



***"Behavioral Heuristics and Market Patterns in a Bertrand-Edgeworth Game"***

Por Daniel Heymann (Universidad de San Andrés & Universidad de Buenos Aires), Enrique Kawamura (Universidad de San Andrés), Roberto Perazzo (Instituto Tecnológico de Buenos Aires & Universidad de San Andrés) y Martín Zimmermann (Universidad de San Andrés).

**D.T.: N° 108**

**Marzo, 2011**

Vito Dumas 284, B1644BID, Victoria, San Fernando, Buenos Aires, Argentina  
Teléfono 4725-7053, Fax 4725-7010  
Email, [economia@udesa.edu.ar](mailto:economia@udesa.edu.ar)

# Behavioral Heuristics and Market Patterns in a Bertrand-Edgeworth Game

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

(<sup>†</sup>)Universidad de San Andrés, Depto. de Economía,  
Vito Dumas 284, (B1644BID) Victoria, Buenos Aires, Argentina,

(<sup>§</sup>)Universidad de Buenos Aires, Depto. de Economía,  
Av. Córdoba 2122, (1120AAQ), Buenos Aires, Argentina,

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

February 16, 2011

## Abstract

This paper studies price-setting behavior in a Bertrand-Edgeworth game, a traditional setup that provides a stylized representation of real-world markets (such retail markets). It explores methodological complementarities between analytical, experimental and agent-based simulation approaches. The game under analysis only admits hard-to-characterize mixed-strategy Nash equilibria. This paper uses families of heuristic rules for individual behavior to analyze the resulting market prices. A set of laboratory experiments finds that price-setting choices of agents are described reasonably well through a sales-based simple pricing rule. Average market prices tend to converge from above to *quasi*-steady states with small individual dispersion, occasionally disturbed by agents who search for profits by raising prices. Salient features of the experimental outcomes can be represented through simulations with artificial agents who apply the sales-based heuristics with parameters calibrated to the experiment.

**JEL: D43, C63, C91**

# 1 Introduction

The Bertrand-Edgeworth (BE) oligopoly, where price-setting, capacity-constrained firms compete for customers, provides a fair approximation of features in different markets, particularly in the retail sector<sup>1</sup>. The basic form of such oligopoly is very simple to describe. In a “Marshallian short-run” (where the number of firms and their physical size are given) each of the suppliers in such markets operates under more or less constant marginal costs, and is subject to a capacity constraint determined by the volume of goods they can produce or keep in store. These firms post prices, which signal their willingness to sell goods (conditional on the capacity limits) to customers who seek to buy at the lowest price available. From the demand side, all potential buyers know all posted prices and can move costlessly between suppliers.

Despite its simplicity, how this BE oligopoly actually works is still imperfectly understood. From the theoretical point of view, the BE oligopoly is a game without Nash equilibria in pure strategies, except for very special cases. The mixed-strategy equilibria are difficult to characterize, and their implementation requires a non-trivial coordination in the individual behavior of firms, which reduces their empirical plausibility as representations of actual market outcomes<sup>2</sup>. This empirical implausibility of mixed-strategy equilibria in BE oligopolies stands out in various experimental results, such as those in Kruse et al. (1994), Fonseca and Norman (2008, 2009) and Buchheit and Feltovich (2010). All these papers reject the hypothesis that the prices observed in the experiments are consistent with the predictions of their respective mixed-strategy Nash equilibria. In some cases, such as Kruse et al (1994) and Fonseca and Norman (2009), experimentally observed prices both fluctuate over time in a way resembling Edgeworth-cycles<sup>3</sup> (i.e. phases where prices rise as if driven by implicit collusion, followed by sharp “competitive” price cuts) and show a tendency to converge from above to the “competitive range”. However such prices do not get close either to alternative theoretical benchmarks like the competitive equilibrium or tacit collusion (see also Davis and Holt, 1994). Thus, the representation of market outcomes and individual behavior in the BE setting remains an open (and relevant) problem.

The comments presented above suggest a radical change in the standard methods to study BE oligopolies. First, it demands a departure from the standard behavior of the optimal mixed-strategy following firms. Instead, our approach starts from the assumption that agents follow much simpler price-setting heuristic decision rules. This assumption, in turn, naturally brings with it the introduction of the agent-based modeling tools into the BE oligopoly framework. This combination of tools implies the use of computer simulations to generate theoretical predictions

---

<sup>1</sup>Ketcham et al. (1984) remark that the generalization of pre-determined prices in retail is a 20th century development, which replaced the traditional practice of “haggling”.

<sup>2</sup>For an analysis of the mixed strategy solutions of these games, see Dasgupta and Maskin (1986), Allen and Hellwig (1986), Maskin and Tirole (1988), the appendix in Kruse et al (1994) and more recently De Francesco and Salvadori (2008). (See also Tirole, 1988, for a textbook treatment). In some special cases computation of mixed-strategy equilibria is easier, but these they correspond to very particular step demand functions (see Holt and Solis-Soberon, 1992, for a discussion).

<sup>3</sup>See Maskin and Tirole (1988) for a theoretical treatment of the Edgeworth cycles.

for each type of rule. However, such computer simulations are not sufficient to determine which heuristic rule is the most empirically relevant. As several authors from the agent-based and experimental literatures argue (such as, e.g., Duffy, 2006) the experimental method does provide such type of empirical validation for each rule.

This paper develops the above-mentioned new approach to study BE oligopolies. It considers a BE market in a simple form, as a benchmark case that may lead to various extensions. In this setting buyers are assumed to purchase a non-storable good and to be perfectly informed about the prices posted by all the firms. Consumers have identical demand functions with price-elasticity equal to one. This is equivalent to assuming that there is a single buyer with a fixed amount of “money” available for purchases, who visits stores in an order of ascending prices. Sellers produce to order, and only decide about their prices. Sales equal the minimum between the realized demand at the pre-determined posted price and the capacity limit. The parameters determining the aggregate demand, the marginal cost and the capacity limit are common knowledge, and firms also know the value of the competitive price (always strictly higher than marginal cost). However, price-setters are assumed to know only their own individual history of prices, volumes of sales and realized profits, but not those of their competitors. This last assumption, corresponding to a case of “dispersed” sellers, reduces the dimension of the set of variables on which agents can condition their price decisions, and also facilitates the search for simple characterizations of behaviors.

As a starting point, this paper considers two families of heuristic pricing rules: i) a set of *sales-based* rules, where individual prices respond in a very simple manner to past demand, by increasing or lowering the firm’s prices by exogenous amounts according to whether the volume of sales in the previous market day reached full capacity or not; and ii) *profit-gradient* rules where prices depend on the individual record of changes in profits following price changes. According to the profit-gradient rule, a firm continues or reverses its ongoing behavior (increasing or decreasing the price) looking at the sign of the change in profits realized in the immediate past.

The main results of this paper are summarized as follows. The simulations show quite different price behaviors between the two pricing rules. In particular, they show that the average price (across firms) according to the sales-based rule does not necessarily converge to the same range in the long run compared to the profit-gradient rule. In terms of the experimental evidence, the basic important result is that the data definitely favors some form of sales-based pricing. Moreover, simulations using a specific form of the sales-based rule (with parameter values calibrated using the experimental results) replicate salient stylized market patterns observed in the experiments.

These results show how extremely simple heuristics represent major aspects of actual subjects’ decisions and generate market consequences consistent with experimental evidence. This is a remarkable feature considering that the experimental agents were both advanced undergraduate and graduate Economics students, who could be expected, by their training, and their understanding of the nature of the game, to apply more sophisticated strategies if they found

them useful to increase profits. When viewing the BE oligopoly as a system of interacting agents, our results suggest that the BE markets have tendencies to “order themselves” in quasi-stationary states for average prices, but where neither individual nor market behavior become frozen. The fluctuations of observed prices also confirm the presence of effects reminiscent of Edgeworth cycles.

Regarding the behavior of individuals through time, almost all the experiments show that, at any stage, agents reduce their prices when selling below full capacity. Also, when they sell at full capacity, experimental subjects either increase or do not change the price (but almost never reduces it) with time-varying frequencies along the course of the experiments. In the initial “market days”, subjects tend to act aggressively, raising prices in practically every instance that they sell at full capacity, while the average price goes well above the competitive level. As the play proceeds, the average of prices and their dispersion decline (this matches with the finding reported by Kruse et al., 1994). In this phase, agents are found more often keeping their price constant after full capacity sales, as if fearing to lose sales otherwise. Eventually, the frequency of upward price adjustments after full-capacity sales slightly fluctuates around a value that is clearly positive but well below one.

Overall, in the last stages of the experiments, after going through the initial transient phase, the behavior of subjects is consistent with a random choice of positive or zero price changes when selling at full-capacity, while maintaining the property of cutting prices following less than full capacity sales. Using such a *probabilistic* sales-based rule to calibrate the parameters of the price-setting algorithms with the averages across the whole experimental population, the artificial-agent simulations of the model yield average long-run prices that approximate the corresponding values observed in experiments. Such simulations also generate price dynamics similar to those present in the experimental data.

The basic probabilistic sales-based rule assumes that both the probabilities of increasing prices or leaving the prices given full capacity sales and the absolute value of the adjustment are constant. Clearly, such rules on pricing decisions are not conditioned on profits. However, the time-varying frequency of zero price changes after a full-capacity sales observed in the experiments show that what those rules assume as constant parameters do in fact evolve through time. To capture this observation we modify the simulation model so that the frequency of price changes becomes a choice parameter for one (profit-seeking) “mutant” individual in a market where the other artificial agents set prices through the original fixed-parameter rule.

The simulated results of the new model indicate that, when the other firms have a high propensity for raising prices after full capacity sales, the mutant agent gains by keeping its price constant more frequently than the rest. When decreasing the propensity for raising prices after full capacity sales of the other firms there is a point where the remaining supplier gains by increasing its frequency of price rises. This result suggests a potential mechanism through which the search for higher profits could be reflected in adaptive changes in the behavioral parameters within the “sales-based family” of pricing rules, in a way that may induce a convergence to states with prices “bunched” not much above the competitive level, but subject to be disturbed

by upward price movements when the system seems to have settled in what may look as a steady state.

This paper has points of contact with the recent analytic-experimental literature on price-setting, such as Brandts and Guillén (2007), Abbink and Brandts (2005), Cason et al. 2005 (see also Cason and Friedman, 1999, 2003), Fonseca and Norman (2008, 2009) and Buchheit and Feltovich (2010). Compared to those references, a major difference of this paper is the stronger simplicity of the setup of this model, closer to the basic BE framework than what the other papers study. (The paragraphs below will make this point more precise).

More specifically, Brandts and Guillén (2007) consider a BE oligopoly with either two or three firms, in a market where buyers purchase up to a maximum amount of goods at a cutoff price, and zero at a higher price (the “box” demand schedule). In their setup, firms “produce for stock”, that is, where they must pre-determine both the posted price and the produced quantity that will be available for sale in the period. Suppliers are informed about the history of prices and sales of their competitors. The results of the paper suggest that prices are set at values substantially higher than the competitive level, and that they actually tend to increase in the course of the game. This last finding contrasts with the results obtained here, and those from Kruse et al. (1994). An interesting open question is whether those differences may be linked to the information provided to agents in each case.

Fonseca and Norman (2008, 2009) and Buchheit and Feltovich (2010) also deal with models where agents are informed of the past prices posted by their competitors, an assumption that is absent in our experiments. Those papers focus on specific features of the market behavior: respectively, the effects of varying the number of firms (from three to two, interpreted as a merger) on the market price, the consequences of exogenous changes in the capacity constraint on the pricing behavior by duopolies and triopolies, and the potential influence of sunk cost on the prices set by duopolies. Instead, this paper concentrates on exploring behavioral regularities in a basic setup, looking for stylized representations of the experimental evidences.

As for Abbink and Brandts (2005) and Cason et al. (2005), they study models where firms do not have capacity constraints. The first of these papers focuses on the behavior of firms with random unit costs, while Cason et al. (2005) consider a case of zero marginal costs and no limits to the quantities produced in a setting where buyers choose search strategies. The latter define a set of potential heuristics (in their case, profit-based rules that use the information about prices posted by other firms) as benchmarks to simulate market outcomes and to compare them with experimental results. This paper, instead, uses a closely related general approach, in the context of another market setup. Our analysis confirms the existence of useful complementarities between analytical, agent-based and experimental approaches to the study of markets of economic interest.

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  $M > N c q^*$ , so that the desired value of spending by buyers exceeds the cost of production of the capacity volume of output.

Realized profits 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:

$$\Pi_i \equiv q_i(p_i - c) \quad (1)$$

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

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

At that price every firm sells all its potential output.

The next claim can be easily verified:

**Claim 1.** *The competitive outcome is not a pure strategy Nash equilibrium.*

*Proof.* Let  $p_i = p^*$ ;  $\forall i$ . The individual profit of a firm is  $\Pi^* = q^*(p^* - c) = (M/N)(1 - c/p^*)$ . 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  $\Pi_j = (p_j - c) q_j = p^* q^* (1 - c/p_j)$ . 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:

**Claim 2.** *no price  $p'$  set by all firms  $p_i = p'$ ;  $\forall i$  can be a pure-strategy, symmetric Nash equilibrium.*

*Proof.* The proof is divided in three cases:

- 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, or capturing the whole “residual demand”  $R = M - (N - 1)p'q^* = [Np^* - (N - 1)p']q^*$  (which is positive in this case) after all the other firms have sold at their capacity levels, by raising its price as to cut production costs to near zero. The condition for preferring a lower price  $p' - \varepsilon$  (for  $\varepsilon \rightarrow 0$ ) is:

$$\Pi_{p'-\varepsilon} \equiv (p' - c)q^* > [Np^* - (N - 1)p']q^* \equiv \Pi_\infty \quad (3)$$

equivalent to:

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

- 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

$$p_{edge} \equiv p^* + c/N \quad (5)$$

□

The results above illustrate a general feature of the problem: the lack of Nash equilibria in pure strategies, a feature well-known in the oligopoly literature. This problem suggests that individual agents may be involved in possibly inconclusive searches for profitable strategies. The profit implications of alternative price decisions will be affected by properties of the state of the system, such as the following:

- The range of prices above the competitive equilibrium such that, if  $N - 1$  firms choose a price in the set, the  $N$ -th one would find profitable a “large” upwards deviation, gets smaller as the number of firms increases, since it is determined by the interval  $[p^*, p^* + \frac{c}{N}]$ . For a large  $N$ , a small rise in the average price set by competitors may change substantially the advantageous move of the remaining supplier, from a sharp price increase to a small reduction.



- If  $N - 1$  firms set a price  $p' < p^* + \frac{\varepsilon}{N}$  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  $\varepsilon$ , would lead to lower profits, while profits 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 by very little<sup>4</sup>. A behavior of this type appeared to be relevant in market experiments (see Section 4 below).

### 3 Price behavior based on two simple heuristics

The previous discussion has shown that the profit opportunities for individual price setters vary dramatically with the decisions of the other firms, so a Nash-equilibrium in pure strategies does not exist. In contrast to this standard analysis of oligopolies, a natural alternative conjecture is that firms set prices through a simple decision rule (a heuristics). This section explores the specification and properties of two types of heuristics: one which conditions prices on past volumes of sales, and another where prices depend on the individual price-profit 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 either lower the price  $p_i$  whenever the realized volume does not reach full capacity  $q^*$ , or increase the price if sales are at their maximum feasible amount. For each agent  $i$ , this pattern of behavior can be described as:

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

where  $t$  denotes the “market day” and  $(\gamma_+, \gamma_-) \gg 0$  are behavioral parameters which determine the magnitude of upwards and downwards price changes<sup>5</sup>.

---

<sup>4</sup>This 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 1991 for a macro application where the effect is relevant).

<sup>5</sup>It is possible to use random adjustments with price-shifts  $\gamma_+$  and  $\gamma_-$  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 this discussion we treat these parameters as fixed.

A system of firms that follow this pricing rule would have an aggregate steady state, with constant average total sales  $\bar{q}$  and selling price  $\bar{p}$ . Consider the case where all agents share the same behavior parameters, and a situation where  $N_s$  firms sell at full capacity, and therefore raise their price, while  $N - N_s$  reduce theirs after realizing sales below capacity. For sufficiently large  $N$  so that “small number” effects can be ignored (see the discussion in Section 5), the steady state would be characterized by a similar value of total price increases and decreases:  $\gamma_+ N_s \approx \gamma_- (N - N_s)$ , which implies that the fraction of agents selling  $q^*$  would be  $N_s/N \approx \gamma_-/(\gamma_+ + \gamma_-) = 1/(1 + \gamma_+/\gamma_-)$ . Therefore the average sales  $\bar{q} = q^* N_s/N$  (for large  $N$ ):

$$\frac{\bar{q}}{q^*} \approx \frac{1}{1 + \gamma_+/\gamma_-} \quad (7)$$

The average sale price  $\bar{p}$  in that aggregate steady state can be expressed as a function of the competitive equilibrium price  $p^*$ , from the condition  $Np^*q^* = M \approx N_s\bar{p}\bar{q}$ :

$$\frac{\bar{p}}{p^*} \approx 1 + \gamma_+/\gamma_- \quad (8)$$

The properties of this state are intuitively simple. Take for example the limit where  $\gamma_+ < \gamma_-$ . Individual prices change as if placed in a “conveyor belt”: an individual increases its price in small steps of size  $\gamma_+$  until a point is reached where sales fall below full capacity, leading to a price reduction  $\gamma_-$ . In an aggregate steady state, the number of such steps in the cycle would be equal to  $\mu = \frac{\gamma_-}{\gamma_+}$ . Then, on average, every agent does not sell once every  $\mu$  steps, so that average sales would be  $\bar{q} \approx (1 - 1/\mu)q^* = (1 - \gamma_+/\gamma_-)q^*$ , a number of an order of magnitude as that implied by (7). As another example, with  $\gamma_+ = \gamma_-$ , the average price becomes  $\bar{p} = 2p^*$ , and on average agents sell all their capacity output during half of the time. Therefore their average profit  $(p^* - c/2)q^*$  is greater than the one that could be obtained in a steady competitive equilibrium  $(p^* - c)q^*$ . The heuristics would induce some sort of implicit collusion among sellers, which would be sustained because individuals do not undercut their competitors by setting a price slightly below  $\bar{p}$ , which would generate full sales without sizable sacrifice in terms of markup.

Numerical simulations show that, starting from randomly chosen initial prices, a system of agents following the sales-based rule does tend towards a steady state similar to that described by (8). Figure 1 shows a representative time series of the average price of the system, together with the minimum and maximum price at each time step, and the time pattern of the price of a single firm displaying the previously described “conveyor belt” dynamics. The initial condition for individual prices was drawn from a uniform distribution in the interval  $[p^*, 2p^*]$ , implying, in this case, that the average initial price was above  $\bar{p}$ . It can be seen that there is a gradual convergence towards the steady state, also observed when the average initial price is smaller than  $\bar{p}$ . Thus, this simple heuristics generates a movement to a “dynamical equilibrium” where the system is permanently changing its detailed configuration as individual agents vary their prices, while the average price and volume of sales gravitates towards a steady state. It may be noted that the average steady state price does not depend directly on the number  $N$  of firms, but only on the relative size of the adjustment parameters. However, the competitive behavior of




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:  $N = 10, p^* = 1, c = 0.75, \gamma_+ = 0.02, \gamma_- = 0.10$ .

agents may be incorporated in their choice of parameters, or in their use of a more sophisticated decision rule; an example is discussed below.

### 3.2 A rule based on price-profit gradients

This section introduces a price-setting rule which considers explicitly the observed performance of individual past profits and allows for price adjustments of variable size. Here, firms adapt their posted price by maintaining the direction of price change if profits rose in the previous period, and reversing course if profits fell or decreases. Allowing agents to have a two-day memory of sales and prices, the rule would make prices vary according to the one-period gradient of the profit function. A random element allows for a behavior of search if profits have not varied, in a way biased towards price reductions if the constancy of profits derives from zero sales in this period and the past one:

$$p_i(t+1) = p_i(t) + \sigma \Delta(\Pi_i(t)) \text{sign}(\Delta(p_i(t))) + \eta \quad (9)$$

Here  $t$  labels the 'market day',  $\sigma > 0$  is an adjustment parameter,  $\Delta(X(t)) \equiv X(t) - X(t-1)$  and  $\eta$  is a random variable, uniformly distributed in either (a)  $[-\epsilon, +\epsilon]$  if sales were positive in

$t$  or in  $t - 1$ , or (b) uniformly distributed in  $[-\epsilon, 0]$  if there were no sales in  $t$  and in  $t - 1$ . The parameter  $\epsilon$  is a small number which determines the “search space” for the agent.




Figure 2: Maximum-profit-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  $p_{edge} = p^* + c/N$ . Notice that the average price fluctuates around  $p_{edge}$ . Parameters:  $N = 10, p^* = 1, c = 0.75, \sigma = 0.4, \epsilon = 0.05$ .

This rule generates an asymmetric behavior in price increases or decreases. Whenever the firm realizes higher profits by augmenting prices, the response is to gradually raise the price, where the size of the increase is proportional to the relatively small profit increase resulting from the higher markup. By contrast, when the price becomes so high that sales go from full-capacity to zero, the abrupt change in profits results in a large price reduction. This would be analogous to the “conveyor belt” feature of the sales-based rule with larger parameters of price decrease ( $\gamma_-$ ) than increase ( $\gamma_+$ ).<sup>6</sup> An additional feature of this pricing strategy (not present in the sales-based rule discussed before) is the trial-and-error search behavior concerning the sign of price changes implicit in the additive noise term  $\eta$  in (9), whenever the change in profit is small:  $\Delta(\Pi_i(t)) \lesssim \epsilon/\sigma$ . Thus, firms may test price reductions before ‘hitting’ the no-sales region.

Numerical simulations using this rule display a convergent dynamics where the average price converges (from above or below depending on initial conditions) to a steady state, which,

<sup>6</sup>However, the analogy is not complete. A built-in feature of the profit- gradient heuristics, which makes it differ substantially from the sales-based rule, is that, after a drastic price reduction after zero sales which makes sales recover, the next step would be once again a downward shift.

however, cannot be characterized easily as in the case of the sales- based heuristic level with this rule is not known. Figure 2 shows a simulated time series starting from random initial conditions in the range  $[p^*, 2p^*]$ , where after a short transient, the system settles in a “dynamic equilibrium” where now individual firms show a “conveyor belt” behavior disturbed by noisy “searches”.

The simulations with the profit-gradient rule result in average market prices that fall when the number of firms  $N$  becomes larger. The behavior of agents induces here price choices that tend to diminish over time the quantity of firms with zero sales, thus tending to get the system in states where the average price does not exceed  $(N/N - 1)p^*$ . Also, as indicated previously, price increases by an individual firm would cause a reduction in profits when the prices of competitors are concentrated around a level higher than  $p_{edge}$ , where  $p_{edge}$  is defined by

$$p_{edge} \equiv p^* + c/N \quad (10)$$

The simulated results indicate that the average price typically converges to a level below  $p^* \cdot (\frac{N}{N-1})$ , implying that all firms make a positive sale, while mean prices do not fall under  $p_{edge}$ . 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 in the simpler sales-based rule with given price-adjustment parameters. Numerical simulations with few firms ( $2 < N < 6$ ) show as a more or less common feature instances where more than one firm simultaneously increases prices, starting an implicit collusion phase which lasts for a few time steps, resembling an Edgeworth cycle. The amplitude of the excursion diminishes as the number of players increases.

## 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 using computer simulations. However, determining which of the two families of rules is the empirically most relevant is not obvious. To do this, the criterion used in this paper is to run experiments where actual players determined prices and competed in the game in search for profits. This section introduces the experimental setup and reports the main results from those experiments. Also, 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 webpage<sup>7</sup>. Except for two experiments, the subjects in all the others were either graduate or advanced undergraduate students in

---

<sup>7</sup>Visit the home page: <http://www.elautomataeconomico.com.ar/>

Economics who knew beforehand the type of game that they would be playing. The instructions for the game are reproduced in Appendix 1. 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 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 subjects posted prices through the interactive webpage, after which 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 profits in the last five market days. Of the nine sessions that we report here, in five of them participants received pecuniary rewards contingent on their performance in generating realized profits. Three different set of parameters were used:  $E_A : c = 0.9, p^* = 2.22, q^* = 3.0$  or  $E_B : c = 2.15, p^* = 2.45, q^* = 3.0$  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 A comparison of simple heuristic rules

We present here some basic statistical tests that allow to check the experimental relevance of the simple decision rules postulated before. Table 1 summarizes the results obtained in regressions run for the price changes of all agents in all the experiments. The sales- based rule is represented directly by estimates of the average price variation decided by a player contingent on full- capacity sales ( $\gamma_+$ ), and the average price variation when sales are lower than  $q^*$  ( $\gamma_-$ ). It was verified that, for most agents, there is no significant difference in behavior when  $q = 0$  and when  $0 < q < q^*$ . Practically all agents show a positive average price response after realizing full-capacity sales and a negative reaction to less than full sales (Columns I and II). For a substantial majority of individuals, those parameters are significantly different from zero (Columns III and IV). Regressions over the whole population of agents in each game produce highly significant coefficients of the appropriate sign (Table 6). By contrast, in the analysis corresponding to the tests of the profit-based heuristics, only a small minority of agents behave in a way that can be approximated by the formula in a statistically significant way. These results clearly favor the sales- based heuristics over the one that uses explicitly the profit gradient in the previous step of play.

The qualitative patterns of the connection between price variations with past sales are similar when either the full sample of plays or the last rounds in the game are considered. However, the

Experiment	N	$\gamma_+$ (I)	$\gamma_-$ (II)	$\gamma_+^*$ (III)	$\gamma_-^*$ (IV)	<i>MaxProfit</i> (V)
13	12	0.91	1.00	0.58	0.83	0.25
19	11	1.00	1.00	0.82	0.91	0.36
22	7	1.00	1.00	0.57	0.71	0.00
23	7	1.00	1.00	0.71	0.71	0.14
16	6	1.00	1.00	1.00	1.00	0.00
18	6	1.00	1.00	0.66	0.83	0.00
21	6	1.00	1.00	0.66	1.00	0.00
17	4	1.00	1.00	1.00	0.50	0.00
20	4	1.00	1.00	0.75	0.75	0.00
14	3	1.00	1.00	1.00	1.00	0.33

Table 1: OLS estimation summary for each experiment. Columns (I) and (II) report the fraction of agents in each game with appropriate signs (positive for  $\gamma_+$  and negative for  $\gamma_-$ ) irrespective of statistical significance. Columns (III) and (IV) report the fraction of agents in each game with estimated coefficients significant at 10% and appropriate signs for each specification of the pricing behavior strategy. The dependent variable is the change in price in the next iteration; the independent variable is a constant. Therefore, the estimations are  $\Delta p_{t+1} = \gamma_+$  over the subsample that sold completely ( $q = q^*$ ) and  $\Delta p_{t+1} = \gamma_-$  over the subsample with partial sales ( $q < q^*$ ). For the profit-based algorithm we estimated for each agent  $\Delta p_{t+1} = \alpha + \Delta b_{t-1} \text{sgn}(\Delta p_{t-1})$ , where the regressors are a constant  $\alpha$ , and an interaction between lagged change in profit,  $\Delta b_{t-1}$  and the sign of lagged change in price  $\text{sgn}(\Delta p_{t-1})$ . Column (V) shows the fraction of agent with coefficients significant at 10%.

quantitative value of the parameters shows substantial differences in both cases (Table 5 and Table 6). Both the sizes of the adjustment coefficients and the ratios of the responses in upwards and downwards directions (the absolute value of  $\gamma_+/\gamma_-$ ) tend to fall as the game progresses which suggests a more “conservative” attitude over time. This behavior can be identified also in regressions where price adjustments given full sales or less than full sales are represented as a function of “time” (rounds of play), which show declining intensities of response.

### 4.3 Behavioral strategies and market outcomes

Given that the results in table 1 discards the profit-gradient heuristics in favor of the sales-based heuristics model, it is now important to analyze whether the simplest version of that model presented in Section 3 completely described the behavior observed in the experimental data, or else whether some extension of it does this job. This analysis is carried over through the computation of empirical conditional probabilities  $P(\Delta p_{t+1}; q_t)$ , which in fact measure the frequency of price changes of a certain sign ( $\Delta p_{t+1}$ ) given the volume of sales in the previous period ( $q_t$ ). The cells in the transition matrix presented below measure the frequency 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  $\Delta p_{t+1} < 0$  in round of play  $t + 1$ , 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 2 shows such conditional probability for an experiment. A salient characteristic of the matrix is that the experimental behavior reproduces an important feature of the sales-based rule, as is the almost unanimity of price reductions following less than full capacity sales. However, a feature that now appears, and remained hidden in the regressions reported before, is that, in contrast with the very elementary rule originally proposed, full sales do not systematically lead to price increases: in a sizable number of cases, agents maintain their price, or sometimes even decide price reductions. In what follows, define 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 the 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. To check the relevance of including sales in the previous round  $q_{t-1}$ , we measure the probability of changing prices conditional on  $q_t$  and  $q_{t-1}$ .

Table 3 shows a summary of experimental results, indicating for each session the number of agents  $N$ , the number of market days ( $T$ ) and the ratio  $p_{edge}/p^*$  defined by the parameters of the experiment ( $M, N, q^*c$ ), together with the observed behavioral parameters (frequencies of conditional price responses and modules of price changes:  $\xi, h_+, h_-, h_0, \gamma_+, \gamma_-$ ). The table also indicates the value of the posted price averaged over agents and rounds of play,  $\langle p \rangle$ , relative to  $p_{edge} = p^* + c/N$ . All the behavioral parameters reported are based on observations for the second half of market days in the experiment, to focus on the performance after the system has gone through the initial transitional phase. Given the statistical significance of all the



	$\Delta p_{t+1} < 0$	$\Delta p_{t+1} = 0$	$\Delta p_{t+1} > 0$	# 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

Table 2: 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 zeros  $q_t = 0$ . The figures correspond to the last 50 rounds of experiment #18 with  $N = 6$  players and  $T = 100$  periods of play.

Exp. #	N	T	$\xi$	$h_-$	$h_0$	$h_+$	$\langle \gamma_+ \rangle$	$\langle \gamma_- \rangle$	$\langle \gamma_+ / \gamma_- \rangle$	$\langle p \rangle / p^*$	$p_{edge} / p^*$
13 : $E_A$	12	50	0.92	0.07	0.50	0.43	0.043	-0.079	0.684	1.209	1.034
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
22 : $E_B$	7	50	1.0	0.04	0.19	0.77	0.016	-0.039	0.442	1.380	1.125
23 : $E_B$	7	45	0.96	0.05	0.45	0.50	0.291	-0.752	0.763	1.146	1.125

Table 3: Main behavioral parameters and price estimation comparison. Notation of the Table:  $\xi = P(\Delta p_{t+1} < 0; 0 \leq q_t < q^*)$ ,  $h_- = P(\Delta p_{t+1} = 0; q_t = q^*)$ ,  $h_0 = P(\Delta p_{t+1} = 0; q_t = q^*)$  and  $h_+ = P(\Delta p_{t+1} > 0; q_t = q^*)$ , and  $\langle \gamma_+ \rangle$ ,  $\langle \gamma_- \rangle$ : parameters computed taking into consideration the last half rounds of play.

estimated frequencies, the first conclusion is that the behavior of experimental subjects is well approximated by a Markovian process, without a visible effect of sales before the current period.

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

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

As shown in Table 2, the conditional probability of price decreases after less than full-capacity sales is statistically not different than one. This value corresponds very closely to one of the two major features of the sales-based heuristics discussed in the previous section: if an agent does not sell  $q^*$  the rule indicates to decrease the price with certainty. In contrast, full-capacity sales

tend to induce either price increases in a sizable proportion of instances (as in that decision rule), or no change in prices, both happening with strictly positive frequencies (see columns  $h_0$  and  $h_+$  in Table 3). There were some cases (around 13% averaging all periods, agents and experiments) in which prices were reduced following full-capacity sales.




Figure 3: Average, maximum and minimum offer price (rescaled with  $p^*$ ) time series. Also shown is the Edgeworth price  $p_{\text{edge}} = p^* + c/N$ . Experiment #21.

#### 4.3.2 Time-series patterns of key variables-parameters

The model with sales-based heuristics assumes that the absolute value of the price increase after selling at full capacity and that of the price decrease if selling less than at full capacity are both constant. Also, the transition matrix computed in the last subsection showed frequencies of price changes that are not statistically different from constants. However, the time-series data generated in the experiments also reflect much richer dynamics for variables like these two than




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 #18.

what the analysis done so far can reveal. This subsection focuses on some time-series properties for these variables.

**Sizes of price changes ( $\gamma_+, \gamma_-$ ).** Measuring the average price change  $\langle |p_t - p_{t-1}| \rangle$  time series over all players shows a rapid decrease in the average module, approaching a small value, although with fluctuations around it. This property 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 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 6).

In all but one of the experiments, the average module of price increases was smaller than that of price decreases:  $\langle \gamma_+ / \gamma_- \rangle < 1$  (see also Table 3). 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  $0.5 < \langle \gamma_+ / \gamma_- \rangle < 1$ . An interesting regularity was that, for all agents and all experiments,  $\langle \gamma_+ / \gamma_- \rangle < 1$  in the last 25 rounds of play. Also, the ratio  $\langle \gamma_+ / \gamma_- \rangle$  averaged across players falls substantially when comparing early and late stages.




Figure 5: 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.

**Frequencies of price increases following full sales.** Figure 6 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  $\hat{h}_+ \approx 1$ ,  $\hat{h}_0 \approx 0$ , and become more cautious later on:  $\hat{h}_0$  increases steadily, while  $\hat{h}_+$  decreases correspondingly (see Figure 6). After this initial trend,  $\hat{h}_0$  and  $\hat{h}_+$  tend to show fluctuations around a level, suggesting that agents do not stop searching for better strategies in a context where there is no simple best response rule.

#### 4.3.3 Market-level variables: average prices tend to decrease and reduce variability through time, but do not properly converge to a steady state.

The interaction between experimental subjects yields a average market price dynamics with identifiable patterns of organization, and a convergence to a state where individual and average prices show small variations. However, this "approximate" steady state is not an actual position of rest for the system, but is subject to disturbances. The dynamics observed in several experi-




Figure 6: (top) Time series of  $\hat{h}_-(t)$ ,  $\hat{h}_0(t)$ ,  $\hat{h}_+(t)$ , measured by aggregating and averaging the last 5 time steps for each  $t$  over those agents having full sales on Experiment #21.

ments (and similarly to what was found by Kruse et al [18]), initially shows high average prices and a large dispersion between the offers posted by different agents, with a gradual downwards trend and a substantial decrease in price dispersion. This feature corresponds to a behavior marked by a first period of intensive learning with agents who start with wide differences in their strategies and react quite strongly to the outcomes they observe, although they share the knowledge of the basic structure of the game. In that situation, the capacity constraints of low-priced suppliers allows high-priced competitors to make positive sales for some time 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 moves forward, 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 “upward excursions”. These perturbations did occur in seven out of nine experimental sessions, and tended to start with average sale prices above the value  $p_{edge}$ , that is, average prices above the suggested critical level defined in 10.

On the other hand, in seven out of nine 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  $p_{edge}$ . In this respect, it would be interesting to perform longer experiments to measure statistics on the “price excursions”, although this is clearly left for future research.

The idiosyncratic behavior of some agents can be experimentally relevant, particularly in small groups. For example, in experiment #20 reported in the Figure 5, the aggregate price




Figure 7: 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:  $N = 10, p^* = 1, c = 0.75, \gamma_+ = 0.02, \gamma_- = 0.10, h_0 = 0.45, h_- = 0.0$ .

was decisively influenced by the aggressive price increases of one subject who was able to obtain high profits. This experiment was special in the sense 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 so.

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 the matter.

## 5 Matching the experimental evidence: a *probabilistic* sales-based model.

The previous section showed that the simplest sales-based rule presented in section 3 does not perfectly fit the empirical transition matrix computed out of experimental data. However, that

simplest rule can be easily extended to do that job. This section introduces the extension in some detail, and comments on the performance of the simulations of the model compared with the experiments.

## 5.1 The model and its properties from the simulations

The central change to the simplest 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:

$$\begin{aligned} p_i(t+1) &= p_i(t) - \gamma_- \quad \text{if } q_i < q^* \\ p_i(t+1) &= p_i(t) + \begin{cases} \gamma_+ & \text{with prob. } h_+ \\ 0 & \text{with prob. } h_0 \\ -\gamma_- & \text{with prob. } h_- \end{cases} \quad \text{if } q_i = q^* \end{aligned} \quad (11)$$

( $h_- + h_0 + h_+ = 1$ ), where  $h_- \ll h_+, h_0$  (as seen from the experimental measurements). The dynamics of the model will then depend on the behavioral parameters ( $\gamma_+, \gamma_-, h_0, h_+$ ), and the parameters of the game ( $N, M, c, q^*$ ). As with the other versions, the analysis of this model can only be done through computer simulations, starting from a random initial condition across agents<sup>8</sup>, who share the same behavioral parameters.

Figure 7 shows a time series of the average posted price,  $p_{\text{ave}}$ , and the price range ( $p_{\min}, p_{\max}$ ) together with the price of one 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-capacity sales.

Starting with the initial condition indicated above, numerical simulations of the model show that: (a) the initial price range decreases steadily, (b) there is a gradual decrease in average posted and selling prices, and, (c) the system approaches what looks an aggregate steady state. Qualitatively (a) and (b) are two features observed in the experiments, and are well captured by the probabilistic sales-based model. All nine 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. To generate this market behavior, it matters that the agents have widely dispersed initial prices. While the cheapest suppliers raise their offers, the higher-priced agents find that they still are unable to sell, until the range of posted prices has an order of magnitude similar to the parameter of price reductions ( $\approx \gamma_-$ )<sup>9</sup>. Then the condition of not selling is more or less uniformly distributed among agents.

It is now possible to figure out the average posted price of the market, in a similar fashion as in Section 3. Consider  $\Delta P_+$  as the aggregate value of the changes in prices decided by

<sup>8</sup>The price initial conditions were computed from a distribution  $p_i(t=0) = p^*(1 + \xi)$ , where  $\xi$  is a uniform random variable with range  $[0, 1]$ , constraining the initial price range to  $p_i(0) \in [p^*, p^*(1 + \xi)]$ .

<sup>9</sup>Fig. 7 shows an average price range for the last half rounds of  $0.107 \approx \gamma_- = 0.10$ .




Figure 8: Comparison of the average price estimate (16) in the probabilistic sales-based model against numerical simulations as a function of  $h_0$ . Parameters:  $N = 20$ ,  $\gamma_- = 0.10$ ,  $\gamma_+ = 0.02$ ,  $h_- = 0$ .

firms with full-capacity sales in the previous period. This expression 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 the capacity volume,  $N_{ns} = N - N_s$  will reduce their price, and accumulate an aggregate value of price reductions  $\Delta P_- \approx N_{ns}\gamma_-$ . At steady state one expects that the average posted price would correspond to a condition  $\Delta P_+ \approx \Delta P_-$ , assuming that  $\Delta P_+ > 0$ , so  $h_- \ll \gamma_+h_+/h_-$ . Taking that as given, we have:

$$(N - N_{ns})/N_{ns} = N_s/(N - N_s) = \gamma_-/(h_+\gamma_+ - h_-\gamma_-) \quad (12)$$

Solving (12) for  $N_{ns}$ ,

$$N_{ns} = N \left( 1 - \frac{1}{1 + \frac{h_+\gamma_+ - h_-\gamma_-}{\gamma_-}} \right) \quad (13)$$

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

$$\langle p \rangle = M/(q^*N_s) = p^*N/(N - N_{ns}) \quad (14)$$

Thus,

$$\langle p \rangle \approx p_{est} \approx p^*(1 - h_- + h_+\frac{\gamma_+}{\gamma_-}) \quad (15)$$

would be a good approximation for large  $N$ .

For a given  $h_-$ , the average price decreases with  $h_0$  (increases with  $h_+$ ) meaning that *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-capacity sales would never raise their price as expected.

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,

$$p_{est} = p^* \frac{N}{N - \lfloor N_{ns} \rfloor} \quad (16)$$

where  $\lfloor N_{ns} \rfloor$  is the integer value of the number of agents with less than full sales derived from (13).

The simulated results with artificial agents following the probabilistic sales based model (9) validate the approximation (16). Iterating over 1000 time steps, taking averages over the last 500 steps, and averaging over 10 random initial conditions, Figure 8 shows that the numerical simulations, closely follow (16) 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. 8). Using (16) and (13), and assuming  $h_-$  to be negligible, one can solve for the critical  $h_0^*$ , where  $N_{ns} = 1$ ,

$$h_0^* = 1 - \frac{1}{N-1} \frac{\gamma_-}{\gamma_+} \quad (17)$$

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

$$N^* = 1 + \gamma_-/\gamma_+ \quad (18)$$

For the parameters of Fig. 8,  $N^* = 6$ .

## 5.2 Probabilistic sales-based heuristics and the experimental evidence

The results from the simulations discussed in the last subsection can now be contrasted with the experimental results obtained in Section 4. Two dimensions of analysis are relevant in the model-experiment correspondence: i) the extent to which the average posted prices shown by the experiments after an initial transient can be approximated by a sales-based rule, 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 6.

**Average prices.** We have estimated mean values of the behavioral parameters  $h$ 's and  $\gamma$ 's for each experiment by averaging over the last 30 rounds and across players. Figure 9 compares




Figure 9: Comparison of the average price estimate (16) in the probabilistic sales-based model against the experimental results. Also the values of  $h_0$  are shown. Only the results for experiment #20 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.

the (rescaled) average market price of the last 30 rounds with the one estimated from (16). For most of the observations the estimated price is below, but not far from, the experimental market price. 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.

**Changes in behavioral parameters during the play of the game: a model with a mutant player.** The experimental evidence presented in Figure 6 suggests that: (a) when adjusting behavioral parameters, players seem to be more active varying the frequencies  $h$ 's rather than the magnitudes of price adjustment  $\gamma$ '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. The probabilistic sales-based heuristics model, instead, assumes that the probability of increasing the price or maintaining it after a full-capacity sales is constant. Thus, that version cannot really capture the time evolution for such behavior of probabilities. However, a variant of the probabilistic model may explain this dynamics.

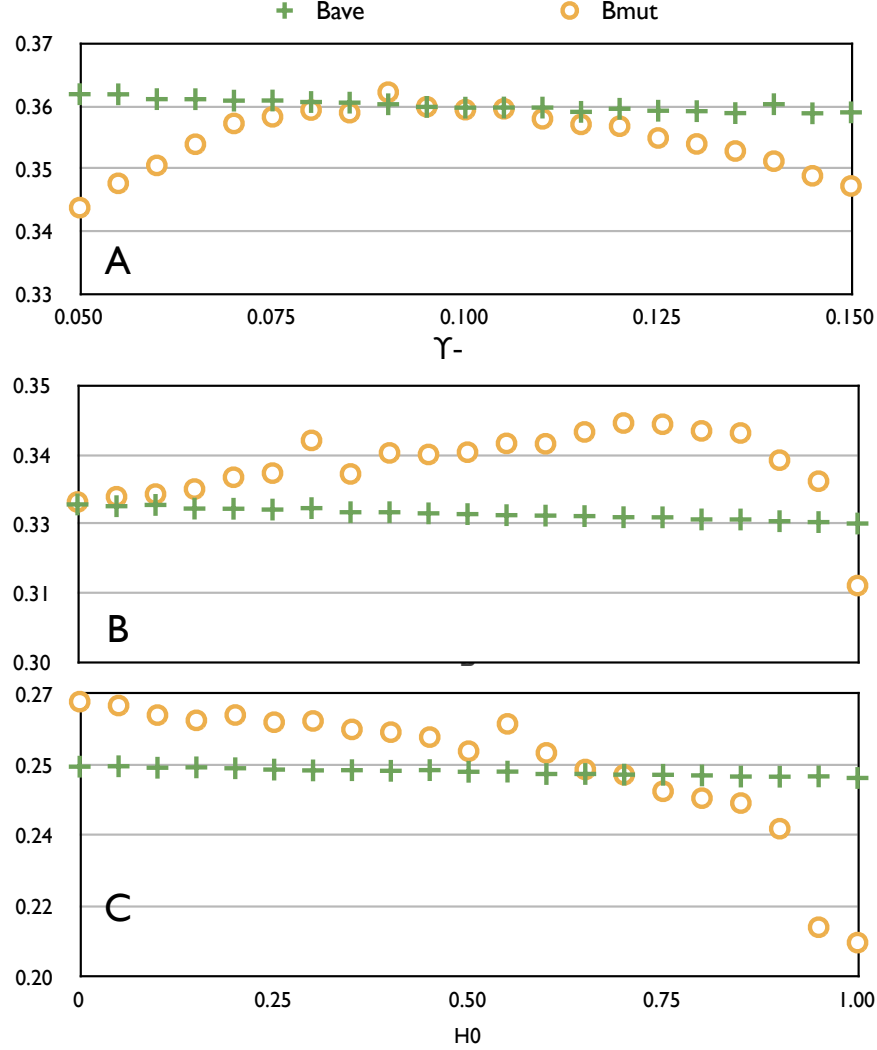


Figure 10: Comparison of probabilistic sales based model average benefit ( $B_{ave}$ ), against the average benefit of single mutant player ( $B_{mut}$ ) which varies: (panel A)  $\gamma_-$ , (panel B, C)  $h_0$ . Panel A shows that the mutant player cannot improve the average benefit above the population's by changing  $\gamma_-$ . 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).

Consider a homogenous population of players with default behavioral parameters  $\gamma_+^{pop}, \gamma_-^{pop}$ ,

$h_-^{\text{pop}}$ ,  $h_0^{\text{pop}}$ ,  $h_+^{\text{pop}}$ , and a single “mutant” player which is allowed to vary one of its behavioral parameters (either  $\gamma_-$  or  $h_0$ ). Consider how the profit realized by the mutant player compares with the average profit of the homogenous set of remaining players. Numerical simulations were performed, averaging 50 simulations, to measure the average profit of the mutant  $\Pi^{\text{mut}}$  and that of the rest of the population  $\Pi^{\text{pop}}$ , and repeated the exercise for various values of the population behavioral parameters.

Figure 10A shows the comparison of  $\Pi^{\text{mut}}$  vs  $\Pi^{\text{pop}}$  using as deviating parameter  $\gamma_-$  in the range  $(0.05, 0.15)$ . The change made by the mutant player does not seem noticeably beneficial. However, analogous numerical simulations but using as deviating parameter  $h_0 \in (0, 1)$  shows the following result. Figure 10B depicts a population of players with  $h_0^{\text{pop}} = 0$ . In this case, the mutant agent would find it profitable to increase  $h_0$  and react less often to realized capacity sales, which is precisely the observed pattern of behavior in experiments. This result suggests that in principle players *may find beneficial to act more cautiously (i.e. increase  $h_0$ ) whenever the population is reacting strongly with full sales ( $h_0^{\text{pop}} \approx 1$ )*.

If all players follow this pattern, the population’s  $h_0^{\text{pop}}$  would also increase. We now study the case of a parameter deviation starting from  $h_0^{\text{pop}} > 0$ . The simulations show that the pattern shown in Figure 10B is maintained for increasing  $h_0^{\text{pop}} > 0$ , up to a point close to  $h_0^{\text{pop}} = 0.70$ . Outcomes for this population parameter are displayed in Figure 10. Here, it is seen that the mutant player would increase its profits by *reducing  $h_0 < h_0^{\text{pop}}$* .

This observation corresponds with the analysis of Section 2 on price behavior. As the population parameter  $h_0^{\text{pop}}$  was increased, the market would settle down at a lower average price. Given the set of parameters, for a value  $h_0^{\text{pop}} = 0.70$ , the average market price coincides with  $p_{\text{edge}}$ . Thus, an individual player would increase profits by becoming more aggressive, and raising its prices as the average population price approaches  $p_{\text{edge}}$ . This may be achieved by reducing the parameter  $h_0$  of the agent, which implies a more frequent response of increasing prices after experiencing full-capacity sales.

The perturbation analysis highlights the possibility of making the behavioral parameters subject to adaptive, profit-seeking changes for all agents. Indeed, a ‘slow’ drift in the behavioral parameters is clearly observed in the experiments (see Fig. 6). The dynamics in this extended system with varying parameters would not rest at the steady states found in the previous simulations, which assumed constant parameters. In particular, when the average price is below  $p_{\text{edge}}$ , one would expect a drop the average value of  $h_0$ , which would then increase the market price. The value  $p_{\text{edge}}$  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 Bertrand-Edgeworth markets, which have features found in relevant real- world counterparts. The game lacks deterministic Nash equilibria, while the equilibria in mixed strate-

gies are difficult to identify and to characterize. After sketching some basic analytical properties of the model, we considered two families of price-setting heuristics, and simulated the market outcomes that they would give rise to. These results were then contrasted with experimental evidence, which provided support to the relevance of some simple pricing rules, and also led to revisions of the starting, elementary assumptions about price- setting.

In an environment which does not allow for readily defined maximizing strategies, experimental agents turn to relatively simple heuristics, but also adapt their behavior in a profit-seeking manner. Specifically, in most instances, agents who do not realize full capacity sales in a period react by lowering prices; when sales reach capacity, both positive and zero price increases are often observed. Experimental subjects initially tend to post prices above the competitive equilibrium level (even when they have been explicitly informed about the value of this price). Then, in an “exploratory” stages of play, players make aggressive moves (with relatively large modules of price change, and a high likelihood of raising the price when sales reach capacity). Further on, agents gradually shift to more conservative strategies (with smaller modules of price changes and higher chances of keeping the price constant). This pattern of behavior induces a *slow* average price decrease, characteristic of the convergence from above to the “competitive range” reported also in Kruse et al. (1994), and which can be simulated in a model with artificial agents. It is also possible to approximate the average price analytically, using parameters obtained from the experiments.

As established in Section 2, when prices approach the competitive level and get close to the critical value  $p_{edge}$  defined in equation (10), there is room for profitable deviations by agents who raise their price with sufficient intensity. The experimental results actually show upwards “excursions” of the prices set by some agents in markets which previously were in a more or less tranquil condition. Thus, the image would be one where “competitive forces” put limits to the unit margins of participating firms at a price level close to  $p_{edge}$ , but where “intermittent shakeouts” push up market prices for some periods (in the experiments, often at starting levels above  $p_{edge}$ ). Thus, market prices do not actually “freeze” in stationary states, either at the individual or the aggregate levels. Rather, the observed downwards price convergence itself generates favorable conditions for disturbances to the emergent “ordered” state.

The analysis can potentially be extended in several ways. Further experimental exercises could inform about the role of  $p_{edge}$  as a critical value, and about the likelihood of large price excursions as a function of the number of players. It would also 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. Allowing for search by buyers and varying the information structure can enricher the understanding on how individual and aggregate performance depend on specific market conditions, and provide additional bridges with the previous literature. 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 a building block for more elaborate agent-based models or, perhaps, apt to be confronted with actual data.

## **Acknowledgements**

Comments by Omar Chisari, Nicolás Aragón and participants in seminars at the Universidad de San Andrés, Universidad de Buenos Aires, Asociación Argentina de Economía Política, Instituto Balseiro and TANDAR are gratefully acknowledged. The usual caveat applies.

# Appendix 1

## Instructions for players

The players in the experiments were senior-level undergraduates or graduate students in Economics. They were presented with the instructions which are shown below in English translation. Clarifying questions were allowed.

- You will participate in the experiment as the manager of one of the  $N$  firms that produce and sell a good (the number of players is not known to any of the participants). The market will open during a sequence of  $T$  "days", numbered  $t = 1, 2, \dots, T$ . The experimenter will announce publicly when the play ends.
- Your cost of producing each unit of the good is  $c$  "credits". You will only produce goods "on demand", that is, once they have been sold. The production, and sales, of your firm in each market day cannot exceed a capacity level of  $y^*$  units. The cost per unit,  $c$ , and the maximum production  $y^*$  are equal for all the firms in the market, and all (as yourself) have been informed of the fact.
- At the beginning of each market day, you (and all the other players) will be asked to write down a price at which you are willing to sell the good. You will see an indication on the screen asking you to write down a number determining the price that you have chosen for the day. No other player will know your price and you will not be informed of the prices set by others. Your information will consist of the prices that you posted, the quantities of goods that you have sold and the profits that you have made in each day.
- Your payoff in the game, and your monetary reward, will depend on the total profits realized by your firm over the  $T$  days of the game. The benefits in a given day will be the product of the margin of the posted price over production costs (that is:  $p - c$ ) and the quantity effectively sold  $y$ . The two players with the highest total benefits will receive a prize of \$xx.
- Your sales will depend on the price that you have set, on the prices of the other participants, and on market demand. The quantities sold by each firm will be determined in the following way:

Once all players have informed their price for day  $t$ , the program will proceed to rank them in an increasing order (from lowest to highest):  $1, 2, \dots, N$ . The market demand is determined by an amount of money,  $M$ , available to buy the good. The program operates by directing that demand first to the firm with the lowest price, 1.

If it happens that the total amount of money  $M$  is insufficient to buy the maximum possible production  $y^*$  of that firm, the quantity sold by 1 is  $y_1 = M/y_1$ , and there is no money left to buy from other firms, whose sales are therefore equal to zero.

If that is not the case, and there is money still available for purchases after buying the whole production of 1, that is, if  $M$  exceeds  $p_1 y^*$ , the remaining available demand  $M - p_1 y^*$  is directed to the second-cheapest firm, 2, and the procedure is repeated.

Firms are "visited" in the order of increasing prices, until one of three things happens:

- a) the demand capacity is exhausted, so that the whole quantity of money  $M$  has been spent buying from firms  $1, 2, \dots, n$ , and then all firms with higher prices (numbers  $n + 1, \dots, N$  in the increasing price order) will make zero sales; or
  - b) the amount of money available for purchases is exhausted and all  $N$  firms make positive sales, possibly with the most expensive firm (the  $N$ th) realizing a volume of sales less than its full capacity  $y^*$ ; or
  - c) all firms  $1, 2, \dots, N$  have been visited, all have made sales at full capacity  $y^*$  and yet demand has not been satisfied (so that  $M$  is higher than the sum of the amounts  $p_i y^*$  spent in firms  $1, 2, \dots, N$ ), but no production capacity remains. In this case, the money which could not be used in buying the good because no supply was available "goes elsewhere" (that is: is not kept for future periods), and in the next day demand is again determined by amount  $M$ .
- When the program has performed these operations, players are informed of their individual results: you will see in the screen the data for your quantity sold  $y$  and your profits  $(p - c)y$  in day  $t$ . This being done, the program initiates day  $t + 1$ , as before, by asking participants to determine their price for that day.
  - As a possible reference for your decisions, the "competitive price" (that is, the price which, if posted by all firms, leads to sales exactly equal to full- capacity output by all of them:  $p^* = M/Ny^*$ ) is .... credits

## Appendix 2

For each agent in each game the behavioral parameters were estimated through OLS regressions for the sales-based pricing heuristics. The dependent variable is change in price in the next iteration and the independent variable is a constant. In order to identify the proposed coefficients we estimated  $\Delta p_{t+1} = \gamma_+$  over the subsample that sold completely ( $q = q^*$ ) and  $\Delta p_{t+1} = \gamma_-$  over the subsample with partial sales ( $q < q^*$ ) for each agent.

Table 4 reports the parameters estimation for each agent in a single experiment. Notice that in at least 82% of the agents, significant estimations were obtained.

In Table 5 we report summary statistics for the sales-based heuristic for each game. Columns (I) and (II) report the mean of the estimated behavioral parameter across agents. Columns (III) report the results of the regression over the whole population of agents in each game. The ratio



Agent	All iterations			Iteration > 25		
	$\gamma_+$	$\gamma_-$	$\gamma_+/\gamma_-$	$\gamma_+$	$\gamma_-$	$\gamma_+/\gamma_-$
1	0.024**	-0.049***	-0.489	0.010**	-0.047**	-0.211
2	0.011**	-0.111**	-0.097	-0.002**	-0.056*	0.045
3	0.010**	-0.020***	-0.498	0.001**	-0.015***	-0.047
4	0.012***	-0.027***	-0.423	0.004**	-0.027***	-0.170
5	0.018***	-0.027***	-0.673	0.003**	-0.019***	-0.159
6	0.033***	-0.038***	-0.872	0.015***	-0.030***	-0.503
7	0.023***	-0.026***	-0.905	0.002***	-0.010***	-0.244
8	0.046***	-0.165***	-0.278	0.020*	-0.096***	-0.214
9	0.056***	-0.067***	-0.848	0.036***	-0.055***	-0.646
10	0.009	-0.058*	-0.161	-0.001	-0.042**	-0.032
11	0.012	-0.022	-0.550	-0.006	-0.013***	0.440

Table 4: OLS estimates of sales-based rule for each agent in Experiment 19. \* represents significance at 10%, \*\* represents significance at 5% and \*\*\* represents significance at 1%.

between the parameters for each game is also reported in Column (IV). Table 6 is identical to Table 5, except that the estimations taken over iterations greater than 25, in order to discard the initial phase of the game.

Experiment	N	$\gamma_+$			$\gamma_-$			$\gamma_+/\gamma_-$	
		Mean (I)	Mean (II)	Aggregate(III)	Mean(I)	Mean (II)	Aggregate (III)	Mean (IV)	
13	12	0.018	0.018	0.015***	-0.057	-0.057	-0.049**	-0.56	
19	11	0.023	0.026	0.016**	-0.056	-0.059	-0.036***	-0.52	
22	7	0.105	0.038	0.080*	-0.267	-0.357	-0.320	-0.50	
23	7	0.021	0.026	0.024***	-0.058	-0.056	-0.060	-0.40	
16	6	0.015	0.015	0.015***	-0.024	-0.024	-0.024***	-0.80	
18	6	0.007	0.004	0.006***	-0.011	-0.011	-0.012***	-0.644	
21	6	0.009	0.009	0.009***	-0.026	-0.026	-0.026***	-0.314	
17	4	0.085	0.085	0.097*	-0.141	-0.044	-0.109	-1.109	
20	4	0.088	0.117	0.062**	-0.094	-0.063	-0.120	-0.593	
14	3	0.016	0.016	0.016**	-0.044	-0.044	-0.037***	-0.48	

Table 5: Columns labeled (I) report the mean of the estimated parameters among players, irrespective of statistical significance. Columns labeled (II) report the mean of the significant (at 10% significance) estimated parameters among players. Columns labeled (III) report the aggregate estimate of the parameter in each game (over the whole population). Finally, Column (IV) reports the ratio of the behavioral parameters in each game.

Experiment	N	$\gamma_+$			$\gamma_-$			$\gamma_+/\gamma_-$	
		Mean (I)	Mean (II)	Aggregate(III)	Mean(I)	Mean (II)	Aggregate (III)	Mean (IV)	
13	12	0.004	0.004	0.004	-0.015	-0.013	-0.013**	-0.23	
19	11	0.008	0.010	0.002**	-0.038	-0.038	-0.009**	-0.15	
22	7	0.003	0.003	0.004**	-0.015	-0.015	-0.019**	-0.25	
23	7	0.004	0.006	0.003**	-0.014	-0.014	-0.018	-0.22	
16	6	0.003	0.003	0.003***	-0.009	-0.009	-0.008**	-0.27	
18	6	0.001	0.001	0.001***	-0.005	-0.004	-0.004*	-0.22	
21	6	0.003	0.001	0.003**	-0.021	-0.017	-0.022**	-0.11	
17	4	0.044	0.044	0.49**	-0.071	-0.088	-0.059	-0.79	
20	4	0.105	0.139	0.076**	-0.57	-0.045	-0.780	-0.14	
14	3	0.002	0.003	0.002**	-0.007	-0.007	-0.009***	-0.18	

Table 6: Identical to previous Table, but only utilizing the subsample post iteration 25.

## References

- [1] Abbink, Klaus and Jordi Brandts (2005). “Price Competition Under Cost Uncertainty: A Laboratory Analysis.” *Economic Inquiry* **43**, pp. 636-648.
- [2] Allen, Beth and Martin Hellwig (1986). “Bertrand-Edgeworth Oligopoly in Large Markets.” *Review of Economic Studies* **53**, pp. 175-204.
- [3] Ball, Laurence and David Romer (1991). “Sticky Prices as Coordination Failures”. *American Economic Review* **81**(3), pp. 539–552.
- [4] Brandts, Jordi and Pablo Guillén (2007). “Collusion and Fights in an Experiment with Price-Setting Firms and Production in Advance.” *Journal of Industrial Economics* **55**, pp. 453-473.
- [5] Buchheit, Steve and Nick Feltovich (2010). “Experimental evidence of a sunk–cost paradox: a study of pricing behavior in Bertrand–Edgeworth duopoly.” *International Economic Review*. Forthcoming.
- [6] Cason, Timothy, Daniel Friedman and Florian Wagener (2005). “The Dynamics of Price Dispersion, or Edgeworth Variations.” *Journal of Economic Dynamics and Control* **29**, pp.801-822.
- [7] Chamberlin, E. (1956). *The Theory of Monopolistic Competition*. Cambridge MA, Harvard University Press.
- [8] d’Aspremont, C. and J.J. Gabszewicz (1985). “Quasi-Monopolies.” *Economica* **52**, pp. 141-151.
- [9] Davis, Douglas and Charles Holt (1994). “Market Power and Mergers in Laboratory Markets with Posted Prices.” *The Rand Journal of Economics* **25**, pp. 467-487.
- [10] De Francesco, Massimo and Neri Salvadori (2008). “Bertrand-Edgeworth games under oligopoly with a complete characterization for the triopoly.” Working Paper.
- [11] Dasgupta, Partha and Eric Maskin (1986). “The Existence of Equilibrium in Discontinuous Economic Games, II: Applications.” *Review of Economic Studies* **53**, pp. 27-41.
- [12] Duffy, John (2006). “Agent-Based Models and Human Subject Experiments.” In Testfatsion, Leigh and Kenneth Judd (eds.) *Handbook of Computational Economics, Volume 2*. Elsevier B.V., pp. 949-1011.
- [13] Edgeworth, Francis Y (1925). “The Pure Theory of Monopoly”, in *Edgeworth, Papers Relating to Political Economy*, New York, Burt Franklin, pp. 111-142.
- [14] Fonseca, Miguel and Hans-Theo Normann (2008). “Mergers, Asymmetries and Collusion: Experimental Evidence.” *The Economic Journal* **118**, pp. 387-400.

- [15] Fonseca, Miguel and Hans-Theo Normann (2009). "Excess Capacity and Pricing in Bertrand-Edgeworth Markets: Experimental Evidence." Working Paper.
- [16] Holt, Charles and Fernando Solis-Soberon (1994). "The Calculation of Mixed-Strategy Equilibria in Posted-Offer Markets." In R.M. Isaac (ed.) *Research in Experimental Economics*, Vol 5. Greenwich: JAI Press.
- [17] Ketcham, Jon, Vernon Smith and Arlington Williams (1984). "A Comparison of Posted-Offer and Double-Auction Pricing Institutions." *Review of Economic Studies* **51**, pp. 595-614.
- [18] Kruse, Brown J., Stephen Rassenti, Stanley Reynolds and Vernon Smith (1994). "Bertrand-Edgeworth Competition in Experimental Markets", *Econometrica* **62**, pp. 343-371.
- [19] Maskin Eric, and Jean Tirole (1988). "A Theory of Dynamic Oligopoly II: Price Competition, Kinked Demand Curves and Edgeworth Cycles", *Econometrica* **56**, pp. 571-599
- [20] Tirole, Jean (1988). *The Theory of Industrial Organization*. Cambridge MA, The MIT Press.