

Behavioral Heuristics and Market Patterns in a Bertrand-Edgeworth Game

D. Heymann(†, §), E. Kawamura(§), R. Perazzo(‡, §)
and M. G. Zimmermann(§)

(†)ECLAC,

Paraguay 1235, (X1233UVW), Buenos Aires, Argentina,

(§)Universidad de San Andrés, Depto. de Economía,

Vito Dumas 284, (B1644BID)Victoria, Provincia de Buenos Aires, Argentina,

(‡)Depto. de Investigación y Desarrollo, Instituto Tecnológico de Buenos Aires,
(C1106ACD) Avda. Madero 399, Buenos Aires, Argentina.

April 15, 2010

Abstract

In this paper we consider a setup of market interactions which may provide a stylized representation of real- world counterparts. We explore complementarities between analytical, agent-based and experimental approaches, as seems to be required by the problem. The Bertrand-Edgeworth game (price- setting with capacity constraints) does not admit Nash equilibria in pure strategies; the mixed strategy equilibria have quite complicated characterizations. We consider two families of heuristic rules which may define frameworks for price-setting: one where prices respond to the past volume of sales; the other where prices are conditioned by the observed correspondence between price changes and profits. The markets inhabited by artificial agents following such rules, with fixed behavioral parameters, show convergence to "dynamic equilibria ", where average prices are more or less constant average prices, but where individual prices keep changing. These states, in addition, can potentially be disturbed by agents who modify adaptively their behavior parameters in a profit- enhancing way. The tendency to reach "quasi steady states " in aggregate prices is also observed typically in experimental results. The price-setting choices of players in experiments can be approximated reasonably with a variant of a simple sales- based rule which governs period-by-period behavior, and with parameters which show a dynamics on a longer time-scale, as if incorporating an adaptive search for profit. Various salient features of the experimental results can be reproduced with simulations of artificial agents who apply the postulated behavioral sales- based strategy with parameter values calibrated from the experiments.

1 Introduction

In this paper we study a price setting game which may provide a schematic counterpart of relevant economic situations and which, despite the extreme simplicity of its specification, poses interesting analytical and practical questions. We carry out the discussion combining the use of agent based simulations and experiments with actual players. The results suggest that a representation with artificial agents acting on the basis of simple behavioral heuristics can reproduce approximately salient features of the experiments, and that those heuristics have reasonable, although not optimal, properties given the characteristics of the game. The qualitative performance of the markets that we observe, and simulate, tends to show a convergence from above of average prices to near the competitive level but, once that state is reached, it can be disturbed by agents seeking higher profits through larger unit margins.

Actual price-setting agents face complicated decision problems in their routine activities. Consider the case of firms such as small retail businesses (grocery shops, drugstores): even assuming away considerations regarding product or location differentiation, and the multiplicity of goods sold, the strategic interaction between the price behavior of firms remains far from trivial. In a "Marshallian short-run" where the size and the number of the firms in the market remains constant, a natural way to represent schematically the setup of that interaction is through a Bertrand-Edgeworth game, where individual price setters with constant marginal costs subject to a capacity constraint meet customers who are fully informed about posted prices and can move costlessly between suppliers.

Since the original formulation by Edgeworth (1925), the literature on this price-setting game has included contributions from theoretical and experimental perspectives. The constraint on the volume of sales of firms implies that, when market demand is sufficiently high to exceed aggregate capacity at a price equal to costs, the standard Bertrand equilibrium with prices at the competitive level does not apply: if $N - 1$ firms set that price, the N th firm would be in the position of a monopolist against the "residual" demand. This observation generalizes to the result that the game does not admit a Nash equilibrium in pure strategies (see Chamberlin, 1956, D'Aspremont, 1985 and Tirole 1988). The mixed strategy equilibria have been characterized, among others, by Dasgupta and Maskin (1986) and Allen and Hellwig (1986). They involve quite elaborate calculations and a subtle coordination in the behavior of agents, which reduce their plausibility as representations of actual market outcomes.

In an experimental setting, Brown Kruse et al. (1994) find no evidence of the emergence of such equilibria, or a definite correspondence with alternatives like the competitive outcome or collusion, although they observe that average prices tend to converge from above to the "competitive range" (see also Smith, 1982). Cason et al. (2005) and Guillén Álvarez (2004) identify a pattern akin to the cycles suggested by Edgeworth, where prices initially fall to near the competitive level, but rebound as some agents try to take advantage of their "market power", and are subsequently driven down again by competitors undercutting their rivals. We will encounter again in what follows the image of a system which orders itself in quasi-stationary states for aggregate market variables, but where neither individual nor market behavior becomes frozen: those "as if" steady states are

compatible with constantly changing actions of particular agents, and they can be disturbed by profitable deviations on the part of some individuals.

The difficulty of defining equilibrium strategies and the lack of correspondence of the standard oligopoly model with experimental results suggest the use of a different approach, where individuals search for profits but without necessarily finding mutually optimal strategies, and where the competitive process is governed by behavior that follows simple decision rules. The agent-based modeling literature has shown that the interaction of behaviorally uncomplicated players can result in non-trivial system dynamics, which have to be analyzed through computer simulations. But any such simulation exercise presumes a previous specification of decision rules that describe the choices of the artificial agents, and there may be quite different ways to define "reasonable" decision criteria. Lacking some form of validation of the postulated decision algorithms, the relevance of the results of a model can be difficult to ascertain. Thus, it seems especially important to combine the analysis of the properties and outcomes of agent-based representations with a look at evidence with which to match the model's behavioral assumptions and its predictions about the performance of the system under study. With this purpose, we contrast the results of simulations with those drawn from experiments.

The behavior of players in variants of the Bertrand-Edgeworth game has been studied analytically and experimentally by Álvarez Guillén (2004), Abbink and Brandts (2005), Cason et al. (2005). In the first of those works, firms must pre-determine both output and prices while, in our setup, they fix prices and "produce for order" instantaneously after observing their own demand: this focuses the analysis on the price-setting decision, isolating it, at this stage, from the choice of production volumes. The other two papers consider scenarios where agents are informed about the history of their competitor's prices, and thus can in principle calculate counterfactual measures of the profits they may have obtained by posting prices at different levels from those effectively announced. We concentrate on the simpler (and not unrealistic) case where firms can observe only their own sales, and must condition behavior on their specific market experience. In that context, we first study the features and operation of two families of heuristic rules parametrized by "behavioral coefficients": one where prices respond directly to the sales history of the individual price-setters, another where prices are function of the individual record of profit changes following price changes. The experimental results provide a quite clear support for the relevance of a modified version of the simple sales-based strategy, which implies falling prices if sales were below capacity, and "alternations" between rising and constant prices (with some instances of reductions) when sales reached the capacity level.

The experiments confirm the presence of market patterns previously discussed in the literature: starting from high average prices with a large individual dispersion, there is a gradual downwards converge. These features can be reproduced with simulations where artificial agents apply the modified, "probabilistic", sales-based decision rule. In addition, we can use the model to make numerical "predictions": an analytical expression for the asymptotic average price level at which the artificial-agent system would settle (quantifying the behavioral parameters with averages over the whole experimental population) provides good lower bound approximations for average market prices observed in the experiments .

The sales- based schemes for price setting do not take into account explicitly the evolution of profits. However, while "everyday" decisions may follow a simple rule, the search for higher profits can be embodied in adaptive changes in the behavioral parameters within the "heuristic family". In fact, the experiments show that behavioral parameters evolve, with some defined patterns, particularly regarding the likelihood of maintaining or increasing prices after capacity sales. Players in experiments tend to start the game with aggressive policies, setting higher prices almost every time they have realized full sales. Later on, as the average price and the dispersion decline, agents are found more often keeping their price constant, as if fearing to lose sales otherwise. Eventually, the frequency of price increases following capacity sales seems to reach a level around which it slowly fluctuates. The model with simulated artificial agents is useful in rationalizing this experimental feature. We can identify "cutoff points" such that, if $N - 1$ agents unanimously use sales- based heuristics with a certain probability of no price change after full sales, the N th individual would profitably choose its parameter with a similar order of magnitude; with a population parameter near zero, the deviating agent would gain by selecting a higher value. Both analytically and experimentally, a price level of particular importance is defined by the value $p_{edge} \equiv p^* + c/N$, (where p^* is the competitive price level, c is the unit cost of the good, and N are the number of players in the market): this marks the limit for the population price below which a deviator can increase profits by raising the price. This critical price level, would mark the point around which Edgeworth cycles are generated. In several experimental sessions, we observe large upwards price excursions after the average price has converged to a value near p_{edge} .

On the whole, the results we have obtained appear to confirm the existence of useful complementarities between analytical, agent- based and experimental approaches to the study of systems of economic interest. The evidence that experimental agents tend to act as if using heuristics belonging to a specific set of simple rules, and the emergence of a patterned order with room for case-to-case variations and for occasional shakeouts seem interesting indications of the actual working of competitive processes.

The rest of the paper is organized as follows. Section 2 presents the Bertrand-Edgeworth model, and discusses the non- existence of deterministic Nash equilibria. Section 3 introduces the heuristic price-setting rules, and studies their properties using numerical simulations. Section 4 presents the experimental results, and their interpretation using the agent- based models; the comparison between models and experiments is developed further in Section 5. The final section presents our conclusions.

2 The Bertrand-Edgeworth price-setting game

The market is populated by N firms who sell a single, homogeneous, non- storable good. Each firm can produce every market day up to a maximum quantity q^* ; for $q \leq q^*$ the unit cost is a constant c . At the beginning of a day, all firms post a selling price p_i ; $i = 1 \dots N$ (it must always be that $p_i > c$), without any exchange of information among them. The announcement implies a commitment to supply the good at that price, up to the limit determined by the production capacity. Subject to

that limit, the firm produces a quantity equal to the orders it has received in the day.

The information available to potential buyers prior to their purchase decision is assumed here to include the whole set of posted prices. There are no search or mobility costs. This implies that buyers will visit firms in an order of increasing prices, and buy from them until their demand for the good is satisfied or the supply capacity becomes exhausted. For all practical purposes, the demand side of the market acts as if it consisted of a single "representative" buyer who acts non-strategically according to the market demand function. We specify that function with a unit price elasticity, corresponding to planned purchases with a fixed value M per day, and concentrate on the case where $M > Nc q^*$, so that the desired value of spending by buyers exceeds the cost of production of the capacity volume of output.

Realized benefits for a firm i who sets price p_i and sells a volume q_i of the good (constrained by $0 \leq q_i \leq q^*$) are given by:

$$B_i = q_i(p_i - c) \quad (1)$$

The *competitive price* that would equilibrate the market is:

$$p^* = \frac{M}{Nq^*} \quad (2)$$

At that price every firm sells all its potential output. It can be easily verified that the competitive outcome is not a Nash equilibrium. Let $p_i = p^*; \forall i$. The individual benefit of a firm is $B^* = q^*(p^* - c) = (M/N)(1 - c/p^*)$. A single deviating agent who posts a price $p_j > p^*$ would realize a value of sales $p_j q_j = M - (N - 1)p^* q^* = p^* q^*$, independent of its price, while profits would be $B_j = (p_j - c) q_j = p^* q^* (1 - c/p_j)$. It is clear that this agent has an incentive to set a *higher* price $p_j > p^*$, and reduce production costs for a given value of sales. The maximum profit would be attained choosing $p_j \rightarrow \infty$.

More generally, it can be seen that no price p' set by all firms $p_i = p'; \forall i$ can be a Nash equilibrium.

- If $p' > \frac{N}{N-1}p^*$ so that the demand would be exhausted with less than $N - 1$ firms selling at full capacity, there is one supplier with zero sales, who thus has an incentive to reduce prices.
- If $\frac{N}{N-1}p^* > p' > p^* + \frac{c}{N}$, the N th firm gains by setting a price slightly below p' . Here, the N th firm effectively faces a choice between "guaranteeing" for itself full capacity sales, by lowering its price relative to its competitors, and capturing the whole "residual demand" (which is positive, in this case) after all the other firms have sold at their capacity levels, and raise its price as to cut production costs to near zero. The condition for preferring a lower price $p' - \varepsilon$ (for $\varepsilon \rightarrow 0$) is:

$$B_{p'-\varepsilon} = (p' - c)q^* > Np q^* - (N - 1)p' q^* = B_\infty \quad (3)$$

or:

$$p' - p > \frac{c}{N} \quad (4)$$

- If $p' < p^* + \frac{c}{N}$ for $N - 1$ agents, the N th firm has an incentive to set a very high price. The price level $p^* + \frac{c}{N}$ thus defines a boundary with important implications for the incentives of agents. We refer to that boundary as

$$p_{edge} \equiv p^* + c/N \tag{5}$$

The results above illustrate a general feature of the problem, the lack of Nash equilibria in pure strategies. This suggests that individual agents will be involved in possibly inconclusive searches for beneficial strategies. The profit implications of alternative price decisions will be affected by properties of the environment, such as the following:

- The range of prices above the competitive equilibrium such that, if all firms choose a price in the set, would make profitable a “large” upwards deviation by one of them, clearly diminishes with the number of firms N , since it is determined by the interval $[p^*, p^* + \frac{c}{N}]$. For a large N , a small increase in the average price set by competitors may change substantially the advantageous move of the remaining supplier, from a large raise in its price to a small reduction.
- If $N - 1$ firms set a price $p' < p^* + \frac{c}{N}$ the profits of the N th supplier would vary non-monotonically with the size of a positive deviation of the price, to $p' + \varepsilon$. It can be seen that a “local” shift, given by a small value of ε , would lead to lower profits, while benefits would increase for a “sizeable” price jump.
- In states where prices are concentrated around a level $p' > p^* \frac{N}{N-1}$, non-selling agents would have incentives to undercut their competitors. However, the downwards evolution of aggregate prices could be very gradual. If the variance of individual prices is small, the high-priced firms can become competitive by lowering their price very little. Suppose that n firms thus reduce prices by a small amount ε in a given market period (the number n would be such that, when $p' \sim p^* \frac{N}{N-j}$, $n \sim j$, that is, those firms which have realized zero sales cut prices); the aggregate effect on the average price would be $\delta \bar{p} = \frac{n\varepsilon}{N}$. The proportional change could then be approximated by: $\frac{\delta \bar{p}}{p'} \sim (1 - \frac{p^*}{p'}) \frac{\varepsilon}{p'}$ which, for a given deviation of the average price from the competitive level, would be a number independent of N , and of a small magnitude if the size of the individual price adjustment ε is small¹. A behavior of this type appeared to be relevant in market experiments (see section 4 below).

¹This feature evokes the property found in imperfect competition models that the price of individual firms depends strongly on the average price (rather than respond to the level of market demand) when the goods that they produce are close substitutes (see for example Ball and Romer 1989 for a macro application where the effect is relevant).

3 Price behavior based on some simple heuristics

The previous discussion has shown that the profit opportunities for individual price setters vary considerably with the environment which results from the whole set of decisions of the firms. Since there is no obvious way to determine an optimal strategy, a natural conjecture would be that prices are set through some simple and more or less robust decision rule. We explore here the specification and the properties of two types of heuristics: one which conditions prices on past volumes of sales, and another where prices depend on the price- benefit history. The results of this section prepare the ground for the analysis of experimental evidence in the next section, by providing a "lense" through which to look at the data.

3.1 A simple sales-based pricing rule

A particularly simple scheme would be for the agent to watch sales q_i , and lower the price p_i if the realized volume does not reach full capacity q^* , and increase the price if sales are at their maximum feasible amount. This pattern of behavior can be described as:

$$p_i(t+1) = p_i(t) + \begin{cases} \gamma_+ & \text{if } q_i = q^* \\ -\gamma_- & \text{if } q_i < q^* \end{cases} \quad (6)$$

where t denotes the "market day" and $\gamma_{+/-}$ are behavioral parameters which determine the magnitude of upwards and downwards price changes². At this point, we consider a model with agents which share the same behavioral parameters.

Price setters following this pricing rule would reach an aggregate steady state, with constant average total sales and selling price, whenever N_s firms selling at full capacity raise price ($\approx \gamma_+$) and $N - N_s$ suppliers lower prices ($\approx \gamma_-$). If N is sufficiently high so that "small number" effects can be ignored (see the discussion in Section 5), the steady state would be given by a similar value of price increases and decreases: $\gamma_+ N_s \approx \gamma_- (N - N_s)$. The intuition is simple assuming that $\mu = \frac{\gamma_-}{\gamma_+}$ is an integer. It would take μ upward movements to go from the minimum to the maximum level of the price distribution, while the opposite change would take place in one jump (a similar argument applies if μ is any rational number). At any given moment agents can be placed in $1 + \mu$ levels according to their prices. Individual prices change as if placed in a "conveyor belt" : they grow in small steps of order of magnitude γ_+ until the moment in which they are forced to fall by an amount in the order of γ_- . In the steady state, the ratio of the number of firms raising and cutting their prices would be μ .

The average sale price \bar{p} can be related to the competitive equilibrium price p^* . Since $p^* = M/Nq^*$ the average sale price $\bar{p} \approx p^* N/N_s$, and the average sales as $\bar{q} = q^* N_s/N$. We therefore get

²It is possible to use random adjustments with price-shifts γ_+ and γ_- assumed, for example, to be drawn from a uniform distribution in the intervals $\gamma_+ \in [0, \Delta^+]$ and $\gamma_- \in [0, \Delta^-]$. For the sake of clarity, in tyhis discusion we treat these parameters as fixed.

(for large N),

$$\frac{\bar{p}}{p^*} \approx 1 + \gamma_+/\gamma_-$$

$$\frac{\bar{q}}{q^*} \approx \frac{1}{1 + \gamma_+/\gamma_-}$$

For instance, when $\gamma_+ = \gamma_-$, $\bar{p}/p^* = 2$. If settled in this dynamic equilibrium agents, on the average, sell all their capacity output during half of the time. This is done at an average price that is twice p^* and therefore their average benefit is greater than the one that could be obtained in a steady competitive equilibrium. The heuristics would induce some sort of implicit collusion among sellers, who nevertheless would be clearly acting suboptimally, since for any of them, fixing a price slightly below \bar{p} would ensure full sales without a price sacrifice.

Using this price adjustment rule the average prices and quantities would depend only on the parameters γ , and not directly on N , although the number of competitors may operate indirectly on price behavior by changing the modules of price increases and decreases chosen by agents.

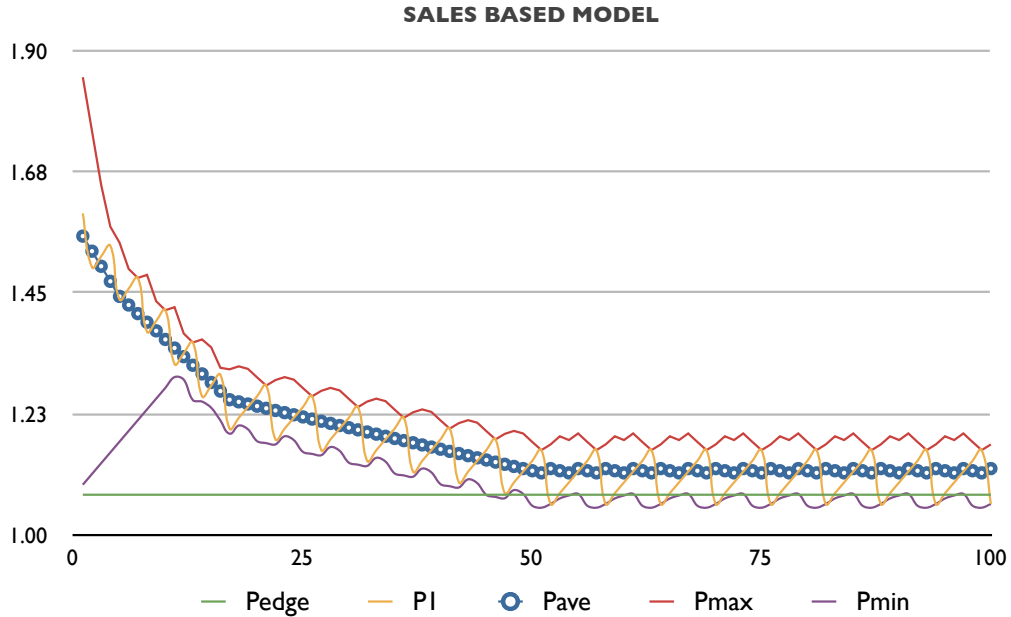


Figure 1: Sales based model price time series. Average (p_{ave}), maximum (p_{max}) and minimum (p_{min}) price as a function of time, together with a sample trajectory (p_1) of an agent and $p_{edge} = p^* + c/N$. Note that $\bar{p} \simeq 1.20 \simeq (1 + \gamma_+/\gamma_-)$. Parameters: $N = 10, p^* = 1, c = 0.75, \gamma_+ = 0.02, \gamma_- = 0.10$.

Numerical simulations of the system show that, homogenous agents starting from randomly chosen initial prices, after a transient, the average price does tend to the steady state $p/p^* = 1 + \gamma_+/\gamma_-$ derived above (for large N). The price of each individual agent fluctuates; however, the average price of the N agents remains constant in time. The system is permanently changing its detailed configuration because the agents continuously vary their prices, and their relative position in the price distribution. Firms selling at full capacity increase their prices until they are undercut by competitors (who previously were high-priced non-sellers) and led to revise prices downwards. Thus, agents are continuously exchanging roles while the aggregate features of the system remain constant.

3.2 A rule based on price-profit gradients

In this section we study a price-setting rule which considers explicitly the observed performance of profits, and allows for price adjustments of variable size. Here, agents adapt their posted price following the direction (of price increases or decreases) that has been associated in the past with rising benefits, reversing course, if benefits have fallen. Allowing agents to have a two-day memory of sales and prices, the rule would make prices vary according to the gradient of the benefit function:

$$p_i(t+1) = p_i(t) + \sigma \delta(B_i(t)) \text{sign}(\delta(p_i(t))) + \eta \quad (7)$$

Here t labels the market day, σ is an adjustment parameter, $\delta(X(t)) \equiv X(t) - X(t-1)$ and η is a random number, uniformly distributed in $[-\epsilon, +\epsilon]$ if there were sales in t or in $t-1$, and uniformly distributed in $[-\epsilon, 0]$ if there were no sales in t and in $t-1$; with ϵ a small number which determines the “search space” for the agent.

The pricing strategy determines an increase in the price proportional to the absolute value of the change in benefits in the past period if either the benefit has increased with an increasing price or it has diminished with a lower price. A symmetric behavior applies in the other cases. Due to its small magnitude, the additive term η in (7) is effective in case the benefit has not changed: it represents a random search for a better situation in the immediate neighborhood of the current price, making sure that after two periods with no sales, the search is towards lower prices.

This decision rule, meant to mimic a change along the gradient of the benefit function, generates asymmetric sizes of price increases and decreases. The feature results from the discontinuity of the benefit gradient, as profits of an individual firm change abruptly from positive (and potentially high) values to zero following small adjustments which make the firm unable to sell³. The behavior thus generated makes prices eventually fall, and sales become positive, when the sales of the firm have been zero. Thus, in contrast with the case of the sales-based heuristics discussed in the previous section, here the system would not be expected to settle in a steady state where more than one firm has zero sales, that is a state where the average selling price is higher than $\frac{N}{N-2}p^*$. In addition, as

³This profits-based rule makes prices fall when, after a price reduction, the volume of sales goes from zero to full capacity. With constant prices of the other firms, that move cannot be convenient for the agent. The points underscore the “suboptimality” of simple decision procedures.

indicated previously, small price increases by an individual firm would cause a reduction in profits when the prices of competitors are concentrated around a level higher than $p_{edge} = p^* + c/N$. Those bounds above p^* depend negatively on the number of firms; thus, a larger number of competitors N can be associated here with lower average prices, which was not the case with the simple rule that made prices jump every time sales were at full capacity.

In simulations with agents using this decision rule, prices converge to a “dynamic” equilibrium which is different from that with the sales- based algorithm (6), although here there is also a “conveyor belt” pattern of gradually increasing prices (as a response to small increments in benefits), and sudden falls when additional price raises make sales and profits abruptly go to zero. In cases where the number of agents is small, sellers following this rule occasionally coincide for some periods in increasing their price, as in an implicit collusion. The system undergoes big fluctuations (resembling Edgeworth cycles), which are not long- lasting. Also, the simulated systems tend to converge to states where the average sale price is below $p^* \frac{N}{N-1}$, implying that, in general (that is, leaving aside ”excursions” that modify transitorily the configuration), all firms make positive sales, while p_{edge} appears like a lower bound. The simulations with the profit- gradient rule thus result in market prices that fall with N .

4 Experimental results

The previous section introduced two families of price setting rules, studied their properties and their implications for the performance of the market. Here we report the results of experiments where actual players determined prices and competed in the game in search for profits. The experimental outcomes are compared with the the assumptions and implications of the models based on the simple heuristics, a comparison which is carried out further in Section 5.

4.1 The experimental setup

The experiments were conducted using a specially designed webpage⁴; most players were graduate students in Economics who knew beforehand the general nature of the game. The parameters of the experiments were given by N , the number of participants in a game (since the play was ”silently” carried out in individual computers, on some occasions simultaneous games were ran simultaneously),: M , the magnitude determining the level of market demand, c , the unitary cost, and q^* , the maximum individual capacity. Participants were not informed of the number of ”competitors” they would be playing with. The experimenters announced publicly the following parameters: the (common) unitary cost c , the maximum individual capacity q^* , and the competitive price level p^* . In each session, the system always suggested as the initial price, the competitive price level p^* .

In each market round, all players posted prices through the interactive webpage, after which

⁴Visit the home page: <http://www.elautomataeconomico.com.ar/>

the server computed the amount sold by each agent. The server buys sequentially from "firms", starting with the cheapest seller, and visiting all other participants in an increasing price order. The server buys up to the seller's capacity constraint, until either the total amount of money M is spent or there are no more participants to visit. Finally, each player receives from the server an indication that a new round of play is open, together with her realized levels of price, quantity sold, and benefits in the last five market days. Of the eight sessions that we report here, in three of them participants received pecuniary rewards contingent on their performance in generating realized profits. Three different set of parameters were used: $E_A : c = 0.9, p^* = 2.22, q^* = 3.0$, $E_B : c = 2.15, p^* = 2.45, q^* = 3.0$ or $E_C : c = 0.9, p^* = 22.2, q^* = 3.0$, with M adjusted so as to satisfy: $M = p^* q^* N$. The number of players in the experiments, N , varied in the range (3, 12).

4.2 Behavioral strategies and market outcomes

In order to characterize the average behavioral strategies followed by agents, we computed for each experiment a realized mean (across all players) of the conditional probability $P(\Delta p_{t+1}; q_t)$, which measures the frequency of price changes of a certain sign (Δp_{t+1}) given the volume of sales in the previous period (q_t). Thus, the cells count the number of events in which players either increased the price $\Delta p_{t+1} > 0$, maintained the posted value, $\Delta p_{t+1} = 0$, or decreased the price $\Delta p_{t+1} < 0$ on the next round, given that in the previous period they had either (a) sold all their capacity output $q_t = q^*$, (b) made positive but not "full" sales, $0 < q_t < q^*$, or (c) sold a zero quantity $q_t = 0$. Table S1 shows such conditional probability for an experiment.

Table S1

Average over players of conditional probability $P(\Delta p_{t+1}; q_t)$ of next round's change in price $\Delta p_{t+1} = p_{t+1} - p_t$ given that the player sold the whole production $q_t = q^*$, a partial amount, $0 < q_t < q^*$, or zero $q_t = 0$. The figures correspond to the last 50 rounds of experiment #18 with $N = 6$ players and $T = 100$ periods of play.

	$\Delta p_{t+1} < 0$	$\Delta p_{t+1} = 0$	$\Delta p_{t+1} > 0$	# of events
$q_t = q^*$	0.13	0.42	0.45	200
$0 < q_t < q^*$	0.97	0.0	0.03	57
$q_t = 0$	0.94	0.02	0.04	43

The transition matrix $S1$ shows that the experimental behavior reproduces an important feature of the sales-based rule discussed in the previous section, as is the almost unanimity of price reductions following less than full capacity sales. However, in contrast with that very elementary rule, full sales do not systematically lead to price increases: in a sizeable number of cases, agents maintain their price, or even lower them sometimes. In what follows, we will refer to the parameters $h_- \equiv P(\Delta p_{t+1} < 0; q_t = q^*)$, $h_0 \equiv P(\Delta p_{t+1} = 0; q_t = q^*)$, $h_+ \equiv P(\Delta p_{t+1} > 0; q_t = q^*)$, as the measures of frequencies of price falls, non-variations and increases, respectively, after full sales.

The number $\xi \equiv P(\Delta p_{t+1} < 0; 0 \leq q_t < q^*)$ is the frequency of price decreases conditional on sales below capacity. In order to check the relevance of including sales in the previous round q_{t-1} , we measured the probability of changing prices conditional on q_t and q_{t-1} ; the result, which was consistent in all experiments, was that the behavior was well approximated by a Markovian process, without a visible effect of sales before the current period.

Table S2 shows a summary of results, showing for each experiment the number of agents N , the number of market days (T) and the ratio p_{edge}/p^* , together with the observed behavioral parameters (frequencies of conditional price responses, modules of price changes: $\xi, h_+, h_-, h_0, \gamma_+, \gamma_-$; also shown is the average posted price $\langle p \rangle$ relative to p_{edge} . All values were averaged using the last half of market days.

Table S2

Most relevant parameters and results for all the conducted experiments. Probabilities $\xi = P(\Delta p_{t+1} < 0; 0 \leq q_t < q^*)$, $h_- = P(\Delta p_{t+1} = 0; q_t = q^*)$, $h_0 = P(\Delta p_{t+1} = 0; q_t = q^*)$ and $h_+ = P(\Delta p_{t+1} > 0; q_t = q^*)$, and $\langle \gamma_+ \rangle$, $\langle \gamma_- \rangle$ where computed taking into consideration only the last half of the rounds.

Exp. #	N	T	ξ	h_-	h_0	h_+	$\langle \gamma_+ \rangle$	$\langle \gamma_- \rangle$	$\langle \gamma_+ / \gamma_- \rangle$	$\langle p \rangle / p^*$	p_{edge} / p^*
13 : E_A	12	50	0.92	0.07	0.50	0.43	0.043	-0.079	0.684	1.209	1.034
14 : E_C	3	50	0.79	0.04	0.54	0.42	0.111	-0.282	0.443	1.237	1.014
16 : E_A	6	100	0.94	0.10	0.35	0.55	0.012	-0.019	0.633	1.354	1.068
18 : E_B	6	100	0.97	0.13	0.42	0.45	0.019	-0.023	0.734	1.237	1.146
17 : E_A	4	100	0.90	0.01	0.18	0.81	0.147	-0.203	0.857	1.395	1.101
19 : E_A	11	100	0.98	0.11	0.52	0.37	0.048	-0.081	0.650	1.155	1.036
20 : E_B	4	100	0.61	0.01	0.75	0.24	0.589	-1.283	1.039	2.310	1.220
21 : E_B	6	80	0.84	0.10	0.66	0.24	0.064	-0.091	0.636	1.150	1.146

The experimental results show diversity in the behavior of agents and markets, but also some reasonably clear stylized features of the decision rules of players and the consequent market outcomes. These regularities are described next.

4.2.1 Price strategies can be approximated by a modified sales- based adjustment rule

As shown in Table S2, the conditional probability of price decreases after less than full sales was very close to one. This corresponds very closely to the mayor characteristic of the sales-based heuristics discussed in the previous section. In contrast, full sales tend to induce either price increases in a sizeable proportion of instances (as in that decision rule), but there are also frequent occasions where

prices remain unchanged (see columns h_0 and h_+ in Table S2). There were some cases (around 13% averaging all periods, agents and experiments) in which prices were reduced following full sales.

Simulations performed using the profit-gradient rule (keeping a similar number of events and of agents), were used to measure the conditional probabilities $P(\Delta p_{t+1}; q_t)$. Table S3 shows the results of this computation in one numerical experiment, and indicates that the salient differences with Table S1 are that: (a) the profit-gradient rule either increases or decreases prices with full sales, with very few cases of constant prices and, (b) when sales are positive, but less than full capacity, $0 < q_t < q^*$, the rule either increases or decreases prices, while the experiments nearly always decreases prices. Thus, using the conditional probability $P(\Delta p_{t+1}; q_t)$ to discriminate between the two heuristics as potential representations of behavior in experiments, a modified sales- based rule appears to do a better job in a simple way.

Table S3

Average over players of conditional probability $P(\Delta p_{t+1}; q_t)$ of next rounds change in price $\Delta p_{t+1} = p_{t+1} - p_t$ given the player sold the whole production $q_t = q^*$, part of it $0 < q_t < q^*$, or none at all $q_t = 0$. The figures are computed from a profit gradient rule using the following parameters: $N = 10$, $c = 0.75$, $p^* = 1$, $\gamma = 0.2$, $\epsilon = 0.05$, for a total number of $T = 100$ rounds.

Only the last 50 rounds were considered in the statistics.

	$\Delta p_{t+1} < 0$	$\Delta p_{t+1} = 0$	$\Delta p_{t+1} > 0$	# of events
$q_t = q^*$	0.49	0.05	0.41	415
$0 < q_t < q^*$	0.52	0.0	0.48	50
$q_t = 0$	0.88	0.0	0.12	25

4.2.2 Average prices tend to converge from above, but do not rest in a steady state

The interaction between price setters in the experiment results in market dynamics which show identifiable patterns of organization, but which do not relax into a sustained steady state, either at the individual or the aggregate levels. A typical evolution, observed in several experiments (and with analogies with what was found by Kruse et al [4], initially shows high average prices and a large dispersion, with a gradual downwards trend and a convergence of individual prices. This corresponds to a behavior marked by a first period of intensive learning, where agents (who, it may be recalled, are informed about the structure of the game) start with wide differences in their strategies and react quite strongly to the outcomes they observe; due to the capacity constraints, the existence of low- priced suppliers allows high- priced firms to make positive sales for some time




Figure 2: Average, maximum and minimum offer price (rescaled with p^*) time series. Also shown is the Edgeworth price $p_{edge} = p^* + c/N$. Experiment #21.

by capturing the residual demand. Subsequently, the competitive responses of agents lead to a drop in price dispersion, along a declining trend for the average (connections between market features and individual strategies are discussed in more detail in the next section).

As the play goes on, the system goes through a phase of slow price movements. However, the market does not settle in a state with steady mean prices. As the average price falls, some agents tend to perceive incentives to search for higher profits through price increases, which can generate market-level “excursions”. These perturbations did occur in five out of eight experimental sessions and most of them started when the average price was $\bar{p} > p_{edge}$; that is above the suggested critical level discussed in Section 2.

In general, in six out of eight experiments, the average selling price was above the bound $p^*N/(N - 1)$ at which only $N - 1$ agents would realize full sales. Also, even after 100 ‘market days’, only in two experimental sessions, the average selling price became lower than the limit $p_{edge} = p^* + \frac{c}{N}$. In this respect, it would be interesting to perform longer experiments to measure statistics on the “price excursions”.

The idiosyncratic behavior of some agents can be experimentally relevant, particularly in small groups. For example, in experiment #20 reported in the Figure 4, the aggregate price was decisively influenced by the aggressive price increases of a subject, who was able to generate high profits. This




Figure 3: Average, maximum and minimum offer price (rescaled with p^*) time series. Also shown is the Edgeworth price $p^* + c/N$ and p_1^* . Experiment #18.

experiment was special in that the subjects who participated had already played in a previous session; this experience possibly made the market develop without undergoing the first stage of “wide search” with high mean prices and a slow decline in price averages. Here, the mean price soon reached a level close to $\bar{p} \approx p_{edge}$, and players engaged in several large excursions, as they had incentives to do.

Our experiments do not show a clear pattern linking the average market price with the number of players N (other parameters constant), although the limited quantity of experiments performed allows no definite conclusion to be drawn on this.

4.2.3 Behavioral parameters show some patterns, and change over time

Sizes of price changes (γ_+, γ_-) Measuring the average price change $|\delta p(t)|$ over all players shows in general a rapid decrease in the average module, reaching a small value, with oscillations around it. This results in a narrowing of the range between the maximum and the minimum price, which takes place on a time scale of some 30 ‘market days’. There are large variations on the average size of price adjustments in different experiments, but the magnitude of price changes tends to be




Figure 4: Average, maximum and minimum offer price (rescaled with p^*) time series. Also shown is the Edgeworth price $p^* + c/N$ and p_1^* . Experiment #20.

relatively high in the first rounds of play, when the variance of individual prices is large, and to diminish as the range of price dispersion falls (see Figure 5).

In all but one of the experiments, the average module of price increases was smaller than that of price reductions: $\langle \gamma_+ / \gamma_- \rangle < 1$ (see Table S2). For individual players, when N is small (< 6), there is a tendency towards roughly similar modules $\langle \gamma_+ / \gamma_- \rangle \approx 1$, while for larger number of competitors, agents appear to be more cautious, in the sense that $0.5 < \langle \gamma_+ / \gamma_- \rangle < 1$.

Frequencies of price increases following full sales Figure 5 shows a time series of an estimation of $\hat{h}_-(t)$, $\hat{h}_0(t)$, $\hat{h}_+(t)$ obtained by aggregating those events where players sold at full capacity, and computing frequencies over the last five steps for each 'market day' t and each agent.

A general pattern observed in all experiments was that players start the experiment with an 'aggressive' policy of increasing prices if sales reach capacity $h_+ \approx 1$, $h_0 \approx 0$, and become more cautious later on: h_0 increases steadily, while h_+ decreases correspondingly (see Figure 5). After this initial trend, h_0 and h_+ tend to show oscillations, suggesting that agents do not stop searching for better strategies in a context where there is no simple best response rule.




Figure 5: (top) Time series of $\hat{h}_-(t)$, $\hat{h}_-(t)$, $\hat{h}_-(t)$, measured by aggregating and averaging the last 5 time steps for each t over those agents having full sales on Experiment #21.

5 The revised sales-based model and the experimental evidence

The previous section showed that the sales-based rule can be readily extended to accommodate salient features of the experimental results like the transition matrix shown in Table S1. Here, we carry out the extension in some detail, and comment on the performance of the model compared with the experiments.

The central change to the simpler heuristics (6) discussed in section 3 is the introduction of a probabilistic choice of price changes following full sales, parametrized by probabilities h_- , h_0 , h_+ of decreasing, keeping constant, or raising the price. The price rule for agent i would then be:

$$\begin{aligned}
 p_i(t+1) &= p_i(t) - \gamma_- \quad \text{if } q_i < q^* \\
 p_i(t+1) &= p_i(t) + \begin{cases} \gamma_+ & \text{with prob. } h_+ \\ 0 & \text{with prob. } h_0 \\ -\gamma_- & \text{with prob. } h_- \end{cases} \quad \text{if } q_i = q^*
 \end{aligned} \tag{8}$$

($h_- + h_0 + h_+ = 1$). The dynamics of the model will then depend on the behavioral parameters (γ_+ , γ_- , h_0 , h_+), and the number of players N . We have performed numerical simulations of the above model, starting from a random initial condition across agents⁵, who share the same behavioral parameters.

In Figure 6 a time series of the average price p_{ave} and the price range (p_{min} , p_{max}) is shown

⁵The price initial conditions were computed from a distribution $p_i(t=0) = p^*(1+\xi)$, where ξ is a uniform random variable with range $[0, 1]$, constraining the initial price range to $p_i(0) \in [p^*, p^*(1+\xi)]$.

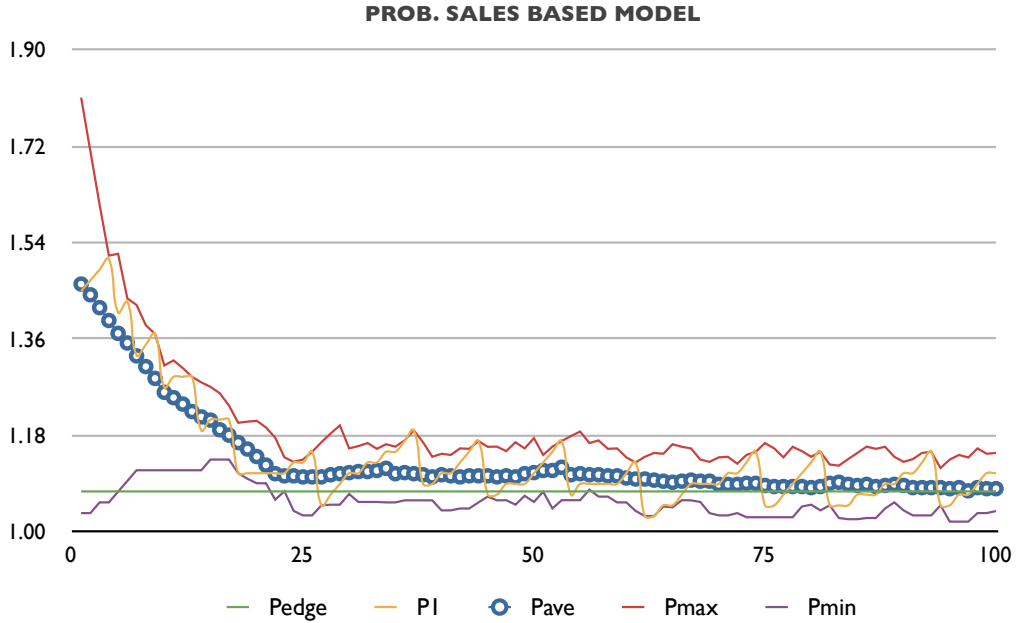


Figure 6: Probabilistic sales-based model time series. Average (p_{ave}), maximum (p_{max}) and minimum (p_{min}) price as a function of time, together with a sample trajectory (p_1) of an agent and $p_{edge} = p^* + c/N$. Parameters: $N = 10, p^* = 1, c = 0.75, \gamma_+ = 0.02, \gamma_- = 0.10, h_0 = 0.45, h_- = 0.0$.

together with the price of an individual agent. Qualitatively, the average price behavior is analogous to the one shown in Fig. 1.; however, at the individual level, the stochastic element is noticeable, with occasional drops in prices after full sales.

Starting with the above initial condition numerical simulations of the model show (a) the initial price range decreases steadily, (b) there is a gradual decrease in the average prices, and finally (c) it reaches a steady state. Qualitatively (a) and (b) are two features observed in the experiments, and are nicely captured by the probabilistic sales-based model. All 8 sessions except one (Exp #20; we will later discuss this exception) start with a large dispersion of individual prices, and a gradual convergence (which takes about 30 market rounds) around an average that initially shows a downwards trend.

In order to generate this market behavior, it matters that the agents have widely dispersed initial prices; while the cheaper suppliers agents increase steadily with no downturn, the higher priced agents find that they on average become more frequently not selling than the lower price partners. This occurs until the average price has an average price range of the order of magnitude as

a downturn ($\approx \gamma_-$)⁶ and the average frequency of not selling across agents is uniformly distributed.




Figure 7: Comparison of the average price estimate (13) in the probabilistic sales-based model against numerical simulations as a function of h_0 . Parameters: $N = 20$, $\gamma_- = 0.10$, $\gamma_+ = 0.02$, $h_- = 0$.

We can figure out the average posted price of the market, in a similar fashion as in Section 3. Consider δP_+ as the number given by the aggregation of the price changes of firms having realized full sales in the previous period. This can be approximated by: $\delta P_+ \approx N_s(h_+\gamma_+ - h_-\gamma_-)$, where N_s is the number of sellers who sell the capacity quantity (remember that, with probability h_0 the price is kept constant). On their side, the agents who sold less than capacity, $N_{ns} = N - N_s$ will reduce their price, and accumulate an aggregate price change $\delta P_- \approx N_{ns}\gamma_-$. A steady state in the average posted price would correspond to $\delta P_+ \approx \delta P_-$. Then, in that steady state

$$(N - N_{ns})/N_{ns} = N_s/(N - N_s) = \gamma_-/(h_+\gamma_+ - h_-\gamma_-) \quad (9)$$

Solving (9) for N_{ns} ,

$$N_{ns} = N \left(1 - \frac{1}{1 + \frac{h_+\gamma_+ - h_-\gamma_-}{\gamma_-}} \right) \quad (10)$$

The average posted price is related to N_{ns} by,

$$\langle p \rangle = M/(q^*N_s) = p^*N/(N - N_{ns}) \quad (11)$$

Thus,

$$\langle p \rangle \approx p_{est} \approx p^*(1 - h_- + h_+\frac{\gamma_+}{\gamma_-}) \quad (12)$$

is a good approximation for large N .

⁶Fig. 6 shows an average price range for the last half rounds of $0.107 \approx \gamma_- = 0.10$.

For a given h_- , the average price decreases with h_0 (increases with h_+): more cautious players make the average price go closer to the competitive market equilibrium. In the limit $h_0 = 1$, clearly $p_{est} = p^*$, since agents with full sales would never raise their price as expected when agents never increase the price.

The integer constraint for the numbers of agents (N_s, N_{ns}) becomes relevant for not-too-large N . Then, the average market price would be approximated by,

$$p_{est} = p^* \frac{N}{N - \lfloor N_{ns} \rfloor} \quad (13)$$

where $\lfloor N_{ns} \rfloor$ is the integer value of the number of agents with less than full sales derived from (10).

The numerical simulations of the probabilistic sales based model (8) validate the approximation (13). Iterating over 1000 time steps, taking averages over the last 500 steps, and averaging over 10 random initial conditions, Figure 7 shows that the numerical simulations, closely follow (13) as a function of h_0 . The price jumps occur at the discrete changes of N_{ns} and are of order $1/N$.

It is interesting to notice that there is a critical h_0^* , above which $p_{est} = p^*$ (see Fig. 7). Using (13) and (10), and taking h_- negligible, one can solve for the critical h_0^* , where $N_{ns} = 1$,

$$h_0^* = 1 - \frac{1}{N-1} \frac{\gamma_-}{\gamma_+} \quad (14)$$

For small N , this critical h_0^* crosses zero, so there is value of $N = N^*$ below which the only steady state corresponds to $p_{est} = p^*$,

$$N^* = 1 + \gamma_-/\gamma_+ \quad (15)$$

For the parameters of Fig. 7, $N^* = 6$.

5.1 Comparison with the experiments

The above results can now be used to compare with the experimental results obtained in Section 4. We study the model- experiment correspondence in two dimensions: i) the extent to which the average prices shown by the experiments after an initial transient can be approximated by a model-based formula, and ii) if the model can shed some light about the evolution of behavioral parameters during the course of the experiments and, in particular, about the pattern of frequencies of price responses (parameters h_-, h_0, h_+) found experimentally and shown in Figure 5.

Average prices We have estimated the values of the behavioral parameters h 's and γ 's averaging over the last 30 rounds and across players. In Figure 8 we compare the (rescaled) average market price of the last 30 rounds with the one estimated from (13). In general the estimated price is below the experimental market price, but not far from it. The larger departures from the p_{est} correspond to experiments #17 and #20, with 4 players each and large price variations across the whole session.

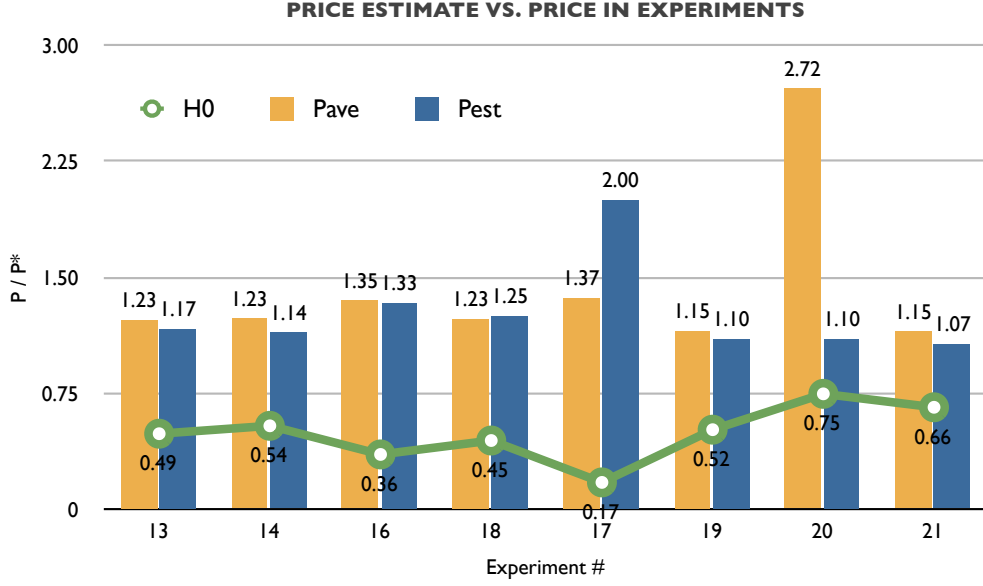


Figure 8: Comparison of the average price estimate (13) in the probabilistic sales-based model against the experimental results. Also the values of h_0 are shown. Only the results for experiment #21 and #17 are outliers. The $h_{+,-}$ and $\gamma_{+,-}$ for each experiment were estimated taking the average over the last 30 rounds. The average price of the experiment was also for the last 30 rounds. Notice that shown prices are rescaled with p^* of the experiment.

Changes in behavioral parameters during the play of the game The evidence presented in Figure 5 suggests that: (a) players seem to be more active varying the frequencies h 's than the magnitudes of price adjustment γ 's, once the initial transient is over, (b) the average probability h_0 starts close to zero and increases steadily until it reaches a level around which it fluctuates. In order to study these results, we consider a perturbation on the probabilistic sales based rules, and look at the consequences for an individual's profits of a deviation from the prevailing behavioral parameters. Given a population of players with a set of default behavioral parameters $\gamma_-^{pop}, h_0^{pop}$, we let a single 'mutant' agent play with a different γ_- or h_0 . In order to estimate the 'normal' change in performance deriving from the deviation, we aggregate over 50 numerical realizations the average benefit of the 'mutant' b_{mut} and measure this figure against the average benefit of the whole population b_{pop} .

Figure 9A shows the comparison of b_{mut} vs b_{pop} for a deviating parameter γ_- in the range $(0.05, 0.15)$. It is clear that the change would not have much effect: a 'mutant' agent would be more or less indifferent to the alternative of γ_- from γ_-^{pop} .

However, with a similar numerical simulation, average benefits are quite sensitive to changes of

the frequency $h_0 \in (0, 1)$, (see Figure 9B): here, given that the population maintains its behavioral parameters, the mutant agent will prefer to react less often to full sales, and thus raise its h_0 . The simulation suggests that the search for profits in this way can explain that in actual experiments agents change more readily h_0 rather than γ_- , and that, starting from an initial $h_0(0) \approx 0$, $h_+(0) \approx 1$, players would tend to increase their probability of not varying prices, h_0 , which would be reflected in a decline of the population average of the parameter.

This result holds when the exercise starts $h_0^{pop} = 0$, but can change when the population has a higher value of h_0^{pop} . Figure 9C shows a perturbation analysis with $h_0^{pop} = 0.70$. In this case, the mutant agent would increase its benefits by reducing h_0 . The observation corresponds with the analysis of Section 2. Here, in the initial state, the average market price with the assumed behavioral parameters would be below $p_{edge}^{pop} = p^* + c/N$; as discussed earlier, agents trying to maximize benefit would find it to their advantage to increase their price and, in this context, to react more strongly to the realization of full sales by reducing h_0 .

The perturbation analysis highlights the possibility of making the behavioral parameters subject to adaptive, profit-seeking changes. Indeed, a 'slow' drift in the behavioral parameters is clearly observed in the experiments (see Fig. 5), where in general h 's change in the order of 10-20 iterations. The dynamics in this extended system would never settle in the steady states found in the previous simulations, since the parameters would be continuously changing. In particular, when the average price is below p_{edge} , one would expect a drop the average value of h_0 , which would then increase the market price. The value p_{edge} appears then as a "barrier" below which prices would tend to "rebound", but above which agents would have an incentive to undercut their competitors (cf. Section 2)

6 Conclusions

We have studied the decentralized operation of the process of market competition in a multi-agent game which has potential counterparts in concrete economic environments and which raises interesting analytical questions, particularly because it lacks deterministic Nash equilibria and because the equilibria in mixed strategies are hard to identify and to characterize. From this perspective, we looked for means to represent the self-organization of a system that does not feature intuitively obvious states of rest, either for individual actions or for aggregate outcomes. In order to deal with the problem, we explored in a particular instance the complementarities between analytic arguments, computer simulations of the interactions of behaviorally simple artificial agents, and laboratory experiments with actual players. The results confirm the usefulness of combining those approaches, and they offer reasonably specific suggestions regarding the behavior of agents and the market outcomes in the setting under consideration.

In an environment where there is no definite or salient best strategy that agents may use systematically to maximize profits, it seems natural that they act according to some heuristic procedure with feedback from past experience. The simulations and the experiments permit a

stylized, but quite precise, specification of features of the adaptive behavior of agents. We have found that a very elementary rule (reduce the price when sales are below capacity in the previous period, increase or keep constant the price when "full sales" were realized) describes with a good approximation the price choices of experimental agents. This implies that price decisions are not conditioned directly on the history of benefits. However, within that general framework, the search for profits can be embodied in adaptive changes of the parameters of the price- setting function: the sizes of upwards and downwards price adjustments and the frequency (or probability) of price rises following full sales. Although we have not identified clear regularities in the experimental determination of these parameters, there is evidence that agents modify their behavior in the course of the game, with more aggressive moves (larger modules of price changes, higher likelihood of price rises when sales reach capacity) at the first "exploratory" stages of play, and less intense responses as individual prices converge.

In the setup we have considered, the competitive responses of agents create recognizable patterns of organization in the markets, but these do not "freeze" in stationary states. A characteristic feature in several experiments is a convergence of individual prices around a declining average. Interestingly, this qualitative market performance can be reproduced in a simulation with artificial agents using the sales- based heuristics, and parameters that vary in the form mentioned in the previous paragraph; the gradual aggregate movement results from a persistent oscillation in the prices of particular agents, who "alternate" between selling at full capacity with a rising price until a level is reached where the volume of sales drops, and leads to a fall in the price.

In any case, the downwards price convergence itself generates the conditions for a disturbance to the emergent "ordered" state. As discussed in section 3, when prices approach the competitive level, there is room for profitable deviations by agents who raise their price sufficiently. The experimental results suggest that by groping for profits, individuals can hit upon those opportunities and generate "excursions" in markets which previously were in a more or less tranquil condition. Thus, the image would one where "competitive forces" put limits to the unit margins of participating firms, but where "intermittent shakeouts" push up market prices for some periods.

The analysis can potentially be extended in several ways. Experimentally, it can be interesting to study the behavior of prices when some market parameter (like the value of M , which defines the position of aggregate demand), is shocked, or subject to a gradual change. Also, an extension to the case of durable goods, where firms post prices, decide production in advance and manage inventories, would be a step towards a representation of behavior in real markets. Regarding simulations, a natural next step would be to consider explicitly the adaptation of the quantitative characteristics of sales- based pricing rules through some profit- motivated criterion. Ultimately, this work may help to identify pricing decision rules with sufficient validation which can be used as elements of more elaborate agent- based models or, perhaps, apt to be confronted with actual data.

References

- [1] Allen, Beth and Martin Hellwig (1986). "Bertrand-Edgeworth Oligopoly in Large Markets", *Review of Economic Studies* 53, pp.175-204.
- [2] Arthur, Brian (1994). "Inductive Reasoning and Bounded Rationality." *The American Economic Review* 84 Papers and Proceedings, pp. 406-411.
- [3] Abbink, Klaus and Jordi Brandts (2005). "Price Competition Under Cost Uncertainty: A Laboratory Analysis." *Economic Inquiry* 43, 636-648.
- [4] Brown Kruse, Jamie, Stephen Rassenti, Stanley Reynolds and Vernon Smith (1994). "Bertrand-Edgeworth Competition in Experimental Markets", *Econometrica* 62, pp. 343-371.
- [5] Burgos, E. , Horacio Ceva, and Roberto P. J. Perazzo (2002), "Thermal treatment of the minority game" *Phys. Rev.* **E65** 036711.
- [6] Chamberlin, E. (1956). *The Theory of Monopolistic Competition*. Cambridge MA, Harvard University Press.
- [7] d'Aspremont, C. and J.J. Gabszewicz (1985). "Quasi-Monopolies." *Economica* 52, pp. 141-151.
- [8] Dasgupta, Partha and Eric Maskin (1986). "The Existence of Equilibrium in Discontinuous Economic Games, II: Applications." *Review of Economic Studies* 53, pp.27-41.
- [9] Edgeworth, Francis Y (1925). "The Pure Theory of Monopoly", in *Edgeworth, Papers Relating to Political Economy*, New York, Burt Franklin, pp. 111-142.
- [10] Haar, Ter (YEAR) "—"el libro de macánica estadística
- [11] Heymann, Daniel, Roberto Perazzo and A. R. Schuschny. (2004) "Learning And Imitation: Transitional Dynamics In Variants Of The BAM." *Advances in Complex Systems* 7, pp. 21-38.
- [12] Levin, Simon A. (1999) *Fragile Dominion. Complexity and the commons*, Helix Books.
- [13] Osborne, Martin and Ariel Rubinstein (1994). *A Course in Game Theory*. Cambridge MA, The MIT Press.
- [14] Smith, J. Maynard (1972) "Game Theory and the Evolution of Fighting" in: *On evolution* Edimburgh University Press
- [15] Maskin E., and J. Tirole (1988). "A Theory of Dynamic Oligopoly II: Price Competition, Kinked Demand Curves and Edgeworth Cycles", *Econometrica*, **56**, pp. 3.
- [16] Smith, V. A. (1982) "Microeconomic Systems as an Experimental Science." *American Economic Review* 72, pp. 923-955.
- [17] Tirole, Jean (1988). *The Theory of Industrial Organization*. Cambridge MA, The MIT Press.

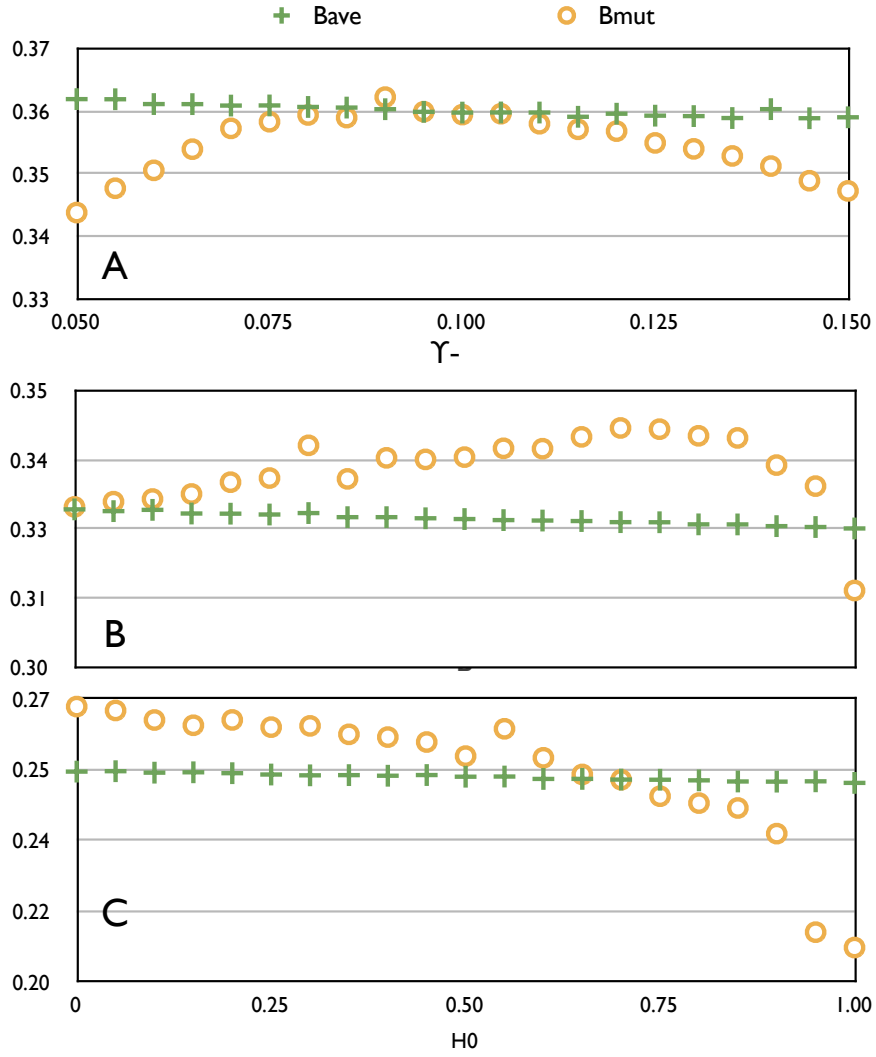


Figure 9: Comparison of probabilistic sales based model average benefit (B_{ave}), against the average benefit of single mutant player (B_{mut}) which varies: (panel A) γ_- , (panel B, C) h_0 . Panel A shows that the mutant player cannot improve the average benefit above the population's by changing γ_- . On panel B, the mutant player improves by increasing h_0 , while on panel C, the mutant improves by *decreasing* h_0 , thus limiting the values of h_0 by which a single agent can optimize its benefit. Parameters: $N = 10, c = 0.75, p^* = 1, q^* = 1$. Behavioral parameters for all players except one $\gamma_+^{pop} = 0.02, h_-^{pop} = 0, h_0^{pop} = 0$ (Panel A), $\gamma_-^{pop} = 0.10$ (Panel B, C).