Grid adjustment to Nash equilibrium

An experimental re-interpretation of the Bertrand paradox

Enrique Fatas[1]

LINEEX - University of Valencia
and University of Texas at Dallas

Ernan Haruvy

University of Texas at Dallas

Antonio J. Morales

LINEEX- University of Malaga

Abstract

The Bertrand paradox is that oligopolisitic firms are not generally observed to play Nash equilibrium in Bertrand competition. This finding has been observed in the lab in intermediate horizons. In a somewhat longer horizon, we show that firms do appear to nearly converge to Nash equilibrium. However, convergence does not appear consistent with naive adjustment, but rather with sophisticated players. We examine the data and find that in addition to naïve adaptation, subjects exhibit behaviour consistent with Nash equilibrium beliefs involving coarser bidding grids. In addition, equilibrium behaviour appears to follow an adjustment process as well, where the bidding grid is gradually refined. Existing empirical works on Bertrand competition have found evidence for boundedly rational models. We find that such models are useful in organizing behaviour in early stages of the game, but less so in later stages.

KEYWORDS: Bertrand, Price Competition, Experiments, Learning.

JEL classification numbers: C91, D43, L13

[1] Corresponding author: *efatas@utdallas.edu*

1. Motivation

A primary concern with respect to oligopolistic markets is the relationship between market structure and competitiveness. And in this regards, theoretical models yield quite different results depending on whether firms are assumed to compete in quantities (à la Cournot) or in prices (à la Bertrand). In fact, one of the striking features of price competition is that equilibrium prices are equal to marginal costs regardless of the number of competitors. As few as two firms competing on prices will bring about competitive behaviour and the ensuing zero profit condition.

This sharp implication is in fact quite rarely observed in real life situations as price dispersion with prices well above marginal costs is the norm (see Wellford, 2002; Baye and Morgan, 1999). Theoretical explanations for this so called *Bertrand paradox* have come by relaxing some assumptions of the basic model: transactions costs, product differentiation, cost heterogeneity, incomplete information, etc (see Baye and Kovenock, 2008). However, the difficulty in obtaining information about these quantities makes difficult to ascertain which of these explanations, if any, underlies the Bertrand paradox in empirical studies.[2]

Given the significance of price competition in economics it is hardly surprising that one of the seminal uses of the experimental methods in economics is the analysis of oligopolistic markets, the classical reference being Fouraker and Siegel (1963). In the lab, researchers can overcome the confounds posed by the empirical data by constructing experimental models that replicate most of the assumptions of formal models, thus testing the theory in its own domain. To the surprise of some, price dispersion persists in the lab after controlling for those explanations offered by many to account for the Bertrand paradox (see Holt, 1995 and Plott 1982, 1989 for reviews of the literature).

Moreover, the most basic version of the Bertrand model, with fixed demand and constant returns, has been shown to fail in the lab. Recently, Dufwenberg and Gneezy (2000, hereafter DG), in a paper aimed at testing the influence of market

[2] Some of these limitations are tempered using studies on competition among internet sites. However, price dispersion is also found. (See for example Baye and Morgan, 2001)

concentration on price competition, conclude that "the Bertrand solution does not predict well when the number of competitors is two, but (after some opportunities for learning) predicts well when the number of competitors is three or four."

This result called for finding alternative explanations and DG offers one in terms of boundedly rational players. Since then, a number of explanations has been proposed, including collusion (Suetens and Potters, 2007), boundedly rational equilibrium models such as quantal response equilibrium and epsilon equilibrium (Baye and Morgan, 2004), and step-level reasoning with altruism (Gneezy, 2005), to name a few.

When we double the length of the DG horizon we find that behaviour actually appears to nearly converge to Nash equilibrium in duopolies, suggesting that Nash equilibrium predictions are more informative than implied from past research. While the converged outcomes correspond to Nash equilibrium, the adaptive process that leads to such an outcome is not as straightforward. Models with adaptive gradual movement towards best response to the past—dynamic process whose variations include well-known models such as EWA (Camerer and Ho, 1999) and reinforcement learning (Erev and Roth, 1998)—do not appear to explain the data well. Instead, it appears that subjects behave as if they have Nash equilibrium beliefs with coarse bidding grids that are refined over time.

The remainder of the paper is as follows. In Section 2 we describe the experimental design. In section 3, we describe the general accepted models for such settings, and why they may or may not explain the patterns. In section 4, we describe the data and formally test the models. In section 5, we draw conclusions.

2. Our experiment

2.1 The games

We consider two duopoly price competition models with homogeneous product, no capacity constraints, inelastic demand q and complete information. For given prices, the demand is served by the firm charging the smallest price, the tie-breaking rule being that demand is equally shared. The only difference between the two models is the cost function: linear versus quadratic. In the linear cost game, the cost function for given production level *q* is

$$c_i(q) = c\,q \text{ for i} = 1{,}2 \quad (1)$$

whereas in the quadratic case the cost function is

$$c_i(q) = c\,q^2 \text{ for i} = 1{,}2 \quad (2)$$

where *c>0*. Note that in the linear case, the marginal cost is equal to c, while in the quadratic model, marginal costs are no longer constant. For the present experiments, c=5, q=20 for the lowest price firm, q=10 for each firm in case of a tie, and q=0 for the higher price firm.

2.2 Experimental design and procedures

The experimental design consisted of two treatments. Each treatment was examined separately with a different group of subjects. Treatments differ only in the cost structure.[3] The main elements of the experimental design are summarized in table 1 below.

[Table 1 around here]

Our study reports the results of computerized experiments conducted at LINEEX, the experimental laboratory at the University of Valencia, using z-Tree (Fischbacher, 2007). Subjects were recruited electronically with the LINEEX web based system. None of them had ever participated in a similar experiment before, even when some of them had already participated in different experiments. By participating in the experiment a participant made an average of US$28, and experiments took less than one hour to run.

---

[3] A set of instructions translated from the Spanish is available upon request.

After the experiment, subjects were debriefed by an on screen questionnaire including some questions about their strategies (just to double check that they did understand the structure of incentives). Given their replies and the procedure, we are confident that both the tasks and the incentives were understood by subjects.

Participants were assigned to one treatment when recruited. When they reached the laboratory, they were randomly allocated to a cubicle choosing a numbered chip (with cubicles' numbers) from a basket. At the beginning of every session they were randomly allocated to a group of two and remained together throughout the whole experiment.

Following standard experimental procedures, instructions are read aloud and subjects play 20 trial rounds with the computer (it uniformly randomizes prices). Subjects are not paid for this trial rounds. Once the trial rounds are over, every subject states a price from 0 to 500 (one decimal was allowed, but very rarely used). The subject with the lower of the two prices wins the demand of 20 and a profit of the winning price times 20. The other subject wins zero demand and zero profit. In case of ties, demand and profits are equally divided. After each round, subjects receive information about prices and profits in their market. They also have a history table in their screen, including all past values of their performance. At the end of the experiment, all subjects were privately paid using sealed envelopes with the number of their cubicle typed in the outside. It was impossible to link names and decisions.[4]

3. Models and theoretical development

3.1. Static Models

Static models ignore the evolution of prices over time. Static models may be applied to the initial period play (step-level thinking, Gneezy, 2005), the end-period (Nash equilibrium), or the distribution of choice over the pooled periods (Baye and Morgan, 2004).

[4] Privacy and anonymity were strictly guaranteed. Subjects were never asked to introduce any personal information in the system. The full protocol is available from authors upon request.

3.1.1 Nash equilibrium prediction

Nash equilibrium requires the highest level of rationality of all models in that it requires each player to best respond to the other's price. The Nash prediction for the linear Bertrand game is that prices equal marginal cost 5, whereas for the quadratic game, the Nash prediction gives the following interval [50, 150]. We state these results as a proposition with the proof in the appendix.[5]

Proposition 1. The set of Nash equilibria is

(i) For the linear game, $p^* = c$

(ii) For the quadratic game, $p^* \in \left[\frac{1}{2}cQ, \frac{3}{2}cQ\right]$

3.1.2. k Nash model -- Nash prediction with a coarse grid

Due to cognitive costs, players may not evaluate strategies in terms of the smallest possible increment. Assume that subjects when competing in the linear cost Bertrand simplify the continuum by using a common grid of size k, i.e. they only consider prices which are multiples of k. In such a setting, it can be shown that together with a price equals to the marginal cost, the linear game has two additional *Nash equilibria*: k and 2k. It is important to note that 3k, 4k, and higher multiples of k on are not. Hence, our theory does not predict any symmetric profile based on prominent numbers to be equilibrium; instead, only some of them will be stable solutions for this kind of games.

In the quadratic cost case, we can state the following proposition (proof in appendix) regarding equilibria for a common grid of size k.

[5] Proposition 1 assumes that the strategy space is a continuum. DG has already pointed out that in experimental settings, subjects are usually restricted to choose integers and that this restriction alters the set of Nash equilibria of the game. A generalized version of this idea is analysed under the name k-Nash Model in the next subsection.

Proposition 2. *Let n^- (n^+) be the minimum (maximum) multiple of k belonging to the interval* $[50{,}150]$. *Then, the set of Nash equilibria of the k-game is {(p,p) with p = kn for n = n^- ,..., n^+, n^++1, n^++2}.*

Rubinstein (1998) considers the tendency to simplify problems as one of the fundamental sources of bounded rationality. Also, it has received strong psychological support (e.g., Payne et al., 1988). It may well be the case that people simplify problems by approximating an interval with a coarser grid which contains fewer elements and is therefore easier to contemplate. From this point of view, discretization is a conscious decision of subjects on the basis of complexity considerations.[6] Note that this phenomenon can be viewed as an example of categorisation, which has been identified by Gigerenzer et al. (1999) as a building block of bounded rationality. Quite recently, there is a strand in the literature analyzing the impact of categorization on game theoretic settings. See for example Jehiel (2005).

Reducing the number of points considered on a grid could also be due to uneven prominence of some numbers relative to others. Albers (1999) models the construction of numerical responses and the perception of numerical stimuli in the decimal system. Experimental evidence abounds, Neugebauer and Selten (2003) noted, in an experiment of first-price sealed-bid auctions, that 58 percent of the observed bids in the first round were in multiples of five; similarly, Whynes et al. (2005) showed that in studies designed to elicit the willingness to pay for medical treatments in the interval [0,200], 60% of the answers were in the set {10, 20, 50, 100}.

3.1.3. Quantal response equilibrium

The quantal response equilibrium (QRE) model was first proposed by McKelvey and Palfrey (1995). Anderson, Goere and Holt (1998) expanded the technique for

[6] This last motivation gets reinforced when we get into the realm of computer science and become aware of the computational complexity of the Nash equilibriums of games. In fact, the problem remains essentially open in the case of n-person games, and as Papadimitriou (2001) states "Together with factoring, the complexity of finding a Nash equilibrium is (…) the most important concrete open question on the boundary of P today." There are however some recent advances which have been done by using discrete constraint solving and interval constraint solving techniques.

games with a continuous action space. Baye and Morgan (2004) applied QRE to Bertrand competition data from Dufwenberg and Gneezy (2000) and found that it greatly outperforms the pure Nash equilibrium model. No error term was permitted around the Nash prediction in that comparison. Hence the finding that Nash predictions perform badly may or may not be surprising. Models with a distribution over choices should be expected to outperform models with a single point prediction when applied to empirical data. Indeed, when Baye and Morgan (2004) compared QRE to epsilon-equilibrium-- another bounded rationality concept with a distribution over choices, the comparison did not yield a clear winner.

The QRE model assigns players probabilistic, typically logistic, decision rules that map expected payoffs to choices. Expected payoffs are derived based on beliefs about the strategies of others and QRE imposes that all players hold correct beliefs about the probability distributions of all other players' strategies.

This can be a complex problem to solve analytically but computationally it is fairly straightforward. Denote expected payoff to a price of k by Π_k and denote the observed empirical probability of event E by P(E). We first compute expected payoffs for each of the 500 possible prices based on the observed empirical frequency in the game as follows:

$$E\Pi_k = 20kP(price > k) + 10kP(price = k) \tag{3}$$

We then compute a logit best response function as follows:

$$\hat{P}(price = k) = \frac{\exp(\lambda \Pi_k)}{\sum_{j=1}^{500} \exp(\lambda \Pi_j)} \tag{4}$$

$\hat{P}(price = k)$ is the estimated probability of observing price k. We next substitute the computed logit probabilities from the last stage for the empirical probabilities. That is, we plug $\hat{P}$'s for the P's on the Leith-hand side. We recompute the logit probabilities and repeat the procedure. We stop when the probabilities become stable. Note that this function involves the estimation of a single parameter λ, which is estimated to maximize the joint likelihood of observed prices.

3.1.4 Step-level reasoning

Gneezy (2005) reported a Bertrand price competition experiment with 10 price points from 1 to 10. In that setting, a naïve player might conjecture that other players are so uninformed that are equally likely to choose any of the 10 price points. Then the best response is to submit a price of 5. Such naïve thinking corresponds to the level-1 of Stahl and Wilson (1994, 1995). A slightly more sophisticated type (level-2) may anticipate others to be of level-1, and therefore the best response price is 4. In a similar manner, a level-3 player would bid 3, level-4 would bid 2, and level 5 and higher would bid 1.

Normally players with levels of rationality greater than 2 or 3 are rarely observed, with the exception of Nash (Bosch-Domènech et al., 2002). While Gneezy (2005) rejected a simple model of step-level thinking (using the model Camerer, Ho and Chong, 2004), he found evidence for such behavior with the addition of a cooperative type.

In contrast to Gneezy (2005), we have a range of 5001 possible choices (ten per integer between 0 and 500, plus the limit value). In that range, for any reasonable error term, the difference between level-1 and level-2 is too small to be of any statistical significance, in contrast to a range of 10. Likewise, levels 3, 4, and higher are really not distinguishable. We are therefore left with only a level-1 player as a proxy for the step-level model. Luckily, level-1 is generally the largest segment of players in any such estimation, so any predictive power of the step-level model should be reflected in the predictive success of level-1.

3.2. A simple adaptive model with a Nash equilibrium player type

Typically, adaptive models would imply some gradual adjustment towards recently observed market prices (as long as market prices are above 5 in the linear case and above 100 in the quadratic case). In a reinforcement learning framework

(Roth and Erev, 1995; Erev and Roth, 1998), the higher price competitor would get a reinforcement of zero for that period, which would push him away (discourage him) from the bid he placed in that period. The winner in the period, in contrast, would get a significant positive reinforcement, encouraging him to repeat that bid.

In a belief learning framework or a hybrid learning framework (Camerer and Ho, 1999), the loser would be aware of the foregone payoffs involved with bids not chosen. In other words, the bidder knows what he could have gotten had he selected a lower price. The highest foregone payoff would be at the market price, and therefore the propensity to select the market price would increase the most.

Here we attempt to be agnostic about the most appropriate learning model. Instead, we pursue a gradual adjustment model (Crawford, 1995) which does not involve utility and propensity specifications and instead merely follows a gradual adjustment towards the recently observed market price.

$$bid(t) = (1-\beta)bid(t-1) + \beta P(t-1) \qquad (5)$$

The beta parameter is an inertia parameter. The above model would not account for non-myopic expectations, for expectations about others, or for the relative payoff differences between possible strategies. As such, it is hardly expected to have a better fit than more sophisticated models.

Alternative modes of behaviour here are Nash equilibrium and the k Nash equilibrium discussed earlier. The k Nash type is implemented by computing the bin size of the grid based on observed bid increments. It then computes the Nash equilibrium corresponding to the estimated grid and selects the equilibrium price. Subjects are assumed to gradually adjust towards Nash equilibrium using the same inertia parameter used by the adjustment to market price described earlier. *NE* denotes either pure Nash or k Nash, the two models we will contrast.

$$bid(t) = (1-\beta)bid(t-1) + \beta NE(t-1) \qquad (6)$$

Note that the inertia parameter nests static notions of Nash equilibrium in the model, as well as dynamic notions of Nash as an adaptive concept. If the inertia parameter is 1, then the pure Nash prediction does not change from one period to the next. Likewise, for an inertia parameter of 0, a participant's bid remains at its initial value and never changes. With an inertia parameter between 0 and 1, the prediction for this type follows a gradual adjustment from the initial bid to Nash equilibrium.

For the k Nash model, there is another dynamic that is not as obvious from the above equation-- the dynamic of the bidding increments. As participants adjust towards Nash equilibrium, their bidding increments naturally shrink. This by construction shrinks the bin size and therefore brings the k Nash equilibrium closer to pure Nash equilibrium over time. In summary, the k Nash dynamic allows for adjustment towards Nash equilibrium based on past prices as well as past increments. Finally, we use a mixture model framework to allow for multiple subpopulations, each following a different dynamic.

4. Results

4.1 Descriptive statistics

Table 2 below presents some basic statistics about prices in our two experimental treatments (Linear and Quadratic). Table 2 includes both average posted and market prices for the first and last period, and all 20 rounds of the experiment. In line with previous experiments, prices in the first period are relatively far from the equilibrium predictions. In the linear (quadratic) treatment, the average posted price is 161.41 (206.58), while the average market price is 95.00 (154.52).

However, a remarkable convergence process leads prices to the predicted outcomes, especially in the case of the quadratic treatment. In the linear (quadratic) treatment, the average posted price in the last period is 55.00 (159.05), while the average market price is 46.30 (149.00). Note that the average market price in the last round is only 1 price unit away from the Pareto dominant equilibrium of the quadratic game (150). Given that subjects were able to choose any price from 0.0 to 500.0, we find this convergence striking.

[Table 2 around here]

More interestingly, Table 2 shows that this convergence is strong enough to get prices very close to the predictions for the 20 rounds, in average. Posted (market) prices go down to 84.94 (66.61) in the linear treatment and get really close to the efficient equilibrium in the quadratic, in which the average posted price is 165.72 and the average market price is 148.82. Figure 1 presents the convergence across rounds for both treatments and posted and market prices.[7]

[Figure 1 around here]

4.2 Econometric analysis: Static Models

We first examine prices in periods 1-5. A histogram of posted and market prices (Figures 2 and 3) shows the data in periods 1 to 5 is fairly diffused in the interval. From early on, there is some concentration of choices near the absolute Nash equilibrium (5 in the linear case; 50-150 in the quadratic case). However, most of the data falls to the right of the Nash equilibrium.

[Figure 2 around here]

[Figure 3 around here]

Figure 4 below shows the cumulative distribution functions for the three models we tested—Quantal Response Equilibrium (QRE), the level-1 model and the Nash model. The maximum likelihoods are -1365, -1368 and -1307 for QRE, Level 1 and Nash, respectively, in the first five periods of Linear Bertrand case and -2525, -2589

[7] Graphs present for every period the median price (central horizontal bar) inside a whiskers box (containing prices in the 25%-75% percentile). The vertical lines capture price range from upper to lower adjacent values and the hollow circles account for outside values.

and -2362 for QRE, Level 1 and Nash, respectively, in the first five periods of the Quadratic Bertrand case. Level-1 and QRE are nearly indistinguishable in the Linear Bertrand case, whereas QRE slightly outperforms Level-1 in the Quadratic Bertrand case. The Nash model greatly outperforms the other two in both scenarios. We see a similar pattern in the last five periods of both treatments. However, the magnitude of the bias of the boundedly rational models is larger.

[Figure 4 around here]

4.3. Adaptive Behaviour

Our examination in section 4.1 of price patterns over time showed that prices declined over time in both treatments and that while bids and market prices began diffused, they appear to converge to Nash equilibrium over time. This by itself suggests some dynamic to the data. Models proposed in the past to deal with Bertrand competition over time have not been adaptive. Baye and Morgan (2004), for example, pooled data over periods, whereas Gneezy (2005) looked at one-shot behavior. Showing that subjects are adaptive suggests looking at dynamic models.

An important question pertains to the manner by which subjects adjust behaviour. Naïve myopic subjects with short memories are likely to be pulled towards each other's prices, where losers would lower their prices somewhat and winners would raise it somewhat. This is not behaviour that lends itself to rapid convergence to equilibrium.

We first test whether we can find evidence for naïve myopic adjustment. We classify bidders into winners and losers, excluding ties. By construction, there are equal numbers of winners and losers. We then look at how what proportion of winners raise their prices following a win, and compare this to the proportion of losers that do the same. We further look at the proportion of each group keeping their prices the same. The reason for looking at both of these variables is that belief-based models would suggest some partial adjustment by the winner towards the losing price and by the loser towards the winning price, whereas reinforcement based models would suggest repeat behaviour by the winners. Table 3 below shows that approximately half of winners in both treatments raise their prices, whereas the vast majority of losers lower their prices. This is consistent with very naïve adjustment.

[Table 3 around here]

Now that we showed evidence in favour of naïve myopic adjustment, we will show the evidence *against* naïve myopic adjustment. Specifically, we show that in addition to adjustment durations, the data has downward price jumps. We define *adjustment durations* as consisting of consecutive periods where either of the two competitors raises their bids. To prevent ties from disrupting such durations, we liberally (in this case only) consider price increases as compared to bids in the previous three periods, as opposed to one previous round. Ties can still disrupt the continuity of adjustment durations, but only with a three-period convergence. Clearly, naïve adjustment would suggest very long, if not permanent, adjustment durations. Table 4 below shows that, in contrast to the prediction of fully naïve adjustment, there are also durations (consecutive periods) of downward movements. That means that there are durations where both competitors lower their prices in consecutive periods. This would only be consistent with somewhat forward looking non-myopic subjects. On average, adjustment durations last 5 to 7 periods. A downward movement, in contrast, lasts on average 1 to 4 periods (see breakdown by treatment).

[Table 4 around here]

The evidence for and against naïve myopic adjustment suggests that potentially two subpopulations co-exist in the data. As such, a dynamic mixture model may be appropriate here. As discussed in section 3.2, we allow for two subpopulations. The first subpopulation consists of naïve myopic adapters. Participants of this type adjust each period towards the recently observed market price. A second subpopulation adjusts towards Nash equilibrium. The model involves three parameters. Table 5 below present the main results of this model. The first parameter is the proportion of participants classified as fitting under the Nash adjustment model. The second parameter is the standard deviation of the error term around the prediction, where the error is assumed to be distributed normally. The last parameter is the inertia parameter, where inertia of 1 implies no adjustment and inertia of 0 implies full instant adjustment.

[Table 5 around here]

The estimated log likelihood for this model is -1386.33. In comparison, a model with only an adaptive type (no Nash) yields a maximized log likelihood of -1421.19. A model with a pure Nash type yields an estimated zero proportion of players falling under pure Nash type (all classified as naïve adaptive) and the same maximized likelihood as a single segment adaptive model. Hence, the adjustment model implies a dynamic involving k Nash choice with coarse grids.

5. Conclusions

Bertrand price competition is one of the most useful models of markets and one that naturally lends itself to laboratory experiments. The conventional wisdom on the Nash prediction in such a setting had been that it was not very useful. We show that for a 20-period run with the standard two-player Bertrand game, Nash equilibrium is predictive. Moreover, for a quadratic-cost case, the Nash prediction is within 1 of both the mean market price in the last round and in all 20 rounds. This is a good predictive power.

With that finding, we seek to re-examine past models that have been used to capture data for somewhat shorter horizons. We find that these models, QRE and level-1 in particular, are not as good in fit as Nash equilibrium with a normally distributed error term at capturing choice patterns. Past comparison to Nash equilibrium have been somewhat unfair in that regard since Nash equilibrium was the only model considered with no error term around the prediction.

In examining a dynamic adjustment model of the data, gradual adjustment towards the pure Nash equilibrium is not found to be a useful model of adaptive behavior. However, an alternative approach—extending Nash thinking to a coarse grid model—appears to configure well into the dynamic adjustment model and capture the data well. Under that adaptive approach, subjects adjust the grid they use over time to capture the transition from the early game to the late game. We find that this model outperforms naïve adjustment models.

6. References

Albers, W. (1999), "Prominence Theory as a Tool to Model Boundedly Rational Decisions", in *Bounded Rationality, The Adaptive Toolbox*, ed. By G. Gigerenzer and R. Selten, 297-318, MIT Press, Cambridge.

Anderson, S., J. Goeree and C. A. Holt (1998), "Rent Seeking with Bounded Rationality: An Analysis of the All-Pay Auction", *Journal of Political Economy* 106, 828-853.

Baye, M. R. and J. Morgan (2004), "Price Dispersion in the Lab and the Internet: Theory and Evidence", *RAND Journal of Economics*, 35 (3), 449-466.

Baye, M. R. and J. Morgan, (1999), "A folk theorem for one–shot Bertrand games", *Economics Letters,* 65, 59–65.

Baye, M. R. and D. Kovenock (2008), "Bertrand Competition," in *The New Palgrave Dictionary of Economics*, 2nd edition, Palgrave Macmillan.

Bosch-Domènech, A., J. G. Montalvo, R. Nagel and A. Satorra (2002), "One, Two, (Three), Infinity, ...: Newspaper and Lab Beauty-Contest Experiments," *American Economic Review*, 92(5), 1687-1701.

Camerer, C., T. Ho and J. K. Chong (2004), A congnitive hierarchy model of games, *Quarterly Journal of Economics* 119 (3), 861-898.

Dufwenberg, M. and U. Gneezy (2000), "Price competition and market concentration:

an experimental study", *International Journal of Industrial Organisation* 18, 7-22.

Erev, I. and A. Roth (1998), "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique Mixed Strategy Equilibria," *American Economic Review* 88, 848-881.

Fouraker, L. and S. Siegel (1963), *Bargaining Behavior*, McGraw-Hill, New York.

Gigerenzer, G., P. M. Todd and ABC Research Group (1999), *Simple Heuristics That Make Us Smart*, Oxford University Press.

Gneezy, U. (2005), Step-level Reasoning and Biddding in Auctions, *Management Science*, 51(11), 1633-1642.

Holt, C. (1995), "Industrial organization: a survey of laboratory research". In: Kagel, J., Roth, A. (Eds.), *Handbook of Experimental Economics*, Princeton University Press.

Jehiel, P. (2005), "Analogy-based expectation equilibrium", *Journal of Economic Theory* 123, 81-104.

McKelvey, R. and T. Palfrey (1995), "Quantal Response Equilibria for Normal Form Games", *Games and Economic Behaviour* 10, 6-38.

Neugebauer, T. and R. Selten (2006), "Individual Behaviour of First-Price Sealed-Bid Auctions: The Importance of Information Feedback in Computerized Experimental Markets", *Games and Economic Behaviour* 54, 183-204.

Papadimitriou, C. (2001), "Algorithms, Games and the Internet", In Proc. Of ACM Symp. On the Theory of Computing (STOC), 749-753.

Payne, J. W., J. R. Battman and E. J. Johnson (1988), "Adaptive Strategy Selection in Decision Making", *Journal of Experimental Psychology: Learning, Memory and Cognition* 14, 534-552.

Plott, C. (1982), "Industrial organisation and experimental economics", *Journal of Economic Literature* 20, 1485–1587.

Plott, C. (1989), "An updated review of industrial organization: applications of experimental economics", In: Schmalensee, R., Willig, R. (Eds.), Handbook of Industrial Organization, Vol. II, North Holland, Amsterdam.

Roth, A. and I. Erev (1995), "Learning in Extensive Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior,* 8, 164-212.

Rubinstein, A. (1988), *Modeling Bounded Rationality*, Cambridge: MIT Press.

Suetens, S. and J. Potters (2007). "Bertrand Colludes More than Cournot," *Experimental Economics* 10, 71-77.

Stahl, D. and P. Wilson (1994), "Experimental Evidence of Players' Models of Other Players," *Journal of Economic Behavior and Organization*, **25**, 309-327.

Stahl, D. and P. Wilson (1995), "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior*, **10**, 218-254.

Wellford, C. P. (2002), "Antitrust: results from the laboratory," in Experiments Investigating Market Power, ed. By R. M. Isaac, and C. A. Holt, pp. 1-215. JAI Elsevier, Amsterdam.

Whynes, D.K., Phillips, Z. and E. Frew (2005), "Think of a number… any number?" *Health Economics* 14, 1191-1195.

7. Tables and Figures

Table 1: Summary of Experimental Design

	Linear Bertrand	Quadratic Bertrand
Subjects	46	84
Sessions	2	2
Experimental rounds	20	20

Table 2: Average Prices

(std dev between brackets)

		Linear Bertrand	Quadratic Bertrand
Posted Prices	Period 1	161.41 (112.07)	206.58 (97.58)
	Period 20	55.00 (78.41)	159.05 (108.41)
	All periods	84.94 (90.04)	165.72 (97.47)
Market Prices	Period 1	95.00 (71.45)	154.52 (60.25)
	Period 20	46.30 (68.36)	149.00 (98.14)
	All periods	66.61 (70.82)	148.82 (87.55)

Table 3: Proportion of participants raising or holding the prices

(participants are classified by whether they won or lost in the previous round)

Treatment	Strategy	Winner in previous round	Loser in previous round
Linear	Raise	0.474 (0.500)	0.080 (0.271)
Linear	Hold	0.114 (0.317)	0.085 (0.280)
Quadratic	Raise	0.531 (0.499)	0.151 (0.358)
Quadratic	Hold	0.118 (0.323)	0.095 (0.293)

Table 4: Means and frequencies of adjustment and downward durations.

Treatment	Average Adjustment Duration	Avg. Downward Duration	# of Adjustment Durations	# of Downward Durations
Linear	5.61	2.08	2.65	2.52
Quadratic	5.88	2.06	2.71	2.31

Table 5: A model of adaptive behaviour as a mixture of naïve myopic adapters and Nash adapters

(Standard errors in parentheses)

Proportion Nash	0.652 (0.359)
Std dev of error term	51.706 (2.426)
Inertia Parameter Nash	0.350 (0.150)
Log Likelihood	-1386.33

Figure 1: Evolution of Posted and Market Prices over time

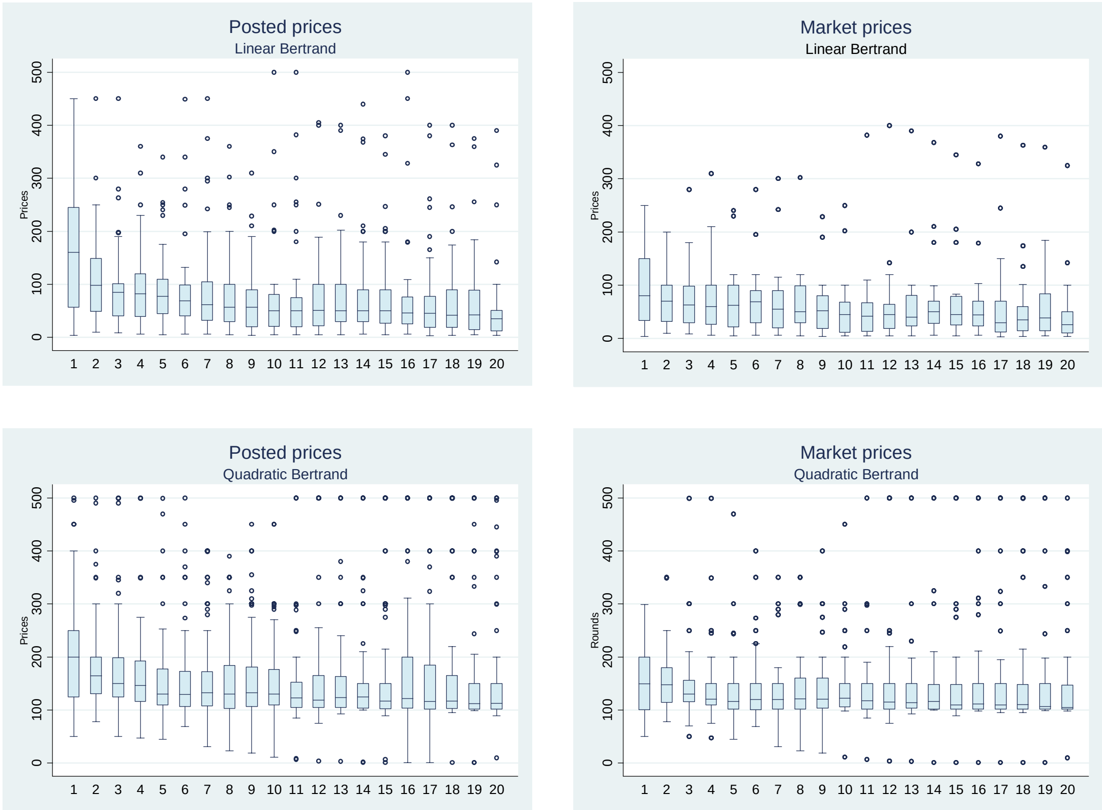

Figure 2: Histograms of posted prices

(periods 1-5 and 16-20 per treatment)

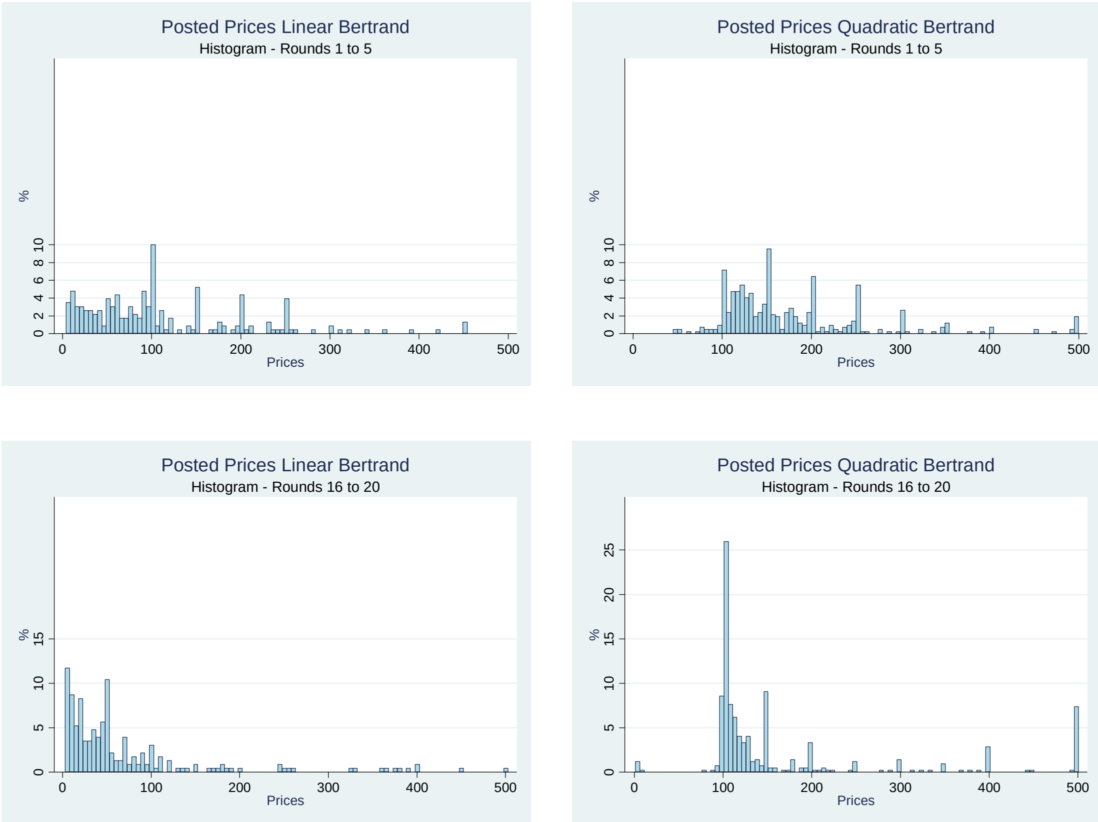

Figure 3: Histograms of market prices

(periods 1-5 and 16-20 per treatment)

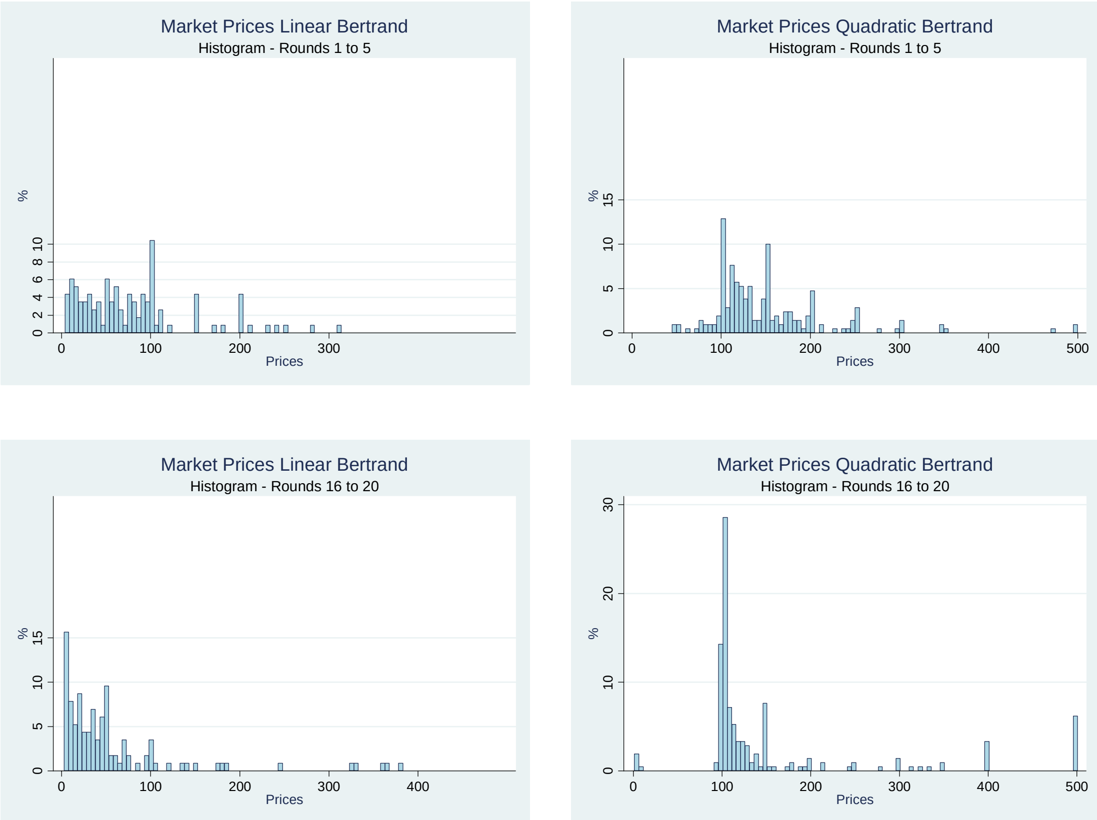

Figure 4: Cumulative distribution functions

Models vs. actual data

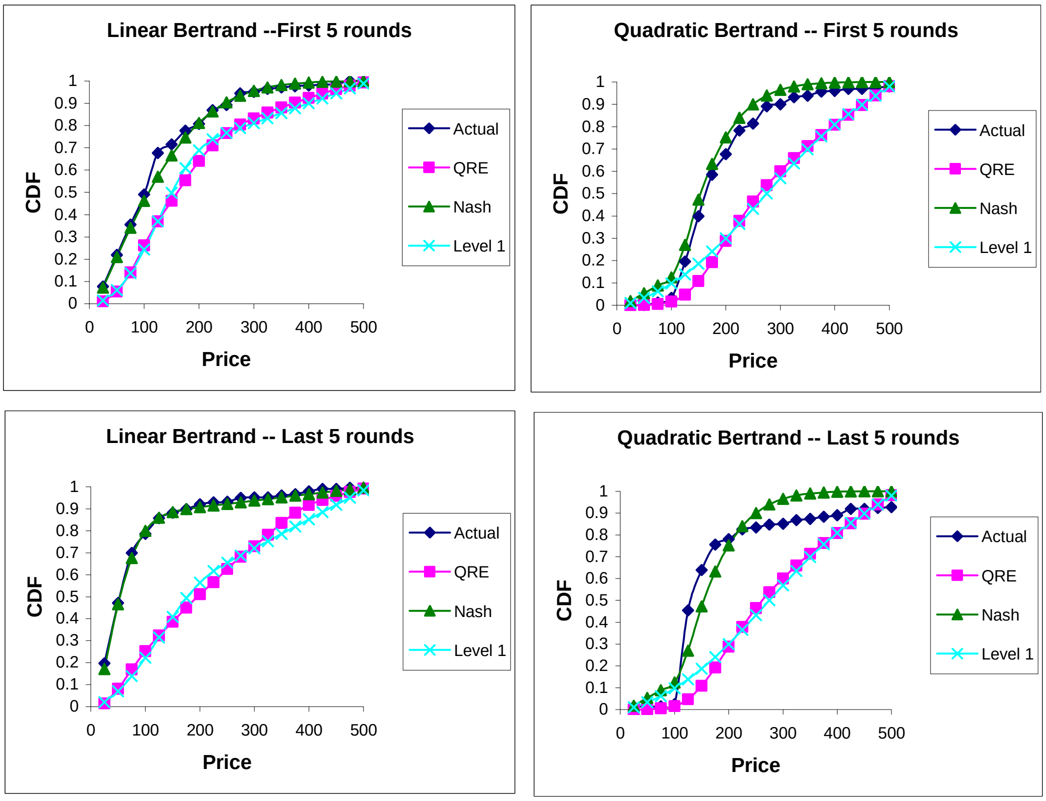

8. Appendix

Proof of Proposition 1.

The linear case is omitted. For the quadratic case, we prove it in three steps.

Step (i). In equilibrium no firm has negative payoffs (note that there are no fixed costs).

Step (ii). There are no asymmetric equilibria. The proof is by contradiction. Let (p_1^*, p_2^*) be any asymmetric Nash equilibrium. Without loss of generality, assume that $p_1^* < p_2^*$. Note that from part (i) above it is true that $\pi(p_1^*, p_2^*) \geq 0$. Then

$$\frac{\partial \pi(p_1^*, p_2^*)}{\partial p_1} = Q > 0 \tag{A.1}$$

Hence, firm 1 can increase profits by increasing its price. This contradicts our original statement that (p_1^*, p_2^*) is Nash equilibrium.

Step (iii). We focus now in symmetric equilibria. The lower limit of the set of Nash equilibria is the one that yields zero profits

$$\underline{p}\frac{Q}{2} - c\left(\frac{Q}{2}\right)^2 = 0 \Rightarrow \underline{p} = \frac{1}{2}cQ \tag{A.2}$$

For given price p larger than $\underline{p}$ the profit is

$$\pi_i(p,p) = p\frac{Q}{2} - c\left(\frac{Q}{2}\right)^2 > 0 \tag{A.3}$$

We now compute the largest price such that there is no profitable deviation. The deviation to consider is to set a price p-ε with ε>0. The profit associated to this deviation is

$$\pi_i(p-\varepsilon,p) = (p-\varepsilon)\frac{Q}{2} - c\left(\frac{Q}{2}\right)^2 \tag{A.4}$$

We finally compute the difference $\pi_i(p-\varepsilon,p) - \pi_i(p,p)$. A bit of algebra shows that

$$\pi_i(p-\varepsilon,p)-\pi_i(p,p)=Q\left(\frac{1}{2}\left(p-\frac{3}{2}cQ\right)-\varepsilon\right) \qquad \text{(A.5)}$$

Hence, for any price lower than $\frac{3}{2}cQ$ there does not exist any profitable deviation. QED

Proof of proposition 2. We consider a family of games indexed by k. All these games will be identical to the 0-game except that the strategy set is S_k={0, k, 2k…,nk ,…} with k>0 and n=1,2,… We proceed in three steps.

Step (i). No firm has negative profits in equilibrium. The proof is trivial again by noticing that there are no fixed costs.

Step (ii). There are no asymmetric equilibria. The proof is indirect. Let (p_1^*,p_2^*) be any asymmetric Nash equilibrium of the k-game. These prices can be written as follows: $p_1^*=n_1k, p_2^*=n_2k$ and without loss of generality we assume that $n_1<n_2$. Note that it is true that $\pi_1(p_1^*,p_2^*)\geq 0$ and $\pi_2(p_1^*,p_2^*)=0$. The proof of lemma 1 applies as long as $n_1 < n_2 +1$ because then n_1+1 would be a profitable deviation for firm 1. Hence, the case we need to analyse is $n_2=n_1+1$.

There are two possibilities:

(a) $\pi_1(p_1^*,p_2^*)=0$. In this case, a profitable deviation for firm 1 is to set the price p_2^* because the payoffs for firm 1 will be positive.

(b) $\pi_1(p_1^*,p_2^*)>0$. In this case, the firm to deviate is firm 2 by setting a price p_1^*. In this case, the profits of firm 2 would be positive.

Whence, we have found a contradiction.

Step (iii). It is plainly true that those multiple of k belonging to the interval $\left[\frac{1}{2}cQ,\frac{3}{2}cQ\right]$ are Nash equilibria of the k-game. We need to show that there exist two additional equilibria with prices larger than $\frac{3}{2}cQ$. Consider a rival's price p=nk larger than $\frac{3}{2}cQ$. It is easy to show that for this price, there does exist a lower bound for an undercutting strategy to be profitable. This lower bound is $\frac{1}{2}\left(p-\frac{3}{2}cQ\right)$. Given that this lower bound is strictly increasing in p, we simple need to compute the largest rival's price for which this lower bound is smaller than k.

Given that n is an integer number; we first compute the rival price $\bar{n}$ for which the lower bound is exactly equal to k. The solution we are looking for is the largest integer smaller than n*.

$$k=\frac{1}{2}\left(\bar{n}k-\frac{3}{2}cQ\right)\Rightarrow \bar{n}=\frac{3}{2}\frac{cQ}{k}+2 \qquad \text{(A.6)}$$

Hence,

$$n^{*}=\text{int}\left(\frac{3}{2}\frac{cQ}{k}\right)+2 \qquad \text{(A.7)}$$

We now make sense of the above expression. Note that $\text{int}\left(\frac{3}{2}\frac{cQ}{k}\right)$ is precisely n^{+}. Hence, the above formula shows that $n^{+}+1$ and $n^{+}+2$ are Nash equilibria of the k-game. QED