



Behavioral heuristics and market patterns in a Bertrand–Edgeworth game



D. Heymann^{a,b}, E. Kawamura^{a,*}, R. Perazzo^{a,c}, M.G. Zimmermann^a

^a Universidad de San Andrés, Economics Department, Argentina

^b IIEP-BAIRES, Universidad de Buenos Aires, Argentina

^c Instituto Tecnológico de Buenos Aires, R&D Department, Argentina

ARTICLE INFO

Article history:

Received 25 July 2011

Received in revised form 6 April 2014

Accepted 26 April 2014

Available online 9 May 2014

JEL classification:

D43

C63

C91

Keywords:

Oligopolies

Capacity constraints

Agent-based models

Experimental testing of heuristic price-setting rules

ABSTRACT

This paper studies Bertrand price-setting behavior when firms face capacity constraints (Bertrand–Edgeworth game). This game is known to lack equilibria in pure strategies, while the mixed-strategy equilibria are hard to characterize. We explore families of heuristic rules for individual price-setting behavior and the resulting market patterns, through simulations of agent-based models and laboratory experiments. Overall, the individual pricing strategies observed experimentally can be represented approximately by a sales-based simple rule. In the experiments, average market prices tend to converge from above and approach a state resembling a steady state, with slow aggregate price variations and low price dispersion around an average near the competitive level. However, that configuration can be disturbed occasionally by excursions triggered by discrete price raises of some agents. Salient features of experimental results can be described by simulations where agents use sales-based heuristics with parameters calibrated from the experiments. The results obtained here suggest the existence of useful complementarities between analytical, experimental and agent-based simulation approaches.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

The Bertrand–Edgeworth (BE) oligopoly, where price-setting, capacity-constrained firms compete for customers, provides a fair approximation of several markets, especially in the retail sector.¹ The basic form of that scenario is easily described. In a “Marshallian short-run”, where the number of suppliers and their respective productive sizes are given, each firm operates under constant marginal costs, up to the bound established by a capacity constraint which determines a maximum volume of goods which can be produced or kept in store. These firms post prices that signal their willingness to sell, conditional on the capacity restriction, to customers who seek to buy at the lowest available price and who, to focus the analysis on the behavior of firms, can be assumed to know all posted prices and move between suppliers at no cost.

Despite the simplicity of this market setup, the resulting pricing behaviors are still imperfectly understood. On the theory side, the BE oligopoly game lacks Nash equilibria in pure strategies, as will be discussed shortly. The mixed-strategy

* Corresponding author at: Universidad de San Andrés, Vito Dumas 284, Victoria (B1644BID), Buenos Aires, Argentina. Tel.: +54 11 4725 7077; fax: +54 11 4725 7010.

E-mail address: kawa@udesa.edu.ar (E. Kawamura).

¹ 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”.

equilibria are hard to characterize, and their implementation requires a non-trivial matching between the decisions of different sellers. Such demanding requirements of coordination limit the plausibility of those equilibria as representations of actual market outcomes.² This implausibility is empirically confirmed in experimental studies such as those in Kruse et al. (1994), Fonseca and Normann (2008, 2013) and Buchheit and Feltovich (2010). The experimental results of these papers reject the hypothesis that behavior corresponds to equilibrium mixed strategies. They do not find, either, evidence for alternative theoretical benchmarks like the perfectly competitive outcome or tacit collusion (see also Davis and Holt, 1994), although they identify some fluctuations resembling “Edgeworth cycles”,³ with alternating phases of price increases and sharp competitive reductions.

The experimental oligopolists studied in the literature appear to face difficulties in specifying mutually compatible optimal plans. This feature reinforces the motivation for studying heuristic procedures that price setters may be using, instead of appealing to the traditional equilibrium notions. The focus on heuristic rules as potential representations of individual behavior facilitates exploring complementarities between experiments and agent-based simulations to analyze market patterns.⁴ This is the aim of the present paper.

We consider the market for a non-storable good. Supply is generated by N firms, subject to a constant unit cost and a fixed upper limit on production. For the sake of the analysis, firms are assumed to be identical in their cost and maximum-output parameters. Sellers post prices before the beginning of the market period, knowing only their own past decisions (prices) and performance (sales, profits). As a simplifying assumption, firms are not informed about the variables pertaining to their competitors.⁵ Also for simplicity, demand is represented through a unit-elastic function such that the representative buyer plans to execute each period a given, constant value of spending in the good. In addition, the model assumes away search or information costs: purchasers know the whole set of posted prices before realizing any transaction. In this setting, the behavior on the demand side is purely mechanical, and it can be assimilated to a sequence of visits to the stores with prices in increasing order, until either the spending plans are satisfied or all firms have sold at full capacity. The paper focuses on the case where, given the capacity constraint, the competitive price is strictly above the unit cost.

With these assumptions, if all firms except one set their prices at the competitive level, the “residual supplier” would maximize its profits by fixing a very high price, even if this leads to a small (in the limit, a vanishing) volume of sales. Another property of this oligopoly is that no uniform price can be a Nash equilibrium and that, if $N - 1$ firms set a “low” price, below a particular threshold, the N th competitor would have an incentive to fix a (much) higher price. These properties confirm the implausibility of a pure strategy Nash equilibrium as a solution concept for this oligopoly. Therefore, as a behavioral alternative, the paper specifies two families of simple pricing criteria which may potentially represent the behavior of sellers, and simulates the market features that would emerge if firms use them. The first of these heuristic rules consists of a straightforward reaction to realized sales, by raising the price if the quantity demanded in the last period has reached the capacity limit, and lowering the price in the opposite case. The other class of pricing functions would follow a profit-gradient strategy, by determining a price change of the same sign than that of the previous round if this led to higher benefits, and reversing course otherwise.⁶

The experimental evidence presented next shows some noticeable patterns. Except for few cases, the experiments show that the observed individual price-setting behavior can be rationalized by a variant of the simple sales-based scheme. This seems an especially remarkable outcome given that most of the experimental subjects were advanced Economics students who, by their training and understanding of the game, could be expected to use more sophisticated strategies looking for higher profits. A second observed feature concerns how experimental markets tend to organize. Posted prices tend to start at relatively high average levels and show large dispersions, with a gradual convergence to a what looks like a quasi-stationary state approaching the competitive price, occasionally disturbed by “upward excursions” where some agents seek to improve profits by substantial but transitory raises in markups. Overall, the experimental results indicate both heterogeneity in behaviors, as could be expected, and price configurations with structures apt to be replicated through agent-based simulations using stylized behavioral heuristics.

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

² 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 (2010). (See also Tirole (1988), for a textbook treatment.) In some special cases computation of mixed-strategy equilibria is easier, but these correspond to very particular step demand functions (see Holt and Solis-Soberon (1994) for a discussion).

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

⁴ See the discussion of such complementarities in Duffy (2006).

⁵ This contrasts with the setups in related works such as Brandts and Guillén (2007), Cason et al. (2005), Fonseca and Normann (2008, 2013), Buchheit and Feltovich (2010). On their side, Abbink and Brandts (2005) and Cason and Friedman (1999, 2003) study experimental markets without capacity constraints.

⁶ These schemes have points in common with two different sets of the literature on learning behavior: the directional learning models introduced by Selten and Stoecker (1986), Selten and Buchta (1999), and treated in experimental contexts in Selten (2004) and Selten et al. (2005); and the reinforcement learning or cased-based schemes in works such as Karandikar et al. (1998), Izquierdo et al. (2007), and the surveys by Bendor et al. (2001) and Gilboa and Schmeidler (2001).

models; the comparison between models and experiments is developed further in Section 5. The final section presents the 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 daily 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 a sequence of increasing prices, and buy from the cheapest available suppliers until either their willingness to buy is exhausted or firms have ran out of goods to sell. To make the argument precise, and to study concrete pricing decisions in experimental contexts, it is useful to represent the behavior of customers through a specific demand schedule. For simplicity, the model assumes a unit-elastic function, as if the agent came to the market ready to spend on the good a certain amount of money (equal for all consumers and fixed across periods), with an aggregate amount of money income M .⁷ In practice, the demand side of the market would act non-strategically, as a single agent with a scaled-up version of the individual demand function. The paper concentrates on the case where $M > Ncq^*$, and the desired value of customer purchases 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 equilibrium price is then:

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

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.*⁸

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 $p_j \rightarrow \infty$. \square

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 $p' > (N/(N - 1))p^*$ so that the demand would be exhausted with less than $N - 1$ firms selling at full capacity. If $N - 1$ suppliers fix that price, firm N would realize a zero value of sales in case $p_N > p'$, and would thus have an incentive to reduce its price slightly below p' .
- If $(N/(N - 1))p^* > p' > p^* + (c/N)$, the N th firm gains by setting a price slightly below p' . Here, the N th firm effectively faces a choice between, either “guaranteeing” for itself full capacity sales by lowering its price relative to its competitors, or acting as a monopolist who captures the “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, and raise its price so as to cut production costs to near zero. The condition for preferring a lower price $p' - \varepsilon$ (for $\varepsilon \rightarrow 0$) is:

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

⁷ Some implications of modifying the demand function, to a linear specification, are briefly discussed below.

⁸ The proposition is proved here using the postulated unit-elastic demand function, but (as well as the non-existence of a Nash equilibrium with uniform prices, stated in Claim 2 below), it also applies to other cases. Consider that of a mass z of “small” customers with linear demand curves $q^d = b(p^m - p)$. Here p^m denotes the price at which the quantity purchased goes to zero (an exogenous parameter). The competitive equilibrium price is then given by $p^* = p^m - (Nq^*)/zb$, which is strictly greater than c by assumption. Assume now that $N - 1$ firms post the competitive price p^* , and that this price is below the monopolistic level in the absence of capacity constraints: $p^{\text{mon}} = (p^m + c)/2$. The residual demand faced by the N th firm is $q_N = ((N - 1)/N)zb(p^m - p_N)$. This leads to maximum profits at price $p_N = p^{\text{mon}}$, which is strictly greater than p^* , and sales below the capacity level, $q_N < q^*$. Here, by contrast with the unit-elastic case, the N th firm has no incentive to post such a high price that the volume of sales tends to zero. Correspondingly, when prices set in the market are so low that an agent can contemplate a profitable “upward excursion”, the magnitude of the jump would be limited by the size of the gap between the competitive and the monopoly prices.

equivalent to:

$$p' - p^* > \frac{c}{N} \tag{4}$$

- If $p' < p^* + (c/N)$ for $N - 1$ suppliers, the N th firm has an incentive to set a very high price.

□

Note that the last part of the proof introduces a critical level of the price of the $N - 1$ th suppliers p' ,

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

When p' becomes smaller than p_{edge} , the best response action of the N th firm suddenly switches from setting a price slightly lower than p' to announcing a price much above that of the $N - 1$ competitors.

The results shown above illustrate a more general feature of this oligopoly: the lack of Nash equilibria in pure strategies, a feature stressed in the literature.⁹ This lack of a set of mutually optimal prices suggests that individual firms 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 such an interval, the N th firm would find profitable a “large” upwards deviation, gets smaller as the number of firms increases, since it is determined by the magnitude $1/N$. 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 $p' < p^* + (c/N)$, the profits of the N th supplier would vary non-monotonically with the size of a positive shift of the price, to $p' + \varepsilon$. A “local” increase, given by a small value of ε , would lead to lower profits, while profits would increase for a “sizeable” price jump.
- In states where prices are concentrated around a level $p' > p^*(N/(N - 1))$, 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. 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 firms have no obvious procedure to find mutually consistent optimal prices. The absence of Nash-equilibrium in pure strategies directs the attention towards possible decision rules that may provide simple, reasonable, guides for price-setting based on the information that firms themselves generate in their own transactions, without requiring intricate conjectures about the behaviors of their competitors. Here, we analyze pricing criteria based only on the performance of the firm in question, and which do not depend on parameters such as the number N of sellers in the market.¹⁰ 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 prepare the ground for the analysis of experimental evidence, by providing a “lens” through which to look at the data.

3.1. A simple sales-based pricing rule

A particularly simple scheme would be for the price-setting 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:

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

where t denotes the “market day” and $(\gamma_+, \gamma_-) \gg 0$ are behavioral parameters which determine the magnitude of upwards and downwards price changes.¹¹

⁹ Several papers show more general proofs of non-existence of pure-strategy, Nash equilibria in Bertrand–Edgeworth oligopolies. See, e.g., d’Aspremont and Gabszewicz (1985) and Dasgupta and Maskin (1986, footnote 3) for different versions of the demonstration. The non-existence proposition does not apply only to the symmetric equilibrium case; Dasgupta and Maskin (1986) prove it for duopolies with every possible price-configuration (covering both the symmetric case as well as the two other possible asymmetric cases). It does not seem too difficult to come up with an extension of that method to the N -firm case, although such a proof is out of the scope of this paper.

¹⁰ Whether this is a plausible assumption seems to be empirical question. Section 4 below does not show a clear connection between the observed pricing behaviors and N , in the range considered.

¹¹ It 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 this discussion we treat these parameters as fixed.




Fig. 1. Price evolution with sales-based model. 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 $\bar{p} \approx 1.20 \approx p^*(1 + \gamma_+/\gamma_-)$. Parameter values: $N = 10$, $p^* = 1$, $c = 0.75$, $\gamma_+ = 0.02$, and $\gamma_- = 0.10$.

A system of firms that follow this pricing rule would have an aggregate steady state, with constant average sales \bar{q} and selling price \bar{p} . Consider the case where all agents share the same behavioral parameters, and a situation where N_s firms sell at full capacity, and therefore raise their price in the following day, 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 $\bar{q} \equiv q^* N_s/N$ (for large N) is equal to:

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

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 $Np^*q^* = M \approx N_s\bar{p}\bar{q}$:

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

The properties of this state are intuitively simple. Take for example the limit where $\gamma_+ \ll \gamma_-$. Individual prices vary as if placed on a “conveyor belt”: an individual increases its price in small steps of size γ_+ up to a value such that sales fall below full capacity, leading to a price reduction γ_- . In an aggregate steady state, the number of such steps in the cycle would be equal to $\mu = \gamma_-/\gamma_+$. Then, on average, every agent does not sell once every μ 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 that obtained in a steady-state competitive equilibrium $(p^* - c)q^*$. The heuristics would induce some kind of implicit collusion among sellers. The latter would be sustained because individuals do not undercut their competitors by setting a price slightly below \bar{p} , which would generate full sales without a sizable sacrifice in terms of the price markup.

Numerical simulations show that, starting from randomly chosen initial prices, a system of agents following the sales-based rule does tend asymptotically towards a steady state price level \bar{p} described by (8): $\bar{p} \approx p^*(1 + \gamma_+/\gamma_-)$. Fig. 1 shows a representative time series of the average price of the market, 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} . The figure describes a gradual convergence from above towards the steady state price level $\bar{p} \approx p^*(1 + \gamma_+/\gamma_-)$. The convergence also takes place starting from an average initial price smaller than the asymptotic level. Simulations with other adjustment parameters γ_+ and γ_- confirm that the corresponding steady state is an attractor for average prices.

Thus, this simple heuristics generates a movement to a “dynamic equilibrium” where the system shows a permanent variability of individual prices, while the average price and volume of sales gravitates towards a constant value. It must 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 agents may be incorporated in their choice of parameters, or in their use of a more sophisticated decision formula.




Fig. 2. Price evolution with profit-based model. 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: $N = 10$, $p^* = 1$, $c = 0.75$, $\sigma = 0.4$, and $\epsilon = 0.05$.

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 have decreased. Prices are assumed to vary by a magnitude proportional to the one-period gradient of profits. The rule incorporates a term allowing for a random search if profits have not varied, in a way biased towards price drops 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’, $\Pi_i(t)$ is the value of profits realized at t , $\sigma > 0$ is an adjustment parameter, $\Delta(X(t)) \equiv X(t) - X(t-1)$ and η 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 ϵ is a small number which determines the “search space” for the price setter.

This rule generates an asymmetric behavior in price increases or decreases. Whenever the firm obtains higher profits by raising prices, the response is to continue with a gradual increment, as the size of the increase is proportional to the (relatively small) profit gain 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 drop. This would be analogous to the “conveyor belt” feature of the sales-based rule with larger parameters of price decrease (γ_-) than increase (γ_+).¹² An additional feature of this pricing strategy (not present in the sales-based rule discussed before) is the trial-and-error search behavior implicit in the additive noise term η in (9), whenever the change in profit is small: $\Delta(\Pi_i(t)) \lesssim \epsilon/\sigma$. Thus, firms may test the consequences of price reductions before ‘hitting’ the no-sales region.

Fig. 2 shows a simulated time series starting from random initial conditions in the range $[p^*, 2p^*]$. After a short transient phase, the system settles on a state resembling a “dynamic equilibrium”, with some aggregate oscillations, and a sizable variability of individual prices.

The simulations with the profit-gradient rule result in average market prices that fall when the number of firms N becomes larger. The pricing behavior tends to reduce over time the number of firms with zero sales, thus moving the system towards values such that the average price does not exceed $p^*(N/(N-1))$. 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} was defined in (5). Table 1 shows simulated asymptotic price levels, which decline with N and have values close to the corresponding p_{edge} .

Also, numerical simulations with few firms ($2 < N < 6$) show, as a fairly regular feature, cases where several firms simultaneously increase prices, starting an implicit collusion phase which lasts for a few time steps, before being reversed. These price excursions resemble Edgeworth cycles. The amplitude of these movements diminishes as the number of players increases. For $N \geq 6$ the excursions cannot be distinguished from small fluctuations of the average price.

¹² 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.

Table 1

Simulation results of profit-based model (9) showing the average price as a function of the number of players N . For each run, the average price of the last 500 rounds was recorded. (\bar{p}) is the average over 50 realizations for each N . Parameters of the simulation: $p^* = 1$, $c = 0.75$, $\sigma = 0.2$, $\epsilon = 0.01$, total time steps 5000.

| N | (\bar{p}) | p_{edge} |
|-----|-------------|-------------------|
| 2 | 1.2801 | 1.375 |
| 4 | 1.1820 | 1.188 |
| 6 | 1.1376 | 1.125 |
| 8 | 1.1141 | 1.094 |
| 10 | 1.0946 | 1.075 |
| 14 | 1.0695 | 1.054 |
| 18 | 1.0525 | 1.042 |
| 25 | 1.0371 | 1.030 |
| 35 | 1.0276 | 1.021 |
| 50 | 1.0193 | 1.015 |

4. Experimental results

The previous section introduced two families of price-setting rules and studied their properties and their implications for the performance of the Bertrand–Edgeworth market, using computer simulations. However, the potential empirical relevance of each of those heuristics remains to be established. This section reports the results of experiments where the participants performed as firms setting prices, and compares the observed individual behaviors with those derived from the simple decision formulas and the simulations just discussed.

4.1. The experimental setup

The experiments were carried out using an especially designed web page.¹³ All subjects were either graduate or advanced undergraduate students in Economics, except in one experiment played by physicists (experiment #17). The instructions for the game are shown in Appendix 1. For each experiment, the computer server generated a unique identifier session number.¹⁴ Since the play was “silently” implemented in individual computers, on some occasions two separate games were ran simultaneously. Before the session started, the experimenters announced publicly the following (common) parameters: the unitary cost c , the maximum individual capacity q^* , and the competitive price level p^* . In each session, the system always showed in the first round of play the competitive price level p^* . Participants were not informed about the number N of “competitors” they would be playing with and therefore could not deduce M , the level of market demand.

In each market round, all subjects posted prices through the interactive web page, after which the server computed the amount sold by each agent. The server sequentially bought from “firms”, starting with the cheapest seller, and visiting all other participants in an increasing price order. In each visit, total purchases were equal to the seller’s capacity constraint, until the total amount of money M was spent or, in case supply ran out before the available cash to spend was exhausted, the resources were discarded, without carrying them to the following round.

At the end of each market day, each player received from the server an indication that a new round of play was open. Information was also provided about the player’s chosen price level and the realized quantity of sales and profits in the last five market days. The players could calculate the one-step profit gradient if they wished to. In two sessions (#29 and #31) the data about that gradient was explicitly supplied by the server.

In this setup, the players were not informed about the value of any variables (prices, sales, profits) pertaining to the other participants. Alternative experimental specifications of the game are clearly possible (as illustrated in the literature). This one is particularly simple, and seems to correspond to real-world cases where, for example, market competitors are small shops in different locations. Out of the eleven sessions reported here, in seven participants received pecuniary rewards contingent on their performance in generating realized profits. Two 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 between 3 and 12, while the total number of market days T was chosen in the range (45, 150) (see Table 8, columns 1–3).

Fig. 3 shows price series in two experimental sessions. As observed in other experiments, the game starts with some prices well above the competitive level (although this value is known by all agents before the play begins), and with a large dispersion. Later, the range of individual prices shrinks persistently, along a declining trend of average prices. After something like thirty “market days”, the cross-player variance of prices becomes quite small. However, this convergence may be occasionally disturbed by abrupt upwards price excursions.

¹³ See the home page: <http://www.elautomataeconomico.com.ar/>.

¹⁴ In order to check the working of the client–server application, we ran some trial sessions, played by ourselves, each of which had its identifier session number. Thus, the reported experimental sessions need not have consecutive session numbers.




Fig. 3. Average, maximum and minimum offer price (rescaled with p^*) empirical time series. Also shown is the Edgeworth price $p_{edge} = p^* + c/N$. (top) Experiment #21, (bottom) Experiment #18.

4.2. A comparison of simple heuristic rules

One of the goals of the experiments was to evaluate whether individual behavior showed patterns analogous to those of the heuristic rules previously discussed. For each agent in each game, we run OLS regressions of the observed price changes on realized sales in the previous market day, as suggested by the sales-based pricing heuristics. To estimate the corresponding coefficients, the paper considers two separate regressions: one specified as $\Delta p_{t+1} = \gamma_+$ over the subsample in which the agent had full capacity sales, ($q = q^*$), and the other $\Delta p_{t+1} = \gamma_-$ over the subsample with partial, or zero, sales for the agent ($q < q^*$). As an example, Table 2 reports the experimental parameters for each agent in a single experiment (#19). The estimates do not change much if the initial 25 market iterations are discarded. Significant coefficients were obtained for most players (9 out of 11 players, 82%), as shown in the table.

The relevance of the profit-based algorithm was analyzed by estimating, for each subject, equations of the type $\Delta p_{t+1} = \alpha + \beta \Delta \Pi_{t-1} \text{sign}(\Delta p_{t-1})$, where the regressors are a constant α and an interaction between the lagged change in

Table 2
OLS estimates of γ_+ and γ_- for sales-based rule for each agent in experiment #19.

| Agent | All iterations | | | Iteration > 25 | | |
|-------|----------------|------------|---------------------|----------------|------------|---------------------|
| | γ_+ | γ_- | γ_+/γ_- | γ_+ | γ_- | γ_+/γ_- |
| 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 |

* Significance at 10%.
 ** Significance at 5%.
 *** Significance at 1%.

Table 3

Summary of regression results for each experiment (discarding the initial 25 iterations), indicated by the corresponding session identifier number. Columns (I) and (II) report the fraction of agents in each game with appropriate signs in the sales-based estimation (positive for γ_+ and negative for γ_-) irrespective of statistical significance. Columns (III) and (IV) report the fraction of agents in each game with estimated coefficients significant at the 10% level, and appropriate signs. Columns (V) and (VI) correspond to the profit-based equations (PB) algorithm. We estimated for each agent $\Delta p_{t+1} = \alpha + \beta \Delta \Pi_{t-1} \text{sign}(\Delta p_{t-1})$, as described in the text. Column (V) shows the fraction of agents with coefficients β significant at the 10% level in estimates without the constant regressor; column (VI) is analogous, in regressions that include the constant.

| Experiment | N | γ_+ (I) | γ_- (II) | γ_+^* (III) | γ_-^* (IV) | PB (V) | PB (VI) |
|------------|----|----------------|-----------------|--------------------|-------------------|--------|---------|
| 13 | 12 | 0.91 | 1.00 | 0.58 | 0.83 | 0.25 | 0.17 |
| 19 | 11 | 1.00 | 1.00 | 0.82 | 0.91 | 0.18 | 0.09 |
| 22 | 7 | 1.00 | 1.00 | 0.57 | 0.71 | 0.57 | 0.57 |
| 23 | 7 | 1.00 | 1.00 | 0.71 | 0.71 | 0.14 | 0.14 |
| 16 | 6 | 1.00 | 1.00 | 1.00 | 1.00 | 0.17 | 0.17 |
| 18 | 6 | 1.00 | 1.00 | 0.66 | 0.83 | 0.00 | 0.00 |
| 21 | 6 | 1.00 | 1.00 | 0.66 | 1.00 | 0.17 | 0.17 |
| 17 | 4 | 1.00 | 1.00 | 1.00 | 0.50 | 0.00 | 0.00 |
| 20 | 4 | 1.00 | 1.00 | 0.75 | 0.75 | 0.25 | 0.25 |
| 29 | 5 | 1.00 | 1.00 | 1.00 | 1.00 | 0.20 | 0.20 |
| 31 | 3 | 1.00 | 1.00 | 1.00 | 0.67 | 0.67 | 0.67 |

Table 4

Columns (I) report the mean among players of the estimated parameters, irrespective of statistical significance. Columns labeled (II) show means of the estimated parameters that were significant (at the 10% level). Columns (III) report the aggregate estimate of the parameter in each game (pooling the whole population). Column (IV) shows the ratio of the behavioral parameters in each game using the means xxx.

| Experiment | N | γ_+ | | | γ_- | | | γ_+/γ_- |
|------------|----|------------|-----------|-----------------|------------|-----------|-----------------|---------------------|
| | | 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 |
| 29 | 5 | 0.047 | 0.047 | 0.044*** | -0.074 | -0.074 | -0.078*** | -0.694 |
| 31 | 3 | 0.068 | 0.068 | 0.069*** | -0.126 | -0.146 | -0.106* | -0.673 |

In columns III, (*) denotes statistically significant estimates at 10% confidence level, (**) means statistical significance at 5%, while (***) implies statistical significance at 1%.

profit $\Delta \Pi_{t-1}$ and the sign of lagged change in price $\text{sign}(\Delta p_{t-1})$. As an alternative, the regressions were also run excluding the constant α .

Table 3 summarizes the results obtained in regressions ran 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, and the average price variation when sales are lower than q^* . Almost 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 of Table 3). For a substantial majority of individuals, those estimates are statistically significantly different from zero (Columns III and IV of Table 3). In contrast, Column V and VI of Table 3 show that in most experiments only a minority of players can be consistently represented using the profit based algorithm. Only in experiments #22 and #31 more than half of the participants had significant estimated parameters.¹⁵

Table 4 summarizes results for the sales-based heuristic for each game. Table 5 is the same except that the regressions were estimated discarding the first 25 market days. The qualitative features of the relation 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 quantitative value of the parameters shows substantial differences in both cases (Tables 4 and 5). Both the sizes of the adjustment coefficients and the ratios of the responses in upwards and downwards directions (the absolute value of γ_+/γ_-) 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.

¹⁵ In a few cases, the coefficients were significant, but with the negative sign. We did not count these as valid cases. The results just commented in the text do not favor the sales-based heuristic as a description of the experimental behavior. A possible conjecture could be that the trouble of calculating the profit gradient discouraged agents from conditioning their prices on that variable. We ran two experiments, #29 and #31, where players were shown explicitly the last-period change in their benefits. In one of them (with a small number of agents), the regression results indicate that players may have used a strategy that looks at the profit gradient, but this would apply to a small minority of agents in the other experiment. See Table 3.

Table 5
Identical to previous table, but only utilizing the subsample post iteration 25.

| Experiment | N | γ_+ | | | γ_- | | | γ_+/γ_- |
|------------|----|------------|-----------|-----------------|------------|-----------|-----------------|---------------------|
| | | 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 |
| 29 | 5 | 0.014 | 0.014 | 0.014*** | -0.030 | -0.030 | -0.029*** | -0.43 |
| 31 | 3 | 0.041 | 0.041 | 0.040*** | -0.073 | -0.091 | -0.062 | -0.71 |

Table 6
Average over players of the conditional probability $P(\Delta p_{t+1} ; q_t)$ of the change in price $\Delta p_{t+1} = p_{t+1} - p_t$ given that in the previous period 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.

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

4.3. Behavioral strategies and market outcomes: some features

4.3.1. Prices and sales: conditional frequencies

This section characterizes more precisely the experimental pricing dynamic behavior and the resulting evolution of the artificial markets. For the first goal, Table 6 shows for one experiment (#18) the conditional frequencies $P(\Delta p_{t+1} ; q_t)$ of price movements of a certain sign in market day $t + 1$, (Δp_{t+1}), given the volume of sales in the previous period (q_t). The cells in the transition matrix measure the frequency of events with price increases $\Delta p_{t+1} > 0$, with null price changes, $\Delta p_{t+1} = 0$, or price drops $\Delta p_{t+1} < 0$ 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 $q_t = 0$.

A salient characteristic of the matrix shown before is that the experimental behavior reproduces an important element of the sales-based rule, as indicated by 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 in the previous section, is that, in contrast with the simplest rule, 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.

Another question refers to the length of the memory implicit in the price decisions of individual agents. To check the relevance of sales in periods before the last, Table 7 shows frequencies of price changes of a certain sign in market round $t + 1$ conditional on q_t and q_{t-1} . The results indicate that sales lagged more than one period, q_{t-1} , do not play an important role in determining current behavior.

4.3.2. Behavioral parameters across experiments

A summary of the results concerning the values of salient behavioral estimated parameters observed in different experiments is shown in Table 8. This presents overall averages with respect to both individual agents and to rounds of play in the second half of market days, after the system has gone through the initial transitional phase. It can be seen that sales

Table 7
Average over players of conditional probability $P(\Delta p_{t+1} ; q_t, q_{t-1})$ of next round's change in price $\Delta p_{t+1} = p_{t+1} - p_t$ given q_t and q_{t-1} . The figures correspond to the last 50 rounds of experiment #18 with $N = 6$ players.

| t | t - 1 | $\Delta p_{t+1} < 0$ | $\Delta p_{t+1} = 0$ | $\Delta p_{t+1} > 0$ | # of events |
|-----------------|---------------------|----------------------|----------------------|----------------------|-------------|
| $q_t = q^*$ | $q_{t-1} = q^*$ | 0.10 | 0.47 | 0.42 | 134 |
| $q_t = q^*$ | $0 < q_{t-1} < q^*$ | 0.23 | 0.46 | 0.31 | 39 |
| $q_t = q^*$ | $q_{t-1} = 0$ | 0.11 | 0.41 | 0.48 | 27 |
| $0 < q_t < q^*$ | $q_{t-1} = q^*$ | 0.97 | 0.00 | 0.03 | 37 |
| $0 < q_t < q^*$ | $0 < q_{t-1} < q^*$ | 0.93 | 0.07 | 0.00 | 14 |
| $0 < q_t < q^*$ | $q_{t-1} = 0$ | 1.00 | 0.00 | 0.00 | 6 |
| $q_t = 0$ | $q_{t-1} = q^*$ | 0.93 | 0.03 | 0.03 | 29 |
| $q_t = 0$ | $0 < q_{t-1} < q^*$ | 1.00 | 0.00 | 0.00 | 4 |
| $q_t = 0$ | $q_{t-1} = 0$ | 1.00 | 0.00 | 0.00 | 10 |

Table 8

Main behavioral parameters and price averages. Notation of the table: $\xi = P(\Delta p_{t+1} < 0; 0 \leq q_t < q^*)$: frequency of price reductions after sales below capacity; $h_- = P(\Delta p_{t+1} = 0; q_t = q^*)$, $h_0 = P(\Delta p_{t+1} = 0; q_t = q^*)$ and $h_+ = P(\Delta p_{t+1} > 0; q_t = q^*)$: frequencies of price increases, no change, and decreases after full-capacity sales, respectively; $\langle \gamma_+ \rangle$, $\langle \gamma_- \rangle$: modules of price increases and decreases; $\langle p \rangle / p^*$: average price relative to the competitive price; for comparison purposes, the last column shows the ratio between $p_{edge} \equiv p^* + c/N$ and the competitive price. The table also indicates for each experiment the number of players, N , and that of market days, T . Experimental parameters are computed taking into consideration the last half of the total rounds of play.

| 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 |
| 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 |
| 29 : E_B | 5 | 90 | 0.79 | 0.02 | 0.40 | 0.58 | 0.009 | -0.020 | 0.475 | 1.261 | 1.176 |
| 31 : E_B | 3 | 150 | 0.77 | 0.04 | 0.36 | 0.60 | 0.054 | -0.060 | 0.896 | 1.320 | 1.293 |

lower than full capacity, or equal to zero, are followed by price drops in the majority of cases. The response of prices after full capacity sales shows few cases of price decreases (column defined by the notation h_-), but a considerable number of no changes (column denoted with h_0), along with a sizable proportion of raises (column denoted with h_+), as postulated in the simple sales-based rule. The absolute value of such price increases is typically smaller than that of downwards price adjustments. Average prices are clearly higher than the competitive levels (in some experiments, with quite substantial spreads), and are also larger than the limit price p_{edge} below which a single deviant agent would have an incentive to produce a large price increment.

4.3.3. Time-series patterns of market variables and behavioral parameters

The simple sales-based heuristics assumes that the absolute values of price increases and reductions are constant during the course of the game. However, the experiments show some recognizable dynamic pattern for such parameter estimates. Although with fluctuations, and differences between experiments, the average over agents of the absolute value of price movements, $\langle |p_t - p_{t-1}| \rangle$, tends to decline as the game proceeds, and reaches small values (Fig. 4A). This is associated with a narrowing of the range between the maximum and the minimum price, which takes place on a time scale of some 30 ‘market days’ (see Fig. 3). In addition, the ratio of up and down adjustment modules $\langle \gamma_+ / \gamma_- \rangle$ averaged across players falls substantially when comparing early and late stages of the game. Actually, for all agents and all experiments, $\langle \gamma_+ / \gamma_- \rangle < 1$ in the last 25 rounds of play. This behavior matches with that observed in the frequency of price increases after full-capacity sales.

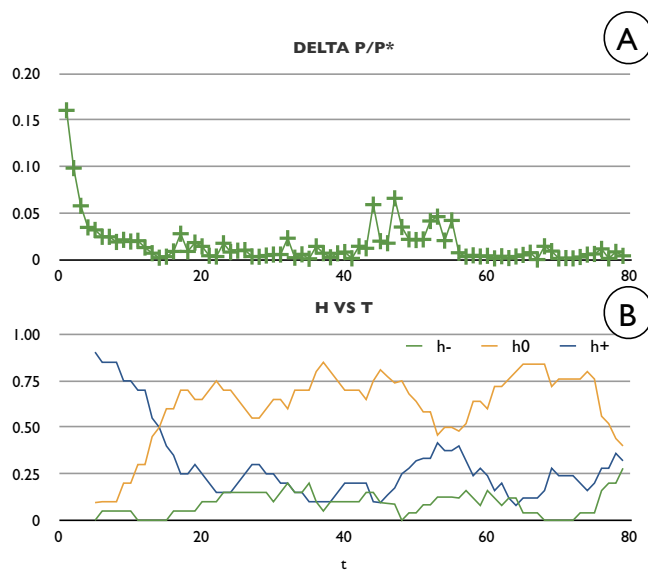


Fig. 4. (A) Time series of mean value over agents of $|p_t - p_{t-1}|/p^*$. (B) Time series of $\hat{h}_-(t)$, $\hat{h}_0(t)$, $\hat{h}_+(t)$, measured by aggregating and averaging the last 5 time steps for each t over those agents having full sales. All data from Experiment #21.

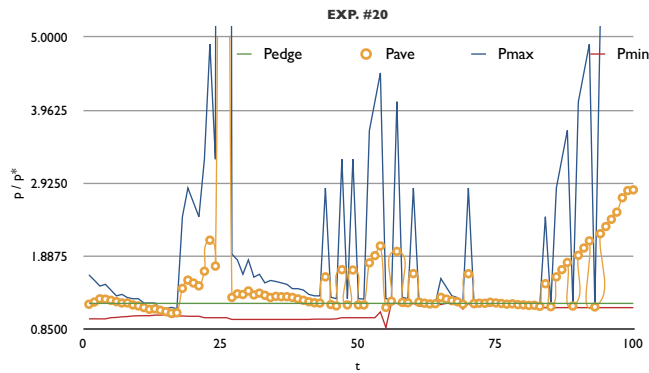


Fig. 5. Evolution of average, maximum and minimum offer price (rescaled with p^*). Also shown are the prices $p_{\text{edge}} = p^* + c/N$ and $p_1^* = p^*(N/(N-1))$. Experiment #20.

Fig. 4B shows measures of the empirical frequencies $\hat{h}_-(t)$, $\hat{h}_0(t)$, $\hat{h}_+(t)$ 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. Allowing for individual specificities, a general pattern observed across the experiments was that players started the game with ‘aggressive’ policies of increasing prices if sales reached capacity $\hat{h}_+ \approx 1$, $\hat{h}_0 \approx 0$, and became more cautious later on: \hat{h}_0 increases steadily, while \hat{h}_+ decreases correspondingly (see Fig. 4A). The speed at which agents become cautious, varies among experiments. Fig. 4A shows an example where after the initial trend, \hat{h}_0 and \hat{h}_+ reach a level with fluctuations, suggesting that agents do not stop searching for better strategies in a context where there is no simple best response rule. Also, it may be noted that the magnitudes of $\hat{h}_-(t)$, $\hat{h}_0(t)$, $\hat{h}_+(t)$ do not depend on the number of players, N .

4.3.4. Price convergence and price excursions

The interaction between experimental subjects generates identifiable patterns of organization, and a convergence from above to a condition where individual and average prices show small variations (resembling the findings of Kruse et al. (1994)). However, this ‘‘approximate’’ steady state is not an actual position of rest for the system, but it is subject to disturbances. The observed dynamics tend to show a first period of experimentation, with players who start with widely different strategies (although they share the knowledge of the basic structure of the game), and who react quite strongly to the outcomes they observe. In that case, 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 price (see again Fig. 3).

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 prices. As the average price falls, some agents seem to perceive incentives to search for higher profits through sudden price increases, which can generate market-level ‘‘upward excursions’’. These perturbations did occur in eight out of eleven experimental sessions. They tended to start with average sale prices above the value p_{edge} , that is, average prices above the suggested critical level defined in (5).

Also, in eight out of eleven experiments, the average selling price remained above the bound $p^*N/(N-1)$ at which only $N-1$ agents would realize full sales. 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’’. Such task is left for future research.

The idiosyncratic behavior of some agents can be experimentally relevant, particularly in games with small groups. For example, in experiment #20 reported in Fig. 5, the aggregate price was definitely 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 allowed the market evolve 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 $\bar{p} \approx p_{\text{edge}}$, and players engaged in several large excursions.

4.3.5. On mixed-strategies Nash equilibria

The previous discussion has focused on the correspondence of the experimental behavior with that generated by simple heuristics. A question that remains is whether the evidence could be consistent with more traditional concepts, particularly the mixed-strategy Nash equilibrium. Kruse et al. (1994) suggested that this could be tested by studying whether prices show an autoregressive structure, in which case the answer would be negative. That paper finds such a property in their experiments. The analysis of our data in this paper points in the same direction (see Table 9).

Table 9

Summary of autoregressive estimates of experimental prices. For each experiment the prices of all agents at a time step t formed vector \bar{p}_t , and the autoregressive model with up to two lagged prices was estimated ($\bar{p}_t = L_0 + L_1\bar{p}_{t-1} + L_2\bar{p}_{t-2}$, L_j regressors). The numbers corresponds to the fraction of agents with significant coefficients (up to the 10% significance level). In most cases, prices were significantly dependent on past values, ruling out independent and identically distributed prices.

| Experiment | N | All iterations | | | Iteration >25 | | |
|------------|----|----------------|-------|-------|---------------|-------|-------|
| | | L_1 | L_2 | L_0 | L_1 | L_2 | L_0 |
| 13 | 12 | 0.92 | 0.58 | 1.00 | 0.67 | 0.33 | 0.50 |
| 16 | 6 | 0.83 | 1.00 | 1.00 | 0.83 | 0.50 | 0.33 |
| 17 | 4 | 0.50 | 0.25 | 1.00 | 0.75 | 0.25 | 1.00 |
| 18 | 6 | 1.00 | 0.83 | 1.00 | 0.83 | 0.50 | 0.17 |
| 19 | 11 | 1.00 | 0.36 | 1.00 | 1.00 | 0.82 | 0.18 |
| 20 | 4 | 1.00 | 0.75 | 1.00 | 1.00 | 0.00 | 0.75 |
| 21 | 6 | 1.00 | 0.67 | 0.67 | 1.00 | 0.33 | 0.50 |
| 22 | 7 | 0.86 | 0.71 | 1.00 | 0.86 | 0.14 | 0.57 |
| 23 | 7 | 0.71 | 0.14 | 0.86 | 0.57 | 0.14 | 0.57 |
| 29 | 5 | 0.60 | 0.40 | 0.80 | 0.80 | 0.40 | 0.60 |
| 31 | 3 | 1.00 | 0.67 | 0.67 | 1.00 | 0.67 | 0.67 |

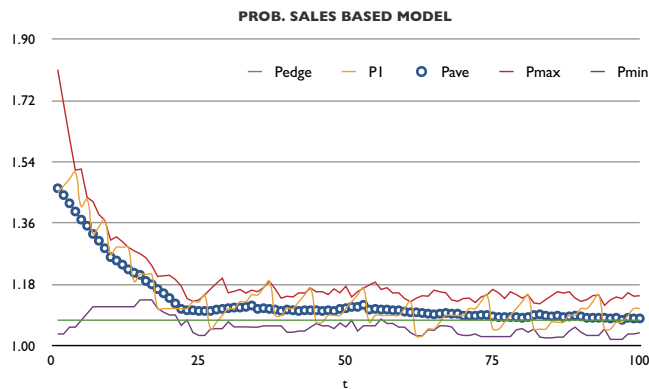


Fig. 6. Simulation of probabilistic sales-based. Average (p_{ave}), maximum (p_{max}) and minimum (p_{min}) price as a function of time, together with a sample trajectory (p_1) of an agent; also shown is $p_{edge} = p^* + c/N$. Parameters: $N = 10$, $p^* = 1$, $c = 0.75$, $\gamma_+ = 0.02$, $\gamma_- = 0.10$, $h_0 = 0.45$, and $h_- = 0.0$. All agents share the same behavioral parameters γ_+ , γ_- , h_0 , and h_- .

5. A modified sales-based model

As shown above, the elementary sales-based pricing heuristics provides a rough description of behavior, and can help to account for some aspects of the evolution of the experimental markets. However, the evidence of price setting in the experiments shows substantial deviations from that simple rule, particularly in the response of prices after the realization of full-capacity sales. This section introduces a relatively small change to the heuristics, and analyzes its performance in simulating market outcomes.

The central change in the rule presented in Section 3 consists in the introduction of a random choice of price changes following full sales, parameterized by probabilities h_- , h_0 , h_+ of decreasing, keeping constant, or raising the price, respectively. The price rule for agent i would then be:

$$\begin{aligned}
 p_i(t+1) &= p_i(t) - \gamma_- \text{ if } q_i < q^* \\
 p_i(t+1) &= p_i(t) + \begin{cases} \gamma_+ & \text{with prob. } h_+ \\ 0 & \text{with prob. } h_0 \\ -\gamma_- & \text{with prob. } h_- \end{cases} \text{ if } q_i = q^*
 \end{aligned}
 \tag{10}$$

($h_- + h_0 + h_+ = 1$), where $h_- \ll h_+$, h_0 (as seen from the experimental observations). The dynamics of the model would thus depend on the behavioral parameters (γ_+ , γ_- , h_0 , h_+), and the parameters of the game (N, M, c, q^*). Fig. 6 shows one simulated evolution of the average posted price, p_{ave} and the price range (p_{min} , p_{max}), together with the price of one individual agent. The simulation starts from a random initial condition of prices across suppliers, assumed to be identical in their pricing




Fig. 7. Comparison of the average price estimate (15) in the probabilistic sales-based model against numerical simulations as a function of h_0 . Parameters: $N=20$, $\gamma_- = 0.10$, $\gamma_+ = 0.02$, and $h_- = 0$.

heuristics.¹⁶ 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.

In the simulated market evolution, the price range decreases steadily, with a gradual decline in the average posted and selling prices. This is consistent with most of the experimental evidence (with the only exception of #20, as mentioned above). Given the dispersion associated with the random initial condition, the cheapest suppliers raise their prices, while the higher-priced agents cannot still sell, until the range of posted prices has an order of magnitude similar to the parameter of price reductions ($\approx \gamma_-$).¹⁷ At that point, then the condition of not selling is more or less uniformly distributed among agents.

Using the parameters of the model, it is possible to generate an analytical approximation of the average posted price to which the simulated market would converge, adapting the procedure used in Section 3. Denote as ΔP_+ the aggregate value of the changes in prices set by firms with full-capacity sales in the previous period. This variable 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 (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 an aggregate steady state, 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 then:

$$\frac{(N - N_{ns})}{N_{ns}} = \frac{N_s}{(N - N_s)} = \frac{\gamma_-}{(h_+\gamma_+ - h_-\gamma_-)} \tag{11}$$

Solving (11) for N_{ns} ,

$$N_{ns} = N \left(1 - \frac{1}{1 + ((h_+\gamma_+ - h_-\gamma_-)/\gamma_-)} \right) \tag{12}$$

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

$$\langle p \rangle = \frac{M}{(q \cdot N_s)} = p^* \frac{N}{(N - N_{ns})} \tag{13}$$

Thus,

$$\langle p \rangle \approx p_{est} \equiv p^* \left(1 - h_- + h_+ \frac{\gamma_+}{\gamma_-} \right) \tag{14}$$

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 would make the average price go closer to the competitive market outcome.

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,

$$p_{est} \approx p^* \frac{N}{N - \lfloor N_{ns} \rfloor} \tag{15}$$

where $\lfloor N_{ns} \rfloor$ is the integer value of the number of agents with less than full sales derived from (12).

The simulated results with artificial agents using the probabilistic sales-based heuristics (10) validate the approximation (15). Iterating over 1000 time steps, taking averages over the last 500 steps, and averaging over 10 random initial conditions (initial agents prices are uniformly distributed in the range $p_i(0) \in [p^*, 2p^*]$), Fig. 7 shows that the numerical simulations closely follow (15) as a function of h_0 . The price steps 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). This number can be calculated analytically. In the case of the parameter values in Fig. 7, then $h_0^* = 0.737$.

¹⁶ The 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 $p_i(0) \in [p^*, 2p^*]$.

¹⁷ In the simulation shown in Fig. 6 the average price range for the last half rounds was of 0.107, very near the size of $\gamma_- = 0.10$.

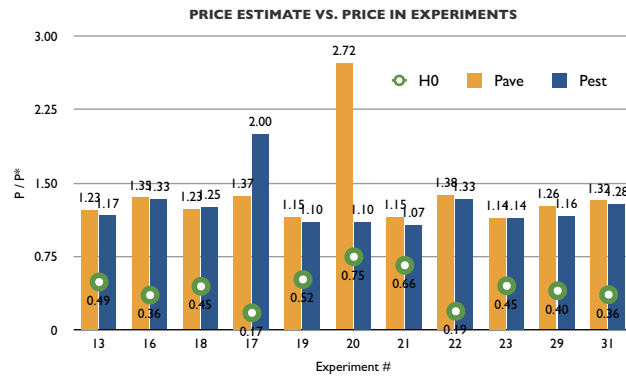


Fig. 8. Comparison of the average price estimate p_{est} (15) in the probabilistic sales-based model against the experimental price average p_{ave} . The experimental values of h_0 are also shown. The behavioral parameters $h_{+,-}$, $\gamma_{+,-}$, and the average price p_{ave} for each experiment were estimated taking averages over the last 30 rounds of play. Notice that the prices shown in the figures are rescaled as ratios with respect to the competitive price p^* .

5.1. Contrasting the simulation-based market average price with experimental evidence

The calculated average prices resulting from the computation of p_{est} correspond to values that would obtain after a process of market convergence has taken place. Consequently, those estimated prices can be compared with those observed experimentally in the last rounds of play. The results of such exercise are shown in Fig. 8. For each experiment, the corresponding p_{est} is calculated using the behavioral parameters h 's and γ 's obtained by averaging across agents and rounds of play the observed frequencies and absolute values of price changes in the last 30 market days. In most cases, the estimated price p_{est} is below, but not far from, the experimental market price. This is another way of showing how the probabilistic sales-based model provides a good approximation of the prices observed in experimental markets, even under the assumption of constant behavioral parameters. As exceptions, two experiments (#17 and #20) stand as outliers, where the estimates performed badly. These were cases with small numbers of players (four in each case), characterized by large changes in individual prices (see Table 8). Only in these experiments the hypothesis of convergence is much less appropriate than in the others.

6. Conclusions

This paper studied price determination in a market setup with potential real-world counterparts. The Bertrand–Edgeworth price-setting games, where suppliers face capacity constraints, do not admit in general deterministic Nash equilibria, while the mixed-strategy equilibria are difficult to characterize. In those conditions, it seems reasonable that agents look for simple, practical heuristics to guide their behavior. We have considered two such families of pricing rules, and analyzed the resulting market outcomes through agent-based. These hypothetical patterns were contrasted with the evidence obtained from laboratory experiments. The observed behavior of experimental agents supported the relevance of some simple pricing criteria, and also led to revisions of the very elementary pricing rules postulated at the beginning of the analysis.

Placed in an environment that does not allow a ready identification of profit maximizing strategies, the behavior of agents in the experiments showed approximate correspondence with some basic heuristics. At the same time, the experimental price-setters did not use those strategies in strictly mechanical ways. The search for improved performances was reflected in changing parameters of response to realized outcomes and occasional “excursions” where some players implemented discrete changes in their reaction functions.

In most instances, agents who do not realize full-capacity sales in a period react by lowering prices; when sales reach full 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” stage 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 decline in average prices, characteristic of the convergence from above to the “competitive range” reported also in Kruse et al. (1994). This pattern can be simulated in a model with artificial agents who use a sales-based pricing rule. It is also possible to approximate the average price analytically, using parameters obtained from the experiments.

The analysis established in Section 2 suggests that, when prices approach the competitive level and get close to the critical value p_{edge} defined in Eq. (5), there is room for profitable deviations by agents who raise their price with sufficient intensity. The experimental results actually show “upward excursions” of the prices set by some agents in markets which previously were in a more or less tranquil condition. Thus, the dynamics 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}).

The analysis may be extended in several directions. An important feature of the experimental design was that the market information available to agents was limited to the history of their own transactions. Thus, the potential decision rules can only condition prices on those variables. The question thus arises of how experimental behaviors would be defined with other information structures, that allow observing data referred to other agents. Analyses and experiments using other specifications of the demand function may also serve as trials for the range of validity of the results obtained in the paper.

Within the general frame of the present setup, 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 parameters such as 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 either 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 can improve the understanding on how individual and aggregate performance depend on specific market conditions, and provide additional links to the previous literature. Regarding simulations, a natural next step would be to consider explicitly the adaptation of the quantitative parameters 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 market data.

Acknowledgements

We thank the co-editor and two anonymous referees for useful and insightful comments that helped to improve the paper. We also acknowledge gratefully the comments received from Dan Friedman, Omar Chisari, and participants in seminars at the Universidad de San Andrés, Universidad de Buenos Aires, Asociación Argentina de Economía Política, Instituto Balseiro and TANDAR. We thank especially the research assistance provided by Nicolás Aragón, Fiorella Benedetti, and Brian Feld at different stages of the paper. The usual caveat applies.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.jebo.2014.04.027>.

References

- Abbink, K., Brandts, J., 2005. Price competition under cost uncertainty: a laboratory analysis. *Econ. Inq.* 43, 636–648.
- Allen, B., Hellwig, M., 1986. Bertrand–Edgeworth oligopoly in large markets. *Rev. Econ. Stud.* 53, 175–204.
- Bendor, J., Mookherjee, D., Ray, D., 2001. Aspiration-based reinforcement learning in repeated interaction games: an overview. *Int. Game Theory Rev.* 3, 159–174.
- Brandts, J., Guillén, P., 2007. Collusion and fights in an experiment with price-setting firms and production in advance. *J. Ind. Econ.* 55, 453–473.
- Buchheit, S., Feltovich, N., 2010. Experimental evidence of a sunk-cost paradox: a study of pricing behavior in Bertrand–Edgeworth duopoly. *Int. Econ. Rev.* (forthcoming).
- Cason, T., Friedman, D., 1999. Learning in a laboratory market with random supply and demand. *Exp. Econ.* 2, 77–98.
- Cason, T., Friedman, D., 2003. Buyer search and price dispersion: a laboratory study. *J. Econ. Theory* 112, 232–260.
- Cason, T., Friedman, D., Wagener, F., 2005. The dynamics of price dispersion, or Edgeworth variations. *J. Econ. Dynam. Control* 29, 801–822.
- Dasgupta, P., Maskin, E., 1986. The existence of equilibrium in discontinuous economic games. II: Applications. *Rev. Econ. Stud.* 53, 27–41.
- d’Aspremont, C., Gabszewicz, J.J., 1985. Quasi-monopolies. *Economica* 52, 141–151.
- Davis, D., Holt, C., 1994. Market power and mergers in laboratory markets with posted prices. *Rand. J. Econ.* 25, 467–487.
- De Francesco, M., Salvadori, N., 2010. Bertrand–Edgeworth games under oligopoly with a complete characterization for the triopoly. Working Paper.
- Duffy, J., 2006. Agent-based models and human subject experiments. In: Testfatsion, L., Judd, K. (Eds.), *Handbook of Computational Economics*, vol. 2. Elsevier B.V., pp. 949–1011.
- Fonseca, M., Normann, H.-T., 2008. Mergers, asymmetries and collusion: experimental evidence. *Econ. J.* 118, 387–400.
- Fonseca, M., Normann, H.-T., 2013. Excess capacity and pricing in Bertrand–Edgeworth markets: experimental evidence. *J. Inst. Theor. Econ.* 169, 199–228.
- Gilboa, I., Schmeidler, D., 2001. *A Theory of Case-Based Decisions*. Cambridge University Press, Cambridge UK.
- Holt, C., Solis-Soberon, F., 1994. The calculation of mixed-strategy equilibria in posted-offer markets. In: Isaac, R.M. (Ed.), *Research in Experimental Economics*, vol. 5. JAI Press, Greenwich.
- Izquierdo, L., Izquierdo, S., Nicholas Gots, J., Gary, P., 2007. Transient and asymptotic dynamics of reinforcement learning in games. *Games Econ. Behav.* 61, 259–276.
- Karandikar, R., Mookherjee, D., Ray, D., Vega-Redondo, F., 1998. Evolving aspirations and cooperation. *J. Econ. Theory* 80, 292–331.
- Ketcham, J., Vernon, S., Arlington, W., 1984. A comparison of posted-offer and double-auction pricing institutions. *Rev. Econ. Stud.* 51, 595–614.
- Kruse, B.J., Stephen, R., Stanley, R., Smith, V., 1994. Bertrand–Edgeworth competition in experimental markets. *Econometrica* 62, 343–371.
- Maskin, E., Tirole, J., 1988. A theory of dynamic oligopoly. II: Price competition, Kinked demand curves and Edgeworth cycles. *Econometrica* 56, 571–599.
- Selten, R., 2004. Learning direction theory and impulse balance equilibrium. In: Friedman, D., Alessandra, C., Reinhard, S. (Eds.), *Economics Lab: An Intensive Course in Experimental Economics*. Routledge, London, UK, pp. 133–140.
- Selten, R., Abbink, K., Cox, R., 2005. Learning direction theory and the winner’s curse. *Exp. Econ.* 8, 5–20.
- Selten, R., Buchta, J., 1999. Experimental sealed-bid first price auctions with directly observed bid functions. In: Budescu, D.V., Ido, E., Rami, Z. (Eds.), *Games and Human Behavior: Essays in Honor of Amnon Rapoport*. Lawrence Erlbaum Associates, Mahwah, NJ, pp. 79–102.
- Selten, R., Stoecker, R., 1986. End behavior in sequences of finite prisoners’ dilemma supergames: a learning theory approach. *J. Econ. Behav. Organ.* 7, 47–70.
- Tirole, J., 1988. *The Theory of Industrial Organization*. The MIT Press, Cambridge, MA.