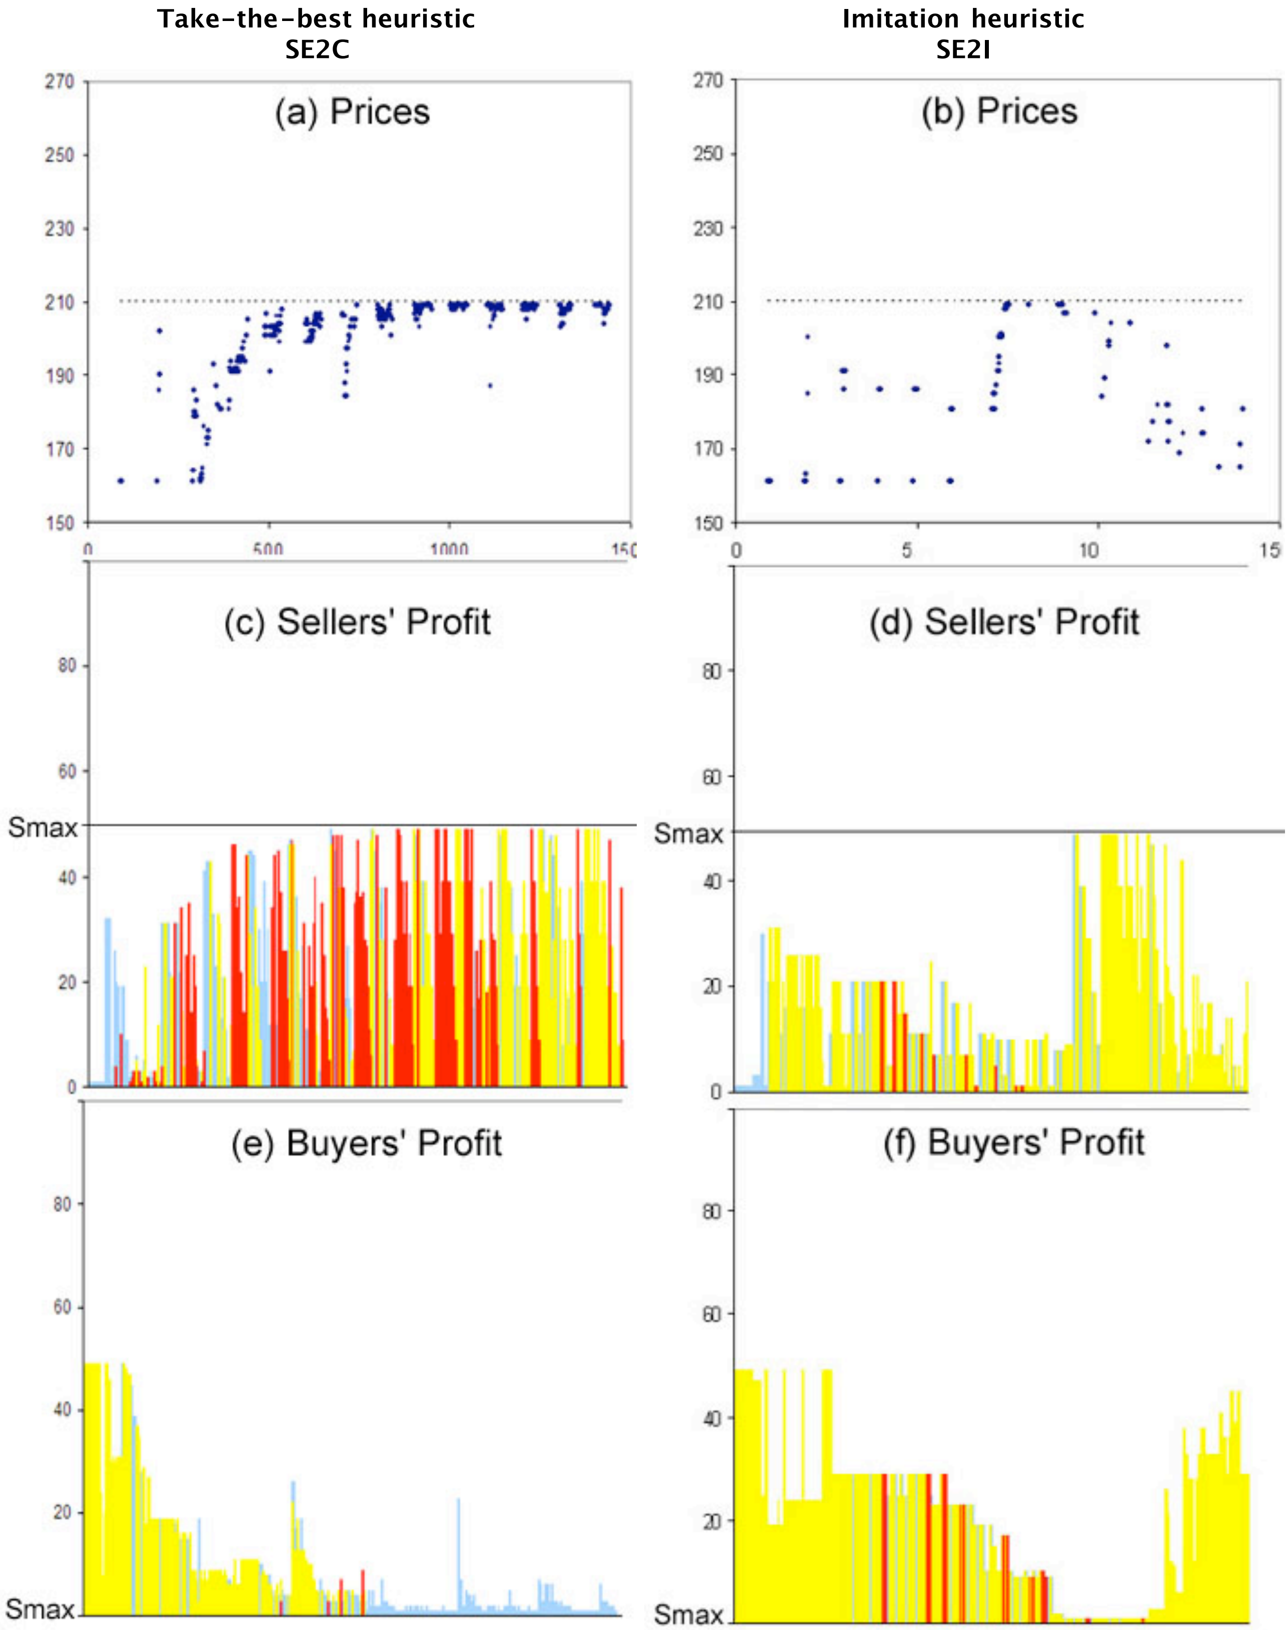

Figure 5. Path of price convergence and agents' surplus in a run of a market populated by an initial

population of 10 K buyers and 10 GD sellers in E2 environment. After some periods, the agents change their bidding strategy. Sellers adopt mainly the K or ZIP bidding strategy to improve their profits (5.c). As many sellers are parasite, the K buyers' profits decrease and buyers, who use the take-the-best heuristic, choose the GD bidding strategy to improve their profits (5.e). As a result, price convergence is achieved when the demand is perfectly elastic (5.a). When agents use the imitation heuristic, price convergence is not achieved because many buyers and sellers are parasite.

Asymmetric environments E3

6.10
The market performance under the E3 asymmetric environment with perfect elastic supply (Figure 6, SE3I and SE3C) is quite different from the market performance under the E1 symmetric environment (Figure 4, SE1I and SE1C).

6.11
In E3 asymmetric environments the prices converge from the market side with lower elasticity as it also happens in E2 asymmetric environments. But now the prices converge from above because the demand is less elastic than the supply

6.12
When the agents use the take-the-best heuristic to choose the bidding strategy, convergence is achieved (Figure 6.a). The α convergence coefficient is always lower than 3 (see fifth column of Table 7: SE3C). As it happens in E2 environment, the K parasitic agents are mainly located in the market side with lower elasticity and GD agents are mainly located in the market side with higher elasticity. But in E3 the market side with lower elasticity is demand side. The sellers' surplus is not zero (Figure 6.c) but it is less than the buyers' surplus (Figure 6.e).

6.13
When the agents use the imitation heuristic to choose the bidding strategy, the market performance gets much worse. Price convergence is not achieved (Figure 6.b). The α convergence coefficient increases (high price volatility) and sometimes the market collapses. The reason is that the imitation heuristic has a bias towards the K parasitic behaviour. This triggers herding and the market collapses.

Take-the-best heuristic SE3C	Imitation heuristic SE3I

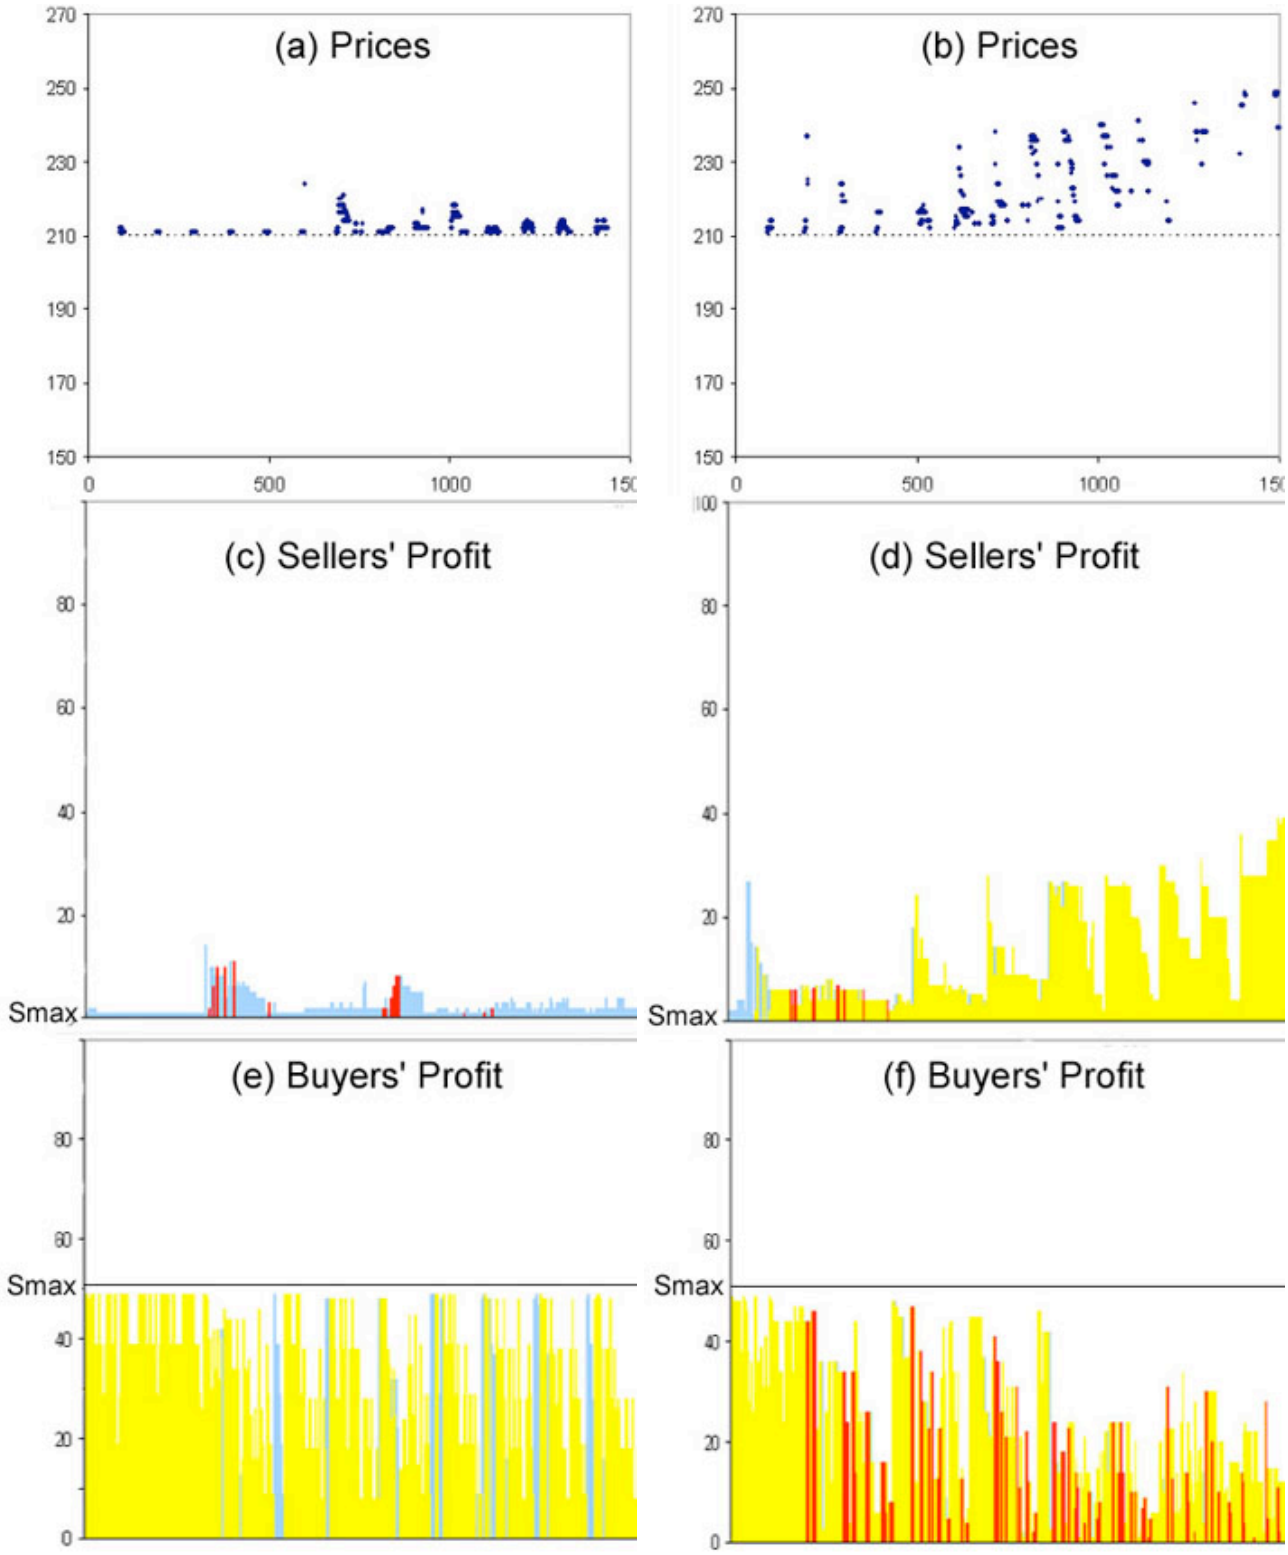

Figure 6. Path of price convergence and agents' surplus in a run of markets populated by an initial population of 10 K buyers and 10 GD sellers in E3 environment (as it was in previous analysis explained in section 6.2). Agents' surplus is close to the maximum surplus (6.c and 6.e). As result there are no changes in bidding strategies when agents learn with the take-the-best. But, when agents change their bidding strategies by imitation, the parasitic K bidding strategy emerges as the main strategy for sellers and buyer. As result, price convergence is not allowed, prices are higher than P0, and the maximum surplus is obtained by sellers (where agents adopt the parasitic behaviour).

Market efficiency

6.14

The market efficiency is an ecological property, thus it is robust against individual learning. One expects that the market efficiency was near 100 percent in all the populations of the bidding strategy space. But it is not so. These results can be improved if we allow the agents to choose the bidding strategy and they use the take-the-best heuristic to choose it.

6.15

Previous to the market performance analysis (in terms of allocative efficency) when agents choose their bidding strategy, we show simulation results when agents can not update their bidding strategy. It is a convenient start point to get an idea of the market efficiency values in all bidding strategy space.

6.16

We report the market efficiency values using a 3D figure where the z-axis is the market efficiency and both x-axis and y-axis are the bidding strategy space (see Figure 7). We represent the bidding strategy space by a two dimensional simplex grid with vertices corresponding to the pure bidding strategies (all ZIP: point a in Figure 7, all GD: point d, or all K: point f). We use a colour code to have a graphical idea of the regions where one bidding strategy is the most popular strategy: ZIP in the red region, GD in the blue one, and K in the yellow one.

6.17

In the regions where most of the agents are ZIP agents (red region of the Figure 7) or GD agents (blue region), market efficiency is near 100 percent and the volatility is very low. In the regions where more than 50 percent of the agents are parasites (yellow region), the market efficiency goes down and it tends to zero when the proportion of parasites tends to 100 percent. This topography is similar in the three alternative environments (E1, E2, and E3) but the volatility under asymmetric environments (E2 and E3) is higher than it is in the symmetric environment (E1). The topography of the market efficiency values has not changed substantially after fifteen trading periods (see Figure 7 for the topography of the average market efficiency[7] in the E1 symmetric environment, see movie).

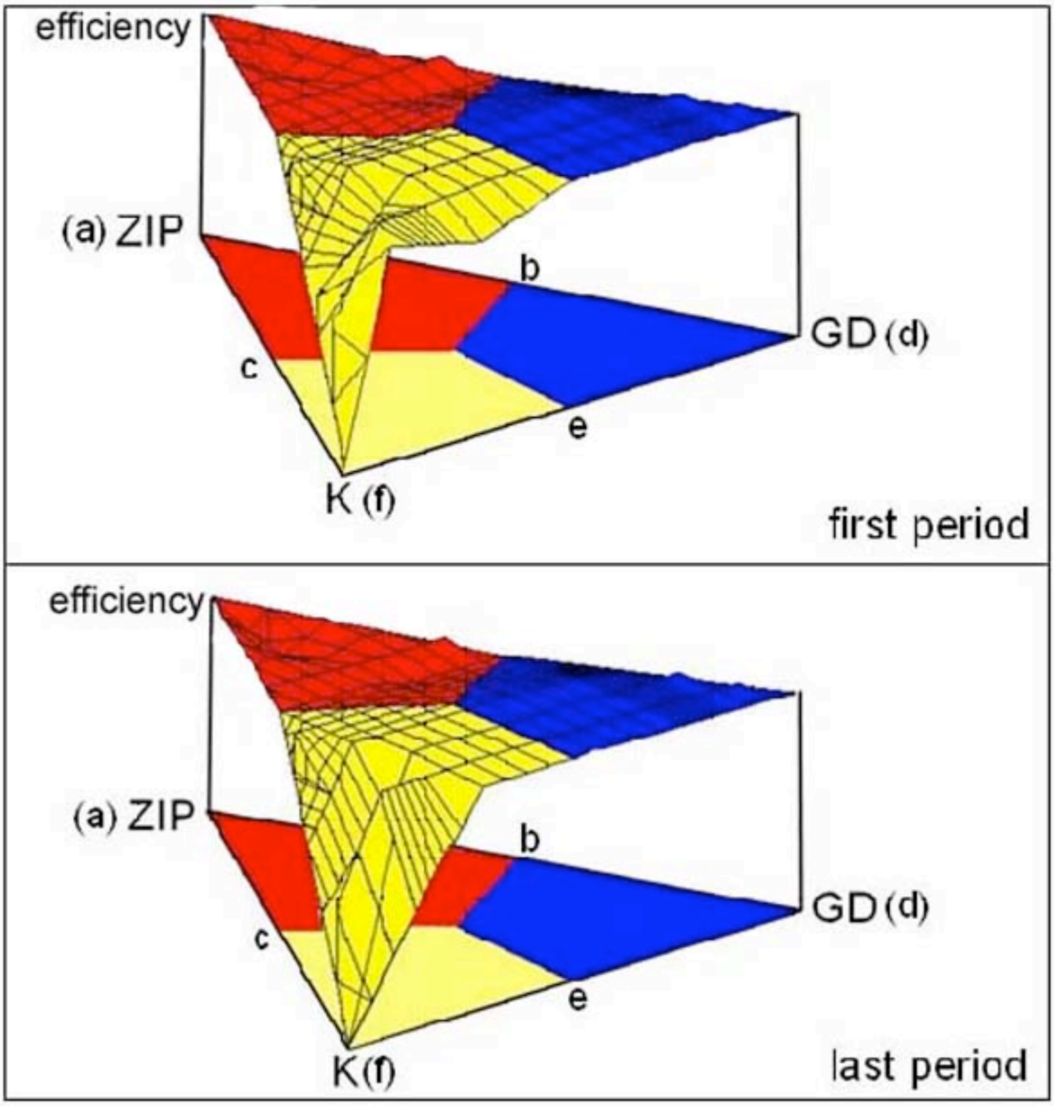

Figure 7. Market efficiency in the bidding strategy space at the first and at the last trading period when the agents have fixed bidding strategies in the E1 symmetric environment

6.18

When the agents use the take-the-best heuristic to choose the bidding strategy, they move to populations that are not in the yellow region for any environment and for any initial population. As a result, the market efficiency increases.

6.19

Nevertheless when the agents use the imitation heuristic to choose the bidding strategy, they move to populations that are in the yellow region for any environments and for any initial population because the agents imitate the parasitic bidding strategy which is the bidding strategy with the highest profit. Therefore, the market efficiency decreases.

Conclusions

7.1

It is conventional wisdom that in complex environments imitative behaviour is more likely to occur. Our results indicate that is not the "rational" choice for the CDA market. We conclude that when the agents with bounded rationality confront a CDA market populated by agents who choose their bidding strategies from a set of the three alternatives (Kaplan or parasitic, Zero intelligent Plus, and GD), the market performance is better when the agents use the take-the-best heuristic to make this choice than when they use the imitation one, as naively one could expect. These results are in agreement with human-subjects experiments by Bosh-Domenech and Vriend (2003) and economic theory research by Dixon et al. (1995). Although market information leads prices to convergence, there is room for anticipation, forward looking and strategic behaviour of the agents. This solves the perfect competition paradox. Price taking behaviour does not imply absence of strategic behaviour that affects both the agents and the functioning of the market.

7.2

Furthermore we conclude that the CDA market performance improves when the agents choose their bidding strategy with the take-the-best (frugal) choice.

7.3

We obtain that, when the market is populated by agents with fixed bidding strategies (Kaplan or parasitic, Zero Intelligent Plus, or GD), price dynamics, market efficiency, and agents' surplus depend on the proportion of the bidding strategies and their distribution on both market sides. In some cases the price convergence is not achieved and the market collapses. We can gear the market to achieve price convergence and allocative efficiency if the agents choose their bidding strategy with an appropriate heuristic choice.

7.4

Since human agents are bounded rational it is reasonable to explore market performance under frugal heuristics. We have used two frugal heuristics choice: imitation and take-the-best. The first one demands more information and has a social learning nature. The second one is individually based and demands more intelligence. The take-the-best heuristic outperforms the imitation one.

7.5

From a methodological point of view, we claim that in order to compare results from artificial agents simulations with the evidence provided by human-subject experiments it is very convenient to build the agent-based market model with a clear specification of the triplet: $I \times E \times A$. (I: institution, E: environment, A: agents' behaviour).

7.6

Human-subject market experiments have established that, for a wide variety of environments, a CDA market guarantees the maximum efficiency (100 percent) and that the transaction prices converge quickly to the competitive equilibrium price. We have demonstrated that they depend mainly on the way agents learn to bid. In particular, our simulations show how agents can modify the conventional patterns of the market. It is due to the choice of the Kaplan parasitic bidding strategy by agents in the supply or the demand side of the market. Thus it seems that the institution matters and so does individual and private behaviour.

Acknowledgements

We are grateful to two anonymous referees for their comments. We are indebted to Professor Cesáreo Hernández for many useful discussions all along the paper. Usual disclaims apply. We acknowledge financial support from the Spanish MEC project nº DPI-2005-05676

Notes

[1]The concept of efficiency is central to financial markets. See Dimson and Mussavian (1998) for a brief introduction to efficiency in financial markets.

[2]In one setting a single agent of one type competes against an homogeneous populations of a different type. In the other setting two types of agents compete in a market where one has a counterpart of the other type.

[3]The model has been programmed in SDML and the source code is available at http://www.insisoc.org/take_the_best/. SDML is a strictly declarative model language freely available from the Centre for Policy Modelling (http: //cfpm.org)

[4]The initial current bid is zero and the initial current ask is the double of the maximum reserve price.

[5]Artificial financial markets that adopt the CDA as institution uses a different approach: agents are endowed with a set of stocks and/or money, and agents can buy/sell at their choice. Our approach is more close to Experimental Economics, where agents behave as buyer or seller. It allows us to validate the results of simulations with real empirical evidence from market experimentation with humans.

[6]The number of combinations with repetition is (k+n-1)!/(n! (k-1)!) where k is the number of bidding strategies from which choose and n is the number to be chosen (i.e., the number of agents).

[7]It is not possible to report in the paper these values for all populations of the bidding strategy space. They are available in http://www.insisoc.org/take_the_best/.

References

BOSH-DOMENECH, A., Vriend, N. (2003) Imitation of successful behaviour in Cournot markets in *The Economic Journal* 113. pp. 495-524.

CHEN, S.-H. (2000) Toward an agent-based computational modelling of bargaining strategies in double auction markets with genetic programming in *Lecture Notes in Computer Science* 1983. pp. 517-531.

CHEN, S.-H., Tai, C.-C. (2003) Trading restrictions, price dynamics and allocative efficiency in double auction markets: analysis based on agent-based modelling and simulations in *Advances in Complex Systems* 6 (3). pp. 283-302

CLIFF, D. (2001) Evolutionary optimization of parameters sets for adaptive software-agent traders in continuous double auction in *Tech. Rep of HP Laboratories* 01-99, Hewlett Packard Laboratories, Bristol, England

CLIFF, D., Bruten, J. (1997) Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets in *Tech. Rep of HP Laboratories* 97-141.

CONTE, R., Paolucci, M. (2001) Intelligent social learning in *Journal of Artificial Societies and Social Simulation*, 4 (1) < http://jasss.soc.surrey.ac.uk/4/1/3.html>

DAS, R., Hanson, J., Kephart, J., Tesauro, G. (2001) Agent-human interactions in the continuous double auction in *Proceedings of the International Joint Conferences on Artificial Intelligence*. pp. 1169-1187.

DIMSON, E., Mussavian, M. (1998) A brief history of market efficiency in *European Financial Management* 4 (1). pp. 91-193.

DIXON, H., Wallis, S., Moss, S. (1995). Axelrod meets Cournot: oligopoly and the evolutionary metaphor, part I. *Discussion Papers in Economics, nº 95/8, Un. Of York. Later on (2002) same title in Computational Economics* 20. pp. 139-156.

EASLEY, D., Ledyard, J. (1992) Theories of Price Formation and Exchange in Double Oral Auctions in Friedman and Rust (eds.), *The double auction markets: Institutions, theories and evidence*. pp. 63-98, Addison-Wesley.

GIGERENZER, G. Goldstein, D.G. (1996) Reasoning the fast and frugal way: Models of bounded rationality in *Psychological Review* 103(4). pp. 650-669.

GIGERENZER, G., Selten, R. (ed.) (2001) *Bounded Rationality: The Adaptive Tool-box*, MIT Press, Cambridge, Mass.

GJERSTAD, S., Dickhaut, J. (1998) Price formation in double auctions in *Games and Economic Behaviour* 22. pp. 1-29.

GODE, D., Sunder, S. (1993) Allocative efficiency of market with zero-intelligent traders: Market as a partial substitute for individual rationality in *Journal of Political Economy* 101. pp. 119-137.

HOLT, C. A. (1995) Industrial Organization: A Survey of Laboratory Research in: J.H. Kagel and A. Roth eds *The Handbook of Experimental Economics*. Princeton University Press. pp. 349-443.

KIRMAN, A., Vriend, N. (2001) Evolving market structure: an ACE model of price dispersion and loyalty in *Journal of Economic Dynamics & Control* 25. pp. 459-502.

MARKS, R (2005) Market Design Using Agent-Based Models in L. Tesfastion and K. Judd eds *Handbook of Computational Economics* 2. North Holland.

MUCHNICK, L., Solomon, S. (2007) Markov Nets and the NatLab platform; Application to Continuous Double Auction in Salzano and Colander (eds.), *Complexity Hints for Economic Policy*, Springer http://shum.huji.ac.il/~sorin/ccs/isi/Markov-Nets.pdf

NELSON, R.R., Winter, S.G. (1982) *An Evolutionary Theory of Economic Change*, Cambridge, MA: Harvard University Press.

POSADA, M., Hernández, C., López, A. (2005) Learning in a Continuous Double Auction Market in *Lecture Notes in Economics and Mathematical Systems* 564. pp. 41-51.

POSADA, M., Hernández, C., López, A. (2006) Strategic Behaviour in a Continuous Double Auction Market in *Lecture Notes in Economics and Mathematical Systems* 584. pp. 31-43.

POSADA, M. (2006) Strategic Software Agents in Continuous Double Auction under Dynamic Environments in *Lecture Notes in Computer Science* 4224. pp. 1223- 1233.

POSADA, M., Hernández, C., López, A. (2007a) An Artificial Economics View of the Walrasian and Marshallian Stability in *Lecture Notes in Economics and Mathematical Systems* 599. pp. 101-111.

POSADA, M., Hernández, C., López, A. (2007b) Testing Marshallian and Walrasian Instability with an Agent Based Model in *Lecture Notes in Computer Science, forthcoming*.

PREIST, C., van Tol, M. (1998). Adaptive agents in a persistent shout double auction in *Proceedings of the First International Conference on Information and Computation Economics. pp*. 11-18.

ROTH, A., Erev, I. (1995) Learning in extensive- form games: experimental data and simple dynamic models in the intermediate term in *Games and Economic Behaviour*, 8. pp. 164-212.

RUST, J., Miller, J., Palmer, R. (1993) Behaviour of trading automata in computerized double auctions in Friedman and Rust (eds.), *The double auction markets: Institutions, theories and evidence* Addison-Wesley. pp. 155-198.

SMITH, V. (1962) An Experimental Study of Competitive Market Behavior in *Journal of Political Economy*, 70. pp. 111-137.

SMITH, V. (1976) Bidding and Auctioning Institutions: Experimental Results in *Papers in Experimental Economics*, Cambridge University Press. pp. 106-127.

SMITH, V. (1989) Theory, Experiment and Economics in *Journal of Economic Perspectives*, Winter. pp. 783-801.

SMITH,V. (2003) Constructivist and Ecological Rationality in Economics in *The American Economic Review*, 93(3). pp. 465-508.

TESAURO, G., Das, R. (2001) High performance bidding agents for the continuous double auction in *Proceedings of the Third ACM Conference on Electronic Commerce. pp*. 206-209.

VRIEND, N. (2000) An illustration of the essential difference between individual and social learning, and its consequences for computational analyses in *Journal of Economic Dynamics & Control* 24. pp. 1-19.

WALSH, W., Das, G., Tesauro, G., Kephart, J. (2002) Analysing complex strategic interactions in multi-agent games in *Workshop on Game Theoretic and Decision Theoretic Agents*. pp. 109-118.

WEIBULL, J. (1995) *Evolutionary Game Theory*, Cambridge, MA, MIT Press.

Return to Contents of this issue